An Ontogenetic Perspective to Scaling Sensorimotor Intelligence

Cynthia B. Ferrell and **Charles C. Kemp**
MIT Artificial Intelligence Lab
545 Technology Square #819
Cambridge, MA 02139
vox: (617) 253-7007
fax: (617) 253-0039
e-mail contact: ferrell@ai.mit.edu

Introduction

Much of the work in autonomous agent design has been heavily inspired and influenced by biological systems. Researchers strive to build artificial systems that emulate the robust, flexible, and adaptive behavior of animals. In particular, the behavior-based approach (Brooks 1991) has adopted many of its key models, metaphors, and concepts from the fields of ethology, neuroscience, and evolutionary theory to name a few. Looking to animals to glean insight into how to build intelligent autonomous systems, the field has brought the issues of embodiment, situatedness, emergence, task-based decomposition, and environmental complexity to the foreground (Brooks & Steels 1995). Indeed, the treatment of these issues characterizes various biologically-inspired approaches, several of which have been successfully applied to the control of autonomous robots—behavior based systems (Brooks 1986), evolution-based systems (Cliff *et al.* 1993), and connectionist systems (Beer 1990), for instance. To date, the behavior of these systems mirrors that of simpler organisms, such as insects, situated in relatively complex environments.

Wishing to move beyond the behavior of insects and other simple animals, researchers question the scalability of these approaches and heavily debate how far systems using these approaches can be pushed. Often the question of scaling is posed in the context of increasing either environmental complexity or task complexity. Implicit in this question is that of cognitive scaling, i.e., can the intelligence of these physically embodied systems approach that of higher-level animals such as dogs, dolphins, primates, and humans? If this is the case, what roles could embodiment, situatedness, emergence, environmental complexity, or having a task-based view play in scaling cognitive and behavioral capabilities of artificial and natural systems?

In contrast to many autonomous agent programs that are shaped by an evolutionary or ethological perspective, we propose a research program that pursues an ontogenetic path. We argue that the field of developmental psychology gives fresh insight into the issues of cognitive and behavioral scaling. It offers a viable alternate viewpoint from which to study embodiment, situatedness, emergence, environmental complexity, and goal-oriented action in the world, illuminating the critical roles they may play in bootstrapping a child to more sophisticated levels of mental competence. Our contention is that a developmental perspective to autonomous agent design profoundly affects the implemented control architecture, mechanisms, subsystems, and internal structures in ways which distinguish an ontogenetic system from other learning-based systems.

A Piagetian View of Cognitive Growth

Developmental psychology is a large and active discipline (Carey & Gelman 1991), (Thelen & Smith 1994), (Michel & Moore 1995). Jean Piaget founded this field upon the premise that clues and insights into the nature of the mind and general intelligence could be discovered by studying the behavior of children (Piaget 1952). Through scientific study, he found that the physical and mental skills of children evolve throughout infancy and childhood on route to the mature adult state. Over the course of development, children become ever more sophisticated in thought and behavior, and thereby more effective in their environment.

Piaget hypothesized that adult-level skills and competencies are not innate. Instead, the necessary cognitive structures are invented, discovered, and constructed by the child. In this view, cognitive and behavioral development is the process of incrementally changing these mental structures of mind, to accumulate new more powerful structures. By modifying and creating these mental structures, a child fundamentally changes the way it thinks. This evolution is made possible by the child's interaction with its environment and a set of biologically specified developmental processes.

The course of development is influenced by past learning experience and by biological endowments, however the broad outline is remarkably consistent across cultures and ages. The consistency is so striking that Piaget characterized the developmental process as a progression through four stages (sensorimo-

tor, pre-operational, concrete operations, and formal operations), each stage being qualitatively distinct in the sophistication of the skills exhibited by the child. Indeed, children are examples of systems that scale par excellence—they start off in a relatively primitive sensorimotor state where their behavior is governed by direct sensory stimulation and reflexive motor actions, and progress throughout several successive stages to eventually reach adult-level cognitive and behavioral abilities.

Developmental versus an Evolutionary Approach

We contend that the popular evolutionary metaphor for scaling system complexity fails to provide a good path for researchers to follow and instead encourages counter productive methods of design. The evolutionary perspective argues that new systems build on previous systems by using similar computational and morphological structures and organizations. This corresponds to traveling down a branch of the evolutionary tree to more sophisticated organisms. However, when addressing the issue of evolving cognitive systems, one does not travel down a single branch of the evolutionary tree but instead jumps between and down branches—i.e. the progression from insects to reptiles to rats to dogs to dolphins to chimpanzees to humans is not a continuous progression. There are huge evolutionary jumps in between these species which are difficult to emulate. In engineering practice, after a system has been painstakingly constructed, it is unclear how the researcher should scale the system to the next level of mental complexity. Evolution covers its tracks well, making its path very difficult for a researcher to follow.

In contrast, the development of cognition and behavior in children is fluid and well documented. Piaget's stages not only provide clues as to what changes are occurring within children at different times in development, but also provides clues as to how these changes bootstrap off of each other to allow the child to progress to higher levels of competence. Experiments attempt to pry out what kinds of computational processes and representational structures are available to the child at each stage. We use the description of these stages, and supporting experimental results, as a road map to guide the design of our ontogenetic system, each successive stage acting as a goal that a system must achieve through development. Using this road map and the vast library of literature specifically devoted to the computational mechanisms of human intelligence, we propose a distinctive research agenda—to design an ontogenetic system which progresses through the first few stages of development such that system reaches successive stages by using only what has developed in the previous stages. It is important that the ontogenetic system not skip stages in the developmental process, as this would imply more prior knowledge than necessary and would decrease the likelihood that the system would be able to scale to later stages.

While these two perspectives are not inherently contradictory, they do promote very different methods of design. The evolutionary approach encourages the researcher to explicitly scale an old design by deliberately adding new components and modifying old ones. While, from the start, the developmental perspective forces one to consider how to design an integrated system which scales itself. However, within the context of our ontogenetic approach evolutionary theory is useful.

Important Insights from Ontogeny

This section tersely presents a few of the many insights which can be gained from an ontogenetic perspective. The two most important themes of this section relate to how the body, environment, and mind work together to simplify learning and facilitate scaling. The first theme emphasizes the way the body, environment, and mind work together to direct development and simplify learning through constraints and biases. A good constraint removes extraneous possibilities, thereby reducing the size of the input or output space and subsequently assisting learning. An input bias makes a particular input easier to acquire or more likely to occur. Similarly, an output bias makes a specific action easier to execute or more likely to be made. The second theme highlights the way in which an ontogenetic system's goals, morphology, environment, and cognitive abilities all grow in complexity. These intimately related paths give continuity and coherence to the challenges the system is presented with, making it possible for the system to incrementally increase its abilities by building on previous competencies.

Goal-Oriented Action in the World

Innate and self-created goals form the engine behind an ontogenetic system's life long development. A child's efforts to meet its goals in a variety of situations result in failures which push the child to modify its mental structures. In general, the inability of a child to meet its expectations of success implies that the child needs a better model of reality, improved skills, or modified goals. For example, the high level motivation to accomplish a goal puts pressure on a child to focus mental resources on related tasks and to attempt these tasks more often. This effort promotes the emergence of new tools and resources which help the child accomplish the goal with higher probability over a larger set of situations. Furthermore, as the child becomes more proficient at accomplishing this goal, other goals can use it as a sub-goal.

Goals must become more complex and elaborated as the system develops, so that they will remain commensurate with the system's capabilities and continue to promote incremental development. Internal and external factors work to maintain a useful level of goal complexity, which lies somewhere between goals of stifling intricacy and goals of unchallenging simplicity.

For example, unreachable, unreasonable, and overly simplistic goals which waste the ontogenetic system's resources, tend to be filtered out by boredom and frustration. On the other hand, good goals stimulate the cognitive growth of the system by presenting challenges which require incremental progress to meet, thus providing an addictive level of positive feedback.

Also, goals help the system to self-organize by supporting the development of task-dependent skills and promoting coherent activity. A system's active goals implicitly specify a task which sub-systems can use to organize themselves. These active goals also focus the ontogenetic system's resources on a particular task, thus giving coherence to the system's behavior. Through this coherent activity, sub-systems learn how to work together.

Situatedness and Embodiment

For an ontogenetic system, the body plays an essential role in constructing the system's cognitive foundation. Human intelligence relies on a huge amount of knowledge about the world. The majority of this common sense knowledge is stored in lower level systems and tends to be very specific to each input and output modality. Spoon feeding a large percentage of this necessary knowledge to a developing system is both inefficient and infeasible. In order for an ontogenetic system to truly scale, it must be able to acquire and organize this knowledge on its own. Consequently, the system must have an interface to the world which allows it to easily compose and administer relevant queries of the environment. The body is an interface optimized to build this common sense low-level knowledge through body-oriented activity. Hence the body is necessary, since without the representations gained through these ostensibly simple interactions with the world, abstract reasoning would have no meaningful representations with which to build tools and resources.

It is important to recognize that a child's growing and changing body usefully constrains and biases its interaction with the environment. At any temporal point in development, the system's developing morphology provides an interface to the world which is well matched to a child's level of cognitive abilities. For example, an infant's input and output is limited by several means including poor visual acuity, poor visual accommodation, muscular weakness, and short limbs. Most of these factors reduce the volume of space with which the child can interact, thereby constraining the input and output spaces in reasonable ways. As the child matures this volume of space increases, allowing the child to encounter more difficult situations when it is prepared to deal with them.

Environmental Complexity

A sufficiently rich environment aids development by challenging the system. In contrast, a limited environment causes two significant problems. First, at best a system can only learn about the contents of the environment. Second, an ontogenetic system forced to develop in a sparse environment will not benefit from the synergy of simultaneously learning from many related environments.

An overly complex environment is no better for an ontogenetic system than a very simple environment. What is important is that the environment's complexity be controlled in beneficial ways. In particular, the complexity of the environment should be tailored to the ontogenetic system's stage of cognitive development. Several things help regulate the complexity of the environment, including internal goals and predispositions, physical limitations, and dependency on a caretaker.

Fundamentally, the environment is limited by the types of situations the developing system is capable (likely) to get into. For example, infants can assume close interaction with a caretaker, and use his or her presence to aid development. The environment is further simplified by predispositions which attract humans to environments of comprehensible complexity and push them to avoid environments of overwhelming complexity. Infants cry for help if they are placed in an environment of undue complexity and over stimulation. Boredom and frustration also help keep a human in environments of appropriate complexity. Also, limited ability to locomote curtails the variety of environments the child can explore in a short time frame.

Emergence

An ontogenetic system is a complex system that exhibits emergent behavior. The the child's mind is the product of a intricate, dynamical system consisting of many interacting sub-systems (such as memory, attention, motor control, perception, emotions, etc.) and is shaped by the child's experience in its environment. Experiments in developmental psychology have revealed that many of these systems undergo a developmental progression, becoming more elaborated and sophisticated over time. For example, a child's memory becomes richer, more easily accessible, and more flexible. These sub-systems develop concurrently and assist in the mutual development of each other. Throughout development, the system grows and increases in complexity. Self-organization is critical to manage this complexity so that the system remains effective in its environment. Self-organization, increasing complexity, and the interaction among many constituents are typical of complex systems that exhibit emergent behavior.

The Epigenesis of an Artificial Mind

We view the design of our artificial system as a meta-design problem, where the challenge is to design a system which can build itself through its interactions with a rich environment and the use of ontogenetic processes. Thus, the system incrementally builds new and more powerful mental and behavioral structures

which are inherently dynamic, possessing both declarative and procedural characteristics. The system must acquire new resources and learn how to use its current resources efficiently and effectively, otherwise the system cannot scale and will become bogged down in its own intricacy. The system should be highly opportunistic and creative, reusing and modifying lower level structures for higher level function, and thus maximizing the return on the considerable time and energy used to construct the lower levels. Inherent in this process is that learning be pervasive, existing at all levels of the system so that every level can grow and refine. This is coupled with the ability of the system to limit what it represents in the environment. An ontogenetic system has control over its representations of the input and output spaces, and thus has the power to internally constrain and bias them. Over the course of development these input and output representations increase in complexity as they become more refined, segmented, and specialized. In this way, the input and output representations are yet another tool with which the ontogenetic system can avoid being overwhelmed with confusing detail or bored by unchallenging simplicity.

Research Interests

Thus far we have argued for the merits of an ontogenetic approach to scaling up cognitive and behavioral skills in artificial systems. We have articulated several important insights this perspective provides for the importance of embodiment, situatedness, environmental complexity, and goal-directed behavior in scaling up cognition and behavior.

Our research program pursues the ontogenetic path, and we are in the process of exploring design and implementation issues on a humanoid robot (Brooks & Stein 1994). Although the robot is immobile, it has hardware to support visual, auditory, tactile, orientation, and manipulation abilities. This allows us to investigate developmental phenomena in multiple domains such as manipulation and social skills. We believe it is important to do so, as it allows us to study knowledge and skill transfer across domains, which is a critical property of a successful ontogenetic system.

Our long term research goal is to build a system whose ontogeny parallels that of the sensorimotor period of children. Although one could dedicate a lifetime to this objective, we believe that our approach will lead to interesting and useful results in tractable lengths of time.

The sensorimotor period spans the first two years of life and is well studied and documented. Piaget decomposed the sensorimotor period into six stages. Although the actual progression is more fluid than a stage decomposition implies, Piaget's characterization of the sensorimotor period is a good overview of the types of changes that children undergo during this time. (Drescher 1991) provides a nice synopsis of the sensorimotor period, and (Diamond 1991) gives a detailed account of the development of reaching skills over this time period. A breif synopsis of the sensorimotor period is provided below:

- *Stage 1—0 to 1 month*: Adaptive reflexes. The infant's behavior is characterized by reflexive motor responses, either to appropriate sensory stimuli or to spontaneous activation. As the infant begins to discriminate between different perceptual states, the reflexes adapt to suit the particular circumstance. For example, the child adapts its sucking reflex so that it sucks on fingers differently from how it sucks on a bottle.
- *Stage 2—1 to 4 months*: Primary circular reactions. The infant begins to chain its adaptive reflexes together, often tending toward repetition. If the first reaction produces a familiar signal that triggers a second reaction, then a functional coordination results. For example, bringing the hand to the mouth combines with the sucking behavior. These functional coordinations can repeat either spontaneously or when they have an immediate affect on the body. These reactions begin to bridge the gap between different sensorimotor modalities. For example, a milestone of this stage is the appearance of visually guided reaching.
- *Stage 3—4 to 7 months*: Secondary circular reactions. The infant begins to chain primary circular reactions together, often to reproduce fortuitously discoverd effects on objects. An example of this is batting a toy hung over head to see it move and make noise. Until now, the child's behavior has been very solopist in nature, but now she begins to attend to her own actions which have a noticible effect on the external world. In particular, she attends to cause and effect relations. Her behavior becomes more directed, but is still heavily pulled by the external world.
- *Stage 4—7 to 9 months*: Coordination of secondary reactions. The child begins to exhibit means-end behavior, i.e., taking intermediate actions toward a desired goal. For instance, the child may uncover a hidden toy in order to grasp it. Now the child's behavior indicates intentionality and creativity, applying familiar means to achieve new goals. It should be noted, however, that the child does not invent new actions or vary existing actions; instead, she uses known actions in novel ways.
- *Stage 5—9 to 15 months*: Tertiary circular reactions. "The little scientist". Now the child has an active interest in pursuing novelty for its own sake. Her behavior indicates genuine curiosity, and she varies her actions in new ways in different contexts to perform "experiments", often to see if she can repeat interesting outcomes. These experiments are not hypothesis driven, but are better characterized as trial and er-

ror explorations with the goal of simply seing what happens.

- *Stage 6—15 to 24 months*: Simulation of events. For instance, a child may suddenly invent the behavior of using a stick to acquire a toy that is out of reach without engaging in a lot of prior groping. Piaget argues that the child engages in internal reenactment of physical activity to do so. There is an explosion of intellectual skill which is a precurser of the Pre-operational period of development. Language is coming on-line as well.

We aspire to design and implement a system which exhibits a similar progression as outlined above. Whereas children undergo skill development in many different domains; we concentrate our efforts in two areas: manipulation and social interaction. The development of manipulation skills is a good choice because it is richly documented, it is an important skill which enables the child to purposefully explore its environment, and it allows the child to exhibit competence in cognitive as well as behavioral tasks. The social realm is also important to study as evidenced by the crucial role that parents (mentors) play in the development of children. No child grows up in isolation, and the constraints and biases that adults provide when interacting with them have a large influence on the course of development.

Designing and implementing a system which allows us to explore developmental issues in multiple domains is no small task. We contend that it requires a holistic approach, building an integrated system consisting of multiple dynamic interacting subsystems (motivations, attention, emotion, memory, perception, motor, etc.), each undergoing growth, differentiation, diversification, and refinement. The power of the ontogenetic perspective is that it is a synthesis of many ideas and concepts thought to be important to designing autonomous systems that scale. It unites these concepts in a common framework, binding them and inter-relating them in specific ways. Although we are in our early design stages, we can already see the profound effect this perspective has had in shaping our architecture, mechanisms, and representations.

References

Beer, R. 1990. *Intelligence as Adaptive Behavior: An experiment in computational neuroethology.* New York: Academic Press.

Brooks, R., and Steels, L. 1995. *The Artificial Life Route to Artificial Intelligence*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, publishers.

Brooks, R., and Stein, L. 1994. Building brains for bodies. *Autonomous Robots* 1:1:7–25.

Brooks, R. 1986. Achieving artificial intelligence through building robots. Memo 899, Massachusetts Institute of TechnologyArtificial Intelligence Lab, Cambridge, Massachusetts.

Brooks, R. 1991. New approaches to robotics. *Science* 253:1227–1232.

Carey, S., and Gelman, R. 1991. *The Epigenesis of Mind: Essays on biology and cognition.* Hillsdale, New Jersey: Lawrence Erlbaum Associates, publishers.

Cliff, D.; Harvey, I.; ; and Husbands, P. 1993. Explorations in evolutionary robotics. *Adaptive Behavior* 2(1):73–110.

Diamond, A. 1991. Neuropsychological insights into the meaning of object concept development. In Carey, S., and Gelman, R., eds., *The Epigenesis of Mind: Essays on biology and cognition.* Hillsdale, New Jersey: Lawrence Erlbaum Associates, publishers. 67–110.

Drescher, G. 1991. *Made-Up Minds: A constructivist approach to artificial intelligence.* Cambridge, Massachusetts: MIT Press.

Michel, G., and Moore, C. 1995. *Developmental Psychobiology: An interdisciplinary science.* Cambridge, Massachusetts: MIT Press.

Piaget, J. 1952. *The Origins of Intelligence in Children.* New York: Norton.

Thelen, E., and Smith, L. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action.* Cambridge, MA: MIT Press.