

Evolving agents as a metaphor for the developing child

Matthew Schlesinger

Department of Psychology, Southern Illinois University Carbondale, USA

Abstract

The emerging field of Evolutionary Computation (EC), inspired by neo-Darwinian principles (e.g. natural selection, mutation, etc.), offers developmental psychologists a wide array of mathematical tools for simulating ontogenetic processes. In this brief review, I begin by highlighting three of the approaches that EC researchers employ (Artificial Life, evolutionary robotics and comparative stochastic optimization). I then focus on the advantages of using comparative stochastic optimization as a method for studying development. As a concrete example, I illustrate the design and implementation of an EC model that simulates the development of reaching in young infants.

Introduction

Although psychologists long ago dismissed Haeckel's proposition that *ontogeny recapitulates phylogeny* (i.e. that development of the individual mirrors the evolution of the species), the use of evolution as a metaphor for changes on the developmental timescale remains a popular theme in psychology (e.g. Edelman, 1987; Siegler, 1994). These metaphors often stir debate over whether developmental mechanisms are unique, or if instead, some aspects of development are best interpreted in a neo-Darwinian framework (e.g. neural Darwinism, socio-biology, etc., see Dawkins, 1976).

In this paper I make the (hopefully) uncontroversial claim that while many aspects of child development need not be interpreted as obeying neo-Darwinian principles, modeling these phenomena from an evolutionary perspective can still offer a variety of insights to developmental researchers. In particular, I propose a 'global evolutionary analogy' (Siegler, 1984) that employs a set of evolutionary concepts (e.g. variation-and-selection, competing behavioral strategies, etc.) as a computational framework for simulating developmental changes during infancy and childhood.

In the next section, I introduce the notion of a *population of evolving agents*, and describe three different ways that such a population can be used to study learning and development. I then focus on one of these three approaches (i.e. comparative stochastic optimization), and highlight some of the unique advantages that it provides for modeling development. In the final section, I illustrate how to apply these ideas by describing a

recent model of motor-skill development in young infants (Schlesinger & Parisi, 2001a; Schlesinger, Parisi & Langer, 2000).

A brief overview of Evolutionary Computation

Evolutionary Computation (EC) is a broad mathematical approach for investigating adaptive systems, including a wide array of mathematical techniques (e.g. cellular automata, classifier systems, fuzzy logic, genetic algorithms and genetic programming, neural networks, etc.). While EC methods are widespread in the fields of engineering and computer science, they have only recently emerged in cognitive science (e.g. special section on EC and cognitive science in *IEEE Transactions on Evolutionary Computation*, Wiles & Hallinan, 2001).

In contrast to conventional connectionist models (e.g. backprop nets), which often take an individual-based approach to learning, EC models employ a population-based approach. There are three key elements of EC models.

First, the population in the EC model is composed of *autonomous agents* or organisms, whose phenotypic characteristics (e.g. physical form, behavior, etc.) are the product of a genetic code that is expressed in an artificial or real environment. Note that the notion of an agent is intended to be quite flexible and inclusive, spanning the continuum from chemical or cellular agents (e.g. molecules, viruses, amoeba, etc.) to complex biological organisms (e.g. insects, mammals, primates, etc.). Second, EC

agents are normally *embodied*, that is, they occupy space and in the case of biological agents, possess sensorimotor systems that allow them to both sense and respond to their physical environment. Third, these populations of agents gradually change and adapt to their environment over time as a function of neo-Darwinian processes (i.e. competition and selective reproduction; variation via genetic recombination, etc.).

Note that there are not only numerous domains to which an EC model can be applied (e.g. linguistic, cognitive, social, etc.), but also numerous ways in which a model can be interpreted. Indeed, there is a broad range of interpretations, spanning the continuum from strict homology of the model (to a particular biological organism or system) to loose analogy or metaphor. I describe here three examples of how EC models are applied and interpreted.

1. **Artificial Life.** The general goal of *artificial life* (ALife) research is to simulate all biological phenomena, both real and possible (Langton, 1995; Parisi, 1996). Insofar as evolution is a critical aspect of biological processes, EC methods play a prominent role in many ALife models. ALife researchers often stress the correspondence between evolutionary processes in their models, and those that occur in the natural world. In other words EC methods are not just a metaphor for evolutionary change, but rather a compact set of techniques for investigating neo-Darwinian processes on the computer.

Thus, EC methods provide a platform for investigating a wide array of questions posed by evolutionary biologists. Possible topics include modularity of neural structure and behavior, cooperative social strategies, communication (e.g. signaling systems, language, etc.), parasitism and mutualism, navigation and landmark use, predator–prey relations, and so on. Perhaps more relevant to developmental researchers, the ALife framework has also been used to examine potential interactions or feedback between evolution and development (Hinton & Nolan, 1987; Parisi, Nolfi & Ceconi, 1992).

2. **Evolutionary robotics.** While some ALife researchers use robots to address evolutionary questions (e.g. Husbands, Harvey, Cliff & Miller, 1997; Nolfi & Marocco, 2001), a more general goal of *evolutionary robotics* is the study of adaptive behavior in machines and robotic systems. In this context, the use of EC methods represents a practical or pragmatic approach, in which the researcher asks: ‘How can we build a robot that learns to solve this problem?’

In contrast to traditional AI methods (e.g. expert- or knowledge-based systems), evolutionary robotics

often begins with a ‘naïve’ robot (or population of simulated robots) that learns by exploring and interacting with its environment. There are two key advantages to this approach. First, in many task domains the solution may not be known in advance, or it may be too difficult to compute directly (e.g. scheduling of 20 elevators in a 100-story building, controlling a multiple degree-of-freedom robot arm and vision head, etc.). Instead of supervising or directing the robot’s behavior, the researcher need only evaluate its performance with respect to a fitness or performance function, and selectively ‘reproduce’ those robots that perform best. (Note that reinforcement learning, a set of algorithms for simulating secondary or conditioned reinforcers, offers a related approach; for a detailed introduction, see Sutton & Barto, 1998.) Second, EC methods provide a powerful form of parallel search, where 100 or perhaps even 1000 robots simultaneously compete to solve the problem.

3. **Comparative stochastic optimization.** As I have characterized them, ALife and evolutionary robotics often represent two ends of the continuum of how EC methods are implemented and interpreted (i.e. as either homologous to, or simply inspired by biological evolution). Somewhere between these two extremes is a more epistemological approach, in which the researcher asks: ‘How does organism X learn to solve this problem, and how optimal is the solution?’

The approach of *comparative stochastic optimization* is not a particular modeling technique or framework, but rather a metatheoretical perspective on the modeling process. In the domain of EC models, variation-producing functions (e.g. sexual reproduction, mutation, etc.) are interpreted as stochastic processes, analogous to the role of exploratory behavior in an individual. The goals of this approach are to identify the range of solutions that exist for a given task, and to investigate alternative metrics for measuring the adaptive value of these strategies (e.g. spatial navigation and the development of shortest-route behavior; see Cornell & Heth, 1983; Rieser & Heiman, 1982).

When applied to learning and developmental processes, EC methods offer a comparative approach that addresses a variety of questions. For example, what behavioral strategies are available to infants or children on a given task? Do these strategies differ by a given measure of optimality (e.g. minimum effort or time, maximum reward, etc.)? What environmental and biological constraints affect the availability of these strategies? Do certain strategies emerge in a regular sequence, and if so, how do earlier strategies limit or direct the emergence of later ones?

Comparative stochastic optimization as a model of development

As I noted earlier, my goal is not to suggest a strong correspondence between evolution and development, but rather to highlight a few particularly relevant similarities, and to use those similarities as a basis for constructing a computational framework that simulates developmental processes. From this perspective, the evolution of a population of autonomous agents provides a metaphor onto which individual development can be mapped and investigated. There are several features shared by evolution and development that help support this metaphor.

First, trial-and-error problem-solving is a pervasive feature of development, spanning a wide range of ages and knowledge domains (e.g. motor skill development, tool-use, spatial search, etc.; for a broad theoretical and empirical review, see Siegler, 1986). On a longer time-scale, natural selection exploits a comparable strategy, in which multiple variants (i.e. genes, organisms, etc.) are generated, compared and selected according to their fitness. Second, like evolution, development is not ‘supervised’, that is, many developmental milestones are reached without explicit direction or instruction (e.g. learning to walk). Although development and evolution are (normally) progressive, each may be the result of successive adaptations rather than a teleological process with a predetermined end (Piaget, 1952). Finally, both evolutionary and developmental processes are recursive and hierarchical, insofar as the basis for future structures or behaviors is provided by existing structures or behaviors.

What advantages over existing modeling techniques, if any, does such a metaphor offer to developmental researchers? What insights might be provided by modeling development as comparative stochastic optimization, and in particular, by applying EC methods? I highlight here three unique features of the EC approach to modeling development, and present in the next section a concrete example that illustrates these features.

1. *Individual differences.* One advantage to modeling a particular problem or task with an evolutionary framework is that *populations of strategies* are revealed. Note the versatility of this approach, which provides not only a normative description of behavior (e.g. the most optimal strategy, the modal or most common strategy, etc.), but also a population-based level of description in which several distinct ‘families’ of behaviors may emerge. As in other modeling paradigms, these patterns depend both on initial conditions (e.g. random connection weights in an artificial neural net-

work) and also on the individual organism’s particular sequence of experiences. Therefore, EC methods provide an ideal framework for investigating the development of behaviors that differ from individual to individual (but that also appear to cluster into organized behavioral profiles or response patterns), both as a function of biological and environmental factors. A few developmental phenomena amenable to this approach include attachment, perceptual-processing or cognitive styles, gender roles, and so on.

2. *Autonomy.* A second feature of EC models that complements developmental research is that the organisms or agents in these models inhabit real or artificial environments, which they explore and interact with autonomously. Surprisingly, autonomous agents do not sample sensory data from their environment uniformly, but instead tend to focus much of their initial learning experiences in specific sub-regions of the environment. For example, as we will see below, a population of agents learns to limit their range of eye and arm movements as they evolve the ability to reach for objects (Schlesinger *et al.*, 2000). This phenomenon of *sensory self-selection* (Nolfi & Parisi, 1993) emphasizes the role of self-produced activity during learning and development (e.g. Bertenthal, Campos & Kermoian, 1994; Held & Hein, 1963).
3. *Sensorimotor cognition.* A third advantage of EC models is that when they simulate embodied organisms, they often reveal a wide array of sensorimotor or body-based strategies for not only encoding and representing information from the environment (in particular, spatial information, e.g. the location of a landmark, the trajectory of a moving object, etc.), but also exploiting regularities in the environment as a source of ‘external’ representations (Ballard, Hayhoe, Pook & Rao, 1997; O’Regan, 1992; Schlesinger & Parisi, 2001b). Thus, such models provide an ideal platform for investigating how more abstract levels of representation and knowledge (e.g. mental symbols, language, analogies and metaphors, metacognition, etc.) may evolve and/or develop out of basic forms of sensorimotor activity (e.g. Brooks, 1991; Johnson, 1987; Piaget, 1952).

From theory to practice: an EC model of reaching in infants

In this final section, I address the question: How does one build and study an evolutionary model that simulates developmental processes? There are literally dozens if not hundreds of different ways to address this question, and so instead of cataloging the wide array of learning

algorithms and modeling architectures that are available, I present here a detailed description of a relatively flexible set of techniques employed in Artificial Life research that combine genetic algorithms and artificial neural networks (i.e. 'econets' or Artificial Life neural networks; see Parisi, Cecconi & Nolfi, 1990). However, where appropriate, I mention alternative methods or techniques that might also be implemented. As a concrete example, I focus on the development of reaching in young infants.

By about age 3 months, young infants begin to reach for nearby objects (Clifton, Muir, Ashmead & Clarkson, 1993; Hofsten, 1979; Thelen, Corbetta, Kamm, Spencer, Schneider & Zernicke, 1993). This deceptively simple behavior is in fact a major developmental milestone, not only because it affords new opportunities for exploring and manipulating the world, but also because learning to reach is a complex problem with many levels that must be simultaneously coordinated (e.g. muscle activity patterns, arm and eye movements, etc.). For example, how do infants choose among numerous alternative reaching patterns (e.g. arm configurations or postures, reaching paths, etc.)? In addition, what types of sensory feedback (e.g. visual, proprioceptive, tactile) do infants use to guide their reaching movements?

In order to address these questions using an EC approach, four elements of a model need to be designed and implemented: (1) a body, including a 'brain' or nervous system (i.e. an information-processing device that receives sensory signals and generates motor behaviors), (2) a genetic code (including a genotype-to-phenotype mapping), (3) an environment and (4) a learning algorithm.

1. The body. Reaching requires, minimally, one arm and one eye. Consequently, a compact design is a 'one-eyed' agent with a two degree-of-freedom (DOF) arm (i.e. an arm with shoulder and elbow joints). Note this relatively simple agent might be simulated on the computer (e.g. Schlesinger & Parisi, 2001a; Schlesinger *et al.*, 2000), or implemented with a robotic arm and video camera.

Next, we can provide the agent with sensory input from its eye, arm (i.e. proprioceptive signals) and 'hand' (i.e. a tactile or haptic sensor at the arm endpoint). In addition, there are a set of simple 'muscles' for the eye and arm, which cause rotation of the eye, and extensions and flexions of the shoulder and elbow joints, respectively. Finally, we need a motor control system that converts these sensory inputs into motor signals.

One advantage of choosing an artificial neural network as our 'brain' or motor control system – aside from the fact that neural networks employ a brain-like style of computing – is that virtually every aspect

of the network's structure (e.g. the architecture or pattern of connectivity, the number of hidden units, activity thresholds for individual units, etc.) can be parameterized, encoded into the genome and consequently tuned by the learning algorithm. This is an especially useful feature for developmental researchers who would like to contrast innate or maturational (i.e. built-in) with evolvable or developmental components of the neural network, by either predetermining and holding constant a given parameter, or conversely, allowing it to be modified by the learning algorithm.

2. The genetic code. Assuming that we have chosen a neural network as the motor control system, there are then a variety of ways to design a corresponding genetic code. One straightforward method would be to take each modifiable parameter in the network (e.g. each connections weight), and encode it as a binary string. Thus, all of the connection weights in the network can be encoded in a 'genetic' form as one large string of 1s and 0s, in which each segment of 8 bits represents a particular weight. Other parameters in the model would be encoded in the same manner. Of course, part of the appeal of such an encoding is that like strands of DNA, binary strings are produced and copied from one end to the other, and are therefore amenable to modification by 'mutation' (i.e. copying errors), crossover and other related processes.

It is important to note that not only the artificial neural network, but in fact any quantitative characteristic of the agent can be encoded this way. In the reaching model, for example, we might choose to parameterize the length of the upper and lower arm segments, to create different-sized agents. Alternatively, we might make the width of the visual field modifiable.

Finally, we must design a genotype-to-phenotype mapping. In the simplest case, an agent's genotype is used to create an 'instantaneous' realization of the phenotype. For example, at 'birth' a fully mature neural network is automatically generated, with connections weights as encoded by the genotype. Of course, there are more interesting and complicated mappings. For example, the genotype might encode the initial location of units in a neural network, but instead of predetermining the connection pattern between units, allow it instead to develop as a function of the agent's experience (e.g. Cangelosi, Parisi & Nolfi, 1994; see also Raijmakers & Molenaar, this issue).

3. The environment. Conventional models of development often conceptualize the environment as a set of training patterns from which the agent samples randomly. In contrast, EC models take a more holistic approach: from this perspective, the agent is situated

or embedded in, and coupled dynamically with the environment. Sensory patterns do not occur randomly, but instead vary systematically as a consequence of the agent's behavior (e.g. head and eye movements, locomotion, object manipulation, vocalizations, etc.) in a structured world.

For our reaching example, a suitable environment (i.e. reaching workspace) might be a two-dimensional surface, level with and no wider than the length of the agent's arm. As a target for reaching, we can place a simple object anywhere on this surface, either in or out of the agent's visual field. It is important to remember that one need not design a complex, 3-D world filled with color, sound, etc. Indeed, the point of modeling is to limit study to just the essential biological and physical features. Nevertheless, the degree of realism across models can vary widely, from mobile robots in the real world to amoeba-like creatures on a computer (e.g. Brooks, 1991; Parisi *et al.*, 1990; Weng, McClelland, Pentland, Sporns, Stockman, Sur & Thelen, 2001).

4. The learning algorithm. The last component of the model is the learning algorithm, or in this case, the set of computations that modify the genome over successive generations. A common strategy is to first design a *fitness function*, which increases the agent's chances of reproduction each time it accomplishes the task (e.g. reaches the target, finds food, wins a match, etc.). For the reaching example, we might place a population of agents in the reaching workspace, and increment each agent's fitness by 1 when it succeeds in touching the target. Note that by modifying the fitness function appropriately (e.g. imposing a cost term that decrements fitness), we can bias the agents toward a *minimum-time solution*. For example, a bias toward producing relatively quick, direct reaches might be accomplished by reducing fitness by 0.1 after each movement, until the target is reached.

Each generation of agents is then sorted according to fitness, with the most fit individuals selected for reproduction; alternatively, the selection process might include a stochastic routine in which fitness values are first mapped to a probability, making individuals with high fitness more likely to be selected (e.g. Holland, 1975). If reproduction is asexual, each offspring is a 'copy' of its parent. Mutations are accomplished by an occasional 'copy error' (i.e. switching a bit in the genetic string during replication). A more complex model might employ sexual reproduction, in which each offspring is the combination of genetic strings from two parents. In this diploid scenario, additional variation mechanisms (e.g. crossover, co-dominance, etc.) are available that may have consequences independent of the effects of mutation.

Summary of the model

To briefly summarize the reaching model, a population of agents is designed and built. In the initial population, each agent has a completely random binary string as its genotype. The genotype is mapped to a phenotype, which includes connection weights in the agent's neural network; other characteristics may also be genetically encoded. Each generation of agents competes on the task of reaching for a nearby object, with each agent incrementing its fitness when it succeeds in contacting the target over several trials. The most successful agents are then selected and used to produce the next generation. This process of evaluation and reproduction continues until a predetermined performance criterion is reached (e.g. optimal reaching, a fitness ceiling or asymptote, etc.).

Overview of modeling results

How might such a model inform the study of infant motor control? To address this question, Domenico Parisi and I designed and investigated an EC model like the one described above that simulates the development of reaching in young infants. I highlight here three important findings from the model (Schlesinger & Parisi, 2001a; Schlesinger *et al.*, 2000).

1. *Redundant joints are 'frozen'*. Like human infants, the reaching model has more joints than it needs to reach the target. In particular, the body axis and shoulder joints are redundant: any target can be reached by rotating one of these joints – while holding the other stationary – in combination with rotation of the elbow joint. This is exactly what the model learns to do. In ten out of ten simulations, the shoulder joint was 'frozen' or held in place while the body axis and elbow joints were used to control the position of the hand (Schlesinger *et al.*, 2000). This solution of freezing redundant joints is not only a common strategy for adults as they learn new motor skills (e.g. skiing, in Vereijken, van Emmerik, Whiting & Newell, 1992) but has also recently been observed in infants near the onset of reaching (Berthier, Clifton, McCall & Robin, 1999).
2. *Movement slowdown during 'prereaching' movements*. As they learn to reach, young infants generate ballistic movements of the hand toward nearby targets, or prereaches (Hofsten & Ronnqvist, 1993). However, these prereaching movements temporarily decline, and then appear to reorganize as true reaches a few weeks later (Hofsten, 1984). The reaching model not only replicates this developmental pattern, but also suggests that it may be part of a general movement

slowdown, in which both arm and eye movements become slower. In particular, analysis of the vision and proprioception systems in the model indicates that this brief slowdown helps to coordinate felt position of the arm with vision of the target (Schlesinger & Parisi, 2001a).

3. *Touch links vision with proprioception.* A recent debate focuses on whether infants learn to coordinate vision and proprioception by watching the movements of their hand as it approaches the target. This is unlikely, as infants begin reaching for luminous objects in the dark at about the same age they reach in the light (Clifton, Muir, Ashmead & Clarkson, 1993). Another possibility is that a third sensory signal helps 'triangulate' sight of the target with felt position of the arm. For example, tactile sensory information when the target is reached can provide this feedback, which is particularly useful because it also serves as a reinforcement signal (i.e. contact with the target). Indeed, removing the sense of touch in the reaching model severely disrupted learning to reach (Schlesinger & Parisi, 2001a). In addition, we also found that tactile feedback accelerates relearning of the vision-proprioception mapping during prismatic displacement of the target.

Summary

My goal has been to not only highlight a few key developmental processes that correspond to evolutionary change (e.g. trial-and-error learning, individual differences, optimal behavior strategies, etc.), but also advocate EC methods as a novel set of techniques for investigating these phenomena. EC models offer developmental model-builders a variety of choices. First, they are flexible with respect to how the correspondence between development and evolution is interpreted. Second, they are theoretically neutral, yet powerful and general enough to allow the model-builder to contrast the relative roles of maturational and environmental factors. Third, and most importantly, they are also compatible with a wide variety of modeling architectures and learning algorithms, including the other techniques presented in this special section (ART, autoassociator networks and Hebbian networks).

Acknowledgement

This work was supported by a grant from the National Institute of Child Health and Human Development (1R03 HD40780-02).

References

- Ballard, D.H., Hayhoe, M.M., Pook, P.K., & Rao, R.P.N. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, **20**, 723–767.
- Bertenthal, B.I., Campos, J.J., & Kermoian, R. (1994). An epigenetic perspective on the development of self-produced locomotion and its consequences. *Current Directions in Psychological Science*, **3**, 140–145.
- Berthier, R.K., Clifton, N.E., McCall, D.D., & Robin, D.J. (1999). Proximodistal structure of initial reaching in human infants. *Experimental Brain Research*, **127**, 259–269.
- Brooks, R.A. (1991). Intelligence without representation. *Artificial Intelligence*, **47**, 139–159.
- Cangelosi, A., Parisi, D., & Nolfi, S. (1994). Cell division and migration in a 'genotype' for neural networks. *Network*, **5**, 497–515.
- Clifton, R.K., Muir, D.W., Ashmead, D.H., & Clarkson, M.G. (1993). Is visually guided reaching in early infancy a myth? *Child Development*, **64**, 1099–1110.
- Cornell, E.H., & Heth, C.D. (1983). Spatial cognition: gathering strategies used by preschool children. *Journal of Experimental Child Psychology*, **35**, 93–110.
- Dawkins, R. (1976). *The selfish gene*. New York: Oxford.
- Edelman, G.M. (1987). *Neural Darwinism*. New York: Basic Books.
- Held, R., & Hein, A. (1963). Movement produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, **56**, 872–876.
- Hinton, G.E., & Nolan, S.J. (1987). How learning can guide evolution. *Complex Systems*, **1**, 495–502.
- Hofsten, C. von (1979). Development of visually directed reaching: the approach phase. *Journal of Movement Studies*, **5**, 160–178.
- Hofsten, C. von (1984). Developmental changes in the organization of prereaching movements. *Developmental Psychology*, **20**, 378–388.
- Hofsten, C. von, & Ronnqvist, L. (1993). The structuring of neonatal arm movements. *Child Development*, **64**, 1046–1057.
- Holland, J.H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor, MI: University of Michigan Press.
- Husbands, P., Harvey, I., Cliff, D., & Miller, G. (1997). Artificial evolution: a new path for artificial intelligence? *Brain and Cognition*, **34**, 130–159.
- Johnson, M. (1987). *The body in the mind: The bodily basis of imagination, reason, and meaning*. Chicago: University of Chicago Press.
- Langton, C.G. (Ed.) (1995). *Artificial life: An overview*. Cambridge, MA: MIT Press.
- Nolfi, S., & Marocco, D. (2001). Evolving robots able to integrate sensory-motor information over time. *Theory in Biosciences*, **120**, 287–310.
- Nolfi, S., & Parisi, D. (1993). Self-selection of input stimuli for improving performance. In G.A. Bekey & K.Y. Goldberg (Eds.), *Neural networks in robotics* (pp. 403–418). Boston: Kluwer.

- Nolfi, S., & Parisi, D. (1999). Exploiting the power of sensory-motor coordination. In D. Floreano, J.-D. Nicoud & F. Mondada (Eds.), *Advances in Artificial Life: Proceedings of the Fifth European Conference on Artificial Life* (pp. 173–182). Berlin: Springer-Verlag.
- O'Regan, J.K. (1992). Solving the 'real' mysteries of visual perception: the world as an outside memory. *Canadian Journal of Psychology*, **46**, 461–488.
- Parisi, D. (1996). Computational models of developmental mechanisms. In R. Gelman & T. Au (Eds.), *Perceptual and cognitive development* (pp. 373–412). San Diego: Academic Press.
- Parisi, D., Cecconi, F., & Nolfi, S. (1990). Econets: neural networks that learn in an environment. *Network*, **1**, 149–168.
- Parisi, D., Nolfi, S., & Cecconi, F. (1992). Learning, behavior, and evolution. In F. Varela & P. Bourguin (Eds.), *Toward a practice of autonomous systems* (pp. 207–216). Cambridge, MA: MIT Press.
- Piaget, J. (1952). *The origins of intelligence in children*. New York: International Universities Press.
- Reiser, J.J., & Heiman, M.L. (1982). Spatial self-reference and shortest-route behavior in toddlers. *Child Development*, **53**, 524–533.
- Rosenbaum, D.A., Loukopoulos, L.D., Meulenbroek, R.G.J., Vaughan, J., & Engelbrecht, S.E. (1995). Planning reaches by evaluating stored postures. *Psychological Review*, **102**, 28–67.
- Schlesinger, M., & Parisi, D. (2001a). Multimodal control of reaching: the role of tactile feedback. *IEEE Transactions on Evolutionary Computation: Special Section on Evolutionary Computation and Cognitive Science*, **5**, 122–128.
- Schlesinger, M., & Parisi, D. (2001b). The agent-based approach: a new direction for computational models of development. *Developmental Review*, **21**, 121–146.
- Schlesinger, M., Parisi, D., & Langer, J. (2000). Learning to reach by constraining the movement search space. *Developmental Science*, **3**, 67–80.
- Siegler, R.S. (1984). Mechanisms of cognitive development: variation and selection. In R.J. Sternberg (Ed.), *Mechanisms of cognitive development* (pp. 141–162). New York: W.H. Freeman.
- Siegler, R.S. (1986). *Children's thinking*. New Jersey: Prentice Hall.
- Siegler, R.S. (1994). Cognitive variability: a key to understanding cognitive development. *Current Directions in Psychological Science*, **3**, 1–5.
- Sutton, R.S., & Barto, A.G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Thelen, E., Corbetta, D., Kamm, K., Spencer, J.P., Schneider, K., & Zernicke, R.F. (1993). The transition to reaching: mapping intention and intrinsic dynamics. *Child Development*, **64**, 1058–1098.
- Vereijken, B., Emmerik, R.E.A. van, Whiting, H.T.A., & Newell, K.M. (1992). Free(z)ing degrees of freedom in skill acquisition. *Journal of Motor Behavior*, **24**, 133–142.
- Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M., & Thelen, E. (2001). Artificial intelligence: autonomous mental development by robots and animals. *Science*, **291**, 599–600.
- Wiles, J., & Hallinan, J. (Eds.) (2001). Evolutionary computation and cognitive science: modeling evolution and evolving models (Special section, pp. 122–128). *IEEE Transactions on Evolutionary Computation*, **5**.