

Chapter 8

Evo-devo

Tony J. Prescott¹ and Leah Krubitzer²

¹ Sheffield Robotics and Department of Computer Science, University of Sheffield, UK

² Centre for Neuroscience, University of California, Davis, USA

Technology in the areas of artificial intelligence, cognitive systems, and robotics has already made many significant advances. We have constructed bipedal robots that walk and run, quadrupeds that negotiate uneven ground, cars that drive themselves, and micro-scale hover-bots that fly in formation. We have developed sensing systems that read hand-writing, recognize faces, and parse human speech; and we have devised planning, reasoning, and inference systems that integrate terabytes of information to coordinate traffic, handle the logistics of large organizations, optimize complex financial transactions, and mine scientific data sets. In some areas we are close to mimicking the achievements of biological systems, in others we have already outstripped them. But we are reaching a limit in the design of these systems and this is reflected in their brittleness in the face of unexpected challenges and the increasing difficulty and cost of updating these systems to keep pace with the changing world. “Evolvability” is therefore recognized as an important challenge for the design of complex artifacts (Lehman and Belady 1985; Mannaert et al. 2012) and technologists have been inspired by biological systems to make artificial systems that are more dynamic and evolvable (Fortuna et al. 2011; Le Goues et al., 2010).

But what can we learn from nature about evolvability? In the biological sciences, evolution has traditionally been considered through the lens of selection—preferring organisms that thrive over ones that fail. Selection is inarguably one of the most powerful mechanisms operating in biology that has made possible the evolution of complex organisms with rich repertoires of behavior. However, selection can only work if there is suitable variation within the population from which to select. Put it another way, whilst selection is the mechanism through which evolution operates it is not what makes organisms evolvable. The study of biological evolvability points to the sources of natural variability (e.g. Figure 8.1) as critical to understanding and replicating the power of natural evolution (Carroll 2012; Kirschner and Gerhart 2006).

A general definition of biological evolvability (e.g. Wagner and Altenberg 1996) might consider two sources of variability in the population. The first is concerned with how gene sequences vary within a population to give rise to variations in phenotype. Mechanisms that affect this form of variability include processes that directly alter chromosomal DNA such as mutation, crossover, and recombination, as well as population effects such as migration, translocation, and so forth. A second important source of variability, less widely known, derives from the developmental processes through which this genetic infrastructure produces a specific phenotype, as beautifully illustrated by the differences in patterning of the butterfly wings in Figure 8.1. This kind of phenotypic diversity does *not* require variability in the underlying gene sequences but instead relates to variation in the way that genes are expressed during development.




Figure 8.1 Variability in natural systems. Understanding how evolution operates requires that we go beyond the processes that allow *selection* of genetic material and decode the mechanisms that give rise to *variation*. For example, the stunning variety of patterns in these butterfly wings is thought to arise through changes in the way that regulatory gene networks operate without requiring any changes in the underlying gene sequences.

Reproduced from *Proceedings of the National Academy of Sciences of the United States of America*, Emerging principles of regulatory evolution, 104 (Supplement 1), pp. 8605–8612, Figure 1, doi:10.1073/pnas.0700488104, Benjamin Prud'homme, Nicolas Gompel, and Sean B. Carroll, Copyright (2007) National Academy of Sciences, USA.

To understand variability due to development we first have to appreciate that genes do not specify the organism directly—rather, there is a rich set of mechanisms operating inside the cell and across the embryo that, amongst other things, determine how genes are transcribed into proteins and messenger molecules, and the probability with which they are transcribed. The operation of these *epigenetic mechanisms* determines whether any given cell becomes a neuron or a white blood cell, and they resolve the mystery of how bodies can form that have a myriad of different parts despite the fact that all cells share the same DNA. Developmental biologists have now dissected many of the underlying molecular mechanisms that regulate gene expression. Interestingly, we are now realizing that these mechanisms can also be affected by the environment, within and outside the embryo, in a manner that can produce non-genetic (epigenetic), but heritable, trans-generational change.

In recent decades the *evo-devo*—evolutionary and developmental—approach (see, e.g. Carroll 2012; Muller 2007) has become central to our current understanding of biological evolvability. A key starting point was the discovery that similar networks of regulatory genes underlie body patterning in all modern multi-celled animals from sea urchins to humans (see Figure 8.2 and Raff 1996; Swalla 2006). This surprising finding demonstrated the flexibility of these networks to generate different body plans and also drove the realization that the manner in which genes determine phenotypic outcomes depends on multiple interactions across different organizational levels—the adult organism is the outcome of a series of genetic cascades modulated in time and space by the wider embryological, bodily, and environmental context. Selection can operate on phenotypic variability that arises at any stage in this process.

What we have also learned since these remarkable discoveries is that the toolbox of the genetic-developmental system is both sensitive to small changes and yet remarkably robust. For example, slight alterations in the spatial and temporal patterning of gene expression can lead to large changes in the developmental outcomes, as we illustrate with respect to nervous system evolution below.

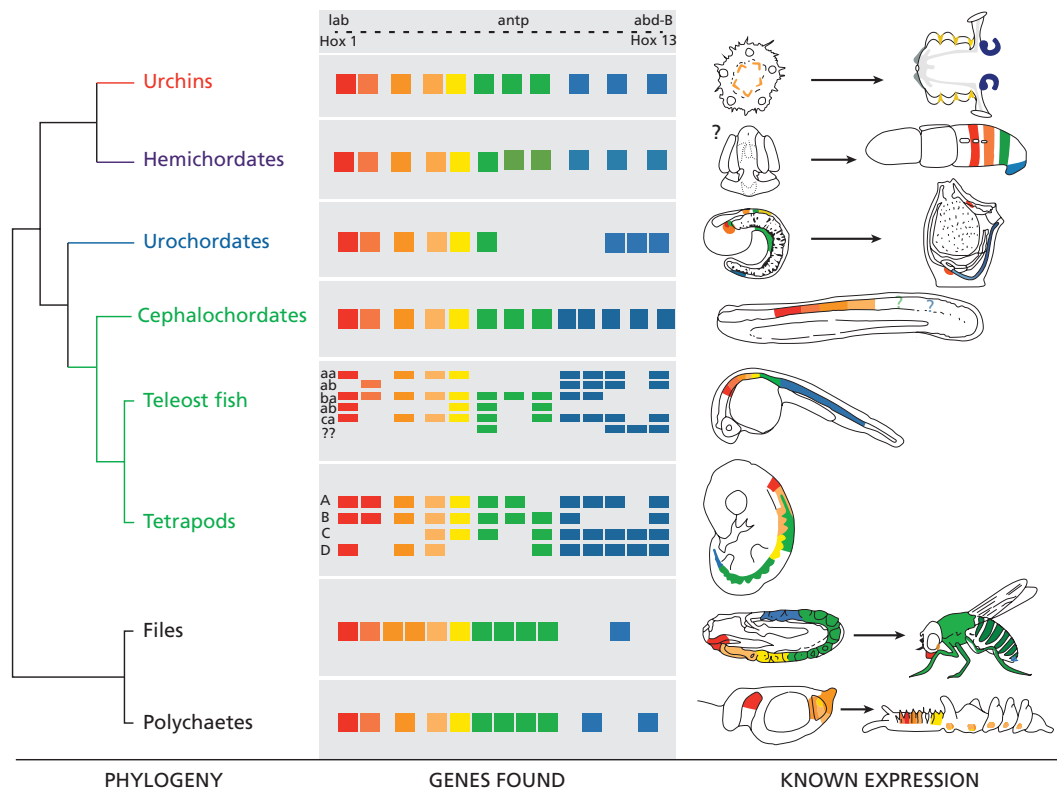


Figure 8.2 The evo-devo of biological systems. In the first half of the twentieth century evolutionary biology was dominated by a reductionist emphasis on the role of genes. A renewed interest in the contribution of development began with the discovery that the homeotic gene clusters, involved in body patterning in *Drosophila*, were strongly conserved across all multi-celled animals including vertebrates, despite obvious differences in the morphological outcomes they produced. This finding forced a re-examination of the role of developmental processes in specifying body and brain designs. This figure shows how the same regulatory gene cluster (Hox) is involved in specifying body patterning in all the different groups of modern multi-celled animals (our own species is a member of the class of tetrapods).

Reprinted by permission from Macmillan Publishers Ltd: *Heredity*, 97(3), Building divergent body plans with similar genetic pathways, B J Swalla, pp. 235–243, Figure 6, doi:10.1038/sj.hdy.6800872 Copyright © 2006, Nature Publishing Group.

However, at the same time, the wider system is able to respond flexibly to these shifting patterns in a manner that still results in a viable organism. This takes us far from traditional notions of evolution as exploring random changes in form. Understanding how developmental dynamics resists damaging perturbations and promotes convergence to viable outcomes could allow us to crack the puzzle of biological evolvability and uncover principles that can be usefully applied to artificial systems.

Our particular interest is in understanding the evolution and development of one specific complex biological system—the mammalian brain and nervous system—and in the possibility of designing brain-like adaptable architectures that could control biomimetic robots. In the next section, we give some additional consideration to particular mechanisms involved in brain evo-devo and to the evolutionary history of brains including our own. In the third part of the chapter we then explore how the principles underlying natural evo-devo have inspired, and are continuing to inspire, the design of artificial living machines.

Evo-devo of the mammalian nervous system

As a control system becomes more complex, interconnected, and integrated, it becomes more likely that any structural variation will result in degraded function. So how did animal nervous systems retain the capacity to evolve as they became larger and more sophisticated? One answer is by avoiding over-specification of the system in the genome and by exploiting developmental mechanisms that, as we have discussed above, promote compensation for disruptive change (Deacon 2010; Katz 2011). More generally, evo-devo research in neurobiology is now beginning to identify the toolbox of developmental processes that can give rise to useful forms of variability in nervous system organization, recognizing that this set of mechanisms has itself been the target of strong selection pressure during evolution (Charvet et al. 2011).

The toolbox for brain evolvability

So what's inside the toolbox for brain evolvability? First, like other developing systems the complex wiring of the brain is the result of a much simpler (though still very complex!) set of growth rules, implemented through genetic cascades, and interacting within a self-assembling structure, first over-specifying and then using intrinsic mechanisms, such as activity-based pruning, to sculpt out useful network topographies. One set of tools thus relates to the *generative*, or self-organizing, mechanisms that permit the brain to wire during early development by tuning in to internal neural signals (see also Wilson, Chapter 5, this volume). The system is first directed to incoming sensory signals, and then to the correlations between inputs from different modalities (see e.g. Krubitzer and Kaas 2005).

Second, there is a set of tools that instantiate *adaptive* mechanisms (which in this chapter we use to mean those involved in learning as opposed to general adaptation). Learning operates to promote the selection of circuits that support behavioral capacities that are well matched to the environment in which the individual develops. The development of cortical areas, for instance, is driven in significant part by specific patterns of incoming sensory inputs that can support tasks such as object recognition and detection, decision-making, path-planning, and motor control. These circuits are continually modified and honed by experience in the world, using different forms of internal and externally-mediated feedback (see Herreros, Chapter 26, this volume).

A third set of tools relates to systems components and architectural principles as discussed in Sections Four (building blocks) and Five (capabilities) of this book. For example, oscillators that can create patterned rhythms, decision circuits that can resolve competitions, and sensorimotor circuits that can implement predictive control. Architectural principles include *layered control* whereby lower tier pathways that link sensing with action can operate in the absence of higher ones (Prescott et al. 1999), and *redundancy*, where multiple substrates provide alternative means to support a given function (Deacon 2010). These architectural features create robustness and thus provide protection against catastrophic local change—smoothing out the fitness landscape by providing more options for exploring useful variability (Kauffman 1990). An important feature that we see in brain evolution, and that can underlie layering and redundancy, is the duplication of existing structures. Making copies of working sub-circuits can lead to system components that are under-used and that can then adjust to take on new functions (Deacon 2010; Whitacre and Bender 2010).

From phenotypic variability to hopeful monsters

Biologists have long been puzzled by the difficulty of reconciling rapid changes in animal lineages with the notion of gradual change brought about by genetic operators such as mutation. In the 1940s, the biologist Richard Goldschmidt coined the phrase “*hopeful monster*” (Goldschmidt 1940) to capture

the idea of an organism that represented a large leap in design-space (also termed macroevolution). This idea was deemed controversial by many biologists of the day who reasoned that any change in the genome that brought about a radical reorganization of the organism would be likely to be lethal. Nevertheless, evolution does seem to have proceeded by taking a number of larger steps in addition to many smaller ones. How might this be possible? Evolvability again appears to be key and the new understanding engendered by the evo-devo approach is resurrecting Goldschmidt's hopeful monster hypothesis. Specifically, genetic mutations, or consistent non-genetic departures from normality, that bring about radical changes can be compensated for by pattern-forming processes that steer development towards building a viable animal. Moreover, where a single population is affected by a significant disruptive event, such as abrupt environmental change, similar changes might be triggered in multiple individuals. Where this happens successful large strides within design space could be made.

We can illustrate this idea using studies of brain evolution that show a number of dramatic changes in the size and organization of the mammalian cortex through evolutionary history, perhaps most remarkably, in terms of speed of change, in the sequence of primate and hominid forms leading to modern humans. This is illustrated in Figure 8.3 which shows approximate shape and areal structure of the cortex in different groups of modern mammals.

Experimental manipulations in brain development are beginning to suggest how such changes may have come about. For example, Chenn and Walsh (2002) have shown that the impact of a single




Figure 8.3 Cortical evolution in mammals. The major mammalian groups and their family relationships illustrating how the shape of the cortical mantle, and of some of its primary sensory areas, vary relative to an inferred common ancestor.

Reproduced from *Proceedings of the National Academy of Sciences of the United States of America*, Cortical evolution in mammals: The bane and beauty of phenotypic variability, 109 (Supplement 1), pp. 10647–10654, Figure 1, doi: 10.1073/pnas.1201891109, Leah A. Krubitzer and Adele M. H. Seelke, Copyright (2012) National Academy of Sciences, U.S.A.

change in the sequence of developmental events involved in early neurogenesis in the mouse can lead to a change from a smooth rodent-like neocortex to a convoluted primate-like one, bringing about an expanded cortical surface without an increase in thickness. Although these particular mice have brain abnormalities that make them non-viable, what is impressive is that the overproduction of neural progenitor cells triggers changes later in the developmental cascade that at least partially compensate for this perturbation (see Rakic (2009) for further discussion). This study demonstrates that changes that occur early in development can have particularly profound effects (thus larger strides). Other experiments, summarized by Krubitzer and Seelke (2012), have used knock-out mice to explore the role of different transcription factors in regulating arealization—the size, shape, and position of different cortical areas. As illustrated in Figure 8.4, eliminating any one of the four factors *Emx2*, *COUP-TF1*, *Pax6*, or *Sp8* radically alters the size and shape of multiple fields—enlarging some, reducing others, whilst preserving their overall topological relationships.

In a final illustrative experiment, by Fukuchi-Shimogori and Grove (2001), an electrically induced overexpression of a growth factor in a specific locus in the mouse forebrain caused the animal to develop a second complete barrel field (a representation of the facial whiskers) that was a mirror-image of the normal one. This result indicates the relative ease with which developmental processes can generate useful redundancy in the form of substantial and well-organized structures that can then be co-opted to serve new roles.

Krubitzer and Seelke (2012) summarize the main differences in mammalian cortex that can be observed by comparing across species, as shown in Figure 8.3. These include (i) the size of the cortical sheet, (ii) the relative amount of space allocated to different sensory domains (and more generally, to the different cortical fields), (iii) cortical magnification of behaviorally relevant body parts, (iv) the addition of new modules, (v) the overall number of cortical fields, and (vi) the connections between cortical fields. Interestingly, many of the same changes can be seen when comparing between individuals of a given species but to a much less dramatic extent than when comparing across species. This illustrates that the phenotypic variability required to allow selection to direct cortical evolution in any of these different directions may be present in most species most of the time.

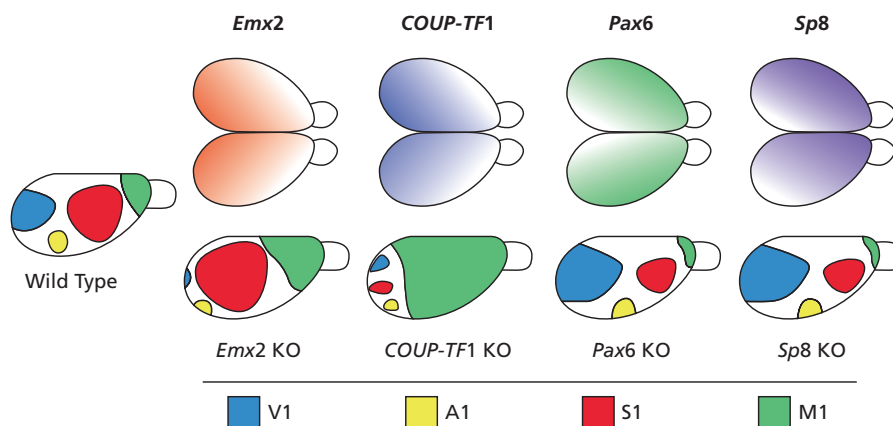


Figure 8.4 The effect of removing specific transcription factors (via genetically-modified ‘knock-out’ mice) on arealization in mouse cortex. The size and shape of major cortical areas—primary visual (v1), auditory (a1), somatosensory (s1), and motor (m1) cortices—is radically disrupted by each knock-out whilst the overall topology is maintained.

Reprinted from *Current Opinion in Neurobiology*, 18 (1), Dennis D. O’Leary and Setsuko Sahara, Genetic regulation of arealization of the neocortex, pp. 90–100, Figure 6, doi: 10.1016/j.conb.2008.05.011, Copyright © 2008, Elsevier Ltd., with permission from Elsevier.

Conservation of basic design principles

Whilst exploring the paths to effective change through biological evolution it is also important to recognize that some things have stayed the same. Indeed, the story of biological evolution can be regarded as one of conservation of basic design principles, that were discovered early on, and that instantiated evolvable control architectures (Kirschner and Gerhart 2006). For example, fossil evidence shows a remarkable explosion of animal forms during the Cambrian period (~541–485 million years ago) in which all of the major bilaterian phyla were represented, despite being absent, for the most part, from the fossil record at the end of the previous Ediacaran period (see also Prescott, Chapter 38, this volume). Current hypotheses suggest that the last common ancestor of all modern bilaterian animals, the *Urbilateria*, evolved some time before the Cambrian boundary, and is thought to have possessed the “essential bilaterian toolbox” (Erwin and Davidson 2002) of regulatory genes, including the homeotic gene clusters that underlie cell differentiation and body patterning in all modern bilaterians (as in Figure 8.2). Vertebrate-like animals also appear earlier than was once thought—finds from Chengjiang in China (the Chinese “Burgess shale”) show the presence of fish-like creatures (craniates) in the early Cambrian (Mallatt and Chen 2003). This evidence implies a rapid evolution of complex nervous systems, as part of the general evolution of new body plans (Gabor Miklos et al. 1994).

Analysis of comparative brain architecture indicates some startling conservation of nervous system structure across all vertebrate classes. Notably, all vertebrate brains comprise a layered architecture with spinal, hindbrain, midbrain, and forebrain components; an integrative core, comprised of the medial reticular formation and the basal ganglia; and a number of specialized learning/memory systems—hippocampus, striatum, cortex, and cerebellum (Prescott et al. 1999). Of course there have been considerable changes in the size, shape, number of sub-divisions (parcellation), and microstructural organization of many brain regions, as well as the addition of new cell types, and the migration of cell clusters within the brain. Connectional changes have included axonal invasion of new areas and increased differentiation of local areas through selective connection loss. Nevertheless, comparative neurobiology indicates that the scope for alteration of the basic vertebrate brain plan may be remarkably limited (Charvet et al. 2011). Vertebrates have a control architecture that is evolvable but only within a set of constraints that were already laid down within the nervous system of a 500 million-year-old fish. If we could better understand this constrained but evolvable architecture—that carried generations of owners from sea to land, from four legs to bipedal walking, into the air, even to explore space—this should make for a powerful model for the design of artificial living machines.

Biomimetic and biohybrid systems

To learn from the evo-devo of natural systems it is clearly necessary to identify, simulate, and abstract the evolutionary and developmental mechanisms that gave rise to their evolvable nature. This approach could result in a toolbox of generative, adaptive, and selective mechanisms that can be used to formulate a design methodology to be applied to complex bio-inspired systems. Such a framework will also require an abstracted understanding of body plans and control architectures since it is clear that evolution acts to select evolvable and flexible control systems and then operates within some limits determined by those designs (and by the laws of physics).

In computational neuroscience, artificial intelligence, and robotics, researchers have explored generative methods, such as self-organization, or selective methods, such as genetic algorithms and genetic programming, and adaptive methods, such as reinforcement and supervised learning, for several decades. A smaller, but still significant, community have investigated how generative, selective, and adaptive mechanisms, or some combination of these, can operate together.

Much of this literature has followed a biomimetic approach. There is insufficient space here to provide a comprehensive review, therefore we briefly survey the history of this field and highlight some examples that were particularly groundbreaking or insightful.

Foundations of artificial evo-devo

The attempt to create artificial systems using the methodologies of evolution, development, and learning is driven by the recognition that engineered systems cannot match either the complexity or the robustness of biological life, and moreover, that nature has invented some remarkable techniques for encapsulating the specification for such complex entities in highly compact codes. For example, the human brain is estimated to have 100 trillion neural connections and yet our genome has less than 30,000 active genes. One of nature's most important tools, then, is this indirect but superbly productive mapping of the genome onto the phenotype.

One analogy that occurs repeatedly is that of a language—with a small vocabulary, and a clutch of grammatical rules saying how words can combine, a language can generate millions of varying but valid sentences. In computer science, the productivity of a formal language, defined as set of rules that operate over an identified set of symbols, has been understood since Chomsky's work in the 1950s, and this general idea was first applied to model biological development by Lindenmayer (1968) to make branching shapes similar in structure to threads of algae (e.g. Figure 8.5). This approach, now known as L-systems (a contraction of Lindenmayer-systems), continues to be widely used. For instance Torben-Nielsen et al. (2008) describe a




Figure 8.5 A plant-like structure generated by an L-system grammar for modeling biological development.

Reprinted from *Journal of Theoretical Biology*, 18(3), Aristid Lindenmayer, Mathematical models for cellular interactions in development I. Filaments with one-sided inputs, pp. 280–99, doi: 10.1016/0022-5193(68)90079-9, Copyright © 1968, Elsevier Ltd., with permission from Elsevier.

modeling system, EVoL-Neuron, that builds models of single neuron arborization, using genetic algorithms to tune these to match experimentally measured cells.

The L-systems approach was also generalized by Kitano (1990) to the construction of artificial neural networks. Whereas Lindenmayer's goal was to show that a small grammar could give rise to differentiated and life-like structures, Kitano sought to demonstrate a specific advantage, in terms of coding efficiency, of a grammar-based system compared to a more direct encoding approach. Moreover, Kitano used L-systems in combination with selective (genetic algorithm) and adaptive (back-propagation) methods, that is, as part of a toolbox approach, and sought to show that this combination of methods could be usefully applied to difficult computational problems (he chose decoding, a classic problem in computer science). Kitano demonstrated that, compared to a direct encoding approach, a grammar-based generative model converged faster and scaled better as the size of the target network increased—a clear win for evo-devo. He also showed that larger networks could be specified without growing the size of the genetic coding—a path to concise and scalable developmental codes for specifying useful artificial complex systems was beginning to emerge.

Whilst grammars beautifully capture some of the power of a compact encoding to recursively express a developmental program, it is not obvious how such a system can map onto the chemical and mechanical systems inside the cell. Historically, two other approaches have provided important leads.

The first begins with Turing's (1952) paper on the chemical basis of morphogenesis (the development of body form). Now seen as a foundational paper for understanding self-organizing systems in general (see Wilson, Chapter 5, this volume), Turing was particularly concerned with the question of how the genes of a fertilized egg enable it to develop into a multicellular organism with asymmetric differentiated structure. Drawing on the idea that development can be guided by diffusing chemicals, that he termed *morphogens*, Turing provided a mathematical model of interacting chemical gradients, termed *reaction-diffusion* processes, that could create elegant patterns such as dappling, stripes, and whorls. Models based on Turing's idea are able to capture some of the more analog and non-local character of processes that happen inside the developing organism. For instance, Lewis (2008) reviewed half a century of work, building on Turing's idea, that has used mathematical models to understand the role of chemical gradients in gene expression, with examples from patterning in plant meristems to the dorsoventral organization of frog embryos. Fleischer and Barr (1994), in following a toolbox approach, combined reaction-diffusion modeling based on differential equations with a grammar-based model of growth through cell-division and a mechanical model of cell-to-cell interactions to create a 2D model of a developing cell sheet. The resulting system produced rich life-like patterns, such as those illustrated in Figure 8.6, that have found applications in areas such as computer graphics. For a contemporary toolbox for synthetic development capable of defining multicellular “soft robots” see Doursat and Sánchez (2014).

A third starting point was the realization that chromosomal DNA, and the chemical processes that surround it, instantiate a rich dynamical system with compressed degrees of freedom and emergent order, and for which it might be possible to define a simpler analog with similar dynamical properties. Kauffman (1969a, 1969b) attempted just this using randomly wired Boolean networks as models of regulatory gene systems. He showed that, given the right degree of connectivity between network nodes, a Boolean network of a size mapped to a specific animal genome will enter dynamically stable states that can predict the number of cell types in that organism, behavior cycles that can predict cell division times, and responses to noise that look like behavior switching. That Boolean networks could be a useful abstraction of regulatory gene systems was further demonstrated by Dellaert and Beer (1996) who embedded one within a model cell, and used the state of the network, as it varied over time, to regulate cycles of cell division and differentiation. Using a genetic algorithm to configure the initial network, they were able to grow




Figure 8.6 Fleischer and Barr (1994) applied their developmental tool-box, inspired by Turing's reaction-diffusion model and Lindenmayer's L-systems, to the challenge of generating computer graphics such as these spheres coated with animal-like scales.

Reproduced from Fleischer, K. W. "A multiple-mechanism developmental model for defining self-organizing geometric structures," Figure 7, PhD thesis, California Institute of Technology, Pasadena, 1995.

multicellular 2D model organisms, with a mix of sensor and actuator cells, that could be evolved to follow curved lines. Dellaert and Beer also built a second model organism more closely modeled on the chemistry of genetic regulatory networks; however, configuring such a model to be evolvable proved to be less tractable than for the simpler Boolean system. The status of Boolean networks as simplified models of cellular regulatory networks is discussed in Bornholdt (2008), who argues that Boolean network dynamics can provide useful insights into how cells "compute." Giacomantonio and Goodhill (2010) have applied Boolean networks to the problem of understanding the genetic regulatory networks underlying arealization in mammalian cerebral cortex, as shown in Figure 8.3, and involving all of the genetic transcription factors shown in Figure 8.4 plus one more (Fgf8). Exhaustively simulating all possible network configurations of these five genes, they found that only 0.1% of possible networks could reproduce experimentally observed expression patterns. Moreover, the networks that worked tended to have certain kinds of within-network interactions and not others, thus giving clues as how these gene networks may operate during brain development.

Reviewing several decades of work, Stanley and Miikkulainen (2003) proposed a taxonomy of model evo-devo, or to use their term *artificial embryogeny*, systems that identified five major dimensions of design choice: (1) cell fate—the mechanism(s) through which cell type is determined; (2) targeting—how cells connect with each other; (3) heterochrony—how the timing of development events is regulated; (4) canalization—how systems are made robust to variation in the genotype; and (5) complexification—how the genome (and the phenotype) become more complex over time.

Various further proposals have been made as to how these different design mechanisms can be abstracted from biological systems to lead to more compact and scaleable codings (and hence to improvements in evolvability). For example, Stanley (2007) has suggested that the core role of local interactions within developing organisms is to provide cells with information about their location in space. If so, a useful abstraction of the developmental process could be to provide an explicit

coordinate frame allowing development at each location to be specified by the composition of a series of parameterized functions. This process can be visualized as a directed graph or “compositional pattern producing network” (CPPN). Removing local interaction from the mix also allows this method to collapse time, that is to compute the final outcome at each location in a single step. Applying artificial evolution to CPPNs shows that they are flexible and evolvable ways of encoding patterns that exhibit phenomena characteristic of natural developing systems such as symmetry and repetition with variation. Other approaches have explored the utility of redundancy, via duplication, of network components (Calabretta et al. 2000); still others have explored the benefits of modularity within evolved controlled systems (e.g. Bongard 2002), and in the emergence of functionally distinct sub-networks in mammalian cortex (Calabretta 2007). Garcia-Bernardo and Eppstein (2015) have recently described an approach to complexification that involves pruning dense model networks to find minimal configurations that retain required functionality, then using these more compact circuits as fixed building blocks within larger systems.

Many model systems have included adaptive mechanisms (learning) as a key step in building a working model organism (e.g. Sendhoff and Kreutz, 1999). Recognized as the “Baldwin effect” (see Wilson, Chapter 5, this volume), a learning capacity can make a system more evolvable by smoothing out the fitness landscape around peaks (if learning can reliably bring the model system to the fitness peak then it should be sufficient for evolution and development to place the system somewhere in the vicinity of the peak). Adding learning capacity is thus another means for reducing code-size for evolving systems whilst making good solutions easier to find (reducing search time). Of course, this comes at the cost of the time needed for lifetime learning and a period of reduced fitness while the system adapts.

Scaling-up

An important goal of research in artificial evo-devo has been to show that model solutions can be scaled to cope with real-world complexity. A useful stepping-stone for this has been the use of simulated 3D worlds that include real-time physics engines adapted from applications in computer graphics and gaming. A well-known example of this approach is Karl Sims (1994) “creatures.” Sims used genetic algorithms to select directed graphs (analogous to an L-system grammar) that specify solutions to the problem of being a digital creature built of collections of blocks, linked by powered flexible joints, and controlled by circuits. Sims evolved a neural network control system for his creatures alongside their physical morphology and designed fitness functions for model aquatic and terrestrial environments that selected for proficiency at tasks such as swimming speed, swimming after a light source, moving across a surface, jumping on a surface, and attempting to possess a square block in contest with another creature—a model evolutionary arms race (see Figure 8.7). These experiments produced a fascinating array of creatures, some with familiar morphologies and behavior reminiscent of that of actual animals such as snakes, tadpoles, or crabs. Others, often equally effective at their tasks, accomplished them with less familiar or bizarre patterns of movement and body-form. The resulting panoply of digital organisms has been compared to the Cambrian explosion of early life.

The success of Sims’ approach depended on the power of a generative encoding to build working model organisms out of parameterized simple elements (blocks), and the ability of genetic algorithms to exploit physics in a manner that simplified the challenge of control. Lipson and Pollack (2000) took this idea closer to physical reality by combining evolution in 3D simulation with additive manufacturing (3D printing) of physical robots for the most successful designs. To perform evo-devo in an actual physical system is more difficult, although some work in reconfigurable robots is moving in this direction. For example, the Eyebot, shown in Figure 8.8, developed by Lichtensteiger and Eggenberger (1999) and modeled on the insect compound eye, applied




Figure 8.7 A simulated evolutionary arms race. Evolved artificial creatures fight for possession of a block in one of the experiments performed by Karl Sims (1994).

Reproduced from Karl Sims, *Evolving virtual creatures*, *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pp. 15–22, doi:10.1145/192161.192167 © 1994, Association for Computing Machinery, Inc. Reprinted by permission.




Figure 8.8 Eyebot (right), a reconfigurable robot inspired by the insect compound eye (left), designed to apply artificial evolution to robot morphology.

A: ©Denis Vesely/Shutterstock.com **B.** Reprinted by permission from Macmillan Publishers Ltd: *Nature*, 406 (6799), Automatic design and manufacture of robotic lifeforms, Hod Lipson and Jordan B. Pollack, pp. 974–978, Figure 5, doi:10.1038/35023115, Copyright © 2000, Nature Publishing Group.

genetic algorithms to tune the position of model ommatidia; Gomez et al. (2004) extended this approach to a robot hand-eye system to explore parallel development of control systems and morphology (degrees of freedom) in a model of human hand-eye coordination. Most recently Vujovic et al. (in press) have used additive manufacturing, together with automated assembly using a robot arm, to evolve and grow robots with variable morphology, testing their ability to locomote across a flat arena.

Future directions

We are reaching a point in time where the widespread availability of super-fast computers, combined with increased sharing of open source toolboxes, could allow a much wider community to get going with artificial evo-devo. On the other hand, a limiting factor is education—current approaches require mastery of a range of different tools and, if mining of biological principles is to be successful, understanding of one of the most daunting scientific literatures—developmental biology—with its complicated vocabulary and detailed data sets. To overcome these problems an emphasis on multi-scale modeling must continue (see Verschure and Prescott, Chapter 2, this volume) alongside interdisciplinary training opportunities for young scientists (see Mura and Prescott, Chapter 64, this volume).

The field of artificial evo-devo also has many territories still to explore. At the most fundamental level the identification of the principles required for living systems—such as being viable and self-sustaining—can be effectively explored through simulation (e.g. Agmon et al. 2016) in parallel with efforts to create new artificial living systems from organic chemistry (Mast et al., Chapter 39, this volume). At the other end of the complexity spectrum, there is also a community seeking to apply this approach to the understanding of human development and cognition (Parisi 1997) through the medium of embodied robotic modeling (e.g. Cangelosi et al. 2015; Metta and Cingolani, Chapter 47, this volume).

One strategy to take these methods into the real world is to try to circumvent some of the bottlenecks that have set limits to natural evolution, such as the need to compress the information required to create new life-forms into a single cell, to limit inheritance to family lines, and to start each experiment afresh with a new organism built from that single cell. For instance, we can explore the retention of acquired adaptations, mimicking Lamarckian evolution (or memetic cultural evolution) (Le et al. 2009), perhaps directly copying useful adaptations from one experiment to the next (much as biologists now do with genetic tools). As already noted, experiments can run in simulated bodies and worlds, where reasonably realistic model systems can now be simulated faster than real time. Model populations can be tuned to optimize evolvability, as theories of this improve, for instance by choosing additional selection criteria that promote phenotypic diversity (Lehman et al. 2016). For physical systems we can repeatedly adapt and re-use modifiable and modular robot hardware whilst taking advantage of auto-fabrication methods. Further, we can more effectively utilize our understanding of biological evolution to start with pre-structured models that are not created from scratch. Nature has identified many “forced moves and good tricks” (Dennett 1995) for the construction of animal bodies and of complex control architectures and conserves these as species evolve (Prescott 2007). We can identify design features that promote robustness and evolvability and either build these in from the beginning or promote them via selection criteria.

To apply the lessons of evo-devo to technology will require a continued effort to identify the evo-devo principles that gave rise to biological life, to determine appropriate levels of abstraction away from the biological detail, and to define design methodologies that allow these principles

be applied to artificial systems. Current toolbox approaches show great promise, although, in the spirit of evo-devo, these systems are now getting sufficiently complex that they themselves could be candidates for automated redesign.

Learning more

Several books have captured the excitement around the fusion of evolutionary and developmental approaches in modern biology—Raff (1996), Carroll (2012), and Kirschner and Gerhart (2006) are all good places to start. The application of the evo-devo approach to neurobiology is still relatively new with most aspects of brain architecture still unexplored. Charvet et al. (2011) discuss some of the key developmental mechanisms in brain evo-devo and consider the balance between constancy and change in vertebrate brain evolution, Deacon (2010) and Katz (2011) discuss a number of neural generative, selective, and adaptive mechanisms that operate to make behavior evolvable, whilst Krubitzer and Seelke (2012) explore the bases for phenotypic diversity in relation to the evolution of mammalian cortex. Much of the research that applies evo-devo methodologies to artifacts has emerged under the umbrella of “Artificial Life,” and a number of classic contributions in this field are provided in Langton (1995); Parisi (1996) also provides a conceptual overview and introduction to much of the foundational work, whilst Stanley and Miikkulainen (2003) provide a useful synthesis and taxonomy. Nolfi and Floreano (2000) review much of the classic work in evolutionary robotics, whilst Pfeifer and Bongard (2006) explore the application of some ideas from evo-devo in robots with particular emphasis on embodiment. Haddow and Tyrrell (2011) provide a review and critical assessment of the application of evo-devo methodologies to electronic systems. Finally, Downing (2015) considers many of the topics introduced in this chapter with an overall emphasis on the emergent nature of biological intelligence, and with similar enthusiasm for re-purposing nature’s evo-devo toolbox to create new kinds of Living Machines.

References

- Agmon, E., Gates, A. J., Churavy, V., and Beer, R. D. (2016). Exploring the space of viable configurations in a model of metabolism–boundary co-construction. *Artificial Life*, 22(2), 153–71. doi:10.1162/ARTL_a_00196
- Bongard, J. (2002). Evolving modular genetic regulatory networks. In: *Proceedings of the 2002 Congress on Evolutionary Computation*, CEC ’02, Honolulu, HI, 2002, pp. 1872–7. doi: 10.1109/CEC.2002.1004528
- Bornholdt, S. (2008). Boolean network models of cellular regulation: prospects and limitations. *Journal of The Royal Society Interface*, 5(Suppl 1), S85–S94. doi:10.1098/rsif.2008.0132.focus
- Calabretta, R. (2007). Genetic interference reduces the evolvability of modular and non-modular visual neural networks. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1479), 403–10. doi:10.1098/rstb.2006.1967
- Calabretta, R., Nolfi, S., Parisi, D., and Wagner, G. P. (2000). Duplication of modules facilitates the evolution of functional specialization. *Artificial Life*, 6(1), 69–84.
- Cangelosi, A., Schlesinger, M., and Smith, L. B. (2015). *Developmental robotics: from babies to robots*. Cambridge, MA: MIT Press.
- Carroll, S. B. (2012). *Endless forms most beautiful: the new science of evo devo and the making of the Animal Kingdom*. London: Quercus.
- Charvet, C. J., Striedter, G. F., and Finlay, B. L. (2011). Evo-devo and brain scaling: candidate developmental mechanisms for variation and constancy in vertebrate brain evolution. *Brain Behavior and Evolution*, 78(3), 248–57.

- Chenn, A., and Walsh, C. A. (2002). Regulation of cerebral cortical size by control of cell cycle exit in neural precursors. *Science*, 297(5580), 365–9. doi:10.1126/science.1074192
- Deacon, T. W. (2010). Colloquium paper: a role for relaxed selection in the evolution of the language capacity. *Proc. Natl Acad. Sci. USA*, 107 (Suppl 2), 9000–06. doi:0914624107 [pii] 10.1073/pnas.0914624107
- Dellaert, F., and Beer, R. D. (1996). A developmental model for the evolution of complete autonomous agents. In: P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack, and S. W. Wilson (eds), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press, pp. 394–401.
- Dennett, D. C. (1995). *Darwin's dangerous idea*. London: Penguin Books.
- Doursat, R., and Sánchez, C. (2014). Growing fine-grained multicellular robots. *Soft Robotics*, 1(2), 110–21.
- Downing, K. L. (2015). *Intelligence emerging: adaptivity and search in evolving neural systems*. Cambridge, MA: MIT Press.
- Erwin, D. H., and Davidson, E. H. (2002). The last common bilaterian ancestor. *Development*, 129(13), 3021–32.
- Fleischer, K. W. (1995). *A multiple-mechanism developmental model for defining self-organizing geometric structures*. PhD thesis, California Institute of Technology, Pasadena.
- Fleischer, K. W., and Barr, A. (1994). A simulation testbed for the study of multicellular development: the multiple mechanisms of morphogenesis. In: C. G. Langton (ed.), *Artificial Life III*. Boston, MA: Addison-Wesley.
- Fortuna, M. A., Bonachela, J. A., and Levin, S. A. (2011). Evolution of a modular software network. *Proceedings of the National Academy of Sciences of the United States of America*, 108(50), 19985–89.
- Fukuchi-Shimogori, T., and Grove, E. A. (2001). Neocortex patterning by the secreted signaling molecule FGF8. *Science*, 294(5544), 1071–4. doi:10.1126/science.1064252
- Gabor Miklos, G. L., Campbell, K. S. W., and Kankel, D. R. (1994). The rapid emergence of bio-electronic novelty, neuronal architectures, and organismal performance. In: R. J. Greenspan and C. P. Kyriacou (eds), *Flexibility and Constraint in Behavioral systems*. New York: John Wiley and Sons, pp. 269–93.
- Garcia-Bernardo, J., and Eppstein, M. J. (2015). Evolving modular genetic regulatory networks with a recursive, top-down approach. *Systems and Synthetic Biology*, 9(4), 179–89. doi:10.1007/s11693-015-9179-5
- Giacomantonio, C. E., and Goodhill, G. J. (2010). A Boolean model of the gene regulatory network underlying mammalian cortical area development. *PLOS Computational Biology*, 6(9), e1000936. doi:10.1371/journal.pcbi.1000936
- Goldschmidt, R. (1940). *The material basis of evolution*. Yale: Yale University Press.
- Gomez, G., Lungarella, M., and Eggenberger, H. (2004). Simulating development in a real robot: on the concurrent increase of sensory, motor, and neural complexity. In: L. Berthouze, et al. (eds), *Proceedings of the Fourth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*. Lund University Cognitive Studies, 117. Lund: LUCS, pp. 119–22. doi:citeulike-article-id:549800
- Haddow, P. C., and Tyrrell, A. M. (2011). Challenges of evolvable hardware: past, present and the path to a promising future. *Genetic Programming and Evolvable Machines*, 12(3), 183–215.
- Katz, P. S. (2011). Neural mechanisms underlying the evolvability of behaviour. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1574), 2086–99.
- Kauffman, S. A. (1969a). Homeostasis and differentiation in random genetic control networks. *Nature*, 224(5215), 177–8.
- Kauffman, S. A. (1969b). Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology*, 22(3), 437–67. doi:http://dx.doi.org/10.1016/0022-5193(69)90015-0

- Kauffman, S. A. (1990). Requirements for evolvability in complex-systems—orderly dynamics and frozen components. *Physica D*, 42(1–3), 135–52.
- Kirschner, M. W., and Gerhart, J. C. (2006). *The plausibility of life: resolving Darwin's dilemma*. Yale: Yale University Press.
- Kitano, H. (1990). Designing neural networks using genetic algorithms with graph generation system. *Complex Systems*, 4, 461–76.
- Krubitzer, L. A., and Kaas, J. (2005). The evolution of the neocortex in mammals: how is phenotypic diversity generated? *Current Opinion in Neurobiology*, 15(4), 444–53. doi: 10.1016/J.Conb.2005.07.003
- Krubitzer, L. A., and Seelke, A. M. H. (2012). Cortical evolution in mammals: the bane and beauty of phenotypic variability. *Proc. Natl Acad. Sci. USA*, 109(Supplement 1), 10647–54. doi:10.1073/pnas.1201891109
- Langton, C. G. (1995). *Artificial Life*. Reading, MA: Addison-Wesley.
- Le Goues, C., Forrest, S., and Weimer, W. (2010). The case for software evolution. In: G.-C. Roman (ed.), *FoSER '10 Proceedings of the FSE/SDP workshop on Future of software engineering research*. New York, NY, USA: ACM.
- Le, M. N., Ong, Y.-S., Jin, Y., and Sendhoff, B. (2009). Lamarckian memetic algorithms: local optimum and connectivity structure analysis. *Memetic Computing*, 1(3), 175. doi:10.1007/s12293-009-0016-9
- Lehman, J., Wilder, B., and Stanley, K. O. (2016). On the critical role of divergent selection in evolvability. *Frontiers in Robotics and AI*, 3, 45.
- Lehman, M. M., and Belady, L. A. (1985). *Program evolution: processes of software change*. San Diego, CA: Academic Press Professional, Inc.
- Lewis, J. (2008). From signals to patterns: space, time, and mathematics in developmental biology. *Science*, 322(5900), 399–403. doi:10.1126/science.1166154
- Lichtensteiger, L., and Eggenberger, P. (1999). Evolving the morphology of a compound eye on a robot. In: *The Third European Workshop on Advanced Mobile Robots, (Eurobot '99)*, Zurich, pp. 127–34. doi: 10.1109/EURBOT.1999.827631
- Lindenmayer, A. (1968). Mathematical models for cellular interaction in development: Parts I and II. *Journal of Theoretical Biology*, 18, 280–315.
- Lipson, H., and Pollack, J. B. (2000). Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799), 974–8. doi:http://www.nature.com/nature/journal/v406/n6799/supinfo/406974a0_S1.html
- Mallatt, J., and Chen, J. Y. (2003). Fossil sister group of craniates: predicted and found. *Journal Of Morphology*, 258(1), 1–31.
- Mannaert, H., Verelst, J., and Ven, K. (2012). Towards evolvable software architectures based on systems theoretic stability. *Software: Practice and Experience*, 42(1), 89–116.
- Muller, G. B. (2007). Evo-devo: extending the evolutionary synthesis. *Nat. Rev. Genet.*, 8(12), 943–9.
- Nolfi, S., and Floreano, D. (2000). *Evolutionary robotics*. Cambridge, MA: MIT Press.
- O'Leary, D. D., and Sahara, S. (2008) Genetic regulation of arealization of the neocortex. *Current Opinion in Neurobiology*, 18(1), 90–100. doi: 10.1016/j.conb.2008.05.011
- Parisi, D. (1996). Computational models of developmental mechanisms. In: R. Gelman and T. K. Au (eds), *Perceptual and cognitive development*. San Diego, CA: Academic Press, pp. 373–412.
- Parisi, D. (1997). Artificial life and higher level cognition. *Brain and Cognition*, 34(1), 160–84.
- Pfeifer, R., and Bongard, J. C. (2006). *How the body shapes the way we think: a new view of intelligence*. Cambridge, MA: MIT Press.
- Pfeifer, R., and Gómez, G. (2009). Morphological computation—connecting brain, body, and environment. In: B. Sendhoff, E. Körner, O. Sporns, H. Ritter, and K. Doya (eds), *Creating brain-like intelligence: from basic principles to complex intelligent systems*. Berlin, Heidelberg: Springer, pp. 66–83.

- Prescott, T. J. (2007). Forced moves or good tricks in design space? Landmarks in the evolution of neural mechanisms for action selection. *Adaptive Behavior*, 15(1), 9–31.
- Prescott, T. J., Redgrave, P., and Gurney, K. N. (1999). Layered control architectures in robots and vertebrates. *Adaptive Behavior*, 7(1), 99–127.
- Prud'homme, B., Gompel, N., and Carroll, S. B. (2007). Emerging principles of regulatory evolution. *Proc. Natl Acad. Sci. USA*, 104(Suppl 1), 8605–12. doi:10.1073/pnas.0700488104
- Raff, R. A. (1996). *The shape of life: genes, development and the evolution of animal form*. Chicago: Chicago University Press.
- Rakic, P. (2009). Evolution of the neocortex: perspective from developmental biology. *Nature Reviews Neuroscience*, 10(10), 724–35. doi:10.1038/nrn2719
- Sendhoff, B., and Kreutz, M. (1999). A model for the dynamic interaction between evolution and learning. *Neural Processing Letters*, 10(3), 181–93.
- Sims, K. (1994). Evolving virtual creatures. In: D. Schweitzer, A. Glassner, and M. Keeler (eds), *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*. Association for Computing Machinery, Inc., pp. 15–22. doi:10.1145/192161.192167
- Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines*, 8(2), 131–62. doi:10.1007/s10710-007-9028-8
- Stanley, K. O., and Miikkulainen, R. (2003). A taxonomy for artificial embryogeny. *Artificial Life*, 9(2), 93–130.
- Swalla, B. J. (2006). Building divergent body plans with similar genetic pathways. *Heredity*, 97(3), 235–43. doi: 10.1038/Sj.Hdy.6800872
- Torben-Nielsen, B., Tuyls, K., and Postma, E. (2008). EvOL-Neuron: Neuronal morphology generation. *Neurocomputing*, 71(4–6), 963–72. doi:http://dx.doi.org/10.1016/j.neucom.2007.02.016
- Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences*, 237(641), 37.
- Vujovic, V., Rosendo, A., Brodbeck, L., and Iida, F. (in press). Evolutionary developmental robotics: improving morphology and control of physical robots. *Artificial Life*.
- Wagner, G. P., and Altenberg, L. (1996). Perspective—complex adaptations and the evolution of evolvability. *Evolution*, 50(3), 967–76.
- Whitacre, J., and Bender, A. (2010). Degeneracy: a design principle for achieving robustness and evolvability. *Journal of Theoretical Biology*, 263(1), 143–53.