

Further clues come from analysing the development of the mind in small children. In recent years, the abilities of the infant mind have been increasingly recognized. Susan Carey, taking the concept of number as a case study, finds that infants and non-human primates share the ability to make simple numerical distinctions. Humans, rather than being born with a mind like a blank sheet, seem to have some basic cognitive skills that are programmed into the genome and available from a very early age. Finally, Steven Pinker traces several aspects of seemingly irrational human behaviour back to their probable roots in early hunter-gatherer societies.

## Chapter 5

### Artificial evolution in virtual and real world settings

*Daniel C. Dennett*

The law of uphill analysis and downhill synthesis (Braitenberg, 1984) states that it is much easier to deduce the behavioural competence of a system whose internal machinery you have synthesized than to analyse the machinery of a black box whose behavioural competence you have observed. This is particularly the case when studying the mind (Dennett, 1994a): downhill synthesis can shed light on issues that are difficult to understand by reverse engineering, i.e., trying to analyse an existing system. Artificial Life uses the bottom-up approach to model cognitive competence by examining how simpler systems could have evolved into more sophisticated ones; Artificial Intelligence, by contrast, starts at the top by attempting to build from scratch machines capable of complex tasks. The two methods produce important differences in the architecture of the systems, partly because of the limited knowledge of the designer.

Artificial Life projects usually depend on a form of genetic algorithm, in which bit strings serve as genomes, formal recipes for performing some task or building some device. These are randomly mutated and set in competition against each other, the winners being permitted to replicate and advance to the next round of competition. I illustrate the power and limitations of this method by comparing three projects: two involve selection in an entirely simulated or virtual environment, the third in a real environment. The comparison demonstrates the importance of making the evolutionary setting as close as possible to the real world, allowing for the myriad physical effects of actual environments. Some of the startling results of these projects confirm the power of bottom-up strategies, particularly as the models

produce solutions to problems that a human engineer would be unlikely to consider. In Artificial Intelligence, the importance of using the real world, rather than a simulated one, as the challenge has been recognized by those who have moved away from 'bedridden' systems (Dennett, 1979) to build robots that have to move and act in the world. The state of robotics is represented here by Cog, a humanoid robot that is being developed using inputs from the environment interacting with a basic structure.

### *Computer-generated models of evolution*

Evolution both of body form and behaviour occurs even in a model where organisms do not interact and neither the genome nor the developmental program receives any input from the environment. In the simulation program Evolved Creatures (Sims, 1994a,b), a set of simple creatures was generated using a series of random numbers as the genetic code. The code determined the phenotype, i.e., the size, number and arrangement of the articulated blocks that made up the body, which had joint sensors and muscles, and the organization of the simple nervous system. These virtual creatures competed for the best locomotory ability either in a virtual liquid representing the sea or on a virtual solid plane that served as a land surface. The winners of these competitions were repetitively bred and mutations were simulated by making random changes in the original genetic code.

In spite of the extreme oversimplification of the creatures and their environments, different body plans emerged and, in later generations, some familiar characteristics, such as symmetry (or near-symmetry) evolved several times. Some individuals were unable to move as a result of a mutation that damaged the nervous system but others evolved unique and unexpected solutions. Some developed the ability to swim when placed in water. Selection was accomplished by an objective, automatic test: each candidate was placed in the simulated liquid space and allowed to behave according to its genotype for a fixed period. The

distances from the origin — if any — covered by all the competing genotypes were measured and the winners advanced to the next generation. Within a few dozen generations, several efficient, graceful swimming forms evolved.

When put on land, various forms of locomotion appeared within 15 generations. The blindly automatic selection process was nicely revealed by a lineage that simply measured the distance, of any body part, from the origin: it was 'born' at the origin as an upright tower of connected blocks and simply fell over! In subsequent generations it was taller and taller at birth and even adjusted its entirely rigid form to execute a sort of somersault when it landed, approximately doubling its distance from the origin. The loophole in the selection regime that allowed this creature to evolve was closed in subsequent competitions, making the selective environment somewhat more realistic.

Selection for other behaviours included jumping, phototaxis and competing for control of a cube, rather like the face-off that starts a hockey game. The computer-generated solutions were often completely different from anything that a computer programmer would have devised. Bizarre shapes that utilized friction or angular momentum in novel ways were among the previously unimagined solutions discovered by this evolutionary process.

This virtual world has several shortcomings. Because the physical forces of the real world are largely absent from the program, many limitations are missing, e.g., energetic demands, resource limitations, wear-and-tear and cost/benefit complexity. Another problem is that neither the genome nor the developmental program receives an input from the environment, so interventions, such as changing the length of the genome, were required to enable new abilities, such as phototaxis, to evolve. Lastly, the system contains no noise, which turns out to be an important lubricant in the evolutionary process (see below).

Some of these limitations are taken into account in another

simulation, ECHO, which is a general platform for exploring artificial life (Holland, 1995). As in Sims' work, the creatures that have evolved in this simulated environment had to find efficient solutions to oversimplified and idealized problems: locomoting, finding 'food' and 'mates' and avoiding predation by other virtual inhabitants of their virtual world. But the world in this simulation is significantly more realistic: the basic building blocks in ECHO are not the body segments used in Evolved Creatures but resources that can be employed in many ways. They may become body parts but can also form part of the genome, represent energy for moving and acting, or be material resources from which offspring can be built.



The genome is in the virtual world, so that both the developmental process and the length of the genome are visible to selective forces, thus opening up opportunities for environmental novelty to intrude spontaneously into the selection process. This means ECHO is open-ended in a way that Evolved Creatures is not. Evolution requires an abundance of undesigned bits and pieces to serve as raw material for incorporation into designed bits and pieces.

Although sexual reproduction was built into ECHO, one of the biological features that has emerged in simulations is an elaboration of mating systems. Varieties of parasitism, symbiosis and mimicry, and even some crude forms of communication, have also appeared. In spite of the simplicity of ECHO, the fecundity of its evolutionary process is striking: the solutions have an ingenuity that is manifestly not imbued by the creators of the system and that combine biological familiarity and novelty. If we discovered life on a distant planet, we would expect to find both convergent evolution, due to underlying deep similarities posed by environmental problems, and entirely novel forms arising from the differences in environment. Both types of solution are abundant in the alien simulated worlds of Artificial Life.

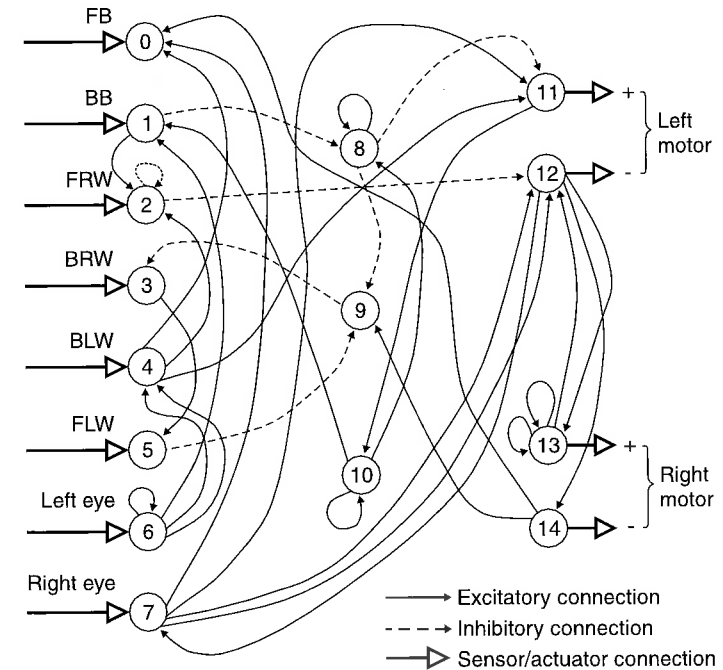
### *Entering the real world*

The most serious limitation of computer simulations is the lack of input from the real world. This is being overcome by using robots, to which real-world physics apply, programmed using genetic algorithms. A popular robot platform for such experiments is the tiny Khepera, designed and built by Francesco Mondada at the Federal Institute of Technology in Lausanne, Switzerland. It is about the size of an ice-hockey puck, with wheels, three eye spots and vibrissae for collision detection. A simple neural net connects the sensors to the wheels (see Fig. 25) and contains a program that allows several individuals to flock and compete. The bit strings that make up the genome completely determine the phenotype. The locations of the eye spots and whether the connections between the sensors and wheels are excitatory or inhibitory is under genetic control. All the connections in the controlling neural net are subject to change over the course of evolution but are not plastic in the individual, unlike the connections in many neural nets that 'learn'.

To enable Khepera-type robots to evolve efficiently, a robot simulator suspended from a gantry has been built (Figs 24, 25; Harvey *et al.*, 1997). It was driven by a developmental program similar to that of Evolved Creatures (Sims 1994*a,b*). The workings of the Khepera were simulated in the controlling computer at the same time that the real sensory and motor interactions with the world were accomplished by the robot-interface (Fig. 24). This obviated the need to accomplish the developmental program specified in the genome by physically rewiring the robots and relocating their eye spots. As each candidate genotype emerged, the computer faithfully built a simulated nervous system following the instructions in the genome, with just three pixels from the built-in TV camera representing the eye spots. Then the computer positioned the gantry at the point of origin in the middle of the table and the genotype was given trials similar to those used by Sims (1994*a,b*) but in real time and space.



**Figure 24.** The gantry robot, a system for testing the effects of phenotypic differences between simple robots in the real world. **a**, the robot consists of an array of sensors carried on a platform that moves on wheels in x and y coordinates. Information from the sensors, including the video camera, is fed to a computer where it provides the input to a simulated nervous system. **b**, detail of the array of sensors: the camera inside the top box points down at an inclined mirror, which can be turned by a stepping motor. The plastic disk suspended from a joystick senses bumps. Modified from Harvey et al., 1997.



**Figure 25.** The three-layer neural network of a Khepera like that used to control the gantry robot. Units on the left are initially designated as input units: BB, FB, back and front bumpers; BRW, BLW, FRW, FLW, back and front, left and right whiskers. Units on the right output to the motors for the left and right wheels. Centre column contains hidden units. Modified from Harvey et al., 1997.

The fitness of the phenotypes specified by the different genotypes was tested in a walled environment, e.g., for phototaxis, they had to learn to avoid a triangle painted on one wall and move towards a square on the other. Each individual was evaluated on three trials and selected on the worst one, in simulation of nature — ‘you are only as good as your worst day’. Those with superior performance were allowed to mate, i.e., their genomes were subjected to crossover and mutation in a step of the genetic algorithm and the resulting genomes became part of the next generation.

The robots evolved in some unexpected ways. The nervous systems of the winners developed complex functions that were difficult to determine using reverse engineering and known design principles. Most complex human artefacts, such as engines, computers, aeroplanes and assembly lines, are composed of many single-purpose subsystems carefully isolated from each other to prevent unintended side effects from interfering with their operations. Like biologically designed systems, these robots largely ignored such principles. When Harvey and his colleagues attempted to give some of their robot genotypes a head start by hand-designing first-generation candidates, evolution frequently discarded their handiwork and replaced it with better, but largely inscrutable, solutions. Although these did not fit human ideas of elegance or efficiency, they were very effective, a good example of Leslie Orgel's Second Rule: "Mother Nature is smarter than you are" (see Dennett, 1995).


Noise was very important for effective evolution as it helped the non-plastic nervous system to work better. Although the robots were initially designed with three eye spots, which was thought to be near the functional limit, in several lineages one eye spot came to 'look' in an irrelevant direction, throwing away its input and introducing more endogenous noise into the system. This noise kept the system from settling into sub-optimal states and so was more valuable to the robot than the external information that could have been obtained through a 'seeing' eye spot.

The power of such a simple set of transducers to exploit information in the environment should not, however, be underestimated. In one evolutionary run, conducted by the computer over several days of continuous trials and generations, diurnal and nocturnal subspecies emerged. The ancestors of the diurnal variety had mainly been tested when daylight from the lab windows contributed to the ambient light, whereas the nocturnal version arose from those tested at night. This entirely unintended and unanticipated outcome is a dramatic demonstration

of the importance of a real world as the selective environment.

### *Cog, a humanoid robot*

Cog, a life-sized humanoid robot (*Fig. 26*), is being developed at Massachusetts Institute of Technology by Rodney Brooks and colleagues as a tool for studying both evolution and the human mind (see Dennett, 1994*b* and the Cog website\*). Cog has two arms, using series elastic joints with a variable spring constant regulated by fast feedback motors. Although it has no legs, it can move its body with the same degrees of freedom as a human body. Four microphones serve as ears and the two eyes have both foveal (narrow angle) and parafoveal (wide angle) cameras. The eyes saccade below the 100msec range at three to four fixations per second, approximately the same rate as in humans. Parts of the body are covered with touch-sensitive plastic skin and detectors including strain gauges and heat sensors approximate to an innate pain system.



**Figure 26.** A portrait of Cog. Reproduced with kind permission of R. Brooks.

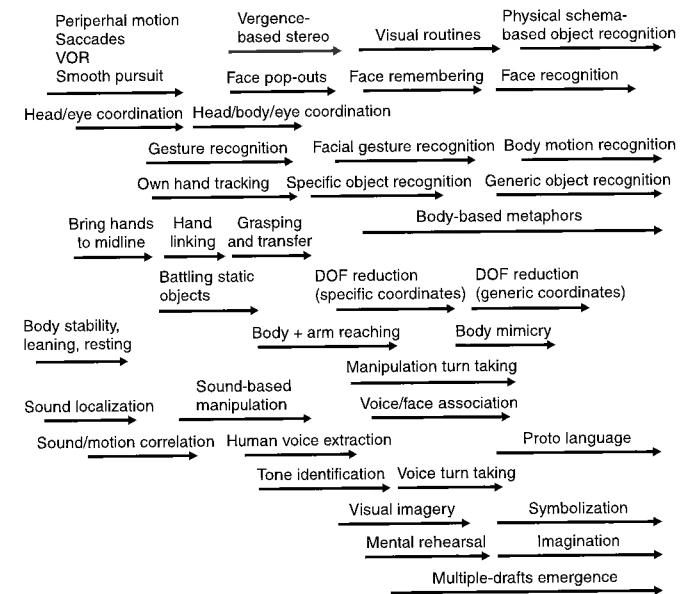
\*Cog website: [www.ai-mit/projects/cog/Text/cog-shop.html](http://www.ai-mit/projects/cog/Text/cog-shop.html)

The original distributed competitive nervous system consisted of 128 Mac2 computer central-processor chips connected in parallel, with no central executive\*. The programs that ran on this hardware, written in a new parallel dialect of Lisp, were partly written by hand and partly evolved using either genetic algorithms or various connectionist learning algorithms. Comparing the efficiency and grace of hand-coded versions of some of these programs with those designed by one of these evolutionary processes encourages the use of the evolutionary methods wherever possible.

Cog's movements depend on feedback and it can perform actions requiring sophisticated motor control, such as manipulating a 'Slinky' toy between its hands, swinging a pendulum back and forth, and hammering a nail into a board, with the joints absorbing much of the shock. It can recalibrate its movements to accommodate for 'growth' and wear-and-tear. For example, the software rapidly compensated for a new set of arms that differed in length and weight from the previous ones and for changes in eye position caused by slippage in the equipment.

Having equipped Cog with hardware and basic software, the intention is to see how it develops from its present 'infancy'. Projected developments are set out in *Figure 27*. The distributed nervous system is provided with separate modules to accomplish different tasks, an organization reminiscent of domain-specific learning (see Gallistel, this volume). Communication between the modules is not designed in advance but emerges with time through their interaction. We hope that the systems that have evolved to solve certain problems will become modified to be available for solving other problems. We anticipate that an intelligent Cog will not have a central processing module but that the task-related modules will develop to access each other to some degree, e.g., to enable a temporal problem to be mapped on to

\* This hardware has been superseded in the summer of 1998 by a second-generation architecture with greater power and ease of use.



*Figure 27.* A scheme of Cog's projected development. Tasks progress from simple on the left and at the top to more complex and abstract on the right and at the bottom. Reproduced with kind permission of R. Brooks.

another dimension, such as the use of an analogue clock to determine time elapsed. Humans achieve this mapping effortlessly and continually use it as a useful crutch for thinking. Cog may discover some mappings of its own that will enable it to exploit its own competences.

The results from Cog may not resemble the very complex behaviours, task performance or consciousness that characterize the functions of the human nervous system but there may be several ways to build a consciousness. In one sense Cog is already a success: according to a Chinese legend, a sage was fishing in the river with a straight pin; this curious news reached the emperor, who was so puzzled that he went to see the man and asked him what he expected to catch; the sage replied, "You, my dear friend". Cog is a preparation for examining a problem, a

stimulus to think about cognitive problems from a different perspective, without assuming *a priori* which are the hard problems.

## Mirror neurons

**Giacomo Rizzolatti**

How do we understand the meanings of the actions of other individuals? Social interaction depends on understanding the behaviour of others, on recognizing that certain behaviours are friendly and others menacing. These capacities are not uniquely human for a monkey understands the social signals emitted by other monkeys and the meaning of actions made by others, such as grasping food. We have discovered neurons with properties that may provide the neuronal substrate for understanding the social signals emitted by others (Rizzolatti *et al.*, 1996; Gallese *et al.*, 1996).

In the rostral part of the ventral premotor cortex (area F5) of the monkey, some neurons were found previously that discharge during goal-directed actions (Rizzolatti *et al.*, 1988). They were classified according to the action that triggers their response: the most common are 'grasping', 'holding' and 'manipulation' neurons. A certain percentage of neurons in F5 discharge both during an action and in response to the presentation of a 3-D object appropriate to the action (Murata *et al.*, 1997). The circuit that mediates the transformation of the intrinsic properties of the object, i.e., size and shape, into hand movements is formed by parietal area AIP and F5 (Jeannerod *et al.*, 1995; see Sakata 1996). Our recent studies have revealed an unexpected class of visuomotor neurons in F5 that discharge both when the monkey grasps or manipulates objects and when it observes the experimenter making a similar gesture. We have called these 'mirror' neurons (Rizzolatti *et al.*, 1996; Gallese *et al.*, 1996).

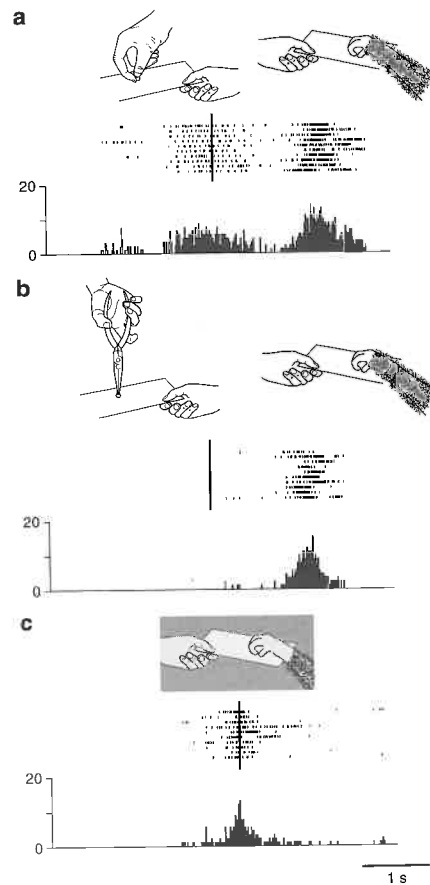
## Properties of mirror neurons

Mirror neurons differ from the neurons in F5 that are involved in object-action transformation because they do not respond to the presentation of objects but are specifically activated by an observed action. Most discharge selectively when the monkey observes a particular action, most commonly grasping (Fig. 28), manipulating, tearing or putting an object on a plate. Some are highly specific, coding not only the aim of the action but also how it is executed. For example, they discharge only during the observation of movements when the object is grasped with the index finger and the thumb. All mirror neurons show visual generalization, i.e., their responses are independent of how far away the agent of the observed action, typically a hand, is. A few respond even when the object is grasped by the mouth.

By definition, mirror neurons have motor properties, typically with a congruence between the effective movement the monkey observes and the one the monkey makes. In some neurons this congruence is extremely strict, the effective motor action coinciding with the observed action that triggers the response. In others it is broader but the motor requirement is usually stricter than the visual; e.g., making a precision grip but observing any type of hand grasping.

## Functional implications

To understand the function of mirror neurons it is important to distinguish between two types of mirror phenomena. Both share the property that some neurons with motor responses are specifically excited, or 'resonate', when the animal observes an action similar to the one coded for by the neuron. But the resonance may have one of two purposes: to allow an individual to repeat movements made by another individual, a mechanism referred to here as first-level mirror behaviour; or to allow an individual to understand actions made by others, which I term second-level mirror behaviour. The latter is the mechanism observed in F5 neurons.



**Figure 28.** Visual and motor responses of a mirror neuron. The experimental situations are represented in the upper part of each panel with the neuron's response shown as an activity raster and histogram below. **a**, the neuron starts to discharge when the experimenter grasps a piece of food from the tray with his hand; it ceases when the food is replaced and the tray is moved toward the monkey; the discharge starts again when the monkey grasps the food. **b**, the neuron does not respond when the experimenter uses a tool to grasp the food. **c**, the monkey grasps food in darkness; the discharge represents the motor activity of the neuron. Vertical line, alignment of rasters and histograms: (**a**, **b**) with the moment when the experimenter grasps the food, (**c**) with the monkey grasping. Modified from Rizzolatti et al., 1996.

Many behaviours in animals can be described as first-level mirror behaviour. The simultaneous flight of an entire flock of birds when its leader starts flapping its wings (Tinbergen, 1953) belongs in this category, as does song learning in some species of songbirds. An understanding of the action is not required, only a fast resonance in the motor system of the observing individual in response to the observed behaviour. It is noteworthy that there is neurophysiological evidence for neurons with mirror properties in songbirds (see Hauser, 1996).

First-level mirror phenomena are also seen in humans, with the best example probably being the capacity of very young infants to imitate buccal and manual gestures (Meltzoff and Moore, 1977). The motor resonance to the observed action seems to be hard wired, because the infant is able to imitate facial gesture even though it has never seen its own face. This is not a simple release of a motor activity, because the response can be artificially delayed (Meltzoff and Moore, 1977), but the aim is to reproduce the observed movement. The learning of syllables by human infants is probably also a first-level mirror behaviour, where the understanding of the action, i.e., the meaning of the utterance, is not required.

The response of neurons in F5, although based on resonance, differs from first-level mirror behaviour in two important ways. First, the motor resonance elicited by observing an action does not determine an overt motor response, i.e., there is no activity in primary motor cortex or the muscles. In one experiment (Rizzolatti et al., 1996), we placed a second monkey in front of the monkey from whom recordings were made and gave the newcomer food. Grasping neurons discharged in the recorded monkey every time the new monkey grasped food. An imitative behaviour would obviously be pointless — the appropriate response would be to chase the newcomer away. The recorded monkey exhibited disappointment with the new situation but never attempted to grasp the food. Even so, the F5 mirror



neurons fired every time the newcomer made a grasping movement.

Second, to be effective the imitation characteristic of first-level mirror behaviour must be precise: infants protruded their tongues in the same direction as the experimenters (see Meltzoff and Moore, 1997). In contrast, the mirror neurons in F5 display enormous generalization, most responding when the grasping hand is seen to move towards or away from the observer, or from above or below the object. In terms of motor organization, these are different movements but all these actions have only one meaning.

I argue that this second-level mirror activity can be interpreted as the mechanism underlying the understanding of motor events. Individuals typically 'know' or predict the consequences of movements they make. This knowledge most probably results from an association between the representation of the motor action, coded in F5 and other motor centres, and the motor consequences of the action. The resonance mechanism does not determine the occurrence of the motor response but evokes neural activity corresponding to the internally generated activity representing a certain action. The meaning of an action can, therefore, be recognized because the two representations are similar.

### *Evolutionary consequences of mirror neurons*

Mirror behaviour seems to be a very interesting example of a mechanism that originated for one purpose, in this case the immediate response of one or more individuals to a signal emitted by another, and has progressively evolved into a complex cognitive system. The evolution of mirror behaviour most probably has not stopped with the recognizing the actions of others. Evidence from comparative anatomy, brain imaging experiments and transcranial magnetic stimulation, as well as from the structure of language, indicate that human language may be derived

from the mirror system (Rizzolatti and Arbib, 1998).

The messages detected by mirror neurons in F5 are not sent voluntarily but are part of the actor's motor repertoire. The receiver, however, can interpret them. This becomes a language when the actor realizes that his/her actions can influence the receiver's behaviour and becomes a sender; conversely, the receiver realizes that the observed behaviour is not unintentional but is precisely directed at her/him. The link between sender and receiver is still, however, the ancient mechanism that makes motor neurons resonate. Although this is only a working hypothesis, the different aspects of the mirror system provide a good illustration of the power of the evolutionary approach for understanding the development even of specifically human capacities, such as spoken language.

## Discussion

Cog can help us to learn about the brain. Cognitive psychology and cognitive neuroscience are the reverse engineering of the mind but a good way to reverse engineer any machine is to try and build it. One discovery of this kind of work is that, when the machine to be emulated is the result of evolution, the principles used for forward engineering need to parallel those found in nature. So Cog's brain does not reflect arbitrary choices made by engineers. Instead, it is designing itself, going through infancy with a history based on real encounters, building up reactions and competences as a child would. This contrasts with top-down models where all the programs are written by hand and the adult emerges, as God designed Adam.

Forward engineering is a good way to identify the hard problems, which are not always obvious — some problems that seem easy turn out to be very hard. There is also the question of what

assumptions can be made. For example, as human babies show mirror behaviour, should some form of mirror neurons be built in as a basic provision or will some analogue of mirror neurons develop spontaneously? The systems Cog develops may not resemble the human nervous system but perhaps there are a number of ways to build a consciousness.

One issue that can be explored by observing Cog is the two opposing ideas about how the cerebral cortex works. The first is that the cortical sheet has a homogeneous structure and uses a general computational algorithm that can be applied equally well to visual perception, motor control and language production. This universal algorithm, which is not yet understood, can cope with any problem the animal encounters, given the correct processing in the input and output pathways. The second view is that the cortex consists of many discrete units, each with a unique algorithm for processing a particular type of information. Although throughout the cortex the structure looks roughly similar under a microscope, this gross similarity does not prove the existence of a universal processing algorithm; indeed molecular and biophysical differences have been demonstrated in several areas. This is also the case with Cog, where the control system looks homogeneous but units differ in the voltages stored in the software.

Proponents of a universal computational algorithm argued that the comparison between the brain and the computer is faulty. In neuronal networks the programme is the structure, which includes the pattern and strengths of the connections. When additional cortical tissue is grafted into a brain, it develops connectivity appropriate to the inputs at that location, not those of its origin; e.g., a piece of visual cortex grafted into the auditory area develops a tonotopic map (Sur *et al.*, 1990). Despite the parcellation of the cortex into many areas with their own specialized functions, many neuroscientists would not expect different pieces of cortex to operate differently at the

abstract computational level, whereas cognitive scientists do.

This disagreement could be a major dilemma for tackling questions about cortical functions. Understanding how complex problems are processed requires additional information about the algorithm(s). Both groups agreed that function is restricted by structure. The debate is whether the restriction is to a single solution or to a set of possible solutions that contains a large number of possible outcomes. Finally, this comes down to the level of analysis: whether the algorithm is determined by overall cortical architecture or by the fine details of connectivity in individual areas. The solution may lie between the two extremes: a basic algorithm that is adapted to solving different problems by small differences in implementation, possibly at the molecular level, that are area specific.

The debate might be assisted by looking beyond the cortex to the cerebellum (Leiner *et al.*, 1991). One idea is that the cerebellum serves as the medium for cortico-cortical pidgins. A pidgin is an *ad hoc* language, developed on the fly by people who need rudimentary communication but do not share a natural language (see Pinker, 1994). Vocabulary items are generally adapted promiscuously from the languages that the speakers know and are put together with a minimum of grammatical nicety, although grammatical patterns soon emerge.

Cortical regions that have developed a domain-specific competence may have difficulty in communicating directly if they execute different kinds of algorithms. A cortico-cortical pidgin would be an *ad hoc* neuronal medium of communication between such regions. It could be that the problem of successfully 'interfacing' such domains is solved by cerebellar intermediaries. The cerebellum and cortex have evolved in parallel, the architecture of the cerebellum is suited to this task and it seems to be involved in precisely timed processes and sequences. As a result of cerebellar lesions, cortical processes that involve timing deteriorate. The neuro-imaging studies that show cerebellar

involvement in relatively difficult and unpractised tasks may be giving us a view of such *ad hoc* pathways of communication being created in the course of cognitive effort.

## Chapter 6

### Evolution and development of brains and cultures: some basic principles and interactions

---

*Paul Rozin*

Novelty can arise in biological evolution, cultural change and cognitive development, although on different time scales. On the long time course of biological evolution, new species with complex, adaptive structures, such as eyes or wings, emerge. Human cultural evolution has a medium time course, with the appearance of social organizations, such as democracy, and new technologies, e.g., tools, printing and electronics. In the individual, cognitive development occurs in a relatively short time, as representational and computational capacities expand in childhood until adult levels of competence are achieved. I believe that similar mechanisms generate novelty on each of these time scales. Small quantitative changes can lead to large qualitative differences in structure or process through mechanisms that include natural selection and preadaptation. Several developmental processes can help to promote large changes with little extra demand on the genome.

Here I discuss the role of preadaptation, probably the most important of these mechanisms. The concept of preadaptation, originally formulated by Charles Darwin, was named by George Gaylord Simpson and developed by Walter Bock (1959) and Ernst Mayr (1960). Contrary to popular opinion, preadaptation is not teleological but indicates a structure or process that is used for a function for which it did not originally evolve; in other words, borrowing rather than starting from scratch (Mayr, 1960). A parallel can be identified in cognitive development, where it is termed accessibility, i.e., the extension of a mental capacity that is highly specialized for solving one type of problem to another type of problem (Rozin, 1976). Similarly in