

Developmental Robotics: Manifesto and Application

BY TERRY ELLIOTT AND NIGEL R. SHADBOLT

*Department of Electronics and Computer Science, University of Southampton,
Highfield, Southampton SO17 1BJ, UK*

We argue that all embodied organisms, whether robots or animals, face the same challenge: of adapting to bodies, brains and environments that undergo constant and inevitable change. After highlighting the evidence for the universal role of a class of molecular factors called neurotrophic factors in the response of animals to this challenge, we suggest that implementing models of neurotrophic interactions on robots may confer on them the adaptability and robustness exhibited by animals. We briefly review a mathematical model of neurotrophic interactions and then discuss its application in a robotic context. Finally, we examine the potential, or otherwise, of our approach to developmental robotics.

Keywords: neuronal development; synaptic plasticity; activity-dependent competition; Koala robot; Khepera robot; computer vision.

1. Of Robots and Men: A Manifesto

The Declaration of Independence asserts that, “[w]e hold these truths to be self-evident, that all men are created equal...” Yet, as any student of developmental biology knows, conspecifics of a given vertebrate species, and to a lesser extent of a given invertebrate species, differ often quite dramatically in their detailed morphologies. And even if all conspecifics were created equal, necessary differences in experience and life would rapidly ensure that they do not remain equal in their detailed morphologies for very long.

All animals, whether vertebrate or invertebrate, undergo an elaborate developmental program during which their bodies and brains are built (and sometimes rebuilt during metamorphosis), with this program often extending into the post-embryonic period. In vertebrates, a quintessential feature of ontogenesis is programmed cell death (Gilbert 2000); cell death also occurs during invertebrate development (Truman 1984), but we shall focus here on vertebrate development. For example, programmed cell death carves out the fingers and toes of the developing foetus. Small differences between the genetic programs and embryonic environments of conspecifics can lead to large differences in the resulting organisms (Edelman 1993), with, for example, skeletal and muscular structures often differing dramatically: students new to the dissection room are frequently surprised by missing or unusual muscles in their cadavers.

If the bodies of conspecifics differ, then their brains must accommodate these differences in order to control them. Extensive cell death occurs during neuronal development too (Purves & Lichtman 1985), with often up to half to three-quarters

of all neurons in a given structure dying during early development. This cell death occurs principally during a phase of neuronal development called target innervation, during which neurons send out axons to make synapses with their intended targets. Their targets may be, for example, muscles or other neurons. Once contact is established, their targets supply molecular factors called neurotrophic factors on which the survival of the innervating neurons depends. If a neuron receives too little factor, it dies. Target innervation and the neuronal death that result therefore constitute an adaptive mechanism by which a brain can learn about its body. The neuronal pool that innervates a muscle, for example, ends up at about the right size for that muscle: had the muscle contained more fibres, more neurons would have survived; fewer fibres, and more neurons would have died.

But the need for the brain to adapt to its body does not end at birth. Vertebrates do not emerge into the world fully grown. Mammals, for example, grow slowly during the pre-pubescent phase and undergo explosive growth during puberty. As any pubescent human male knows to his embarrassment, the growth of the vocal cords is so rapid that strict control is lost. Yet the brain adapts, and eventually the cords are brought back under control.

This kind of “programmed” growth is not the only change that occurs during the post-natal phase, though. Bodies experience wear and tear, minor accidents occur, and perhaps new motor skills are required to survive environmental changes. For example, in response to the loss of a limb or an eye, the neuronal maps representing those parts of the body undergo change, with the neuronal representation of the absent extremity shrinking. The sensory systems of animals also routinely deviate from expected norms. This is well illustrated in the visual system. Animals can suffer from strabismus, in which the eyes do not look at the same point at the same time; they are often afflicted by various degrees of myopia and hyperopia; and astigmatisms are common, in which the cornea has an irregular shape and leads to problems in bringing objects into focus. Experimental studies have revealed that these differences in sensory systems lead to changes in their underlying neuronal representations. To all these day-to-day processes, then, the body adjusts, and the brain must adjust to compensate, to regain precise control or representation of its body.

Myogenesis largely ends before birth, and the changes in muscular size and strength that occur during a life-time take the form of changes in fibre size and physiology. Thus, the pool of neurons that innervates a muscle does not need to change in size during post-natal changes, but the neurons do need to adapt their control of the muscle, changing motor pools and the detailed physiology of the neuromuscular synapses. Similarly, post-natal changes in the sensory representation of the body and in the sensory representation of the environment occur not by cell death, but by synaptic re-arrangement (Purves & Lichtman 1985). Synaptic re-arrangement is frequently competitive in character, and depends on neuronal activity. More active neurons tend to be at an advantage during synaptic re-arrangement and end up with more complex axonal arborisations and greater numbers of synapses than less active neurons innervating the same structure.

During these post-natal changes, then, neuronal birth and death (largely) do not occur, but synaptic birth and death do. Recent experimental evidence, reviewed, for example, in McAllister *et al.* (1999), indicates that neurotrophic factors are also implicated in this later synaptic growth and re-arrangement, in addition to

their role in earlier neuronal death. We are therefore presented with a view of neurotrophic factors in which they are intimately involved during all stages of life in the tuning of a nervous system to its body and to itself as the body and the brain undergo development, growth, wear and tear, injury and decay (Purves 1988, 1994). Development, plasticity, learning and adaptation, on this view, are not separate, distinct processes, but form a continuum on the line of actions of neurotrophic factors at different phases in the life cycle.

Of what relevance are neurotrophic factors to robots? Man may try to create all robots equal, but usually he fails. Differences in motor performance or sensory acuity ensure that no robot is an exact “silicon copy” of another. Robotic neurocontrollers ideally should accommodate and exploit these differences. And even if man could create all robots equal, components would change and degrade at different rates according to different environmental circumstances. Although a robot’s body may not necessarily develop and grow, in other respects the requirements for a robot’s neurocontroller are similar to those for an animal’s brain: it should be able to adapt to wear and tear, changing environmental factors, differing environmental histories, injury and decay. Given these close similarities between robots and men, and given that evolution appears to have found a powerful, adaptive mechanism in neurotrophic support, it appears natural to seek to implement at least a coarse model of competition for neurotrophic support in a robot, in an attempt to confer on it some degree of adaptability both to its own body and to its environment.

This, then, defines our manifesto, and embodies truths that we hold to be self-evident: that robots and men live in the same world; that robots and men must respond to the same challenge of adapting to changing bodies, changing brains and changing environments; and that robots and men should employ the same mechanism to meet this challenge. Indeed in this regard robots and men are each and every one of them unique.

We now put this manifesto into action by first discussing briefly a well-studied mathematical model of competition for neurotrophic support (Elliott & Shadbolt 1998a). We then review the application of this model in two robotic contexts. Finally, we discuss whether or not our manifesto resembles most other manifestos: long on promises but short on delivery.

2. A Neurotrophic Model

Let pre-synaptic or afferent neurons be labelled by letters such as i and j , and post-synaptic or target neurons be labelled by letters such as x and y . Let the number of synapses between afferent cell i and target cell x be denoted by s_{xi} , and let the activity of afferent i be a_i . We assume that target cells release a neurotrophic factor (NTF) that sustains and promotes the growth of these afferent synapses. Let the uptake of NTF by afferent cell i from target cell x be denoted by u_{xi} , and let the average uptake be \bar{u}_{xi} . The central assumption of our neurotrophic model, consistent with experimental data (e.g., Campenot 1982a,b), is that the time average uptake of NTF by afferent i from target x determines the number of synapses projected by i to x ,

$$s_{xi}(t) = \bar{u}_{xi}(t). \quad (2.1)$$

For the time average, we assume that it takes the form of the integral of the function multiplied by a decaying exponential, where the rate of decay is itself time-dependent,

$$\bar{u}_{xi}(t) = \int_{-\infty}^t dt' \epsilon_i(t') u_{xi}(t') \exp\left(-\int_{t'}^t dt'' \epsilon_i(t'')\right). \quad (2.2)$$

Here, $\epsilon_i(t) = 1/\tau_i(t)$, and $\tau_i(t)$ is a time-dependent parameter that sets the instantaneous time-scale for the average for afferent i . From this we obtain the differential equation

$$\frac{ds_{xi}}{dt} = \epsilon_i(t) [u_{xi}(t) - s_{xi}(t)], \quad (2.3)$$

determining the evolution of s_{xi} . Thus, if the instantaneous uptake exceeds average uptake (which sets s_{xi}), then more synapses are grown; otherwise synapses are removed. The rate of evolution is set by ϵ_i , whose form we take to be

$$\epsilon_i(t) = \epsilon(\alpha + \beta a_i(t)), \quad (2.4)$$

where ϵ , α and β are constants. A constant, activity-independent rate is set by α , while activity in afferent pathway i increases the rate of plasticity through the term $\beta a_i(t)$.

As a model for uptake, u_{xi} , we suppose, consistent with experimental data (e.g., Castren *et al.* 1992; Koliatsos *et al.* 1993; Blöchl & Thoenen 1995; Goodman *et al.* 1996), that target cells produce and release NTF in an activity-dependent manner, so that the release from target cell x is given by

$$r_x = T_0 + T_1 \frac{\sum_i s_{xi} a_i}{\sum_i s_{xi}}, \quad (2.5)$$

where T_0 sets an activity-independent level, and T_1 determines a maximum, activity-dependent release of NTF. This released NTF is assumed to diffuse through the target field, so that the amount available for uptake at each target cell following diffusion is given by

$$d_x = \sum_y \Delta_{xy} r_y. \quad (2.6)$$

This is just the convolution of r_x with some characteristic diffusion function Δ_{xy} , which we take for simplicity to be a Gaussian. This diffused NTF is then taken up by afferents according to the equation

$$u_{xi} \propto s_{xi} g(a_i) \rho_i d_x, \quad (2.7)$$

where the constant of proportionality is set by the requirement that $\sum_i u_{xi} = d_x$, i.e. total uptake by all afferents exactly matches availability. This equation embodies the requirements that the uptake by afferent cell i from target cell x should be proportional to the number of synapses projected from i to x , s_{xi} ; some function of the activity of afferent i , $g(a_i)$, which for simplicity we take to be $g(a_i) = a + a_i$; the number of receptors for the factor supported by afferent i , ρ_i , which we take to be $\rho_i = \bar{a}_i / \sum_x s_{xi}$; and the total amount of NTF available for uptake, d_x .

Putting all these assumptions together, we finally obtain

$$\frac{ds_{xi}}{dt} = \epsilon_i(t) s_{xi} \left[\frac{(a + a_i) \rho_i}{\sum_j s_{xj} (a + a_j) \rho_j} \sum_y \Delta_{xy} \left(T_0 + T_1 \frac{\sum_j s_{yj} a_j}{\sum_j s_{yj}} \right) - 1 \right]. \quad (2.8)$$

For a fuller derivation and justification of the various assumptions behind this model, see Elliott & Shadbolt (1998a). Although this model appears complex, it can be shown to be equivalent to a non-linear Hebbian model of synaptic plasticity with competition between afferent neurons implemented through multiplicative post-synaptic normalisation (Elliott & Shadbolt 2002b), and it can also be shown that the model possesses an underlying spin glass structure, in which the function Δ_{xy} provides the “spin coupling” between target cells (Elliott 2002). We have applied this model to the development of the visual system (Elliott & Shadbolt 1998b, 1999, 2002a) and the neuromuscular junction (Elliott *et al.* 2001) in simulation, and have coupled the model to a silicon retina to provide realistic neuronal input (Elliott & Kramer 2002). Such applications have shown that this neurotrophic model robustly develops neuronal maps, even in the presence of high levels of noise in the afferent pathways.

Analysis shows that a critical parameter defining the behaviour of the model is the quantity $c = T_0/(aT_1)$ (Elliott & Shadbolt 1998a). In the absence of NTF diffusion between target cells, when $c < 1$ activity-dependent, competitive afferent segregation occurs in the model, i.e. all but one of s_{xi} for each target cell x go to zero, leaving just one afferent in control of cell x . When $c > 1$, afferent segregation breaks down: all afferents innervate all target cells equally. In the presence of NTF diffusion, the same results can be derived, but the critical threshold is reduced below unity. The point $c = 1$ can be shown to be a transition point at which the model changes from a classic Hebbian to a classic anti-Hebbian model of synaptic plasticity (Elliott & Shadbolt 2002b). Up to the factor a , the quantity c is just the ratio of the activity-independent to the activity-dependent release of NTF by target cells. Since activity-independent release can equally be thought of as exogenous infusion, $c > 1$ simply states that if exogenous infusion is too great relative to cellular release, then afferent segregation breaks down. Thus, the model is consistent with experimental data indicating that exogenous infusion or excess supply of NTFs does indeed lead to a breakdown of normal developmental processes in the visual system and at the neuromuscular junction (Cabelli *et al.* 1995; Nguyen *et al.* 1998).

3. Application to the Khepera Robot

We discuss a previous application of this neurotrophic model to the development of obstacle avoidance behaviour in the Khepera robot (see figure 1); for a fuller treatment, see Elliott & Shadbolt (2001). For this study, we use two Khepera robots in order to demonstrate the intrinsic variability between allegedly identical, off-the-shelf robots, and to show the virtues of developing a nervous system tuned to a particular robot’s morphology rather than using a generic nervous system.

The robots inhabit an A4-sized box whose walls are lined with either plain white paper or alternating black and white striped paper, the stripe width being 14 mm. A robot, connected to a controlling computer through its serial port via a rotating contact, is placed in this arena. Initially, the artificial nervous system of the robot is undeveloped, so that it cannot avoid crashing into the walls. The robot is equipped with three hard-wired, innate reflexes. The “withdrawal” reflex ensures that the robot stops when its total infra-red (IR) sensor input exceeds some threshold, turns randomly and then moves away. This reflex prevents physical damage to a robot by stopping it from actually hitting a wall and can be thought of



Figure 1. K-Team's Khepera robot.

as a pain avoidance reflex. The “boredom” reflex ensures that if the robot remains stationary for more than about 1 s, it turns randomly and moves away. Finally, the “explore” reflex ensures approximately uniform coverage of the arena by the robot by occasionally randomly changing the robot’s direction of motion when the robot is sufficiently distant from a wall. The robot has an innate disposition to move forwards at constant speed. The nervous system modifies this disposition by retarding the wheel speeds according to the IR sensor inputs. When the reflexes are activated, the output of the nervous system is suppressed so that only the reflexes control the robot.

A robot’s nervous system consists in two sets of neurons. One set of six neurons receives input from the six forward-pointing IR sensors (IRs) of the robot. For convenience, we refer to these neurons as the sensory neurons (SNs). We label these six IRs (and six SNs) according to K-Team’s conventions, from a robot’s left to right, as numbers 0 to 5. The synaptic projections from the IRs to the SNs are initially topographically disordered. The degree of disorder is measured by the parameter $b_t \in [0, 1]$, with $b_t = 0$ representing completely random initial topography and $b_t = 1$ representing perfect initial topographic order. Our neurotrophic model is applied to these synapses in order to develop refined topography, so that a robot’s sensory neuron layer learns the spatial relationships between its forward pointing IRs through the patterns of activation in these IRs. The second set of two neurons receives input from the six SNs and these neurons’ outputs are fed directly to the wheel motors to retard the forward motion of the robot, one neuron for each wheel motor. We refer to these neurons as the motor neurons (MNs). All six SNs initially project to both MNs, but a contralateral-ipsilateral distinction, commonly found in biological organisms, is assumed for these projections. The contralateral SNs project initially a few more synapses to an MN than the ipsilateral SNs (cf. Crair *et al.* 1998). This has the effect of tilting the competition, implemented through the use of the neurotrophic model on these synapses, in favour of the initially slightly dominating inputs. Hence, the contralateral SNs finally end up exclusively driving an MN. When sensory topography is refined, the net result is that activation of the IRs on, say, a robot’s left side retards the contralateral (right) wheel, causing the robot to turn away, thus mediating obstacle avoidance. For a full discussion of




Figure 2. IR response profiles for K1 and K2. The length of each line denotes the size of the response of the IR to the stimulus. The arcs around each IR indicate the maximum possible response. [From Elliott & Shadbolt (2001).]

	Plain	Striped
K1	0.0	0.1
K2	43.3	43.2

Table 1. *Khepera Crash Rates under Calibration*

(Crash rates for K1 and K2 in the plain and striped environments running a standard Braitenberg obstacle avoidance algorithm.)

the parameter selections, etc., in the application of our model to this system, see Elliott & Shadbolt (2001).

A complete run for each robot comprises three phases. First the program implementing the plasticity algorithm and reflexes is run for 10,000 iterations (one iteration being one complete update of the nervous system's synapses). Then we de-activate, in software, one or two of the robot's IRs. In this deprived state, we continue to run the plasticity algorithm for a further 5000 iterations in order to determine whether the robot can restore some of the performance lost due to the inactive IR(s). Finally, we restore the nervous system to its state immediately prior to IR deprivation, and run the robot for 5000 iterations while disabling all synaptic plasticity, in order to assess the robot's deprived performance without the possibility of any recovery. Crash rates (numbers of reflex withdrawals per 1000 iterations) for all three phases are determined when the crash rates stabilise.

Because the two Khepera robots (which we call K1 and K2 for convenience) differ quite markedly in their performances, we determine their sensory acuities by measuring the responses of the robots to a small tube of white paper moved at a fixed distance of 15 mm around each IR. Figure 2 show that compared to K1, K2 is rather myopic in its two front sensors, although K1's left rear sensor barely responds at all to the stimulus. These differences are reflected in the two robots' performances when calibrated using a standard Braitenberg obstacle avoidance algorithm (Braitenberg 1984), the results of which are shown in table 1. The myopic K2 crashes slightly over 43 times per 1000 iterations in both environments, while K1 hardly crashes at all. This difference entirely reflects the differences in IR responses of the two robots.

Figure 3 shows typical examples of the sensorimotor maps developed by each




Figure 3. Sensorimotor maps developed by K1 and K2 in the striped world with $b_t = 0.5$. Black circles denote IRs, grey SNs and white MNs; for clarity the left MN is shown on the right and the right MN on the left. The thickness of a line indicates the number of synapses projected. [From Elliott & Shadbolt (2001).]




Figure 4. Sensorimotor maps developed by K1 and K2 in the striped world with $b_t = 0.5$ following deprivation of one IR. [From Elliott & Shadbolt (2001).]

robot after 10,000 iterations in the striped environment with an intermediate value of $b_t = 0.5$. For K1, maps developed in the plain environment are qualitatively identical, but for K2, some systematic misprojections are observed, discussed more fully in Elliott and Shadbolt (2001). Both robots' sensory maps have developed refined topography, but K1's topography is better than K2's, with the latter's sensory map exhibiting some distortions. For both robots, contralateral SNs dominate MN input; with further iterations, the few remaining ipsilateral synapses would be eliminated. Figure 4 shows the impact of depriving IR number 1 on these maps.

Having shown examples of the developed sensorimotor maps for both robots, we can now discuss the quantitative aspects of our results by discussing the crash rates for the two robots in both environments during all three phases of simulation. Tables 2 and 3 summarise our data for both robots. During deprivation, IR sensor 1 is de-activated in all cases. During the undeprived phase, K1 performs similarly in both environments. Following deprivation without plasticity, K1 performs better in the plain world than in the striped. With continued plasticity, K1 recovers some performance in both environments, with the recovery in the plain world being better than that in the striped. Nevertheless, the difference between the crash rates in the plastic and implastic deprived phases for K1 is not statistically significant.

For K2, table 3 shows that its undeprived performance is actually better than its Braitenberg performance, with crash rates between a third and a half those seen during calibration. However, due to the myopia of this robot, its undeprived performance is worse than that of K1. Following deprivation without plasticity, the crash rates roughly quadruple. With plasticity, they are reduced to approxi-

b_t	Runs	\neg D	D \neg P	DP
Plain				
0.25	5	9.3	15.1	10.5
0.50	5	6.9	8.6	5.4
0.75	5	4.7	7.4	5.5
1.00	5	8.4	10.4	6.7
Striped				
0.25	8	10.2	18.6	11.1
0.50	5	5.6	13.6	9.0
0.75	5	6.5	13.8	6.9
1.00	5	6.9	11.4	9.7

Table 2. *Crash Rates for K1*

(Crash rates for K1 in the two environments. Column 1 indicates the value of b_t ; column 2 the number of runs considered, a “*” next to the number indicating a statistically significant difference (at the 1σ level) between the D \neg P and the DP data after some number of iterations); column 3 the crash rates for the non-deprived phase of learning (\neg D); column 4 the crash rates following deprivation but with plasticity switched off (D \neg P); column 5 the crash rates following deprivation but with continued plasticity (DP).

b_t	Runs	\neg D	D \neg P	DP
Plain				
0.25	5	20.0	65.0	44.6
0.50	5 *	15.2	68.9	38.3
0.75	5	19.5	62.5	47.4
1.00	5 *	11.5	66.5	31.2
Striped				
0.25	5 *	13.6	59.6	33.2
0.50	5 *	14.7	62.1	37.1
0.75	10 *	13.7	67.0	32.1
1.00	5 *	15.8	59.3	36.3

Table 3. *Crash Rates for K2*

(Crash rates for K2. The format of this table is identical to that for table 2.)

mately half those without plasticity. The difference between the deprived plastic and implastic regimes is almost always statistically significant. Even with one IR sensor knocked out, in the plastic phase K2 continues to perform better than its Braitenberg calibration performance.

A statistically significant, robust recovery of performance is observed following single IR sensor deprivation for K2, but not for K1. Because K1’s overall performance is so much better than K2’s, it may be expected that single receptor deprivation would not have much impact on it anyway. We test this idea by instead depriving pairs of adjacent IRs on K1. The results are shown in table 4. Depriving IRs 1 and 2 still produces no statistically significant recovery from deprivation, but depriving IRs 2 and 3 does produce a statistically significant degree of recovery.

Pair	Runs	$\neg D$	$D\neg P$	DP
Plain				
1-2	10	7.6	14.1	6.7
2-3	5 *	7.4	21.9	13.4
Striped				
1-2	20	6.9	13.0	9.2
2-3	5 *	6.3	33.5	16.0

Table 4. *Crash Rates for K1 with Double Deprivation*

(Crash rates for K1 for deprivation of the pair of IRs indicated in column 1. The format of this table is otherwise identical to that for table 2. All data in this table are generated with $b_t = 0.50$.)




Figure 5. K-Team's Koala robot with Videre Design's STH-V2 binocular vision head mounted on-board.

4. Application to the Koala Robot

We now discuss the application of our neurotrophic model to the development of visual maps in the Koala robot equipped with a binocular vision system (see figure 5). While the Khepera study discussed above reveals the role of a robot's idiosyncratic morphology and sensory acuity in map development, the Koala study emphasises the importance of environmental factors in map development.

The Koala robot occupies a well-illuminated office furnished with standard items such as a desk, tables, chairs and filing cabinets. As the Koala essentially provides only a mobile platform for the binocular vision head, the robot is set to move continuously in a circle of diameter approximately 1 m, thereby providing a constantly changing visual scene to the vision system. The robot is linked to a computer through its serial port via a rotating contact. Interlaced video images from the two cameras on the binocular vision head are sent down two spare wires in this tethering




Figure 6. The raw, “photoreceptor” images of instantaneously moving edges captured by the binocular vision system (A) and the result of processing the images to extract ON and OFF data (B).

cable and fed into a framegrabber mounted on the same computer. These images are de-interlaced and corrected for a systematic vertical misalignment in the two cameras. The resulting data consist in two 8-bit monochrome 160×120 images captured from the left and right cameras and are fed directly into a simulation of visual map development. These images are too large for real-time processing, so we extract either a 26×26 or a 16×16 subregion, depending on the visual map that we develop, from the centre of the left image and an identically sized region from the right image. The right subimage is usually the central region, but it can also be subjected to a horizontal offset of 20 pixels with respect to the left subimage. These data can be regarded as the raw, “photoreceptor” images. We also consider data corresponding approximately to the ON and OFF outputs of retinal ganglion cells by threshold subtracting consecutive frames. That is, consecutive pixel intensities are subtracted, and only when the absolute difference exceeds some threshold is the pixel’s intensity unchanged; otherwise it is set to zero. With a threshold of 32, we find that such a procedure generates data in which only moving edges survive processing. Figure 6 shows the result of such processing on a set of moving edges.

Processing the raw data to generate ON and OFF data modifies the autocorrelation function of an image with itself and also modifies the cross correlation function between left and right images. In order to understand some of our results on visual map development, it is necessary to understand how these correlation functions change in response to ON/OFF processing. Figure 7 shows the average cross correlation function between the left and right images for raw data and ON/OFF data in the visual world inhabited by the Koala. Similar, figure 8 shows the average autocorrelation function of the left image for raw and ON/OFF data. Although the cross correlation functions for raw and ON/OFF data are broadly similar in shape, we see that the cross correlation function for ON/OFF data has a smaller spatial extent and considerably lower maximum. The maxima of both cross correlation functions occur at a horizontal offset of approximately 20 pixels, hence our selection of a 20 pixel offset when we offset the right subimage with respect to the left subimage. The autocorrelation functions are also broadly similar for both data sets, although for raw data, the function is roughly symmetric with respect to either vertical or horizontal offsets, while that for ON/OFF data exhibits a considerable bias towards




Figure 7. The average cross correlation function for raw (left) and ON/OFF (right) data. A key is shown on the extreme right. The number associated with each box in the key indicates the value of the function on the outer boundary of the enclosed region.




Figure 8. The average autocorrelation function for raw (left) and ON/OFF (right) data. A key is shown on the extreme right.

the vertical. This bias reflects the horizontal motion of the robot, so that captured moving edges tend to be vertically rather than horizontally orientated.

We consider the development of two types of visual map. The first is the development of a topographic representation of one, say the left, retinal array on a sheet of target cells. These target cells may represent the optic tectum in lower vertebrates, or the lateral geniculate nucleus (LGN) in higher vertebrates. Such a process is analogous to the refinement of the topographic representation of the IR cells on the SN cell layer in the above Khepera work, except that the representations develop over a two-dimensional rather than a one-dimensional array of cells. For reasons of computational speed only, we consider a 16×16 array of pixels from the left camera mapping onto the same sized array of target cells. The topographic bias parameter b_t , defined above for the Khepera maps, is set to $b_t = 0.5$. Our results do not exhibit much dependence on b_t unless b_t is close to zero, representing initially almost completely random topography. To visualise the representation of the retinal sheet on the target sheet, we calculate the centre of mass of retinal

projections to each target cell,

$$\vec{M}_{\vec{x}} = \frac{\sum_{\vec{r}} \vec{r} s_{\vec{x}\vec{r}}}{\sum_{\vec{r}} s_{\vec{x}\vec{r}}}, \quad (4.1)$$

where the vector character of the indices has been made explicit. The positions defined by the vectors $\vec{M}_{\vec{x}}$ and $\vec{M}_{\vec{y}}$ are then connected by a line if, and only if, cells \vec{x} and \vec{y} are nearest neighbours on the target sheet. If topography is perfect, then the resulting pattern of lines forms a regular, square grid; increasing deviation away from this pattern represents increasing disruption of the topographic projection.

The second type of visual map that we consider is the development of ocular dominance columns (ODCs) in the primary visual cortex of higher vertebrates (Hubel & Wiesel 1962). ODCs are formed through a competitive process in which LGN afferent cells representing the left and right eyes initially innervate the visual cortex roughly uniformly, so that cortical cells are driven nearly equally strongly by both eyes. As development proceeds, the afferents segregate at the anatomical (and physiological) level into a mosaic of interdigitated, alternating regions of control approximately 500 microns wide, with one eye controlling any given region and the other eye controlling immediately adjacent regions (LeVay *et al.* 1978, 1980). Experimental evidence suggests that neurotrophic factors may be involved in the formation of ODCs (reviewed in McAllister *et al.* 1999). To model the formation of ODCs, we take two 26×26 arrays of cells, one representing the left camera pixels and the other the right camera pixels, possibly with a 20 pixel horizontal offset for the right subimage. Such an offset coarsely resembles convergent strabismus, because the cross correlations increase. These two arrays innervate a 26×26 array of cells representing the primary visual cortex. The projections from each afferent sheet to the target sheet are established as above for topographic map development, so that each target cells is initially controlled roughly equally by both afferent sheets. However, for reasons of computational tractability, we restrict the spread of each afferent cell's connections on the cortex to a 5×5 patch of topographically appropriate cortex. Thus, we consider only the development of ODCs through an activity-dependent, competitive process and not the simultaneous co-development of ODCs and topography.

In figure 9 we show the development of a topographic projection from the sensory, retinal sheet to the target sheet for raw, "photoreceptor" images. Initially, the topographic map is tightly folded because most target cells' inputs are skewed towards the central retinal field (cf. Goodhill 1993). As development proceeds, the map unfolds so that the final map is almost perfect, except for the presence of edge effects. Because the autocorrelation function is approximately symmetric, the receptive fields of the target cells remain approximately symmetric throughout development, exhibiting no systematic bias towards the horizontal or vertical orientations. However, in figure 10, which shows the development of a topographic map in the presence of ON/OFF data, we see that the map is initially stretched in the horizontal direction. An examination of the receptive field of a representative target cell, shown in figure 11, reveals that it is strongly biased towards vertical orientations during early development. This biasing is a consequence of the autocorrelation function for ON/OFF data and is therefore ultimately a function of the robot's purely horizontal motion, which accentuates vertical rather than horizontal structure in the environment. Despite this early biasing towards the vertical, maps

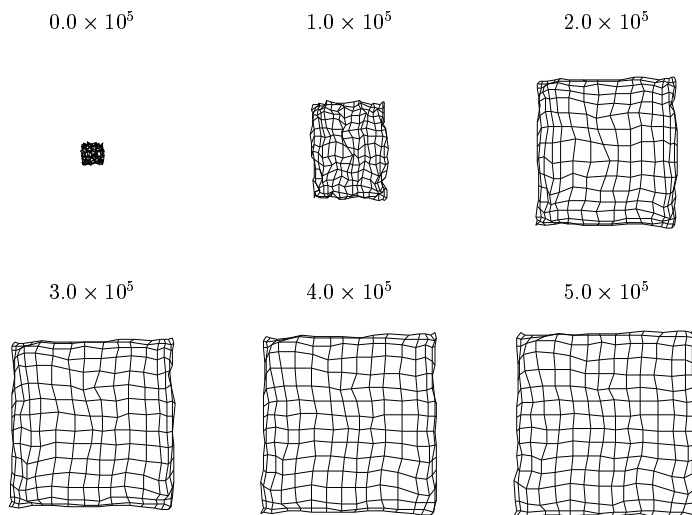


Figure 9. The development of topography for raw, “photoreceptor” data. Each map represents the state of topography at the number of iterations indicated immediately above it.

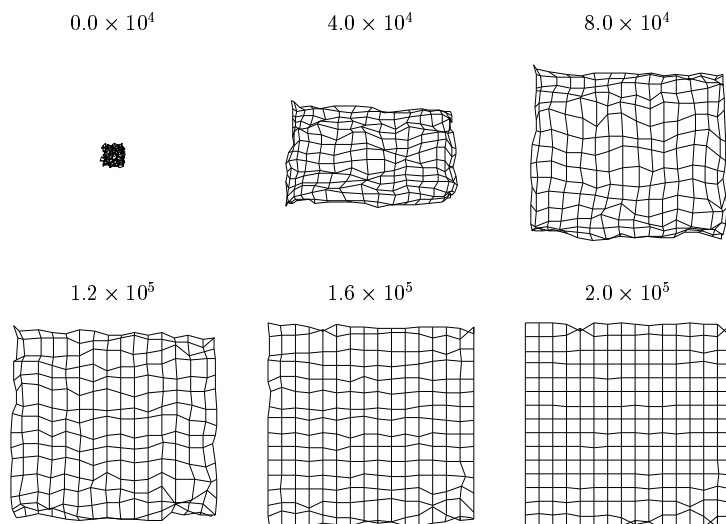


Figure 10. The development of topography for ON/OFF data. The format of this figure is otherwise identical to that for figure 9.

for ON/OFF data develop and mature much more rapidly than those for raw data, so much so that even edge effects are very nearly removed from the final maps.

In figures 12 and 13 we show the final patterns of ocular dominance produced with raw and ON/OFF data, respectively, for either a 0 or a 20 pixel offset of the right subimage with respect to the left subimage. In all cases, we see well segregated patterns of ocular dominance. For raw data, figure 12, we see that a 20 pixel offset,




Figure 11. The development of the receptive field of a target cell near the centre of the sheet shown in figure 10. Each square represents an afferent cell with its shade of grey indicating the number of synapses projected to the target cell. A key is given above the maps.

which induces higher inter-ocular correlations between the left and right images, results in a greater degree of binocularity in the mature maps. Such a result is consistent with developmental work in animals (Hubel & Wiesel 1965; Shatz *et al.* 1977). For ON/OFF data, the 20 pixel offset induces a smaller increase in the inter-ocular correlations, and the resulting map exhibits only slightly greater degrees of remaining binocularity than the 0 pixel offset data. Raw data runs require approximately 250,000 iterations to reach stable, mature maps, while ON/OFF data runs require only 50,000 iterations. As with the development of topography, preprocessing the raw, “photoreceptor” data to generate ON and OFF data therefore dramatically reduces the time required to develop mature visual maps.

5. Discussion

The first application of the neuronal plasticity model to the Khepera robot shows the extent to which we can accommodate to variations in morphology and sensory acuity between robots. It illustrates the ability of our models to develop neurocontrollers tuned to the particular embodiments they encounter.

Using the neurotrophic model, both Khepera robots develop sensorimotor maps that mediate obstacle avoidance behaviour. When we look at K1’s Braitenberg performance, it is near perfect, so it is not surprising that its self-grown map performance lags behind slightly. Nevertheless, K1’s crash rate using a self-grown map can be as low as 4.7 reflex withdrawals per 1000 iterations, and thus is actually pretty good. K2’s self-grown map performance is much better than its Braitenberg performance, with crash rates less than half those seen during calibration. Even




Figure 12. Ocular dominance maps generated by raw, “photoreceptor” data for two different offsets, indicated above the maps. Each square represents a target cell with its shade of grey indicating the degree of control by the left camera. A black (white) square is completely dominated by the left (right) camera. A key is given above the maps.

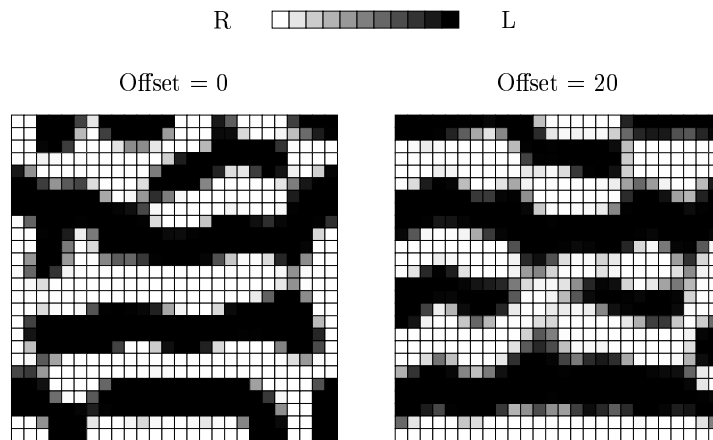


Figure 13. Ocular dominance maps generated for ON/OFF data for two different offsets. The format of this figure is otherwise identical to that for figure 12.

when one IR sensor is knocked out and recovery is permitted, K2’s performance remains better than Braitenberg. For both robots, therefore, self-grown maps perform near or better than a standard Braitenberg algorithm. Why is this? As we argued at the outset, allowing a system to develop its own nervous system using competitive interactions allows it to develop maps tuned to its own particular morphological idiosyncrasies. In the case of K2, its two front IRs are rather myopic, and therefore typically less active than the other forward pointing IRs. Such reduced activity causes these IRs to be disadvantaged during map development, and inputs from neighbouring IRs will compensate. The standard Braitenberg algorithm can

make no allowance for these deficits, and since its performance depends heavily on input from the two front IRs, it is inevitable that it will perform poorly in a myopic robot. Activity-dependent, competitive processes therefore tune a nervous system in a robot-specific fashion without requiring any prior knowledge about the robot in which those processes are instantiated.

Continuing plasticity permits continuing adaptation to ongoing environmental or bodily changes. In particular, when we knock out an IR sensor, the nervous system adjusts by strengthening connections from adjacent IRs. This has the effect of essentially interpolating a missing IR sensor's input from its neighbours' inputs. For both robots, deprivation of one IR followed by recovery restores some degree of lost performance. For K1, this recovery is not statistically significant, but it is noteworthy that the recovered performance in the plain environment matches its undeprived performance. For K2, the level of recovery is statistically significant although it does not match pre-deprivation levels. Despite this, its deprived performance still exceeds Braitenberg performance with all IRs functioning. The failure of K1 to exhibit a significant recovery from single IR deprivation is a consequence of the overall superior performance of this robot. Double sensor deprivation does allow us to see a statistically significant degree of recovery. It is remarkable that even when two IRs are knocked out in K1, whether or not recovery is permitted, its performance exceeds that of K2 running a Braitenberg algorithm without deprivation.

In addition to exhibiting a dependence on a robot's morphology, our results also exhibit a dependence on the particular environment inhabited by a robot, whether plain or striped. K1's undeprived performances are comparable in both worlds, but with deprivation its performance is better in the plain world. This is clearer during double IR deprivation than single IR deprivation. This difference is a consequence of the similarity between the stripe spacing (14 mm) and the IR sensor spacing for the Kheperas (13 mm). In the striped world, when an IR senses a white stripe, its neighbours sense black stripes. Under single IR deprivation, when the deprived IR is near a white stripe, its neighbours cannot interpolate its input because they are near black stripes. But when the deprived IR is near a black stripe, its neighbours are near white stripes and in such a case it is essentially irrelevant that the IR is knocked out, because its output would have been low anyway. Under double IR deprivation, however, one of the two adjacent, active IRs is always near a black stripe, thus always reducing the effectiveness of the interpolation. Hence, K1's deprived performance in the striped world is always worse than that in the plain world, and its double IR deprivation performance is worse than its single IR deprivation performance.

The second application of our neuronal plasticity model to the Koala robot emphasises the importance of environmental factors in map development and shows the extent to which the model can accommodate to variations in sensory input. We have found that both refined retinotopic projections and ODCs emerge robustly under a range of environmental conditions and parameter regimes. We have also shown how the statistics of the visual input play an important role in determining both the nature and rate of development of these structures.

In the Koala experiments, the ON/OFF data attempt to capture something of the temporal processing that occurs in the visual pathway, in particular, how retinal ganglion cells respond to the dynamics of intensity change rather than sim-

ply encoding a raw, “photoreceptor” intensity level. The effect of ON/OFF data extraction is dramatically to reduce the degree of correlation between images compared to the raw, “photoreceptor” data. ON/OFF data also affects the intra-ocular or autocorrelation function. The autocorrelation function is much more narrowly focused than that for raw, “photoreceptor” data and strongly emphasises vertical structure.

Raw and ON/OFF data provide very different statistical characterisations of the visual world. These have a direct impact on our results. Topographic maps and refined receptive fields form more rapidly for ON/OFF data and the receptive fields display a developmental profile that reflects the strong, vertical correlations present in the ON/OFF data. These correlations arise due to the horizontal motion of the robot, so that captured moving edges tend to be vertically rather than horizontally orientated. The weaker horizontal correlations cause the receptive fields to refine more quickly in the horizontal direction. We would predict that providing motion in the vertical plane would lead to more rapidly refined vertical components of the developing receptive fields. Two other effects of the use of ON/OFF data were observed: a greater remaining binocular control of cells at the boundaries of ODCs and much faster ODC segregation.

Researchers have not previously implemented models of ODC and topographic development on real robots with anything like semi-realistic, naturalistic visual input. Our results from the Koala demonstrate that it is important to consider carefully the likely visual statistics that drive neuronal development, and also to consider the sorts of environment in which systems develop.

We have shown that a biologically-inspired model of neuronal development can be transferred from computer simulation to robotic instantiation. In the first application we have shown how the model allows the robot to develop sensorimotor maps tuned to the particular characteristics of the robot. In the second application we demonstrate that the model allows a robot to develop well ordered neuronal maps that correspond to structures found in the visual cortex of higher vertebrates. We have shown how the nature of the visual experience determines important characteristics of these maps.

6. Conclusions and Outlook

We hold this truth to be self-evident: that despite K-Team’s best intentions, no two Khepera or Koala robots are created equal. In light of the stark differences between the sensory acuities of our two Khepera robots, attempts to implement a one-size-fits-all nervous system across a range of such robots, even for such a deceptively simple behaviour as obstacle avoidance, without taking into account their dramatically different capacities, seem doomed to failure. Inter-individual variation is a fact of life with which robotics must come to terms. Precision engineering may seek to reduce this variation to a point at which it can be ignored, but in practice this is infeasible and may for a variety of reasons be undesirable (Smithers 1994).

No two robots experience the same environmental inputs. This may arise from immersion in different environments, variations in acuity or because they are running different early sensory processing algorithms. We need mechanisms that allow neurocontrollers to adapt to their particular morphologies, environmental inputs, contexts and histories. Furthermore, unpredictable parts failure requires either the

discarding of the robot or the adaptation of its neurocontroller to recover acceptable performance. Although throwing away a robot in some circumstances may be an option, in others it may not be, such as a robot explorer on Mars or one operating in an extreme environment. Besides, the mere throwing away of a robot is an “engineering” solution: it does not address the scientific question of how, in principle, do we build organisms that are robust, adaptable and can recover from minor accidents or injury?

This is a challenge that biology faced from the beginning and to which it had to find a solution. Data indicate that its solution does not consist in a mere bag of contingent tricks, one trick for one problem, but may be a universal solution that defines the very manner in which an organism is built and responds to change. Given that biology is a notorious tinkerer, this is perhaps surprising. But when it is recalled that the very first organisms were immediately faced with living in the world, of adapting to it and recovering from damage, perhaps it is not so surprising after all: it is one problem, not a series of problems faced incrementally over time as organisms became more complex. A unitary problem demands a unitary solution. In light of the possible existence of such a solution, it appears foolish to ignore its potential for robotics.

Does the putting into action of our manifesto discussed above support this view? We have only considered coarse grained models running on basic robots inhabiting simple worlds. Nevertheless, we have shown that a biologically-inspired, neurotrophic model facilitates in one case the development of a nervous system that is tuned to the morphological idiosyncrasies of the particular robot in which it is instantiated, and in the other produces well organised maps between various components of a visual system under a wide range of environmental inputs. In the first case the Khepera’s sensorimotor maps perform nearly as well as or much better than what may have been considered to be a roughly optimal algorithm for obstacle avoidance. When the robot has developed a mature, stable map, the latent plasticity inherent in its synapses allows the robot to be ready to accommodate change. In the second application presented here the Koala using the same basic model can accommodate to a wide range of changes and variations in its morphology, experience and environment. These include changing the inter-ocular spacing of cameras, ablating cells in either the afferent or cortical sheets, and varying the extent and nature of temporal preprocessing in the robotic system.

The benefits of such developmental accommodation, of competitive synaptic growth and re-arrangement are very substantial. They permit a nervous system to tune itself to the body in which it finds itself and the environment in which the body resides. The benefit of maintaining plasticity throughout an animal’s lifespan is that it allows the nervous system to adapt to changes that occur throughout a lifetime. Our fundamental equality may reside in the common mechanism of plasticity, a mechanism that can give rise to endless differences and inequalities.

TE thanks the Royal Society for the support of a University Research Fellowship.

References

- Blöchl, A. & Thoenen, H. 1996 Localization of cellular storage compartments and sites of constitutive and activity-dependent release of nerve growth factor

- (NGF) in primary cultures of hippocampal neurons. *Mol. Cell. Neurosci.* **7**, 173–190.
- Braitenberg, V. 1984 *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA: M.I.T. Press.
- Cabelli, R., Hohn, A. & Shatz, C. 1995 Inhibition of ocular dominance column formation by infusion of NT-4/5 or BDNF. *Science* **267**, 1662–1666.
- Campenot, R. 1982a Development of sympathetic neurons in compartmentalized cultures. I. Local control of neurite outgrowth by nerve growth factor. *Dev. Biol.* **93**, 1–12.
- Campenot, R. 1982b Development of sympathetic neurons in compartmentalized cultures. II. Local control of neurite survival by nerve growth factor. *Dev. Biol.* **93**, 13–22.
- Castren, E., Zafra, F., Thoenen, H. & Lindholm, D. 1992 Light regulates expression of brain-derived neurotrophic factor mRNA in rat visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* **89**, 9444–9448.
- Crair, M., Gillespie, D. & Stryker, M. 1998 The role of visual experience in the development of columns in cat striate cortex. *Science* **279**, 566–570.
- Edelman, G. 1993 *Topobiology*. New York: Basic Books.
- Elliott, T. 2002 A spin glass-like Lyapunov function for a neurotrophic model of neuronal development. *Biol. Cybern.* **86**, 473–481.
- Elliott, T. & Kramer, J. 2002 Coupling an aVLSI neuromorphic vision chip to a neurotrophic model of synaptic plasticity: The development of topography. *Neural Comp.* **14**, 2353–2370.
- Elliott, T., Maddison, A. & Shadbolt, N. 2001 Competitive anatomical and physiological plasticity: A neurotrophic bridge. *Biol. Cybern.* **84**, 13–22.
- Elliott, T. & Shadbolt, N. 1998a Competition for neurotrophic factors: Mathematical analysis. *Neural Comp.* **10**, 1939–1981.
- Elliott, T. & Shadbolt, N. 1998b Competition for neurotrophic factors: Ocular dominance columns. *J. Neurosci.* **18**, 5850–5858.
- Elliott, T. & Shadbolt, N. 1999 A neurotrophic model of the development of the retinogeniculocortical pathway induced by spontaneous retinal waves. *J. Neurosci.* **19**, 7951–7970.
- Elliott, T. & Shadbolt, N. 2001 Growth and repair: Instantiating a biologically-inspired model of neuronal development on the Khepera robot. *Robot. Auton. Syst.* **36**, 149–169.
- Elliott, T. & Shadbolt, N. 2002a Dissociating ocular dominance column development and ocular dominance plasticity: A neurotrophic model. *Biol. Cybern.* **86**, 281–292.

- Elliott, T. & Shadbolt, N. 2002b Multiplicative synaptic normalisation and a non-linear Hebb rule underlie a neurotrophic model of competitive synaptic plasticity. *Neural Comp.* **14**, 1311-1322.
- Gilbert, S. 2000 *Developmental Biology*. Sunderland, MA: Sinauer.
- Goodhill, G.J. 1993 Topography and ocular dominance: a model exploring positive correlations. *Biol. Cybern.* **69**, 109-118.
- Goodman, L., Valverde, J., Lim, F., Geschwind, M., Federoff, H., Geller, A. & Hefti, F. 1996 Regulated release and polarized localization of brain-derived neurotrophic factor in hippocampal neurons. *Mol. Cell. Neurosci.* **7**, 222-238.
- Hubel, D.H. & Wiesel, T.N. 1962 Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106-154.
- Hubel, D.H. & Wiesel, T.N. 1965 Binocular interaction in striate cortex of kittens reared with artificial squint. *J. Neurophysiol.* **28**, 1041-1059.
- Koliatsos, V., Clatterbuck, R., Winslow, J., Cayouette, M. & Price, D. 1993 Evidence that brain-derived neurotrophic factor is a trophic factor for motor neurons *in vivo*. *Neuron* **10**, 359-367.
- LeVay, S., Stryker, M.P. & Shatz, C.J. 1978 Ocular dominance columns and their development in layer IV of the cat's visual cortex: A quantitative study. *J. Comp. Neurol.* **179**, 223-244.
- LeVay, S., Wiesel, T.N. & Hubel, D.H. 1980 The development of ocular dominance columns in normal and visually deprived monkeys. *J. Comp. Neurol.* **191**, 1-51.
- McAllister, A., Katz, L. & Lo, D. 1999 Neurotrophins and synaptic plasticity. *Annu. Rev. Neurosci.* **22**, 295-318.
- Nguyen, Q., Parsadanian, A., Snider, W. & Lichtman, J. 1998 Hyperinnervation of neuromuscular junctions caused by GDNF overexpression in muscle. *Science* **279**, 1725-1729.
- Purves, D. 1988 *Body and Brain: A Trophic Theory of Neural Connections*. Cambridge, MA: Harvard University Press.
- Purves, D. 1994 *Neural Activity and the Growth of the Brain*. Cambridge, UK: Cambridge University Press.
- Purves, D. & Lichtman, J. 1985 *Principles of Neural Development*. Sunderland, MA: Sinauer.
- Shatz, C.J., Lindström, S. & Wiesel, T.N. 1977. The distribution of afferents representing the right and left eyes in the cat's visual cortex. *Brain Res.* **131**, 103-116.

- Smithers, T. 1994 On why better robots make it harder. In *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour* (eds. D. Cliff, P. Husbands, J.-A. Meyer & S. Wilson), pp. 64–72. Cambridge, MA: M.I.T. Press.
- Truman, J. 1984 Cell death in invertebrate nervous systems. *Annu. Rev. Neurosci.* **7**, 171–188.