

Chapter 7

Creating and Controlling Complex Biological Brains

Kevin Warwick

Abstract In this contribution, a look is taken at how animal and/or human brain cells can be cultivated (grown) and given a robot physical body (as a controlling brain) in which they can move around and interact with the world. This is realised as a new form of Artificial Intelligence in which the complexity of a highly non-linear biological neural network is employed to uniquely control a real-world robot. The communication/control feedback loop is described and considered in terms of learning, performance, long-term operation and specialisation within the neural structure. Experimental results are presented and philosophical arguments opened up, e.g. can the robot be considered to be a living, conscious entity?

7.1 Introduction

Considerable progress has been made towards hybrid systems in which biological neurons are integrated with electronic components. As an example, Reger et al. [1] demonstrated the use of a lamprey brain to control a small-wheeled robot's movements; meanwhile others were successfully able to send control commands to the nervous system of cockroaches [2] or rats [3] as if they were robots. These studies inform us about information processing and encoding in the brains of living animals [4], however, they do pose ethical questions and can be technically problematic since access to the brain is limited by barriers such as the skin and skull, and data interpretation is complicated due to the number of neurons present in the brain of even the simplest animal. Also, approaches which involve recording the activity of individual neurons or small populations of neurons are limited by their invasive, and hence destructive, nature. As a result, neurons cultured under labo-

K. Warwick (✉)

School of Systems Engineering, University of Reading, Whiteknights, PO Box 225, Reading RG6 6AY, UK

e-mail: k.warwick@reading.ac.uk

ratory conditions on a planar array of non-invasive electrodes provide an attractive alternative with which to probe the operation of biological neuronal networks.

In the past few years, research has focussed on culturing networks of some tens of thousands of brain cells grown *in vitro* [5]. These cultures are created by enzymatically dissociating neurons obtained from foetal rodent cortical tissue and then culturing them in a specialised chamber, in doing so providing suitable environmental conditions and nutrients. An array of electrodes is embedded in the base of the chamber (a multielectrode array; MEA) providing an electrical interface to the neuronal culture [6–9]. The neurons in such cultures spontaneously begin to branch out and within an hour of placement, even without external stimulation, they begin to reconnect with other nearby neurons and commence electrochemical communication. This propensity to spontaneously connect and communicate demonstrates an innate tendency to network. Studies of neural cultures demonstrate distinct periods of development defined by changes in activity which appear to stabilise after 30 days and, in terms of useful responses, last for at least 2–3 months [10–12]. The cultures of neurons form a monolayer on the MEA, making them both amenable to optical microscopy and accessible to physical and chemical manipulation [9].

An aim of the ongoing project described here is to investigate the use of cultured neurons for the control of mobile robots. In order to produce useful processing, the disembodied biological network must develop in the presence of meaningful input/output relationships as part of closed-loop sensory interaction with the environment. This is evidenced by animal and human studies, which show that development in a sensory-deprived environment results in poor or dysfunctional neural circuitry [13, 14]. To this end the overall closed-loop hybrid system involving a primary cortical culture on an MEA and a mobile robot body must exist within a sufficiently rich and reasonably consistent environment. This then constitutes an interesting and novel approach to examining the computational capabilities of biological networks [15].

Typically, *in vitro* neuronal cultures consist of thousands of neurons generating highly variable, multidimensional signals. In order to extract components and features representative of the network's overall state from such data, appropriate pre-processing and dimensionality reduction techniques must be applied. Several schemes have till now been constructed. Shkolnik et al. [16] created a control scheme for a simulated robot body in which two channels of an MEA were selected and an electrical stimulus consisting of a ± 600 mV, 400 μ s biphasic pulse was delivered at varying inter-stimulus intervals. Information coding was formed by testing the effect of electrically induced neuronal excitation with a given time delay termed the inter-probe interval (IPI) between two stimulus probes. This technique gave rise to a characteristic response curve which formed the basis for deciding the robot's direction of movement using basic commands (forward, backward, left and right).

DeMarse and Dockendorf investigated the computational capacity of cultured networks by implementing the control of a 'real-life' problem, namely controlling a simulated aircraft's flight path (e.g. altitude and roll adjustments) [17]. Meanwhile

Shahaf and Marom [18] reported one of the first experiments to achieve desired discrete output computations by applying a simple form of supervised learning to disembodied neuronal cultures. Recently, in [19] a Learning classifier system was used to manipulate culture activity towards a goal level using simple input signals. In both of these latter experiments, the desired result was only achieved in about one-third of the cases, indicating some of the difficulties in achieving repeatability.

This should not really come as a surprise as this is a field of study very much in its infancy. There are bound to be difficulties, however, there is much to be learnt. It is apparent that even at this early stage, such re-embodiments (real or virtual) have an important role to play in the study of biological learning mechanisms and neurological behaviour in general. Our physical embodied robots provide the starting point for creating a proof-of-concept control loop around the neuronal culture and a basic platform for future—more specific—reinforcement learning experiments. The fundamental problem is the coupling of the robot's goals to the culture's input–output mapping, the design of the robot's architecture discussed in this paper therefore emphasises the need for flexibility and the use of machine learning techniques in the search of such coupling.

In the section which follows, the general procedure for laying out the neural culture (the biological component) is described; this is followed by a description of the main elements of the closed-loop control system, including the culture as an important element in the feedback loop. Details of the current system's architecture are given in Sect. 7.3. Section 7.4 includes a description of our initial tests and preliminary results. Section 7.5 meanwhile provides a discussion of possible biological underpinnings providing the motivation and context for the use of the machine learning (ML) methods, and Sect. 7.6 concludes with an overview of current progress. Finally, Sect. 7.7 discusses new ongoing research and planned future extensions.

7.2 Culture Preparation

To realise the cultured neural network, cortical tissue is dissected from the brains of embryonic rats and neuronal cells enzymatically dissociated before seeding onto planar multielectrode arrays (MEAs). The cells are restricted to lie within the recording horizon of the electrode array by means of a template placed on the MEA prior to seeding and removed immediately after cells have settled (~ 1 h). The MEA is also filled with a conventional cell culture medium containing nutrients, growth hormones and antibiotics of which 50 % is replaced twice weekly. Within the first hour after seeding, neurons appear to extend connections to nearby cells (even within the first few minutes this has been observed) and within 24 h, a thick matt of neuronal extensions is visible across the seeded area.

The connectivity between seeded cells increases rapidly over subsequent days. After 7 days, electrical signals are observed in the form of action potentials which, in the 'disembodied culture' (not connected within the closed loop), over the

following week transform into dense bursts of simultaneous electrical activity across the entire network. This bursting feature subsequently continues through to maturity (30 days *in vitro* and onwards). It is not well understood what the bursting actually means and how much it is part of normal neural development. However, such continued behaviour, after this initial development phase, may subsequently be representative of an underlying pathological state resulting from impoverished sensory input and may differ from the activity of a culture developing within a closed loop [20]. This is something which remains to be studied further.

Cultures usually remain active until approximately 3 months of age. During this time, they are sealed with Potter rings [21] to maintain sterility and osmolarity, and are maintained in a humidified, 37 °C, 5 % CO₂ incubator. Recordings are undertaken in a non-humidified 37 °C, 5 % CO₂ incubator for between 30 min and 8 h dependent on environmental humidity and the resulting stability of activity.

7.3 Experimental Platform

The multielectrode array enables voltage fluctuations in the culture (relative to a reference ground electrode outside the network) to be recorded in real-time at 59 sites out of 64 in an '8 × 8' array (Fig. 7.1). This allows for the detection of neuronal action potentials within a 100 µm radius (or more) around an individual electrode. Using spike sorting algorithms [12], it is then possible to separate the firings of multiple individual neurons, or small groups of neurons, as monitored on a single electrode. As a result, multielectrode recordings across the culture permit a picture of the global activity of the entire neuronal network to be formed. It is possible to electrically stimulate via any of the electrodes to induce focussed neural activity. The multielectrode array therefore forms a functional and non-destructive bidirectional interface to the cultured neurons.

Electrically evoked responses and spontaneous activity in the culture (the neuronal network) are coupled to the robot architecture, and thence on to the physical robot, via a machine learning interface, which maps the features of interest to specific actuator commands. Sensory data fed back from the robot is associated with a set of appropriate stimulation protocols and is subsequently delivered to the culture, thereby closing the robot-culture loop. Thus, signal processing can be broken down into two discrete sections (a) 'culture to robot', in which an output machine learning procedure processes live neuronal activity, and (b) 'robot to culture', which involves an input mapping process, from robot sensor to stimulus.

It is important to realise that the overall system employed in this experiment has been designed based on a closed-loop, modular architecture. As neuronal networks exhibit spatiotemporal patterns with millisecond precision [22], processing of these signals necessitates a very rapid response from neurophysiological recording and robot control systems. The software developed for this project runs on Linux-based workstations communicating over the Ethernet via fast server-client modules, thus providing the necessary speed and flexibility required.

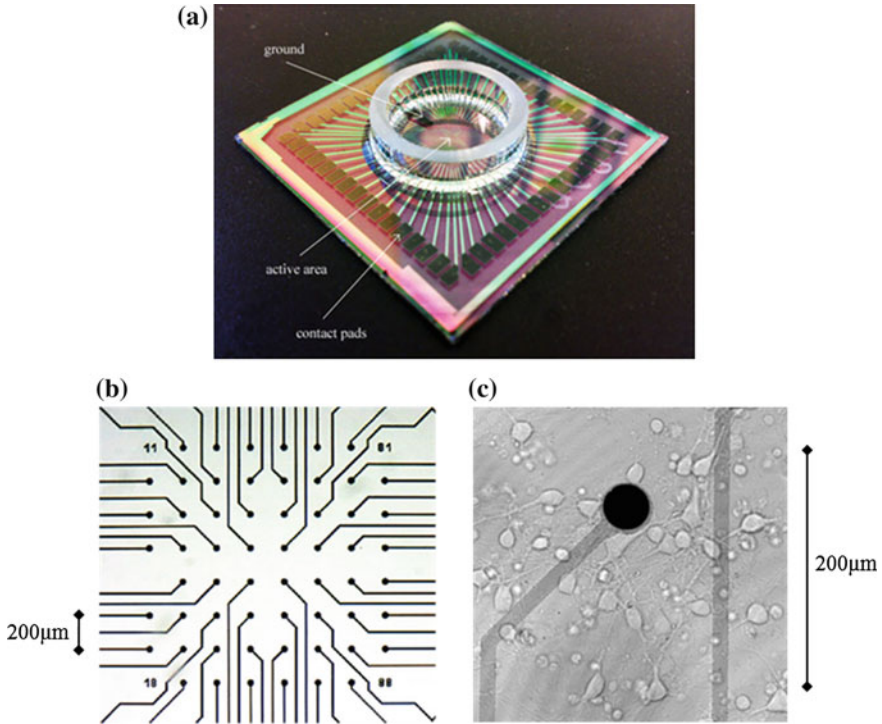



Fig. 7.1 **a** An MEA, showing the 30 µm electrodes which lead to the electrode column–row arrangement, **b** Electrode arrays in the centre of the MEA seen under an optical microscope, **c** An MEA at $\times 40$ magnification, showing neuronal cells in close proximity to an electrode, with visible extensions and interconnections

In recent years, the study of neuronal cultures has been greatly facilitated by commercially available planar MEA systems. These consist of a glass specimen chamber lined with an 8×8 array of electrodes as shown in Fig. 7.1. It is just such one of these MEAs that we have employed in our overall robot system.

A standard MEA (Fig. 7.1a) measures $49 \times 49 \times 1$ mm and its electrodes provide a bidirectional link between the culture and the rest of the system. The associated data acquisition hardware includes a head-stage (MEA connecting interface), 60 channel amplifier ($1200 \times$ gain; 10–3200 Hz band-pass filter), stimulus generator and PC data acquisition card.

To this point, we have successfully created a modular closed-loop system between a (physical) mobile robotic platform and a cultured neuronal network using a multielectrode array, allowing for bidirectional communication between the culture and the robot. It is estimated that the cultures employed in our studies consist of approximately (on average) 100,000 neurons. The actual number in any one specific culture depends on natural density variations in proliferation post-seeding and experimental aim. The spontaneous electrochemical activity of the culture

Fig. 7.2 The Miabot robot with a cultured neural network



realising signals at certain of the electrodes is used as input to the robot's actuators and the robot's (ultrasonic) sensor readings are (proportionally) converted into stimulation signals received by the culture, effectively closing the loop.

We are using a versatile, commercially available, Miabot robot (Fig. 7.2) as our physical platform. This exhibits accurate motor encoder precision (~ 0.5 mm) and has a maximum speed of approximately 3.5 m/s. Hence it can move around quite quickly in real time. Recording and stimulation hardware is controlled via open-source MEABench software [23].

The overall closed-loop system therefore consists of several modules including the Miabot robot, an MEA and stimulating hardware, a directly linked workstation for conducting computationally expensive neuronal data analyses and a separate machine running the robot control interface; a network manager routing signals directly between the culture and the robot body.

The Miabot is wirelessly controlled via Bluetooth. Communication and control are performed through custom C++ server code and TCP/IP sockets and clients running on the acquisition PC which has direct control of the MEA recording and stimulating software. The server sends motor commands and receives sensory data via a virtual serial port over the Bluetooth connection, while the client programs contain the closed-loop code which communicates with and stimulates the MEA

culture. The client code also performs text logging of all important data during an experiment run.

This modular approach to the architecture has resulted in a system with easily reconfigurable components. The obtained closed-loop system can efficiently handle the information-rich data that is streamed via the recording software. A typical sampling frequency of 25 kHz of the culture activity demands large network, processing and storage resources. Consequently, on-the-fly streaming of spike-detected data is the preferred method when investigating real-time closed-loop learning techniques.

7.4 Experimental Results

First an existing appropriate neuronal pathway was identified by searching for strong input/output relationships between pairs of electrodes. Suitable input/output pairs were defined as those electrode combinations in which neurons proximal to one electrode responded to stimulation of the other electrode at which the stimulus was applied (at least one action potential within 100 ms of stimulation) more than 60 % of the time and responded no more than 20 % of the time to stimulation on any other electrode. An input–output response map was then created by cycling through all preselected electrodes individually with a positive-first biphasic stimulating waveform (600 mV; 100 μ s each phase, repeated 16 times). By averaging over 16 attempts, it was ensured that the majority of stimulation events fell outside any inherent culture bursting that might have occurred. In this way, a suitable input/output pair could be chosen, dependent on how the cultures had developed, in order to provide an initial decision making pathway for the robot.

Thus, in the initially developed culture, we found by experimentation a reasonably repeatable pathway in the culture from stimulation to response. We then employed this pathway to achieve desired robot control—for example, for the obstacle avoidance task—if the ultrasonic sensor was active we mapped the culture response onto the robot actuators in order to cause the robot to turn away from the object being located ultrasonically. This ensures constant movement of the robot, without bumping into any obstacles.

The robot followed a forward path within its corral confines until it reached a wall, at which point the front sonar value decreased below a threshold (set at approximately 30 cm), triggering a stimulating pulse as shown in Fig. 7.3. If the responding/output electrode registered activity following the input pulse then the robot turned to avoid the wall. Essentially, activity on the responding electrode was interpreted as a command for the robot to turn in order to avoid the wall. It was apparent that, in fact, the robot turned spontaneously whenever activity was registered on the response/output electrode. The most relevant result for the experiment was the occurrence of the chain of events: wall detection stimulation response. From a philosophical and neurological perspective, it is of course also of interest to




Fig. 7.3 Analysis of the robot's activity during a simple wall detection/right turn experiment

speculate why there was activity on the response electrode when no stimulating pulse had been applied.

The typical behaviour in the cultures studied was generally a period of inactivity (or low-frequency activity) prior to stimulus, followed by heightened network activity induced almost immediately (within few milliseconds) after stimulus, which decayed (typically after ~ 100 ms) to baseline pre-stimulus activity. The study opens up the possibility of investigating response times of different cultures under different conditions and how they might be affected by external influences such as electrical fields and pharmacological stimulants [24]. At any one time we have typically 25 different cultures available and hence such comparative developmental studies are now being conducted.

With the sonar threshold set at approximately 30 cm from a wall, a stimulation pulse was applied to the culture, via its sensory input, each time this threshold was breached—effectively when the robot's position was sufficiently close to a wall. An indication of the robot's typical activity during a simple wall detection/right turn experiment is shown in Fig. 7.3. The green trace indicates the front sonar value. Yellow bars indicate stimulus pulse times and blue/red bars indicate sonar timing/actuator command timing.

As can be witnessed, these response events may occur purely spontaneously or due to electric stimulation as a result of the sensor threshold being breached. Such events are deemed 'meaningful' only in the cases when the delay between stimulation and response is less than 100 ms. In other words, such an event is a strong indicator that the electric stimulation on one electrode caused a neural response on the recording electrode. The red vertical lines indicate the time that a rotation command is sent to the robot. These events are always coupled (the first one starts the right turn rotation and the second simply ends the rotation). Only the second signals of each pair can be clearly seen here as the rotation initiation commands are overlaid by the yellow electrode firing bars (as a result of electrode firing which instantly initiates a rotation command). A 'meaningful' event chain would be for example at 1.95 s, where the sonar value drops below the threshold value (30 cm) and a stimulation-response subsequently occurs.

Table 7.1 Basic statistics from a wall avoidance experiment

Results	Simulation	Live culture
Wall - > Stimulation event	100 %	100 %
Stimulation - > Response event	100 %	67 %
Total closed loop time	0.075 s	0.2–0.5 s
Run time	240 s	140 s
Meaningful turns	41	22
Spontaneous turns	41	16

Table 7.1 shows typical results from a live culture test in comparison with a ‘perfect’ simulation. If the live culture acted ‘perfectly’, making no mistakes, then the two columns would be identical. This of course raises the question as to what a ‘perfect’ response actually is. In this case it could be regarded as a programmed exercise—which some might refer to as ‘machine-like’. In a sense therefore the culture is asserting its own individuality by not being ‘perfect’.

To explain Table 7.1 further—‘Total closed loop time’ refers to the time between wall detection and a response signal witnessed from the culture. ‘Meaningful turns’ refer to the robot turning due to a ‘wall detection-stimulation-response’ chain of events. A ‘wall to stimulation’ event corresponds to the 30 cm threshold being breached on the sensor such that a stimulating pulse is transmitted to the culture. Meanwhile a ‘stimulation to response’ event corresponds to a motor command signal, originating in the culture, being transmitted to the wheels of the robot to cause it to change direction. It follows that for the culture some of the ‘stimulation to response’ events will be in ‘considered’ response to a recent stimulus—termed meaningful, whereas other such events—termed spontaneous—will be either spurious or in ‘considered’ response to some thought in the culture, about which we are unaware.

By totalling the results of a series of such trials carried out (over 100), considerable differences (as typically indicated in Table 7.1) are observed between the ratio of expected and spontaneous turns between the simulation and the live culture. Under the control of the simulation $95 \pm 4 \%$ (Mean \pm SD) meaningful turns were observed whilst the remaining spontaneous turns ($5 \pm 4 \%$) were easily attributable to aspects of thresholding spike activity. In contrast the live culture displayed a relatively low number of meaningful turns ($46 \pm 15 \%$) and a large number of spontaneous turns $54 \pm 19 \%$ as a result of intrinsic neuronal activity. Such a large number of spontaneous turns was perhaps only to be expected in an uncharacterised system and current work aims to both quiet the level of on-going spontaneous, reminiscent of epileptiform, activity present in such cultures and to discover more appropriate input sites and stimulation patterns.

As a follow up closed-loop experiment, the robot’s individual (right and left separately) wheel speeds were controlled by using the spike firing frequency recorded from the two chosen motor/output electrodes. The frequency is actually calculated by means of the following simple principle: A running mean of spike rate

from both the output electrodes was computed from the spike detector. The detected spikes for each electrode were separated and divided by the signal acquisition time to give a frequency value. These frequencies were linearly mapped (from their typical range of 0–100 Hz) to a range of 0–0.2 m/s for the individual wheel linear velocities. Meanwhile, collected sonar information was used to directly control (proportionally) the stimulating frequency of the two sensory/input electrodes. The typical sonar range of 0–100 cm was linearly re-scaled into the range 0.2–0.4 Hz for electrode stimulation frequencies (600 mV voltage pulses).

The overall setup can be likened to a simple Braitenberg model [25]. However, in our case, sensor-to-speed control is mediated by the cultured network acting as the sole decision-making entity within the overall feedback loop. One important aspect being focussed on is the evocation of long-term potentiation (LTP), i.e., directed neural pathway changes in the culture, thereby effecting plasticity between the stimulating-recording electrodes. Although this was not a major initial target in carrying out this part of the experiment, it has been noted elsewhere that a high-frequency burst time can induce plasticity very quickly [27, 28]. As a result, we are now investigating spike timing dependent plasticity based on the coincidence of spike and stimulus.

7.5 Learning

Inherent operating characteristics of the cultured neural network have been taken as a starting point to enable the physical robot body to respond in an appropriate fashion—to get it started. The culture then operates over a period of time within the robot body in its corral area. Experimental duration, e.g. how long the culture is operational within its robot body, is merely down to experimental design. Several experiments can therefore be completed within a day, whether on the same or differing cultures. The physical robot body of course can operate 24/7.

In our studies thus far, learning and memory [26] investigations are at an early stage. However, we were able to observe that the robot appeared to improve its performance over time in terms of its wall avoidance ability. We are currently investigating this and examining whether it can be repeated robustly and subsequently quantified. What we have witnessed could mean that neuronal structures/pathways employ some form of adaptive processes which allow the system to adjust to the closed-loop stimulation environment. In fact, such adaptation may underlie some forms of plasticity and hence learning. Typically, learning is considered to involve creating representations of the relationships between different objects or inputs. This is the so called associative learning paradigm and in its broadest form it includes auto associative and hetero associative paradigms such as classical or Pavlovian conditioning, operant conditioning and many others. However, psychologists also recognise other broad class of learning paradigms which do not involve association (at least not explicitly). Forms of such non-associative learning include sensitisation and habituation. The latter involves gradual

adjustment of responses to repetitive stimuli and has been well documented in the psychological literature. Putative biological mechanisms of habituation have been well characterised in invertebrates.

We believe, that the experiment involving the robot wall avoidance task may constitute a natural setup for a form of habituation to take place. The observed improvement of the robot performance due to repetition of a satisfactory action may then be a result of the strengthening of neural pathways through a process of adaptation–learning due to habit.

This hypothesis opens a plethora of further interesting questions which we are addressing. What processes, structures or pathways in the cortical neural cultures support such habituation?

What is the ‘cognitive’ capacity of such form of learning? The issue involves characterisation of behavioural repertoire that can be supported by habituation. This includes the complexity of behaviours that can be exhibited as well as the number of different behaviours that can be simultaneously acquired by this form of learning. What are the dynamics of habituation? The length of time over which the habituation takes place, the level of robustness of the acquired behaviours, and what are the properties of degradation of them over time, if the behaviour is not maintained.

Another, more practical but far from trivial, problem that needs to be addressed is how to characterise the habituation and elucidate all the problems and issues highlighted above from data. The signals collected from a culture are highly dimensional, noisy and non-stationary. Moreover, operation in the closed loop requires a quick turnaround time which means that, in spite of the potentially long times over which the experiments are run; the data are of finite and short duration. Hence, the issue of reliable recognition and characterisation of neural responses in such circumstances comes to the fore.

We are encouraged by the fact that such habituation processes in animals have been reported on elsewhere, e.g. [29, 31], and experimentation has been carried out to investigate the effects of sensory deprivation on subsequent culture development. In our case we are monitoring changes and attempting to provide a quantitative characterisation relating plasticity to experience and time. The potential number of confounding variables is however, considerable, as the subsequent plasticity process, which occurs over quite a period of time, is (most likely) dependent on such factors as initial seeding and growth near electrodes as well as environmental transients such as feed rate, temperature and humidity.

On completion of these first phases of the infrastructure setup, a significant research contribution, it is felt, lies in the application of machine learning (ML) techniques to the hybrid system’s closed-loop experiments. These techniques may be applied in the spike sorting process (dimensionality reduction of spike data profiles, clustering of neuronal units), the mapping process between sensory data and culture stimulation as well as the mapping between the culture activity and motor commands, and the application of learning techniques on the controlled electrical stimulation of the culture, in an attempt to exploit the cultured networks’ computational capacity.

7.6 Conclusions

We have successfully realised a closed-loop adaptive feedback system involving a (physical) mobile robotic platform and a cultured neuronal network using a multi-electrode array (MEA), which necessitates real-time bidirectional communication between the culture and the robot. A typical culture being employed presently consists of approximately 100,000–150,000 neurons, although at any one time only a small proportion of these neurons are actively firing.

Trial runs have been carried out with the overall robot and comparisons have been made with an ‘ideal’ simulation which responds to stimuli perfectly as required. It has been observed that the culture on many occasions responds as expected, on other occasions, however, it does not, and in some cases it provides a motor signal when it is not expected to do so.

The concept of an ‘ideal’ response is difficult to address because a biological network is involved, and it should not be seen in negative terms when the culture does not achieve (what is in our eyes) an ideal. We know very little about the fundamental neuronal processes that give rise to meaningful behaviours, particularly where learning is involved, we therefore need to retain an open mind as to a culture’s performance.

The culture preparation techniques employed are constantly being refined and have led to stable cultures that exhibit both spontaneous and induced spiking/bursting activity which develops in-line with the findings of other groups, e.g. [15, 32].

A stable robotic infrastructure has been set up, tested and is in place for future culture behaviour and learning experiments. This infrastructure could be easily modified in order to investigate culture-mediated control of a wide array of alternative robotic devices, such as a robot head, an ‘autonomous’ vehicle, robotic arms/grippers, mobile robot swarms and multi-legged walkers.

In terms of robotics, this study and others like it, show that a robot can have a biological brain to make its ‘decisions’. The 100,000–150,000 neuron size is due to present day limitations—clearly this will increase. Indeed it is already the case that 3-dimensional structures are being investigated [20]. Simply increasing the complexity from 2-dimensions to 3-dimensions (on the same basis) realises a figure of 30 million neurons (approximately) for the 3-dimensional case. The whole area of research is therefore a rapidly expanding one as the range of sensory inputs is expanded and the number of cultured neurons encapsulated rises. The potential capabilities of such robots, including the range of tasks they can perform, therefore needs to be investigated.

Understanding neural activity becomes a much more difficult problem as the culture size is increased. Even the present neuron cultures are far too complex at present for us to gain an overall insight. When they are grown to sizes such as 30 million neurons and beyond, clearly the problem is significantly magnified, particularly with regard to neural activity in the centre of a culture volume, which will be (effectively) hidden from view. On top of this, the nature of the neurons may be

diversified. To this point rat neurons are employed in our studies—however, any animal neurons could be used. Human neurons, which are readily available are particularly suitable due to the fact that any results obtained could be more immediately relevant to human studies. The author wishes to record his feelings here that it is important to stress the need for ethical concerns to be listened to in such circumstances.

7.7 Future Research

There are a number of ways in which the current research programme is being taken forward. First the robot is being extended to include additional sensory devices such as extra sonar arrays, audio input, mobile cameras and other range-finding hardware such as an on-board infrared sensor. This will provide an opportunity to investigate sensory fusion in the culture and to perform more complex behavioural experiments, possibly even attempting to demonstrate links between behaviour and culture plasticity, along the lines of [29], as different sensory inputs are integrated.

Provision of a powered-floor for the robot's corral will provide the robot with relative autonomy for a longer period of time while different learning techniques are applied and behavioural responses monitored. For this the Miabot must be adapted to operate on an in-house powered-floor, providing the robot with an unlimited power supply. This feature, which is based on an original design for displays in museums [30] is necessary since learning and culture behaviour tests will be carried out for hours at a time.

Current hardcoded mapping between the robot goals and the culture input/output relationships can be extended using learning techniques to eliminate the need for an a priori choice of the mapping. In particular, reinforcement learning techniques can be applied to various mobile robot tasks such as wall following and maze navigation, in an attempt to provide a formal framework within which the learning capabilities of the neuronal culture will be studied [35, 36].

To increase the effectiveness of culture training beyond the ~30 % success rate seen in previous work, biological experiments are currently being performed to identify physiological features which may play a role in cellular correlates of learning processes. These experiments also investigate possible methods of inducing an appropriate receptive state in the culture that may allow greater control over its processing abilities and the formation of memories [26] involving specific network activity changes, which may allow identification of the function of given network ensembles. In particular, in terms of cholinergic influences, the possible effect of acetylcholine (ACh) [33] in coordinating the contributions of different memory systems is being investigated.

The learning techniques employed and the results obtained from the culture need to be benchmarked. In order to achieve this we are developing a model of the cultured neural network, based on experimental data about culture density and activity. In doing so we hope to gain a better understanding of the contribution of

culture plasticity and learning capacity to the observed control proficiency. At present we are investigating hidden markov models (HMMs) as a technique for uncovering dynamic spatiotemporal patterns emerging from spontaneously active or stimulated neuronal cultures [34, 37]. The use of hidden markov models enables characterisation of multichannel spike trains as a progression of patterns of underlying discrete states of neuronal activity.

Key for the future, however, will be the possible employment of 30 million human neurons as the brain of a robot body. Through technical development, it is envisioned that the number of neurons will subsequently increase further still [38]. As the 100 billion typical number of neurons is approached, other issues then become apparent—as the sheer size of the robot brain tends to approach that of a human—e.g. what, if any, rights should the robot have?

Acknowledgements This work is funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under Grant No. EP/D080134/1.

References

1. Reger, B., Fleming, K., Sanguineti, V., Alford, S., Mussa-Ivaldi, F.: Connecting brains to robots: An artificial body for studying the computational properties of neural tissues. *Artif. Life* **6**, 307–324 (2000)
2. Holzer, R., Shimoyama, I., Miura, H.: Locomotion control of a bio-robotic system via electric stimulation. In: *Proceedings of International Conference on Intelligent Robots and Systems*, Grenoble, France (1997)
3. Talwar, S., Xu, S., Hawley, E., Weiss, S., Moxon, K., Chapin, J.: Rat navigation guided by remote control. *Nature* **417**, 37–38 (2002)
4. Chapin, J., Moxon, K., Markowitz, R., Nicolelis, M.: Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Neuroscience* **2**, 664–670 (1999)
5. Bakkum, D.J., Shkolnik, A., Ben-Ary, G., DeMarse, T., Potter, S.: Removing Some ‘A’ from AI: Embodied Cultured Networks, *Lecture Notes in Computer Science*, pp. 130–145 (2004)
6. Thomas, C., Springer, P., Loeb, G., Berwald-Netter, Y., Okun, L.: A miniature microelectrode array to monitor the bioelectric activity of cultured cells. *Exp. Cell Res.* **74**, 61–66 (1972)
7. Gross, G.: Simultaneous single unit recording in vitro with a photoetched laser deinsulated gold multimicroelectrode surface. *IEEE Trans. Biomed. Eng.* **26**, 273–279 (1979)
8. Pine, J.: Recording action potentials from cultured neurons with extracellular microcircuit electrodes. *J. Neurosci. Methods* **2**, 19–31 (1980)
9. Potter, S., Lukina, N., Longmuir, K., Wu, Y.: Multi-site two-photon imaging of neurons on multi-electrode arrays. *SPIE Proc.* **4262**, 104–110 (2001)
10. Gross, G., Rhoades, B., Kowalski, J.: Dynamics of burst patterns generated by monolayer networks in culture. In: *Neurobionics: An Interdisciplinary Approach to Substitute Impaired Functions of the Human Nervous System*, pp. 89–121 (1993)
11. Kamioka, H., Maeda, E., Jimbo, Y., Robinson, H., Kawana, A.: Spontaneous periodic synchronized bursting during the formation of mature patterns of connections in cortical neurons. *Neurosci. Lett.* **206**, 109–112 (1996)
12. Lewicki, M.: A review of methods for spike sorting: the detection and classification of neural action potentials. *Network (Bristol)* **9**(4), R53 (1998)

13. Saito, S., Kobayashik, S., Ohashio, Y., Igarashi, M., Komiya, Y., Ando, S.: Decreased synaptic density in aged brains and its prevention by rearing under enriched environment as revealed by synaptophysin contents. *J. Neurosci. Res.* **39**, 57–62 (1994)
14. Ramakers, G.J., Corner, M.A., Habets, A.M.: Development in the absence of spontaneous bioelectric activity results in increased stereotyped burst firing in cultures of dissociated cerebral cortex. *Exp. Brain Res.* **79**, 157–166 (1990)
15. Chiappalone, M., Vato, A., Berdondini, L., Koudelka-Hep, M., Martinoia, S.: Network dynamics and synchronous activity in cultured cortical neurons. *Int. J. Neural Syst.*, **17**(2), 87–103 (2007)
16. Shkolnik, A.C.: Neurally controlled simulated robot: applying cultured neurons to handle an approach/avoidance task in real time, and a framework for studying learning in vitro, in mathematics and computer science. Masters thesis, Department of Computer Science, Emory University, Georgia (2003)
17. DeMarse, T.B., Dockendorf, K.P.: Adaptive flight control with living neuronal networks on microelectrode arrays. In: Proceedings of the 2005 IEEE International Joint Conference on Neural Networks, pp.1549–1551 (2005)
18. Shahaf, G., Marom, S.: Learning in networks of cortical neurons. *J. Neurosci.* **21**(22), 8782–8788 (2001)
19. Bull, L., Uruokov, I.: Initial results from the use of learning classifier systems to control in vitro neuronal networks. In: Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation. (GECCO).pp. 369–376. The ACM, London, England (2007)
20. Hammond, M., Marshall, S., Downes, J., Xydas, D., Nasuto, S., Becerra, V., Warwick, K., Whalley, B.J.: Robust methodology for the study of cultured neuronal networks on MEAs. In: Proceedings 6th International Meeting on Substrate-Integrated Micro Electrode Arrays, pp. 293–294 (2008)
21. Potter, S.M., DeMarse, T.B.: A new approach to neural cell culture for long-term studies. *J. Neurosci. Methods* **110**, 17–24 (2001)
22. Rolston, J.D., Wagenaar, D.A., Potter, S.M.: Precisely timed spatiotemporal patterns of neural activity in dissociated cortical cultures. *Neuroscience* **148**, 294–303 (2007)
23. Wagenaar, D.A.D., T.B.; Potter, S.M.: MEABench: a toolset for multi-electrode data acquisition and on-line analysis. In: Proceedings of the 2nd International IEEE EMBS Conference on Neural Engineering, pp. 518–521 (2005)
24. Xydas, D., Warwick, K., Whalley, B., Nasuto, S., Becerra, V., Hammond, M., Downes, J.: Architecture for Living Neuronal Cell Control of a Mobile Robot. In: Proceedings of the European Robotics Symposium EUROS08, pp. 23–31. Prague, CZ, (2008)
25. Hutt, B., Warwick, K., Goodhew, I.: Emergent Behaviour in Autonomous Robots. In: Bryant, J., Atherton, M. and Collins, M. (eds.) Chapter 14 in Information Transfer in Biological Systems. Design in Nature Series, vol. 2, The WIT Press (2005)
26. Hasselmo, M.E.: Acetylcholine and learning in a cortical associative memory source. *Neural Comput. Arch.* **5**, 32–44 (1993)
27. Cozzi, L., Chiappalone, M., Ide, A., Novellino, A., Martinoia, S., Sanguinetti, V.: Coding and Decoding of Information in a Bi-directional Neural Interface. *Neurocomputing* **65**(66), 783–792 (2005)
28. Novellino, A., Cozzi, L., Chiappalone, M., Sanguinetti, V., Martinoia, S.: Connecting neurons to a mobile robot: an in vitro bi-directional neural interface, *Comput. Intell. Neurosci.* (2007)
29. Karniel, A., Kositsky, M., Fleming, K., Chiappalone, M., Sanguinetti, V., Alford, T., Mussa-Ivaldi, A.: Computational Analysis In Vitro: Dynamics and Plasticity of a Neuro-Robotic System. *J. Neural Eng.* **2**, S250–S265 (2005)
30. Hutt, B., Warwick, K.: Museum robots: multi-robot systems for public exhibition. In: Proceedings of the 35th International Symposium on Robotics, p.52. Paris, France (2004)
31. Marks, P.: Rat-brained robots take their first steps. *New Sci.* **199**(2669), 22–23 (2008)
32. DeMarse, T., Wagenaar, D., Blau, A., Potter, S.: The neurally controlled animat: Biological brains acting with simulated bodies. *Auton. Robots* **11**, 305–310 (2001)

33. Chang, Q., Gold, P.: Switching memory systems during learning: Changes in patterns of brain acetylcholine release in the hippocampus and striatum in rats. *J. Neurosci.* **23**, 3001–3005 (2003)
34. Xydias, D., Downes, J., Spencer, M., Hammond, M., Nasuto, S., Whalley, B., Becerra, V., Warwick, K.: Revealing ensemble state transition patterns in multi-electrode neural recordings using hidden Markov models. *IEEE Trans. Neural Syst. Rehabil. Eng.* **19**(4), 345–355 (2011)
35. Warwick, K., Xydias, D., Nasuto, S., Becerra, V., Hammond, M., Downes, J., Marshall, S., Whalley, B.: Controlling a mobile robot with a biological brain. *Def. Sci. J.* **60**(1), 5–14 (2010)
36. Warwick, K., Nasuto, S., Becerra, V., Whalley, B.: Experiments with an in-vitro robot brain. In: Cai Y., (ed.) *Computing with Instinct*, Lecture Notes in Computer Science, vol. 5897, pp. 1–15, Springer, Berlin (2011)
37. Spencer, M., Downes, J., Xydias, D., Hammond, M., Becerra, V., Warwick, K., Whalley, B., Nasuto, S.: Multiscale evolving complex network model of functional connectivity in neuronal cultures, *IEEE Trans. Biomed. Eng.*, **59**(1), 30–34 (2012)
38. Warwick, K.: Implications and consequences of robots with biological brains. *Ethics Inf. Technol.* **12**(3), 223–234 (2010)