

Neuroethology, Computational

Dave Cliff*

School of Cognitive and Computing Sciences,
University of Sussex, BRIGHTON BN1 9QH, U.K.
E-mail: davec@cogs.susx.ac.uk

Technical Report CSRP338

Abstract

In recent years, a number of neural network researchers have used the term *computational neuroethology* to describe a specific approach to *neuroethology*, the intersection of neuroscience (the study of nervous systems) and ethology (the study of animal behavior). Experiments in computational neuroethology can complement classical computational neuroscience studies of neural control of behavior in animals, and can help in the design of artificial autonomous agents. This is a brief introductory review article. (To appear in: M. A. Arbib (editor) *Handbook of Brain Theory and Neural Networks*, MIT Press 1995.)

1 INTRODUCTION

In recent years, a number of neural network researchers have used the term *computational neuroethology* to describe a specific approach to *neuroethology*, the intersection of neuroscience (the study of nervous systems) and ethology (the study of animal behavior). The definition of computational neuroethology is very similar, but is not quite so dependent on studying animals: animals just happen to be biological *autonomous agents*. Non-biological autonomous agents include some robots, and some simulated agents operating in virtual worlds. Here, autonomous agents are self-governing systems capable of operating (i.e. perceiving and acting) in environments which are complex, uncertain, and dynamic. For the sake of brevity in the rest of the text, autonomous agents will be referred to simply as “agents”, and computational neuroethology will be abbreviated to CNE.

CNE can be distinguished from classical computational neuroscience by its increased emphasis on studying the neural control of behavior within the context of neural systems which are both embodied and situated within an environment.

Put most simply, CNE involves the use of computational modelling in trying to understand the neural mechanisms responsible for generating ‘useful’ behaviors in an agent. The word ‘useful’ is rather imprecise: it is more common to talk of *adaptive* behaviors. In the ethology literature, an adaptive behavior is usually defined as a behavior which increases the likelihood that an animal will survive long enough to produce viable offspring. Often implicit in this definition is the assumption that the animal’s environment is sufficiently unforgiving (or hostile) that if the animal does nothing, it will die before it can reproduce. In studying artificial agents, the utility of the behavior is evaluated by different criteria, such as computational or economic efficiency.

Neural networks that generate adaptive behavior should not be confused with *adaptive* neural networks, where connection strengths may alter as a result of experience. Adaptation

*Copyright © 1994 Dave Cliff. All rights reserved.

or plasticity may itself give rise to new or improved adaptive behaviors, but there are many cases of adaptive behaviors which are genetically determined (e.g. “hard-wired” behaviors such as reflexes and instincts).

In the context of adaptive behavior research, it becomes clear that the neural system is one component in the *action-perception cycle*, where actions allow the agent to perceive information concerning its environment, which may lead to changes in the agent’s internal state, which may in turn affect further actions, which affect what information can be perceived, and so on. This is a notion long stressed by Arbib:

“In speaking of human perception, we often talk as if a purely passive process of classification were involved – of being able, when shown an object, to respond by naming it correctly. However, for most of the perception of most animals and much of human behavior, it is more appropriate to say that the animal perceives its environment to the extent that it is *prepared to interact* with that environment in some reasonably structured fashion.” (Arbib, 1972, p.16)

As defined above, CNE may not seem to be particularly distinguishable from most work in neural network research. After all, many people in computational neural network research might argue that their work may, ultimately, lead to understanding of the neural mechanisms underlying the generation of (some) adaptive behaviors. For example, face recognition is an adaptive behavior in humans, and could probably be classed as an adaptive behavior in, say, a security robot. So why can’t a back-propagation network that learns to distinguish between photographs of human faces (for example) be classed as work in CNE?

2 MOTIVATIONS

This section will give only a brief overview of the motivations for the CNE approach. For further discussion of methodological issues, see (Beer, 1990; Cliff, 1990).

Typically, “connectionist” models employ homogeneous groups of highly idealised and simplified neuron models (called *units*), connected in a regular fashion, which exhibit some form of ‘learning’ or adaptation.

The large majority of simplifying connectionist models can be described in essence as mapping or transforming between representations: input data is presented to the network in a particular format, and the network is judged successful when its outputs can be interpreted as a correct representation of the results of performing the desired transformation. In almost all cases, the input and output representation formats are pre-specified by the experimenter (although it should be acknowledged that: this is not entirely true of unsupervised learning; and there are a number of connectionist models which draw inspiration from biological data in their choice of input and output representations). If such networks are to be employed in artificial agents, or are to be of use in understanding biological agents, then this can only be so under the (often unspoken) assumption that, eventually, it will be possible to assemble a ‘pipeline’ of such input-output transducer networks which links sensory inputs to motor outputs, and produces adaptive behavior. The most significant issue here is that there is a dependence on intermediate representations, which may not be justifiable: neural sensory-motor pathways generating adaptive behaviors might not be neatly partitioned into representation-transforming modules; when we “open up the black box” we may not find any patterns of activity identifiable as a representation in the conventional sense, and even if we do, there is no guarantee that they will be in strong accord with representations chosen *a priori* by connectionist modellers.

This should not be mistaken for an argument against representation, nor for a denial of the vital role played by internal states in the generation of adaptive behaviors: it is simply an awareness of the dangers of being misled by *a priori* notions of representation. One of the safest ways of avoiding these dangers is to model, as far as is possible, *entire* sensory-motor pathways (i.e. the complete sequence of neural processing, from sensory input to motor output) involved in the generation of adaptive behavior. This requires that the agent be studied while *embedded in an environment*: most sensory-motor processing for adaptive behavior involves dynamic interaction with the environment; an embedded agent is part of a closed-loop system, in that certain actions can affect subsequent sensory inputs. The sensory-motor pathway should not be viewed as a ‘pipeline’ transforming from a given input representation to a desired output representation, but rather as a link in the action-perception cycle.

When such an approach is adopted, the true nature of the representations and processing necessary for the generation of relatively complex adaptive behaviors is more likely to be revealed, and the validity of any *a priori* assumptions is clarified.

Naturally, it is beyond the state of the art to attempt to model complete sensory-motor pathways in humans or other large mammals, but experimental work in the neuroethology literature provides a wealth of data from less intellectually able animals, such as arthropods (the animal class which includes insects, spiders, and crustacea) and amphibia: it is such animals that are used as the domains of study in most CNE research, as will be seen in the discussion of current CNE projects later in this article.

Before that, illustration of the above arguments can be made by reference to a series of thought-experiments. In his book *Vehicles*, Braitenberg (1984) describes specifications for a series of simple mobile vehicles, operating in a world with simplified kinematics. The series of vehicles starts with an elementary device which performs primitive heat-seeking behavior; it progresses through vehicles that exhibit positive or negative taxes (i.e. orientation towards or away from a directional stimulus), and primitive forms of learning, pattern detection, and movement detection; culminating in vehicles which exhibit chaotic dynamics and predictive behavior. The internal control mechanisms of all the vehicles are rigorously minimal: the simpler vehicles contain nothing more than wires connecting sensors to actuators; while the more advanced ones employ nonlinear threshold devices with delays and pseudo-Hebbian adaptation.

The key point of these thought-experiments is that Braitenberg uses the psychological language indicative of intentional mental states to describe the observed behavior of the vehicles. He ascribes *fear, aggression, love, values and taste, rules, trains of thought, free will, foresight, egotism and optimism* to his vehicles, and demonstrates that while such terms may be useful at the level of description of an external observer, the internal causal mechanisms could be surprisingly simple.

Braitenberg’s vehicles are strongly reminiscent of the simple electromechanical ‘creature’ *Machina speculatrix*, designed and built by Walter (1953): *M. speculatrix* was built from a photo-electric cell, a touch sensor, two electronic ‘neurons’, and assorted gear trains from old clocks and gas-meters (Walter, 1953, p.113, p.244). Nevertheless, Walter notes that the patterns of behavior it produced could be reasonably described as exhibiting *speculation, discernment, and self-recognition* (Walter, 1953, pp.113–114).

While Braitenberg’s vehicles are nothing more than thought-experiments, they provide insight to possible organisational principles in natural and artificial creatures, and demonstrate the limits of applicability of intentional terminology. As such, they are of relevance to the philosophy of cognitive science. In a commentary on Arbib’s work (discussed further below), Lloyd (in (Arbib, 1987, pp.442–443)) notes that there is generally a trade-off between accuracy and completeness in cognitive modelling, and he argues that completeness is more desirable

than accuracy for two reasons highlighted by Arbib's work:

“...Focusing only on components can lead one to overlook emergent effects of cooperative computation; and working with complete models, even if wrong in detail, nonetheless provides us with analytical tools applicable to future data and future, more accurate models... I think the reasons to push for completeness go beyond these, however, and reveal a further source of the value of Arbib's work for cognitive science.

“...[In 1978, Dennett] proposed that one approach the complexity of humans by looking at simpler systems, first solving cognitive problems as they arise in these 'simple minds' and then bootstrapping towards increasingly complex and human like cognizers. Two sorts of systems appealed to Dennett as fruitful stepping stones: living systems and artifactual systems born of engineering imagination, 'Martian three-wheeled iguanas' and the like.” Lloyd, in (Arbib, 1987, p.443).

Thus it can be seen that the CNE approach has some parentage in the philosophy literature. Indeed, the focus in computational neuroethology is on understanding what Dennett, in (Ewert, 1987, p.373), calls the “wise wiring” underlying the generation of adaptive behavior. The presence of such ideas in the philosophy literature is due, at least in part, to prior arguments such as those found in (Arbib, 1972). At present, the focus in most CNE research is on adaptive behaviors which serve the “four F's”: Feeding; Fleeing; Fighting; and reproduction (this is an old joke, of uncertain origin). These behavioral modes can be argued as underlying much of the more complex adaptive behaviors witnessed in 'higher' animals, including the 'intelligent' activities that inspire philosophers to posit the existence of a “language of thought”.

To summarise, research in CNE can be characterised as placing increased emphasis on modelling entire adaptive-behavior-generating sensory-motor pathways in agents embedded in environments which supply sensory-motor feedback. Such an approach lessens the chances of making untenable assumptions concerning issues of representation and processing. In order to study such pathways where there is reliable biological data, it is necessary to focus attention on relatively simple animals such as arthropods or amphibia.

It is important to note that there is a tradition of related work in the artificial neural network literature: research in reinforcement learning for control tasks is most close; see the Handbook article “Reinforcement Learning” by Barto, in this volume.

3 SOME CURRENT RESEARCH PROJECTS

The notes below describe some CNE projects of direct relevance to the arguments summarised above: Arbib's work on visuomotor activity in frogs and toads; Beer's work on locomotion, guidance, and behavioral choice in cockroaches; Cliff's work on visual tracking in hoverflies; and Franceschini's work on equipping autonomous robots with fly-like compound eyes. All of these are ongoing projects, so the descriptions serve as “snapshots” of their current status, rather than as final reviews of completed research programmes. Following the descriptions of these projects, some related work is discussed.

3.1 The Computational Frog

Probably the most advanced project in CNE is the work of Arbib and his students on an evolving family of models of visually mediated behavior in frogs and toads: see (Arbib, 1987) for a review

of the project with peer commentary. Arbib named his simulation model *Rana computatrix*, the computational frog, after Walter's *Machina speculatrix*.

Arbib's computational modelling is accurately based on data from biological experimental work performed by the neuroethologist Ewert (e.g. (Ewert, 1987)). The *R. computatrix* models are faithful to the known biology, and there is an interplay between the experimental and theoretical work: Arbib constructs "an evolving set of model families to mediate flexible cooperation between theory and experiment" (Arbib, 1987, p.407).

Briefly, the 'evolving' nature of Arbib's work is due to his use of *incremental* modelling: his models explore a variety of different connectivities and parameter settings within the overall paradigm of visuomotor brain function in frog and toad. An initial first approximation model was extended and refined in a number of stages, leading to a family of models for *R. computatrix*.

Arbib's approach involves the definition of a number of functional *schemas*: schemas can be modelled by interacting layers of neuronlike elements, or by nets of intermediate-level units; the network models can be related to experimental data concerning neural circuitry, and the development process iterates (Arbib, 1987, p.411 ff.). Further details can be found in Cervantes-Pérez's article "Schema Theory", in this volume.

The primary focus in the *R. computatrix* models has been on how frogs and toads use vision to detect and catch prey, in environments that include obstacles and barriers. Arbib has evolved a series of schema-based models which account for depth perception as interaction between accommodation and binocular clues, and at the lowest level the schemas are plausibly based on known details of the relevant neurological data. For further details, see the article "Visuomotor Coordination in Frog and Toad" by Cervantes-Pérez, in this volume.

One of the more striking results from this work, with reference to Marr's well-known theory of vision, is the indication that (in frogs and toads at least) there are different perceptual mechanisms for different visual stimuli, i.e. the depths to prey and to barriers are extracted from the optic array by different processing channels, and are integrated in the sensorimotor pathways much later than Marr's theory might suggest.

3.2 The Computational Cockroach

Beer's (1990) book contains both methodological arguments for CNE, and also details of experimental work on a computational cockroach, *Periplaneta computatrix*, which is a simulated hexapod agent embedded in an environment, inspired by neuroethological studies of the cockroach *Periplaneta americana*. The real cockroach uses chemotaxis as one of several strategies to locate food sources. If its path along an odour-gradient is blocked by an obstacle, then it performs stereotyped 'edge-following' behavior. The artificial cockroach is controlled by a heterogeneous neural network which was inspired by biological data, and has been used to study issues in locomotion, guidance, and behavioral choice.

The primary external sensory input was simulated chemosensory information: patches of food in the environment gave off odour gradients detectable under an inverse square law relating distance to odour intensity. The neural nets also received mechanosensory input from e.g. proprioceptors in the limbs and tactile sensors which signal the presence of food under the mouth. The simulation model included elementary kinematics: if the artificial cockroach failed to adopt a stable position for a sufficient length of time, it fell down.

Results from the simulation sessions demonstrated behavior in the model that was highly similar to behavior in the real animal, and Beer subsequently performed "lesion" experiments by selectively deleting connections or units from the *P. computatrix* control network. Again, the results from the artificial system were in agreement with the biological data.

P. computatrix was inspired by biological data, but was not intended as a biological model. The various behaviors were generated by heterogeneous neural networks. The neuron model employed by Beer was more faithful to biology than many of the “formal neurons” used in connectionism: the units involved differential equations modelling membrane potentials, which gave his model neural assemblies a rich intrinsic dynamics. For further details, see Beer’s article “Locomotion, invertebrate”, in this volume.

The primary focus was on designing architectures for such units that could act as controllers for the various behaviors that *P. computatrix* should exhibit. Thus there was no treatment of learning in the initial body of work on the cockroach. More recently, Beer has reported on work which extends the original *P. computatrix* simulation model, testing it by allowing it to control walking in a real hexapod robot (Beer et al. 1992).

In the physical implementation, the control network was still simulated (i.e. the units in the neural network were not realised physically) but the sensorimotor connections to the artificial neural network were interfaced to physical sensors and actuators by means of analogue-digital and digital-analogue converters. Beer reports that in all cases, the response of the physical robot was quite similar to that previously observed in simulation. The implementation did however reveal one problem in the controller which had not been examined in the simulation. This problem (involving disturbances in the crossbody phasing of the legs) was easily rectified, but nevertheless this demonstrates that simulation models cannot be trusted as *perfectly* replicating any physical implementation they may ultimately be intended for.

3.3 The Computational Hoverfly

In studying issues in active vision gaze control with spatially variant ‘foveal’ sampling, Cliff (e.g. (Cliff, 1992)) constructed a simulated embedded agent whose environment and optical system were inspired by studies of the hoverfly *Syrirta pipiens*. The computational hoverfly (known as *Syrirta computatrix*, or SYCO for short) was a simulated agent existing in its own ‘virtual reality’.

SYCO replicated (at the behavioral level) the visually guided tracking movements made when male *S. pipiens* pursue conspecific flies in the hope of finding a mate. The network processing model was based on previously untested models proposed in the biology literature by Collett and Land (1975).

In the SYCO simulator, a dynamic 3D world model of the relative positions and orientations of SYCO and a number of ‘target’ flies was used to synthesize visual input, via an accurate model of the optical anatomy of male *S. pipiens*. Within the simulator, the visual input was passed through parallel image-processing networks which effected crude target-identification mechanisms (cf. ‘bug-detectors’); the results of this process were fed to tracking networks based on proposals by Collett and Land. The output of the tracking networks could alter the position and orientation of SYCO within its simulated world. The positions and orientations of the target flies could also vary dynamically, and the positions and orientations of the objects in the model were further varied by perturbations which model noise in effectors and crosswinds or turbulence in the air.

The simulation studies revealed opportunities for correction and extension of the prior models: simulating a proposed model enforces a degree of mechanistic rigour which is highly likely to expose any shortcomings or discrepancies in the model. For example, constructing the SYCO simulation required a more accurate characterisation of both the optical data, and of the fly’s interaction with its environment (Cliff, 1992), than was previously available in the literature. Furthermore, experience with simulating the original proposed model suggested alternative

more parsimonious models which could account for similar behavior.

Further experimentation with real animals would be required in order to establish which of the alternative models comes closer to the mechanisms actually operating in *S. pipiens*. Thus, while (as with Beer's work) SYCO was not intended as a biological model, results from SYCO may inform future studies of the real biological system. The need in CNE for highly coupled modelling, theoretical analysis, and biological experimentation is manifest.

3.4 A Robot with Compound Eyes

The work of Franceschini and his colleagues (see e.g. (Franceschini et al. 1992)) can be considered as applied computational neuroethology. Franceschini's research background is in the neurophysiology of vision in flies (for further details of this topic, see the Handbook article "Visuomotor coordination in fly" by Egelhaaf and Borst, in this volume). He and his students have recently developed a visually-guided autonomous mobile robot that heads for a goal location while simultaneously avoiding nearby obstacles, without the use of 3D world models or explicit representations of the robot's surroundings. The visuo-motor controller for the robot is based on a custom-built massively parallel analog asynchronous network, designed according to principles elucidated in neurophysiological studies of fly vision. The robot uses visual motion information generated by this circuitry (which is essentially a 1D, 360° horizontal compound eye) to 'slalom' through a cluttered environment towards the goal, at a speed of 50 cm·s⁻¹.

However, the robot is not only an engineering endeavor: its development has helped further the understanding of neural processing of visual information for the control of action in flies. In constructing the robotic system, Franceschini and his colleagues were forced to address issues (such as the effects of using low-tolerance components) which lead to further understanding of details of the fly's visuo-motor nervous system (Franceschini, personal communication 1993).

3.5 Related work

Other highly relevant work includes the work of Brooks's group at MIT, who have constructed a number of robots which they refer to as "creatures". Brooks has argued that the study of insect-level behaviors is more likely to reveal the fundamental mechanisms of cognition than is the study of human-level 'intelligent' activities. There is insufficient space here to discuss this work in the detail it deserves: see (Brooks, 1991) for further details and a review of how this work fits into the history of artificial intelligence and cognitive modelling. Typically the creature-robots are autonomous agents which wander around office-style environments (namely, areas of the MIT AI Lab). Brooks and his co-workers have demonstrated that relatively complex adaptive behaviors (such as autonomous navigation by map-building) can be seen to arise from agents whose control systems are organised as layers of *behavior-generating* modules; for discussion of a particular example (a hexapod walking robot), see (Brooks, 1989). Typically, the control architectures for such agents are built from "combinatorial circuits plus a little timing circuitry" (Brooks, 1991). The use of combinatorial circuits does not preclude such work from being classed as CNE: Brooks (1991) uses the term to describe some of his own work (i.e. Brooks, 1989). Further details of such robotics research are given in Arkin's article "Reactive Robotic Systems", in this volume.

Significantly, biologists Altman and Kien (1989) have identified strong similarities between Brooks's control-architecture principles, and recent models of motor control proposed as underlying the generation of behaviors in a number of phylogenetically diverse animals: the similarities are in the rejection of traditional notions of linear hierarchical control of motor output,

with the execution of behavioral outputs governed by a centralized “command center” (e.g. a ‘command neuron’); instead, distributed heterarchical decentralized control systems with inputs and outputs at many levels have been proposed as better accounting for the interaction between sensory input, central pattern generation, and behavioral output, in locusts, cats, and frogs.

4 DISCUSSION

Computational neuroethology studies neural mechanisms which generate adaptive behaviors, and hence requires that agents are studied within the context of their environmental and behavioral niches.

From the above descriptions, some patterns emerge: all of the CNE projects mentioned are dependent on the availability of fairly detailed neuroethological data. Such data invariably comes from invasive *in vivo* experimentation, and the neuroanatomy of ‘lower’ animals such as arthropods is particularly amenable to such techniques: certain neurons performing particular functions are readily locatable in different individual animals of the same species. There are manifest obstacles preventing the collection of such data from human subjects. Furthermore, by definition, any truly *general* principles underlying the neural generation of adaptive behaviors are those which are common to a number of species, so only cross-species studies will help identify general principles (Cliff, 1990, p.37).

Most of the CNE projects have largely eschewed the study of learning (plasticity), postponing study until sufficient knowledge of the architecture of primary sensory-motor pathways is known to clearly understand how plasticity might increase the capacity for generation of adaptive behavior: cf. (Beer, 1990, p.62); so far, the design approach has had much to offer. Nevertheless, it seems reasonable to expect a clean transition from the study of fixed network connectivities to variable connection strengths. Furthermore, the use of genetic algorithms can allow for the study of ‘evolutionary learning’ in networks with non-plastic connectivities: fixed-weight networks can be specified by genotypes; over a number of generations, the average behavior of a population of such networks may improve as a result of the effects of mutation and recombination in reproduction, if coupled with an appropriate selection pressure and fitness evaluation function. For further discussion of such issues, see the articles “Genetic Algorithms and Neural Networks” by Belew and “Genotypes for Neural Networks” by Nolfi and Parisi, both in this volume. Such techniques have been employed to develop useful CNE models (e.g. (Stork et al., 1992)).

Acknowledgements

Thanks to Michael Arbib, Randy Beer, Hilary Tunley, and an anonymous reviewer for helpful comments on earlier versions of this article.

REFERENCES¹

Altman, J. S., and Kien, J., 1989, New models for motor control. *Neural Computation*, 1:173–183.

* Arbib, M. A., 1972, *The Metaphorical Brain: an introduction to cybernetics as artificial intelligence and brain theory*, New York: Wiley-Interscience.

Arbib, M. A., 1987, Levels of modelling of mechanisms of visually guided behavior. *The Behavioral and Brain Sciences*, 10:407–465.

* Beer, R. D., 1990, *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology*, Academic Press.

Beer, R. D., Chiel, H. J., Quinn, R. D., Espenschied, K., and Larsson, P., 1992, A distributed neural network architecture for hexapod robot locomotion. *Neural Computation*, 4(3):356–365.

* Braitenberg, V., 1984, *Vehicles: Experiments in Synthetic Psychology*, Cambridge MA: MIT Press Bradford Books.

Brooks, R. A., 1989, A robot that walks: Emergent behaviors from a carefully evolved network. *Neural Computation*, 1:253–262.

* Brooks, R. A., 1991, Intelligence without reason. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, pp. 139–159.

* Cliff, D., 1990, Computational neuroethology: A provisional manifesto, in *From Animals to Animals: Proceedings of the First International Conference on Simulation of Adaptive Behavior (SAB90)*, (J.-A. Meyer and S. W. Wilson, Eds.), Cambridge MA: MIT Press Bradford Books, pp. 29–39,

Cliff, D., 1992, Neural networks for visual tracking in an artificial fly, in *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91)*, (F. J. Varela and P. Bourguine, Eds.), Cambridge, MA: MIT Press Bradford Books, pp. 78–87.

Collett, T. S., and Land, M. F., 1975, Visual control of flight behavior in the hoverfly, *Syricta pipiens* L. *Journal of Comparative Physiology*, 99:1–66.

Ewert, J.-P., 1987, Neuroethology of releasing mechanisms: Prey-catching in toads. *The Behavioral and Brain Sciences*, 10:337–405.

Franceschini, N., Pichon, J.-M., and Blanes, C., 1992, From insect vision to robot vision. *Philo-*

¹Articles in the *Handbook of Brain Theory and Neural Networks* are limited to no more than 15 references. An asterisk (*) denotes an expository reference.

sophical Transactions of the Royal Society of London B, 337(1281):283–294.

Stork, D., Jackson, B., and Walker, S., 1992, ‘Non-Optimality’ via Pre-adaptation in Simple Neural Systems, in *Artificial Life II*, (C. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, Eds.), Addison Wesley, pp. 409–429.

* Walter, W. G., 1961, *The Living Brain*, Pelican/Penguin Books. First published by Duckworth, 1953.