Brain Anatomy and Artificial Intelligence

L. Andrew Coward

Australian National University, Canberra, ACT 0200, Australia

Abstract. The brain carries out cognitive learning and processing by performing combinations of different types of information processes. Types of information processes are performed by different anatomical structures and implemented in physiology. The information processes performed by different major anatomical structures including the cortex, basal ganglia, thalamus and cerebellum are described, including their implementations in neuron physiology. The implications for the architecture and design of a general intelligence system are discussed.

1 Introduction

For a system that learns many behaviours, theoretical arguments indicate that the need to utilize physical information handling resources as efficiently as possible constrains system architecture [1]. In the limit as the ratio of behaviours to resources becomes very large, the system will tend to adopt the recommendation architecture form. This form separates different types of information processes into physically separate subsystems. Natural selection pressures tend to favour brain architectures which make efficient use of physical information handling resources like neurons and connectivity, and therefore result in recommendation architecture subsystems appearing in biological brains. Any one cognitive process utilizes information processes in most or all of these subsystems. Each subsystem implements its corresponding information processes by physiological mechanisms. The existence of these forms in biological brain is indicated by a range of physiological evidence, including the cognitive effects resulting from damage to different anatomical structures [2]. Analogous theoretical arguments indicate that any system with general intelligence which must learn a complex combination of behaviours will be constrained into the same general architectural forms [3]. The ways in which the various information processes are implemented physiologically can therefore provide guidance for the design of such general intelligence systems. In this paper the information processes performed by different anatomical structures in the brain are defined, and the ways in which these processes are implemented in physiology described. Ways in which physiology can therefore provide guidance for the implementation of a system with general intelligence are discussed.

Fig. 1. The information architecture that results for any system which must learn a complex combination of behaviours with limited information handling resources

2 Brain Information Architecture

Practical considerations result in the physical architecture of the brain being constrained into the architectural form illustrated in figure 1 [1;4;2]. There are two major subsystems, clustering and competition. Clustering defines and detects conditions in the information available to the brain. This information can include sensory inputs indicating the status of the external environment and the body, and information about the internal status of the brain itself. Groups of similar conditions are called receptive fields, with a receptive field being detected if a significant proportion of its conditions are detected. Detections of receptive fields are provided to the competition subsystem, where each such detection is interpreted as a range of recommendations in favour of different behaviours, each recommendation having an individual weight. Competition determines the total weights in favour of each behaviour, and implements the behaviour with the largest weight. Reward feedback following a behaviour changes the weights of recently detected receptive fields in favour of recently accepted behaviours.

To avoid excessive requirements for information handling resources, individual receptive fields must be shared across (i.e. must recommend) many different behaviours. Receptive fields must be initiated and evolved heuristically, but because one receptive field has recommendation strengths in favour of many different behaviours, all such changes to receptive fields must be carefully controlled. To a good first approximation, receptive fields can be expanded by addition of conditions, but existing conditions cannot be deleted or modified [11]. Consequence feedback is

generally derived from one or a few recent behaviours. Use of such consequence feedback from one behaviour to guide changes to existing conditions in a receptive field could result in undesirable side effects on all the other behaviours also influenced by the field. Consequence feedback can only be used to guide changes to recommendation strengths. In other words, consequence feedback can be directly used within competition but not within clustering. The use of consequence feedback within competition means that components within competition must correspond with individual behaviours or types of behaviour, so that changes will only affect one behaviour.

2 Cortex and Clustering

The cerebral cortex makes up the bulk of clustering in the mammal brain. The cortex is a thin sheet of neural matter, crossed by several million columns, each column detecting a different receptive field [5]. The cortex is separated into different areas, with the columns in one area detecting receptive fields within the same body of input information [6]. Column receptive fields within one area can distinguish between circumstances in which different behaviours of a range of types are appropriate, but do not individually correspond exactly with such circumstances. For example, columns in area TE can discriminate between different categories of visual object, but do not individually correspond with object categories [7]. Different areas detect receptive fields appropriate for supporting different types of behaviour. Although most behaviour types could best be managed with a specific group of receptive field types, supporting all such receptive field types for all possible behaviour types would be very expensive in resources. Area receptive fields must therefore be shared across multiple behaviour types. There are perhaps several hundred different cortical areas [8;9;10], and this number reflects a compromise between behavioural effectiveness and resource economy.

Because consequence feedback cannot be used to guide receptive field changes (with some limited exceptions, see [11]), the need to support behavioural discrimination is the only available criterion [1]. An adequate range of behavioural recommendations must be generated in response to every input state, in order to support a high integrity behaviour selection. This means that at least a minimum number of columns in a cortical area must detect their receptive fields whenever that area is presented with inputs. If less than the minimum is detected, some undetected receptive fields must expand slightly so that they are detected. Such expansions mean that the column will still detect its receptive field in the same circumstances as in the past, and the integrity of the recommendation strengths associated with the column will be preserved. However, such expansions must be as small as possible to minimize possible effects on existing recommendation strengths.

Columns are organized into layers of excitatory pyramidal type neurons, conventionally numbered II through IV, layer I being mainly axons. Different types of inhibitory interneurons occur in all layers. Most external input arrives in layer IV, and most output goes from layers V and VI. Pyramidal neurons in layers V/VI therefore indicate detection of the column receptive field. There is a predominant flow of connectivity from IV to II/III and from II/III to V/VI. If there is strong activity in

II/III but no column output from V/VI, the implication is that a relatively slight column receptive field expansion would result in detection. The requirements are therefore to determine when receptive field expansions are needed, to identify which inactive columns will detect their fields with the least expansion, and to drive that expansion. Such requirements could in principle be met by all-to-all connectivity between cortical columns. Greater effectiveness with fewer resources can be achieved using a central resource manager [12].

The hippocampal system detects receptive fields corresponding with large groups of cortical columns that have expanded their receptive fields at similar times in the past, with inputs derived from layers II/III. At each point in time, the hippocampal system receives inputs from all across the cortex, determines whether and where receptive field expansions are required, and sends signals back to the selected columns to drive expansions [13]. Receptive field changes are behaviours which must be recommended by receptive field detections in the clustering part of the hippocampal system, and only implemented if accepted by the competition part of the hippocampal system, of which an important part is the anterior thalamic nucleus [14].

Fig. 2. Organization of synapses to detect conditions, groups of conditions and receptive field defined by groups of groups of conditions

Expansions make use of the long term potentiation mechanism [15]. As illustrated in figure 2, the receptive field of a pyramidal neuron is defined by a group of conditions, with the receptive field being detected if a significant proportion of the conditions is detected. Each condition (such as a_1 , a_2 etc in figure 2) is defined by a group of synapses from other cortical pyramidal neurons (or from the senses via the thalamus), located on one branch of the dendritic tree. Each synapse may have a different weight. If enough synapses on a branch receive action potentials within a

short period of time, the branch injects potential deeper into the dendrite [16]. In other words, the branch condition has been detected. Paths defined by voltage gated ion channels integrate the contributions of different branches [17]. If this integration injects sufficient potential into the soma, an action potential results [18] indicating detection of the neuron receptive field. Such an output action potential is generally combined with an action potential that backpropagates into the dendritic tree [19]. This backpropagated action potential increases the weights of any synapses that have recently received an input action potential [20].

Receptive field expansions occur by increases in the weights of some synapses on a branch. In the extreme case, a branch like b_1 in figure 2 may be made up of silent synapses from all cortical sources. These silent synapses have zero weight [21]. Such a branch defines a provisional condition. The branch will not contribute to neuron receptive field detection in any normal circumstances. However, the branch also has inputs from the hippocampal system. These inputs are active if the column in which the neuron is located is selected for receptive field expansion. If the branch injects potential on the basis of hippocampal activity, and shortly afterwards the neuron fires, the backpropagating action potential will increase the weights of any of the silent synapses that recently received inputs. These synapses now define a condition that could contribute to neuron firing independent of hippocampal activity. In other words, the receptive field of the neuron has expanded slightly by addition of a new condition. Increases of synaptic weights on regular synapses can also contribute to receptive field expansions, although the total weight of any one synapse is limited, preventing excessive contribution by a single source. New silent synapses could also be added to a regular branch.

The cortical sources of silent synapses on a neuron could be selected at random, but resource economies can be achieved by biasing the random selection in favour of sources that have been active in the past at the same time as the target neuron [22]. Such a bias could be achieved by a rapid partial rerun of past experience, with connectivity being guided by this rerun. It has been proposed that REM sleep is the partial rerun supporting creation of provisional connectivity [4;13].

2.1 Indirect Activation of Receptive Fields

Direct activation of a column on the basis of the presence of its conditions in current sensory inputs makes all the recommendations associated with its detection available. However, there can be some inactive columns that have recommendation strengths that are potentially relevant to current circumstances. For example, if a column is inactive but has often been active in the past at the same time as a number of currently active columns, its recommendation strengths may be relevant. Relevant columns may also include columns recently active at the same time as currently active columns and columns that expanded their receptive fields at the same time in the past as currently active columns. Uncontrolled indirect activations would create a chaotic total activation. Hence indirect activation of a column must be a behaviour recommended by columns on the basis of past temporally correlated activity and only implemented if the total recommendation strength is sufficient. Indirect activation recommendation strengths can also be established on the basis of past activity of a column just before or just after activity of another column. A population of columns

directly activated by sensory inputs can generate a secondary indirectly activated column population, which could itself recommend activation of other columns, giving rise to a tertiary population, and so on.

Indirect activation on the basis of frequent past simultaneous activation supports semantic memory, indirect activation on the basis of simultaneous past receptive field expansion supports episodic memory, and indirect activation on the basis of recent simultaneous activity supports priming memory [4;2;23]. Sequences of direct and indirect activations support complex cognitive processing [2;3;11].

3 Basal Ganglia and Behaviour Selection

The basal ganglia and thalamus make up competition in the mammal brain. The connectivity of these nuclei is illustrated in figure 3. There is large input from the cortex going into the striatum, and the output from the basal ganglia comes from mainly from the globus pallidus internal segment and the related substantia nigra pars reticulata. There is strong reciprocal excitatory connectivity between the thalamus and the cortex, and the GPi/SNr produces constant inhibitory output that targets the thalamus. There are two populations of neurons in the striatum: D1 and D2. They support two parallel paths that link the striatum to the GPi/SNr. One path comes from D1 neurons and inhibits the GPi/SNr, (in other words, reduces inhibition of the thalamus) the other path comes from D2 neurons and effectively excites the GPi/SNr. There is a feedback loop from the GPi/SNr back to the striatum via the SNc. The SNc is made up of dopamine neurons. One mode of firing by neurons in the SNc is constant (tonic) at about 4 Hz [24]. This steady firing creates a background concentration of dopamine in the striatum, the dopamine leaks out from the dopamine synapses and permeates the extracellular environment. At the concentrations created by this firing, dopamine affects the relative activity of the D1 and D2 populations.

Individual neurons in the striatum and in the GPi/SNr correspond with very specific individual behaviours [25]. In information terms, the synaptic strength of an input from the cortex to a neuron in the striatum can be viewed as the weight of the recommendation by the receptive field in favour of the corresponding behaviour. An output from the striatum can be viewed as the total current recommendation weight in favour of the corresponding behaviour (or against any other behaviour). D1 striatal neurons target their corresponding behaviour in the GPi/SNr, and encourage that behaviour. D2 striatal neurons target different behaviours in the GPi/SNr, and inhibit those behaviours. There is therefore a competition being performed by the GPi/SNr which selects the most strongly recommended behaviour.

In many cases a behaviour is implemented by release of a cortical information flow. Attention behaviours are releases of information from the senses into the primary sensory areas. Motor behaviours are releases of information from the motor cortex to the spinal cord. Selection of a type of behaviour is release of information from one group of cortical areas that recommend general types of behaviour to other cortical areas that recommend specific behaviours of a general type. Such cortical information flows are coordinated by the thalamus. Reduction in the inhibitory output from the GPi/SNr corresponding with selection of a behaviour results in an increase in thalamic activity generating release of the corresponding cortical information flows.

Fig. 3. Connectivity of cortex-basal ganglia-thalamus-cortex loop

There are two operational problems which must be avoided. One is failure to select any behaviour in circumstances where some behaviour would be appropriate, the other is to try to implement multiple incompatible behaviours. The dopamine feedback loop manages these problems. If there is too much activity in GPi/SNr, there will be no behaviour. This high activity results in an increase in the background dopamine, increasing the chance of a behaviour. If there is too little activity in GPi/SNr, there could be multiple behaviours. This low activity results in a decrease in the background dopamine, reducing the chance of a behaviour. In addition, if high novelty in the environment is detected by the hippocampal system, hippocampal inputs to the striatum increase the tonal firing of dopamine neurons in the SNc [24], resulting in a higher probability of behaviour in novel environments.

3.1 Reward Behaviours

Rewards are behaviours which change the recommendation weights in the striatum in favour of recently selected behaviours. Reward behaviours are recommended by specific striatal neurons that determine total recommendation weights in favour of rewards on the basis of inputs from various cortical areas. Rewards can reinforce behaviours on different levels, ranging from strategic behaviours to specific muscle movements. For example, consider a musician playing a piece of music. Strategic rewards must increase the probability of playing in certain circumstances. Tactical rewards must increase the probability of a certain piece of music being played in certain circumstances. General muscular rewards must increase the probability of certain limb movements relative to the instrument, and detailed muscular rewards must increase the probability of finger movements resulting in notes being played in a specific order and timing.

Reward behaviours are recommended by various cortical areas and selected using basal ganglia structures. However, strategic rewards are managed by the ventral basal ganglia including the ventral striatum (the nucleus accumbens), the VP (corresponding with GPi/SNr) and the VTA (corresponding with the SNc). More specific rewards are managed by the dorsal basal ganglia. Total recommendation strengths in favour of reward behaviours are determined from cortical inputs in the striatum. The competitive selection occurs within the midbrain dopamine neurons (i.e. the SNc and VTA) and selected reward behaviours are implemented back into the striatum. The connectivity arrangements in support of reward behaviours [26] are illustrated in figure 4.

Fig. 4. Spiral of connectivity managing recommendation, selection and implementation of reward behaviours

In addition to the tonal firing mode described earlier, dopamine neurons can fire in a burst mode: 3 – 20 action potentials within a burst, at a frequency of 20 Hz [27]. Striatal neurons can trigger this burst firing by inhibiting interneurons that target midbrain dopamine neurons [26]. The burst firing in turn targets the striatum. Recent firing of striatal neurons results in LTP lasting up to 60 minutes. Burst firing of dopamine can trigger chemical processes that lock in such LTP changes, while a decline in background dopamine can result in LTD [28]. Hence burst firing can implement a positive reward behaviour, while declines in background dopamine can implement a negative reward behaviour.

As illustrated in figure 4, very complex cortical receptive fields recommend strategic rewards to the nucleus accumbens shell, such rewards are implemented in the core. Tactical rewards are recommended by somewhat less complex receptive fields into the core, and implemented in the central striatum. Very specific rewards

are recommended by the premotor or even the motor cortex and implemented into the dorsolateral striatum. This arrangement highlights the issue that reward behaviours must be carefully managed, because the effect of such behaviours on regular behaviours can be very potent, especially in the case of reward behaviours at the strategic level.

4 Thalamus and Cortical Information Management

There are several cortical information flow management issues. One is that receptive field detections by the senses must only be allowed to generate extensive receptive field detections in sensory areas if such generation is a selected behaviour. An example of this is visual attention, where at one point in time, a single visual object is selected to generate higher order receptive field detections over other objects.

A second issue is the need to detect receptive fields simultaneously within multiple different objects, without interference. An example is the need to detect receptive fields within several different visual objects, then to detect receptive fields that combine information from multiple objects in order to develop recommendations appropriate to responding to the group of objects.

These two issues are addressed by related mechanisms. When an action potential spike arrives at a neuron, it injects a postsynaptic potential that decays with a time constant of about 10 milliseconds [29]. For a neuron to detect its receptive field, enough input action potentials must therefore arrive within less than about 10 milliseconds. Spike inputs derived from the senses arrive distributed in time in a relatively random fashion. The rate of spikes with this distribution is too low to generate strong receptive field detections in associative sensory areas and beyond. Visual attention is implemented by a frequency modulation at about 40 Hz (the gamma band of the EEG) imposed on inputs from just the retinal area corresponding with a visual object [11] to the primary visual area. This frequency modulation concentrates spikes within a time slot of less than 10 milliseconds, leaving the rest of the 25 millisecond modulation interval with a relatively low number of spikes. This concentration means that receptive fields will be detected mainly within information derived from just the selected visual object. If different frequency modulations are imposed on inputs from different visual objects, separate populations of receptive fields can be detected within the different objects in different timeslots within the available 25 millisecond interval. This limits the number of different objects to about three, consistent with human working memory capabilities [30].

Each primary thalamic nucleus receives excitatory (glutamatergic) inputs from a number of cortical areas, and provides glutamatergic outputs to one of these areas. The interconnecting axons all pass through the thalamic reticular nucleus (TRN). This nucleus contains only GABAergic interneurons that generate spike trains at the gamma frequency, imposing the gamma band frequency modulation on their targets [31; 32]. At a more detailed level, GABA synapses are generally inhibitory. However, if an action potential arrives at a GABAergic synapse more than about 5 milliseconds before an action potential arrives at a nearby glutamatergic synapse, the result is to increase the excitatory effect of the glutamatergic spike [33]. Hence a train of GABAergic action potentials will tend to result in the outputs generated by a

relatively random distribution of incoming glutamatergic spikes being concentrated in particular time slots relative to the GABAergic inputs.

5 Cerebellum and Behaviour Sequence Management

A sequence of actions is initially managed by detection of receptive fields after each action, and selection of the next action on the basis of the recommendation strengths of those fields. However, if given some initial circumstances a sequence of actions is often performed in the same order, significant gains in speed and accuracy are possible if the sequence is recorded and executed as a whole. The cerebellum performs this function both for voluntary motor functions in which the sequence of actions are muscle movements and for cognitive functions in which the sequence of actions include indirect receptive field activations of various kinds. If the recorded sequences are lost for some reason (e.g. cerebellar damage), the sequences can still be implemented (more slowly and less rapidly) by the full cortex-basal ganglia cycle at each stage.

The connectivity between major anatomical structures that manage voluntary motor control (and other behaviours) is illustrated in figure 5. The regular way to generate an action is for the cortex to detect receptive fields in the available sensory and body position information, the basal ganglia to interpret the receptive field detections into a predominant action, and the action to be implemented by release of cortical receptive fields to, for example, the spinal cord. A sequence of actions will initially be implemented step by step. However, once a frequently used sequence has been defined, greater speed and accuracy can be achieved by automatically implementing each action a short time after the previous action in the sequence.

Fig. 5. Parallel paths through basal ganglia and cerebellum managing behaviour selection and implementation

The cerebellum essentially copies action sequences defined by the full cortex-basal ganglia-thalamus-cortex cycle. It achieves this by establishing connections on to the points that drive actions, and triggering an action when it detects the sensory and motor information that occurred in the past at the time of the preceding action, ensuring correct timing between actions [11]. Hence the cerebellum will tend to carry sequences through to a conclusion, although the full loop could interrupt if necessary.

In the cerebellar cortex there are three layers. The bottom layer is made up of large numbers of granule cells. These have inputs from the pontine nucleus, that in turn gets extensive inputs from the cerebral cortex [34]. Granule cell axons rise to the top layer and bend 90 degrees to run parallel to the surface, forming what are called the parallel fibres. The middle layer is made up of Purkinje cells. These cells have a flat dendritic tree that spans the top layer, perpendicular to the parallel fibres. This dendrite gets single synapses from a very large number of the parallel fibres. Inputs from the inferior olive target Purkinje cell dendrites directly, but each fibre targets just one or sometimes a small group of Purkinje cells, wrapping around the dendrite and therefore called a climbing fibre.

The receptive field of a Purkinje cell (in other words the combination of parallel fibres) can expand by LTP, but an input from the inferior olive triggers LTD [35]. Cerebellar nuclei cells fire tonically, and are excited by the same pontine inputs as the Purkinje cells. They are inhibited by Purkinje outputs. Inferior olive neurons are inhibited by cerebellar nuclei neurons.

A sequence is first implemented step by step by the primary loop through the basal ganglia. If a sequence occurs frequently, different individual cerebellar nuclei neurons establish connections on to groups of neurons driving muscle movement by a coincidence guided LTP type mechanism, on the basis of repetition. Purkinje cells develop receptive fields corresponding with the circumstances just before different actions in a sequence are implemented, again by coincidence guided LTP. Purkinje cells also establish inhibitive connectivity on to cerebellar nuclei neurons, with connection strengths weakened if they tend to be active at the same time.

At this point the cerebellum can take over the sequence. If an action follows the preceding action correctly, the cerebellar nucleus neurons will inhibit the inferior olive, and no LTD will occur in the Purkinje cell. If the action does not occur, there will be LTD which will slightly shift the Purkinje receptive field. Hence the inferior olive tunes the timing of the action sequence [36] and is only active when cerebellar learning is under way [37].

As indicated by the effects of cerebellar damage, the cerebellum manages not only motor sequences including speech generation, but also complex cognitive processes [38]. However, the primary behaviour generation route through the basal ganglia can take back full control of such sequences, with perhaps an undamaged cerebellar region learning the sequences again.

6 Lessons for Design of General Intelligence Systems

The structure and processes of the brain have a number of implications for the design of general intelligence systems. The first is that as the ratio of intelligent behaviours to be learned to available resources increases, the system will be more and more

tightly constrained into the information architecture illustrated in figure 1. If such a system is designed with some different information architecture, either the physical resources will have to be increased, or the system will have a strong tendency towards destruction of older learning when new learning occurs [1].

Secondly, one of the key design decisions is the number of receptive field modules, corresponding with the number of cortical areas in the brain. Each module detects a different type of receptive field. A type is specified by a limited number of sources of the inputs within which fields are defined, and the approximate number of such inputs that contribute to one receptive field. Sources include sensory systems and other receptive field modules. One type will be most effective for recommending a range of behaviour types, and the number of types represents a compromise between the better behavioural support and increased resources as the number of modules increases. A related design decision is the number of different receptive fields supported within a module, and the degree of detailed discrimination that can be provided within one receptive field, analogous with the different pyramidal neurons with potentially slightly different receptive fields that can provide output from one cortical column. The specification of the process by which receptive fields are initially defined is a further important factor. Definitions could be bootstrapped by a design guided initial specification followed by heuristic tuning. This would be analogous with partial genetic specification of types of receptive field in primary sensory areas by guidance to initial connectivity.

A third implication is that a subsystem to manage changes to receptive fields must be provided, including appropriate receptive field modules. Such changes are critical behaviours which must be recommended by receptive field detections and selected in appropriate circumstances by the competitive subsystem. Note that as in the brain, receptive field modules needed for change management may also be useful for other behaviours. Change management must minimize the degree of change to receptive fields, and some mechanism analogous with REM sleep to limit the possible expansions will be required.

A fourth implication is that receptive field modules supporting indirect activations on the basis of past temporally correlated activity must be provided. These modules will contain receptive fields corresponding with groups of receptive fields in other modules which have shown past temporal correlations in activity. Such receptive fields will recommend indirect activation of other receptive fields. Such indirect activations are behaviours which are only implemented if selected by the competition subsystem. A fifth implication is that the management of reward behaviours must be carefully managed, with appropriate receptive field modules to manage rewards for strategic, tactical and detailed actions in an integrated fashion. A sixth implication is that a subsystem to learn and implement frequently utilized sequences of actions will improve the accuracy and speed of the system.

The architecture of the brain is the result of severe natural selection constraints imposed over hundreds of millions of years of evolution favouring species that can learn complex behaviours with the least resources and the least interference between past and new learning. This architecture is therefore a good guide for implementing general intelligence systems that are subject to similar constraints.

References

1. Coward, L.A.: The Recommendation Architecture: lessons from the design of large scale electronic systems for cognitive science. *Journal of Cognitive Systems Research* 2(2), 111–156 (2001)
2. Coward, L.A.: *A System Architecture Approach to the Brain: from Neurons to Consciousness*. Nova Science Publishers, New York (2005b)
3. Coward, L.A., Gedeon, T.D.: Implications of Resource Limitations for a Conscious Machine. *Neurocomputing* 72, 767–788 (2009)
4. Coward, L.A.: *Pattern Thinking*. Praeger, New York (1990)
5. Mountcastle, V.H.: The columnar organization of the neocortex. *Brain* 120, 701–722 (1997)
6. Hilgetag, C.-C., Burns, G.A.P.C., O'Neill, M.A., Scannell, J.W., Young, M.P.: Anatomical connectivity defines the organization of clusters of cortical areas in the macaque monkey and the cat. *Philosophical Transactions of the Royal Society of London B* 355(1393), 91–110 (2000)
7. Tanaka, K.: Columns for complex visual object features in the inferotemporal *cortex: clustering of cells with similar but slightly different stimulus selectivities*. *Cerebral Cortex* 13, 90–99 (2003)
8. Petrides, M., Pandya, D.N.: Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque ventrolateral prefrontal cortex and corticocortical connection patterns. *European Journal of Neuroscience* 11, 1011–1036 (1999)
9. Petrides, M., Pandya, D.N.: Comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns in the monkey. *European Journal of Neuroscience* 16, 291–310 (2002)
10. Morasan, P., Rademacher, J., Schleicher, A., et al.: Human primary auditory cortex: Cytoarchitectonic subdivisions and mapping into a spatial reference system. *NeuroImage* 13, 684–701 (2001)
11. Coward, L.A.: Modelling Memory and Learning Consistently from Psychology to Physiology. In: Cutsuridis, V., Amir, H., Taylor, J. (eds.) *Perception-Action Cycle: Models, Architectures, and Hardware*. Springer Series in Cognitive and Neural Systems, pp. 52–123 (2011)
12. Coward, L.A.: The Hippocampal System as the Cortical Resource Manager: a model connecting psychology, anatomy and physiology. *Advances in Experimental Medicine and Biology* 657, 315–364 (2010)
13. Insausti, R., Amaral, D.G.: Hippocampal Formation. In: Paxinos, G., Mai, J.K. (eds.) *The Human Nervous System*, pp. 871–914. Elsevier, Amsterdam (2004)
14. Aggleton, J.P., Malcolm, W., Brown, M.W.: Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. *Behavioral and Brain Sciences* 22, 425–444 (1999)
15. Bi, G.-q., Poo, M.-m.: Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience* 18, 10464–10472 (1998)
16. Murphy, T.H., Baraban, J.M., Gil Wier, W., Blatter, L.A.: Visualization of quantal synaptic transmission by dendritic calcium imaging. *Science* 263, 529–532 (1994)
17. Gullledge, A.T., Kampa, B.M., Stuart, G.J.: Synaptic Integration in Dendritic Trees. *Journal of Neurobiology* 64(1), 75–90 (2005)
18. Hausser, M., Mel, B.: Dendrites: bug or feature? *Current Opinion in Neurobiology* 13, 372–383 (2003)

19. Stuart, G.J., Sakmann, B.: Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature* 367, 69–72 (1994)
20. Malenka, R.C., Bear, M.F.: LTP and LTD: An Embarrassment of Riches. *Neuron* 44, 5–21 (2004)
21. Malenka, R.C., Nicoll, R.A.: Silent synapses speak up. *Neuron* 19, 473–476 (1997)
22. Coward, L.A.: A Functional Architecture Approach to Neural Systems. *International Journal of Systems Research and Information Systems* 9(2-4), 69–120 (2000)
23. Coward, L.A.: Accounting for episodic, semantic and procedural memory in the recommendation architecture cognitive model. In: *Proceedings of the Ninth Neural Computation and Psychology Workshop: Modelling Language, Cognition, and Action* (2005a)
24. Floresco, S.B., West, A.R., Ash, B., Moore, H., Grace, A.A.: Afferent modulation of dopamine neuron firing differentially regulates tonic and phasic dopamine transmission. *Nature Neuroscience* 6, 968–973 (2003)
25. Alexander, G.E., Crutcher, M.D.: Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends in Neurosciences* 13(7), 266–271 (1990)
26. Haber, S.N., Fudge, J.L., McFarland, N.R.: Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *Journal of Neuroscience* 20, 2369–2382 (2000)
27. Hyland, B.I., Reynolds, J.N.J., Hay, J., Perk, C.G., Miller, R.: Firing modes of midbrain dopamine cells in the freely moving rat. *Neuroscience* 114, 475–492 (2002)
28. Sajikumar, S., Frey, J.U.: Late-associativity, synaptic tagging, and the role of dopamine during LTP and LTD. *Neurobiology of Learning and Memory* 82, 12–25 (2004)
29. Asztely, F., Gustafsson, B.: Dissociation between long-term potentiation and associated changes in field EPSP waveform in the hippocampal CA1 region: an in vitro study in Guinea pig brain slices. *Hippocampus* 4, 148–156 (1994)
30. Cowan, N.: The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behavioral and Brain Sciences* 24, 87–114 (2000)
31. Percheron, G.: Thalamus. In: Paxinos, G., Mai, J.K. (eds.) *The Human Nervous System*, 2nd edn. Elsevier, USA (2004)
32. Macdonald, K.D., Fifkova, E., Jones, M.S., Daniel, S., Barth, D.S.: Focal Stimulation of the Thalamic Reticular Nucleus Induces Focal Gamma Waves in Cortex. *Journal of Neurophysiology* 79, 474–477 (1998)
33. Gulledge, A.T., Stuart, G.J.: Excitatory Actions of GABA in the Cortex. *Neuron* 37, 299–309 (2003)
34. Kelly, R.M., Strick, P.L.: Cerebellar loops with motor cortex and prefrontal cortex of a nonhuman primate. *Journal of Neuroscience* 23, 8432–8444 (2003)
35. Vogt, K.E., Canepari, M.: On the Induction of Postsynaptic Granule Cell–Purkinje Neuron LTP and LTD. *Cerebellum* 9, 284–290 (2010)
36. Yarom, Y., Cohen, D.: The olivocerebellar system as a generator of temporal patterns. *Annals of the New York Academy of Sciences* 978, 122–134 (2002)
37. Kim, J.J., Krupa, D.J., Thompson, R.F.: Inhibitory cerebello-olivary projections and blocking effect in classical conditioning. *Science* 279, 570–573 (1998)
38. Bracke-Tolkmitt, R., Linden, A., Canavan, A.G.M., Rockstroh, B., Scholz, E., Wessel, K., Diener, H.-C.: The Cerebellum Contributes to Mental Skills. *Behavioral Neuroscience* 103(2), 442–446 (1989)