

A Review of Cell Assemblies

Christian R. Huyck · Peter J. Passmore.

Received: date / Accepted: date

Abstract Since the Cell Assembly (CA) was hypothesised, it has gained substantial support and is believed to be the neural basis of psychological concepts. A CA is a relatively small set of connected neurons, that through neural firing can sustain activation without stimulus from outside the CA, and is formed by learning. Extensive evidence from multiple single unit recording and other techniques provides support for the existence of CAs that have these properties, and that their neurons also spike with some degree of synchrony. Since the evidence is so broad and deep, the review concludes that CAs are all but certain. A model of CAs is introduced that is informal, but is broad enough to include, e.g., synfire chains, without including, e.g., holographic reduced representation. CAs are found in most cortical areas and in some sub-cortical areas, they are involved in psychological tasks including categorisation, short-term memory and long-term memory, and are central to other tasks including working memory. There is currently insufficient evidence to conclude that CAs are the neural basis of all concepts. A range of models have been used to simulate CA behaviour including associative memory and more process oriented tasks such as natural language parsing. Questions involving CAs, e.g., memory persistence, CAs' complex interactions with brain waves and learning, remain unanswered. CA research involves a wide range of disciplines including biology and psychology, and this paper reviews literature directly related to the CA, providing a basis of discussion for this interdisciplinary community on this important topic. Hopefully this discussion will lead to more formal and accurate models of CAs that are better linked to neuropsychological data.

Key Words: Cell Assembly, Short-Term Memory, Long-Term Memory, Categorisation, Neural Ensemble, Hebbian Learning, Associative Memory

C. Huyck
Middlesex University. London U.K.
E-mail: c.huyck@mdx.ac.uk

P. Passmore
E-mail: p.passmore@mdx.ac.uk

1 Introduction

The Cell Assembly (CA) hypothesis states that a CA is the neural representation of a concept (Hebb, 1949). The CA hypothesis (and indeed Hebb's related learning rule) has been increasingly supported by biological, theoretical, and simulation data since it was made. The CA can act as a categoriser of sensory stimuli, so the presentation of an object (for example an orange) to an individual may cause a particular CA to become active, allowing the individual to identify the object as an orange. Similarly, CAs can be activated without direct sensory stimuli, so a person's CA for orange will ignite when they think about an orange.

CAs are composed of neurons, so they are an intermediate structure smaller than the entire brain, yet larger than neurons. CAs provide an organising principle for the study of mind and brain, though certainly not the only organising principle. Moreover, CAs cross the boundary between psychology and neurobiology because psychological concepts emerge from the behaviour of neurons.

As the neural representation of a concept, the CA is extremely important, but there is not general agreement on its definition, and the scientific community's understanding of CAs is incomplete. Many in the community support the idea of CAs, but others are not confident that CAs exist. Perhaps this is due to the range of disciplines that make use of the concept. This paper brings together the community's current understanding of this important concept.

This paper reviews evidence supporting the CA hypothesis, and evidence that further explains CA behaviour, function, and structure. While our conclusion does ultimately rely on converging evidence, we have read a significant fraction of the research literature and have actively sought evidence against our conclusion that CAs exist in mammalian brains. A range of psychological, biological and modelling evidence is used to develop a picture of current understanding of CAs. While there is a rich knowledge of CAs, there are many unanswered questions involving CAs, and this paper discusses some of them.

1.1 Hebb's Cell Assembly hypothesis

Hebb's CA hypothesis (Hebb, 1949) states that a CA is a collection of neurons that is the brain's mechanism for representing a concept. This hypothesis ties neurobiology to psychology in that the collection of neurons is a neurobiological entity, and the concept that they support is a psychological entity. A simple example would be that a person's concept of *dog* is neurally implemented by a set of their neurons, their *dog* CA. CAs encode elements of higher cognitive processes like words, mental images and other types of concepts.

The standard model, derived directly from Hebb (1949), is that these neurons have high mutual synaptic strength. When a sufficient number of them fire, due to, for example, mention of the word *dog*, they cause other neurons in the CA to fire, which in turn cause other neurons in the CA to fire, leading to a cascade of neural firing. This neural circuit can persist long after the initial

stimulus has ceased. The psychological correlate of this persistent firing of the neurons in the CA is a short-term memory (STM); without further support, this reverberation can persist on the order of seconds.

The high mutual synaptic strength is a result of Hebbian learning. The neurons have co-fired in response to earlier stimuli, and this has caused the strength of the synapses between them to increase. Thus the CA is typically formed due to repeated presentation of similar stimuli, and thus represents a long-term memory (LTM).

Knowledge of brain physiology and pharmacology has advanced greatly since Hebb's day, but, while taking account of more modern knowledge, the CA hypothesis still rests on three theoretical "principles" about CAs being organized collections of neurons. These principles are:

1. A CA is a relatively small set of neurons, which encodes each concept. With an estimated 10^{11} neurons in the human brain (Smith, 2010), the authors suggest a range of likely CA sizes in terms of potentially constituent neurons, as being 10^3 to 10^7 neurons per CA. At the lower end of the range, we suggest CAs are more likely to be atomic¹, whereas at the upper end of the range we imagine super-CAs composed of many sub-CAs at many levels; e.g. one might consider a cognitive map (Buonomano and Merzenich, 1998; McNaughton et al, 2006) as such a super-CA. Note also, at the individual neuron level, it is assumed that a neuron may be a member of different CAs.
2. The neurons of a CA can show self-sustaining persistent activity. This persistent activity is often called reverberation. The basic idea is that the neurons of a CA fire at an elevated rate after the initial stimulation that ignited the CA. (Please note, we consider that when substantially activated, neurons emit a spike or "fire", whereas activated CAs are said to "ignite" (Braitenberg, 1978), and potentially remain ignited and active for seconds.)
3. A CA's set of neurons, which represents a particular concept, is learned. Typically this learning is assumed to be through synaptic weight change. However, neural death, neural growth, synaptic death, and synaptic growth may all play a part in the growth of the CA, all within a Hebbian learning framework.

These principles are the basis of the standard model, which is further discussed, along with some of its variants, in Section 2.6. Important activities occur in the brain at a range of time scales, and the CA provides coherence on at least two of these scales, seconds for STM, and years for LTM.

1.2 What's in the rest of the paper

Section 2 provides a relatively brief summary of the neural data that supports the existence of CAs in human and other mammalian brains. This section is

¹ An atomic CA is not composed of other CAs.

driven by neural behavioural data linked to psychological behavioural data, and this data provides compelling evidence that CAs do exist in mammalian brains. Though the CA model is informal, this section also introduces a standard model and some of its variants, while discussing one model that is clearly not a CA.

The existence of CAs leads to a series of questions about their location, function, and form. Section 3 reviews current understanding of these questions and their, tentative, answers; this and later sections are necessarily more speculative. One psychological system that is of particular interest is working memory, and Section 4 deals with its relations to CAs in some depth. These questions are more complex than the existence question, and thus their answers depend more on analysis of psychological behaviour, arguments from theory, and simulation. These complex questions, combined with the difficulty of analysing the behaviour of CAs composed of thousands or millions of neurons, currently leave modelling as the best method of addressing these questions about CAs. Section 5 reviews models of CAs, and computational systems that use them.

CAs do not provide all the answers to understanding the brain, nor are they fully understood. In the hope of advancing the understanding of CAs, section 6 discusses some current questions about CAs, particularly some questions involving learning. It also notes that CAs may vary from concept to concept, and within one concept over time. The concluding section, 7, summarizes.

The primary concern of this paper is the functioning of humans both psychologically and neurally. However, this functioning is similar to that of other animals, in particular other primates and other mammals. Consequently, much of the evidence is drawn from animal studies.

Before turning to evidence of CAs in the brain, a final problem we faced building this review is that not all the potentially relevant papers actually mention CAs, or near synonyms such as neuronal ensembles. In this review, we tend to cite papers that are explicitly about CAs, and we can only cite a sample of the papers we have looked at, and these, realistically, are only a fraction of the published literature. The papers that were included were highly cited, seminal, central to our arguments, and usually some combination of the three. We also frequently omitted second papers from authors, as the first provided the link to the body of work.

2 Existence of Cell Assemblies

Hebb proposed CAs as a solution consistent with the physiological and psychological data available in the 1940s. Since that proposal, there has been a vast amount of scientific evidence uncovered supporting his hypothesis that the neural basis of concepts are sets of neurons that remain persistently active and are formed by some sort of Hebbian learning mechanism.

The basic structure of the brain is a loosely coupled net of neurons where neurons are firing constantly at a low rate. Other cells, such as glia, are widely

considered to have little cognitive effect at the level of individual concepts. This neural structure is capable of supporting CAs with the CA igniting when only a small subset of its neurons initially fire (see Section 2.1).

CAs require population coding (principle 1), which is also called ensemble coding. A relatively small set of neurons is the basis of each CA, with neurons able to participate in multiple CAs. The evidence supports the use of population coding throughout the brain (Averbeck et al, 2006a, and see Section 2.2), a review article shows that population coding accurately correlates with a range of cognitive states (Schoenbaum, 1998), and a vast range of specific studies support different aspects of population coding.

There is a large body of evidence for persistent activity (principle 2) (see Section 2.3) and for a fourth principle, synchronous neural firing correlated with cognitive states, (see Section 2.4). This evidence provides very solid support for the existence of CAs. The data related to learning (principle 3) provides further support for the existence of CAs (see Section 2.5). A set of neurons that adhere to these four principles are a CA.

Several articles provide reviews of different sorts of evidence for CAs. A wide range of data is presented showing that neural firing behaviour supports a CA model (Harris, 2005). Spike trains show structure that is not present in the stimulus, and they are not strictly controlled by sensory input; spikes are coordinated, and firing behaviour correlates with internal cognitive states. A range of data, including imaging data such as PET and MEG, but also behavioural, lesioning and event-related potential data, suggests that CAs for words are distributed in brain areas that are semantically specific (Pulvermuller, 1999). For instance, the CAs for action words include neurons in motor cortices. Some of the evidence from the sources in these reviews is recounted in the remainder of this section.

2.1 Anatomical data

When developing his hypothesis, Hebb considered several anatomical theories that still have substantial scientific support. Two in particular are that neurons are constantly active throughout the brain, and that large amounts of cortex can be removed with apparently little effect. Since 1949, there have been many advances in the understanding of neuroanatomy, with, for instance, more modern evidence showing that each neuron tends to fire about once a second (Bevan and Wilson, 1999; Bennett et al, 2000), providing a large range of background activity.

When a neuron fires, it sends activation to neurons to which it is connected, and neurons that have received a great deal of activation fire (Churchland and Sejnowski, 1999). Particular sets of neurons are driven at particular times, leading to firing rates as high as every 5ms (Ylinen et al, 1995). Most neurons lose their activation when they fire, and need more activation to fire again (Krahe and Gabbiani, 2004). Almost all of the synapses of a neuron connect to a unique neuron, and there are rarely as many as three synapses to one

neuron from a given neuron (Braitenberg, 1989). Several neurons need to fire as input to a neuron to make it fire (Churchland and Sejnowski, 1999; Bruno and Sakmann, 2006).

Since many neurons are required to make another neuron fire, and there is a large amount of background activity, a large number of neurons need to fire to make another neuron reliably fire. However, many neurons firing together can make many other neurons fire, so groups of neurons will tend to fire together. Sparse connectivity supports population coding (see Section 2.2). Furthermore, since neurons can be removed with little cognitive impact, the cognitive elements must depend on a large distributed set of neurons.

Similarly, connectivity and firing support persistent activity (Section 2.3), because an initially firing set of neurons can cause a second set to fire, which can in turn cause the first set to fire. Of course, this behaviour need not be limited to two non-overlapping sets of neurons, so the neurons in both sets can fire at an elevated rate.

Broadly speaking, the cortex is an enormous collection of neurons that is not easily separable into smaller structures (Braitenberg, 1989). At a coarse grain, the cortex has a laminar architecture, that is a roughly 2000 cm^2 sheet with six thin layers (Martini, 2001), and this paper mostly considers cortical CAs (but see Section 3.1.4). Neuroscientists have divided the cortex into areas (e.g. Brodmann areas), and when there are feed forward connections from one area of the brain to another, there are almost always connections back from the second to the first (Lamme et al, 1998). These recurrent connections could support CAs that cross brain areas and thus integrate features over a range of complexities and modalities.

2.2 Cell Assemblies and population coding

The first principle of the CA hypothesis is that each concept is represented by a set of neurons. With population coding, a particular concept is coded by a set of neurons that fire at an elevated rate when the concept is perceived or is in STM.

An alternative to population coding is that a single cell represents a concept, commonly known as the grandmother cell (Barlow, 1972). However, a single neuron can not represent a concept because neurons die, and one would lose the concept of one's grandmother if that *grandmother* neuron died. Evidence available to Hebb also eliminates this possibility; the loss of large areas of cortex would remove many concepts, and background neural activity would cause the concepts to pop on and off. More recently, the growing understanding of the prevalence of neural death (Morrison and Hof, 1997) shows that the grandmother cell hypothesis is untenable.

The other extreme alternative to population coding is that all of the neurons code for each concept, and particular firing behaviour determines which concept is active. Hebb argued against this sensory equipotentiality using the effects of learning (see Section 2.5). Subsequently, substantial direct evidence

has supported population coding over equipotentiality, e.g., a review shows population coding is used in a wide range of brain areas (Schoenbaum, 1998).

Much of the evidence for population coding, and other phenomena, is derived from the widely used technique of placing electrodes in or near a neuron to measure its electrical potential. The electrodes can measure depolarisations and thus neural spiking. Direct recording of neural activation levels and spikes is, to a large degree, an ideal way of understanding the behaviour of a neuron because it shows when neurons spike, and it is widely believed that spiking is how most information is passed between neurons. Electrodes have been used for neural recordings for quite some time (e.g. Hubel and Wiesel, 1962; Fuster and Alexander, 1971). The use of electrodes and other single unit recording technology has improved and continues to improve. Still with current technology, it is only possible to measure about 1000 neurons simultaneously. Improved techniques, for instance optical techniques (see Wallace and Kerr, 2010, for a review), are expanding the number of neurons that can be simultaneously recorded. These single unit recording techniques can work *in vitro* or *in vivo*, enabling measurement of neurons in behaving animals, and are leading to a better understanding of CA dynamics. Unfortunately, these measurements are intrusive, so only in rare cases can performing neurons be measured *in vivo* in humans using single unit recording. Moreover, the relatively small number of neurons recorded means that it is difficult to record relatively precise dynamics of a large number of neurons in a living animal.

One piece of evidence for population coding using single unit recording involves neurons in the macaque visual cortex that show population coding of shape (Pasupathy and Connor, 2002) with specific neurons spiking rapidly when specific curves are present at particular angles from the centre of the object. Note that, in this case, firing behaviour is directly linked to environmental stimuli. Several neurons respond to the same feature, but the sampled neurons collectively allow a vast number of objects to be uniquely identified. Population coding allows neurons to respond differentially to a range of stimuli, and collectively they can accurately represent a range of stimuli despite the noise of neural response (Averbeck et al, 2006a).

Population coding supports a diverse range of group behaviour. The simplest behaviour is that a set of neurons all respond to a particular stimulus, this is a group of grandmother cells. In a more complex form, two groups of cells can code for specific features, say lines of two given angles; activation of varying numbers or varying degrees of each set can represent an intermediate feature (Hubel and Wiesel, 1962). The degree of overlap is a topological issue (see Section 3.2) that enables a range of behaviour including continuous valued categories (see Section 3.3).

There is a vast amount of evidence of population coding. Many of the papers described below show population coding (e.g. Sakurai et al, 2004; Nicolelis et al, 1995; Sigala et al, 2008, and many papers mentioned in Section 3.1). These are just some of the thousands of experiments showing sets of neurons with elevated spike rates during specific cognitive events.




Fig. 1 A Reverberating Circuit(reproduced from (Hebb,1949)).

While a range of neurons may respond to a particular stimulus, it is not entirely clear if all responding neurons contribute to the perception. There is evidence (Purushothaman and Bradley, 2005) that only the most active neurons participate. If only some of the firing neurons are involved in the perception, the population coding uses fewer neurons.

Population coding does exist without CAs. For example, neurons in the primary visual cortex use population coding but their elevated activity does not persist when environmental stimulus ceases (Hubel and Wiesel, 1962, though see Section 3.1.2). The connection between these visual cortical neurons may be insufficiently strong to support reverberation. There does seem to be some confusion in the literature about this issue with, in some cases, the term Cell Assembly referring to population coding without reverberation (e.g. Liebenthal et al, 1994). In this case, the authors would prefer the term neuronal ensemble, ensemble coding, or indeed population coding, all of which refer to both reverberating and non-reverberating sets of neurons.

2.3 Persistent activity

The second principle of the CA hypothesis is persistent activity; the CA reverberates (Sakurai, 1998b). A CA can remain active after environmental stimulus ceases, and this reverberation enables the CA to act as an STM (see Section

4). Figure 1 reproduces Hebb's illustration of this reverberation. In the figure, the nodes represent neurons, and the numbers represent order of neural firing, and is meant as an illustration. In measured CAs, this figure is an extreme simplification; hundreds of neurons may spike in any given millisecond, each neuron may spike many times in a second, and any two neurons may not have the same order of spiking in subsequent spikes.

The standard mechanism for showing persistence is neural firing measured by single unit recording, and there is extensive evidence of persistence. One review shows sustained activation in a range of brain areas including premotor and inferior temporal neurons (Funahashi, 2001). For example, prefrontal neurons of monkeys showed elevated firing, associated with particular tasks, which persists during delay portions of those tasks (Assad et al, 2000). Persistent activity has also been shown in rat motor cortex that correlates in task specific ways (Isomura et al, 2009). Rhesus monkeys display persistent neural firing that correlates with duration (Sakurai et al, 2004), so sets of specific task general neurons maintain specific durations.

Firing behaviour has been linked to cognitive states. For example, rhesus monkeys show elevated neural spiking in the prefrontal cortex and medialis dorsalis for times up to a minute while attending and recalling a place for reward (Fuster and Alexander, 1971). Sets of neurons respond to the cue, while other sets respond during the delay, and a third set responds during both cue and delay. As this is not hand or item specific, this suggests sets of neurons that support rehearsal.

There is a range of neuron types, though a discussion is beyond the scope of this paper (but see Klausberger et al, 2003, for example). While typical neurons (e.g. pyramidal) fire relatively regularly given a constant input, other neurons are bursty going through phases when they fire frequently, and then not at all. Bursty neurons oscillate on scales of less than once per second as shown by single unit recording, and this type of persistent activity occurs throughout sensory, motor and association cortices (Steriade et al, 1993). In vitro recordings of rat cortical cells showed persistent firing of stimulus specific cells lasting on the order of 10 seconds (Larimer and Strowbridge, 2009). Many cortical and subcortical regions exhibit this type of behaviour (Seamans et al, 2003). Similarly, spiny neurons change between up and down states at a similar rate, and particular neurons have correlated changes (Stern et al, 1998).

There are a wide range of techniques that, unlike single unit recording, are non-invasive, for example fMRI, MEG, and EEG are all widely used. Unfortunately, none directly measures the behaviour of single neurons; instead they measure the behaviour of a large number of neurons. For example, fMRI measures blood flow in a region of the brain which corresponds to elevated neural firing rates in brain areas showing the aggregate behaviour of many thousands or millions of neurons. As these techniques are non-invasive, humans can be measured relatively easily. fMRI has been used to show persistent activity in human prefrontal cortex that corresponds with the duration of a working memory (Curtis and D'Esposito, 2003). Persistent firing is indicated by fMRI when people locate the source of a sound (Tark and Curtis, 2009). This occurs

in the frontal eye field even though the location is behind the subject, so there is no direct sensory stimulus.

Transcranial magnetic stimulation is a non-invasive procedure that raises activity (neural firing) in specific areas of the brain. If it is done before stimulus, the stimulus is recognized faster, but if it is done during processing, processing is interrupted (Silvanto and Muggleton, 2008). Before stimulus, this is consistent with raising all neural activity in the area which speeds recognition by helping the winning CA ignite faster. During processing, the increased activity causes competing CAs to ignite inhibiting already ignited CAs.

2.4 Correlated activity and functional connectivity

There is evidence that neurons that react to the same stimulus or cause the same action fire synchronously. Modellers have used the saying “neurons that wire together, fire together”. If two isolated neurons fire at an elevated rate, there is no particular reason for a relation between the firing times, but as neurons that are directly connected affect each other’s firing, their mutual firing timings may be correlated. If a large number of neurons are responding to a particular stimulus, they should have correlated firing. Simulation evidence shows that these synchronously firing sets of neurons can be learned using Hebbian learning (Levy et al, 2001).

Correlated activity is not one of Hebb’s principles. von der Malsburg (1981) introduced correlated activity as an extension of the basic CA theory of the day, proposing it as a solution to the binding problem (see Section 5.2.2). Extensive mathematical, simulation, and biological evidence indicate that neurons involved in processing the same percept have correlated activity. This enables synchrony to act as evidence that a particular set of neurons population code a concept (e.g. Bressler, 1995; Sigala et al, 2008); if they are persistent, and the concept has been learned, synchrony is evidence that particular neurons are firing in support of the same cognitive event.

A review article indicates that correlated firing occurs in virtually all brain areas, and can be used to coordinate activity across areas (Bressler, 1995). For example, it has been shown that frontal cortical neurons in monkeys fire synchronously in response to a particular object (Abeles et al, 1993). This synchrony is statistically significant, but not precise to the millisecond level.

Another example shows neurons in the monkey prefrontal cortex coding for task order. Neurons had highly correlated elevated firing rates for particular events in a sequence, and different elements had an orthogonal set of neurons (Sigala et al, 2008). Similarly, there is highly correlated activity in neurons in the prefrontal cortex showing associations between sensory modalities (Fuster et al, 2000). Neural firing in the rat anterior cingulate cortex was measured using multiple single unit recordings, and the neurons tracked each aspect of the task by entering collectively distinct states. Neurons were correlated with specific other neurons in a task and state dependent manner. This indicated that the correlated neurons were members of particular CAs (Lapish et al, 2008)

because they population coded for a concept, persisted, and were learned; that is, CAs were involved in higher order cognitive processing.

Externally measured neural firing in the rat hippocampus can be used to determine the location of the rat, and there is synchronous firing of neurons associated with a particular location. However, firing synchrony is statistically more significant than the correlation of single neurons with location, so the best determinant of a neuron firing is not location but the firing of particular other neurons (Harris et al, 2003).

The evidence for synchronous firing is so significant that it is has been called functional connectivity (e.g. Gazzaley et al, 2004). Though the synapse between two neurons cannot be easily traced, they are functionally connected because they fire synchronously. Synchronous firing has been called a signature of assemblies of cells as components of a coherent code (Singer et al, 1997).

Another link between CAs and synchrony is that percepts can be induced by synchronous presentation (Usher and Donnelly, 1998). When an ambiguous stimulus is presented, flashing a subset of the stimulus at 16ms intervals leads the subject to label the stimulus as an element of the category consistent with the subset. The subjects do not notice the flash and report a constant, ambiguous, stimulus. The visual psychophysics experiments of Usher and Donnelly (1998) are claimed to demonstrate that a global, cortically based, visual binding and segmentation mechanism is sensitive to stimulus asynchronies as low as 16ms. They attribute the effect to synchronous neural activation representing spatially separated stimulus elements. Although Usher and Donnelly admit that further research is "required to test and reveal the nature of the synchrony-binding mechanism", asynchronies of such short duration suggest that if their mechanism is CA based (they do not discuss CAs) then a CA architecture that can comfortably handle differences of a few milliseconds is necessary and the obvious candidate mechanism is the synchronous neural firing of CAs that maintains the genuine asynchrony in the laboratory stimulus.

A third link between synchrony and CAs is shown using voltage sensitive dyes. These dyes respond to the voltage of cell membranes and thus can show the activity of neurons; they are imaged using sensitive fast cameras providing a picture of the membrane potential of many neurons in a few mm square area, evolving rapidly ($< .1\text{ms}$) (Shoham et al, 1999). Experiments using these dyes showed neurons in early visual areas spontaneously and synchronously firing in the absence of visual stimulus; the firing is correlated with firing of stimulus specific response (Grinvald et al, 2003). Note that synchrony in this case is not linked with persistence, but population coding.

Synchrony can also support interactions between CAs. Evidence for CAs bound (see Section 5.2.2) by temporal synchrony has been reviewed (Freiwald et al, 2001); for example, neurons responding to a single moving bar respond in synchrony as a gamma wave (Gray and Singer, 1989). It has been shown that neurons across different brain areas collaborating in a task are synchronized, but are unsynchronized after the task (Roelfsma et al, 1997).

Since binding can induce synchrony, if two neuron, or sets of neurons, are firing synchronously it does not mean they are part of the same CA. However,

it does imply they are in some sense working together. However, the more circumstances that the neurons fire synchronously, the more likely the neurons are to be part of the same CA.

Brain waves come in different forms (e.g. theta waves) and emerge from populations of neurons firing roughly synchronously. These oscillations are related to a long standing theory of brain wave propagation (Beurle, 1956). A review describes how different frequencies have different cognitive effects (Buzaski and Draguhn, 2004), with slower frequencies used to recruit larger groups of neurons. Theta waves are slower than gamma waves, and it has been proposed that they are used in working memory operations (Auseng et al, 2010). Theta waves also support internal state consistency when rat hippocampal ensembles that represent location switch from the representation of one location to another on the theta wave (Jezek et al, 2011).

One way of looking at CAs is that all neurons contribute relatively independently. Another way is that waves of firing create a form of short-term memory. This dynamical CA hypothesis (Fujii et al, 1996) fits in with a synchrony theory (von der Malsburg, 1981).

The variety of waves provide different mechanisms for supporting individual CAs and for linking CAs. These linking mechanisms are not well understood, but it seems these waves are a measurable property of interactions between CAs.

2.5 Learning data

The third principle of the CA hypothesis is that CAs are learned. Hebb used the prevalence of learned response in making decisions to derive this principle. The importance of learning fits in with the long-standing recognition of the importance of learning to the functioning cognitive being (James, 1892). Hebb's famous learning rule was introduced in this context, and subsequently, substantial physical evidence that this type of learning occurs has been found.

Hebb noted the importance of learned behaviour. For example, eye blink to a moving object is a learned response, though it is practically immune to extinction (Riesen, 1947). Later learning is based on earlier learning with early learning tending to be permanent. More recent neurophysiological evidence of neurons becoming tuned to specific categories has been unearthed (see Keri, 2003, for a review).

There is substantial evidence of change in neural response in adult animals that is correlated with cognitive learning. For example, recording of two areas of rat motor cortex showed increasing correlation of particular neurons as a task was learned (Laubach et al, 2000). This enabled prediction of success or failure of the animal performing the task from the behaviour of recorded neurons alone. Another study of adult monkeys shows recorded temporal cortical neurons becoming tuned to specific category relevant features while learning categories, and humans behaved similarly on the tasks (Sigala et al, 2002).

Similarly, gerbil primary auditory cortical neurons become tuned to rising or falling modulated tones (Ohl et al, 2001).

Another type of evidence comes from simultaneous activity, where learning induced functional connectivity persists for at least days in rats that have learned particular tasks (Baeg et al, 2007). This connectivity increases rapidly during early phases of learning, and then approaches an asymptote.

There is evidence that the hippocampus plays a special role in the consolidation of memories (Sutherland and McNaughton, 2000). Also stimulation of *in vitro* hippocampal cells creates temporally separated CAs (Olufsen et al, 2003). This implies that many CAs must include neurons in several cortical areas including (at least in early stages of learning) the hippocampus.

Items learned under high arousal persist longer, but are more difficult to access initially. This is consistent with a CA being formed; the CA formed under high arousal is difficult to access initially since many of its neurons are in a refractory state (Kleinsmith and Kaplan, 1963). One computational model explains this via a CA based simulation (Kaplan et al, 1991).

While an even brief review of the literature on sleep is beyond the scope of this paper, sleep does play a role in learning. For example, hippocampal place cells in rats that are highly active during exploration show higher firing rates during the next sleep cycle (O'Neill et al, 2008). Neurons that co-fire in prefrontal cortex and hippocampus during activity, co-fire again in subsequent sleep and rest periods (Nitz and Cowen, 2008; Maquet, 2001; Sutherland and McNaughton, 2000).

There is a great deal of evidence of learning in adults, but less on earlier learning. Behavioural research suggests that the ability to recognize objects improves throughout childhood and adolescence (Nishimura et al, 2009).

There is neural evidence using fMRI that shows the early visual areas complete their development for retinotopic mapping and contrast sensitivity by age seven, while later areas for place, face and object selection take longer to develop (Grill-Spector et al, 2008). One study uses fMRI to show that the area of cortex devoted to face and place grows from children to adolescents and on to adults while the corresponding cognitive ability also grows (Golari et al, 2007). This does not occur for object recognition implying the child (7 year old) has fully or nearly fully developed object recognition capacity.

Developmental studies have the problem of time. It takes longer to run a developmental study compared to a relatively brief study to learn a particular object or task. Overcoming this problem, one study of visual object recognition of children up to age one uses event related potential as measured by EEG. The study examined brain activity when presenting novel and familiar toys and faces. Statistically significant changes occurred throughout the year, though some types of activity plateaued; e.g., the mid-latency negative component (assumed as an obligatory attentional response) plateaus at eight months (Webb et al, 2005). This indicates some early brain changes that stabilize.

To this point, this section has used a converging evidence approach with neural studies showing that some sets of neurons use population coding, some persist, and some are learned. One study shows a set of neurons that does all

three (Freedman and Assad, 2006). Neurons in Macaque middle temporal and intraparietal areas were recorded showing that neurons responded with population coding to the movement of objects in two particular sets of directions. These neurons persistently fired during a delay. When the task changed so that the monkeys had to respond to new sets of directions, the population of neurons reorganised so that they responded to the new categories. Similarly, primate prefrontal and caudate nucleus neurons population code for expected reward or punishment (Histed et al, 2009). Many fire persistently throughout the delay period, and as the task changes, the neurons' behaviours change.

2.6 The standard model and variants

The authors consider the category of CAs a natural kind, and not one defined by necessary and sufficient conditions. Thus a model can be used to indicate the features that are central to CAs. However, there really is no such thing as the standard CA model, but Hebb's original idea is a good starting point. CAs are typically considered to consist of neurons that maintain persistent activity through strong connections with all neurons being roughly similar, largely as described by Hebb (1949).

A first modification of the model was the introduction of inhibition (Milner, 1957). Hebb did not have solid evidence of inhibitory neurons or synapses, so left inhibition out of the model. With subsequent biological support, inhibition was added to the model, and supports competition between neurons or CAs.

Neurons are firing constantly, and this background activity has been accounted for by Hebb's standard theory. Background activity is not usually sufficient to ignite a CA, but once sufficient additional firing has begun, activity in the CA passes a critical threshold, the CA ignites, and can then persist.

The authors also consider synchronous firing as part of the standard theory. While not part of Hebb's original theory, there is broad evidence for synchronous firing in CAs, and that it occurs is widely agreed in the scientific community. However, the degree to which synchronous firing is used to dynamically bind CAs is not widely agreed.

A neuron that never co-fires with the neurons in a CA is not part of that CA. For that matter, one that is always out of synchrony is also not part of the CA.

In addition to a widely agreed standard model, there are several extensions. There is biological evidence for three other mechanisms for sustaining persistent activity: synfire chains, bistable neurons (Durstewitz et al, 2000), and short-term potentiation (STP). In synfire chains, one set of neurons activates another, then the second set activates a third, and so on, until the initial set is reactivated. These sets of neurons fire synchronously, leading to a chain of synchronously firing neurons (Ikegaya et al, 2004). This differs from the standard model because the CA is broken into independently firing subsets.

Bistable neurons are another less traditional form of persistence. Mentioned in Section 2.3, some neurons move from states of rapid firing (up states), to

states of no firing (Stern et al, 1998; Steriade et al, 1993). In another study, unstimulated neurons moved to up states spontaneously, and synchronously (Cossart et al, 2003); this behaviour is consistent with attractor basins² that can be used for STM. This bistability differs from the standard model because neurons move into states where they fire repeatedly without input; in the standard model the neurons require input to fire. Rested bistable neurons will shift to an up state with sufficient input.

A third alternative to CAs as STM items is storage via STP, rapid synaptic modification that decays (Mongillo et al, 2008; Fusi, 2008). The memory persists as long as the synaptic changes are sufficient to support subsequent reactivation.

To a large extent, synfire chains, bistable neurons, and STP fall within a loose definition of CAs. All use a small set of neurons for a concept, support persistent activity, and can be learned. Of course all add a different flavour to CAs. A CA based on synfire chains would have more cyclic, wave like behaviour; bi-stable neurons require less input to a neuron to maintain persistence; and STP allows CAs to form rapidly compared to the standard model.

There is an alternative to CAs that has some support in the academic community. This model uses a holographic reduced representation so the same active concept can be represented by two entirely separate sets of neurons. A set of neurons represents a concept by its firing pattern (Plate, 1995). This is consistent with population coding, but the firing can pass to another set of neurons, and may never return to the original set. This passing of firing is not entirely consistent with self-sustaining persistence, however it is related to synfire chains.

Proponents (e.g. Eliasmith and Thagard, 2001) of this mechanism show that representations of different concepts can be combined, thus solving the binding problem (see Section 5.2.2). Simulated neurons can be used to represent concepts and the concepts can be bound using these mechanisms. However, the authors are unaware of any biological evidence supporting this type of binding.

A further problem with holographic representations is symbol grounding (Harnad, 1990); how does the neural firing pattern come to represent the particular concept? With CAs, the neurons have been grounded by interaction with the environment (see Section 6). With holographic representations, different sets of neurons need to represent the same concept, and it is not clear how all of these sets might come to represent the same concept. Moreover, if one concept is, for example, associated with a second concept, all of the different sets of neurons in the holographic representation will need to be associated with all the sets of neurons that represent the second concept. In short, learn-

² Statistical mechanics can be used to describe CAs, with a set of firing neurons being an attractor basin. The Hopfield model (Hopfield, 1982, and see Section 5.1.3) is a particularly good example of attractor basins with a particular set of neurons firing and then continuing to fire in response to a particular input. The firing pattern moves from the initial input state, down an energy slope to a new firing state, which has attracted the activity.

ing is the main difference between CAs and holographic representations. It is not clear how holographic representations could be learned.

Holographic representation and associated binding may be compatible with individual neural behaviour. However, the authors feel that it is at best as a complement to the CA model instead of a replacement; it might perhaps be used as a mechanism to create firing patterns of bound CAs in, for example, language processing.

2.7 Cell Assemblies are all but certain

The evidence referenced in this section strongly suggests CAs do exist in the brain. The prevalence of population coding throughout the brain has been verified; voltage sensitive dye techniques show this population coding, as do single cell recordings. Evidence of persistence from single cell recordings includes both persistence of regular spiking, and persistence of up bursty states; other imaging techniques show persistent firing in areas, and transcranial magnetic stimulation can disrupt cognitive behaviour by dynamically modifying neural behaviour in specific small brain regions. Functional connectivity indicates that CAs exist by showing that particular neurons behave together. Finally, though more difficult to analyse, learning data shows the firing of neurons in response to particular cognitive states becomes more correlated over time.

The evidence for CAs is compelling, and many readers will find the support for CAs unsurprising. In many fields, it is already assumed. For example, persistent neural activity is currently being used to indicate working memory (see Section 4).

None the less, the scientific community will be unable to fully confirm the existence of CAs until all the neurons in a functioning brain can be dynamically monitored constantly. Even then it would take years to understand developmental behaviour. Consequently, it is possible to argue that CAs do not exist.

Even though there is substantial evidence for the existence of CAs, the extent and behaviour of CAs is still not well understood. The next section shows the broad extent of CAs and some of their functions.

3 Cell Assembly Details

The existence of CAs does not show that CAs are the basis of all concepts. Of course, the term concept itself is ill defined, roughly meaning a mental unit. Typical concepts include things represented by concrete noun symbols such as *dog*, but the authors would expect animals that did not use symbols, but had categories for dogs (e.g. a macaque or pre-verbal infant familiar with dogs) would have a category for *dog*. Similarly, abstract nouns (like *hope*) and verbs (like *run*) also label categories. Instances of categories (e.g. *Dr Hebb*), and episodes (e.g. *today's lunch*) are also concepts. Procedural tasks (e.g. *hitting*

a tennis ball) might not be concepts, though once labelled, they probably do become concepts with the explicit concept associated with the procedure.

Many questions about CAs remain unanswered, and several are addressed here. The location of CAs in the brain is addressed in Section 3.1. The size and shape of a CA is discussed in Section 3.2. CAs have three primary and complementary functions; they categorise input, they are LTMs, and they are STMs, but CAs are also involved in a range of other processes and some of these processes are discussed in Sections 3.3 and 4.

3.1 Location of CAs

Population coding is common throughout the brain, but persistence is not. That is, neurons may stop firing once the stimulus ceases. However, in higher order areas (most of the cortex), persistence allows cognitive states to persist. Once information is in the brain, it can remain active via CA reverberation.

The cortex can be divided into areas, and one commonly used division is the Brodmann areas (Figures 2 and 3); evidence is provided below regarding the presence of CAs in different cortical areas. In almost all cortical areas, there are neurons that population code for specific features, and that persist without external stimulus, implying that there are CAs in all cortical areas except perhaps primary visual cortex. It is less clear whether there are CAs in subcortical areas, but there is evidence for CAs in some.

3.1.1 Frontal and temporal lobes

Much of the frontal lobe is topologically distant from the sensory and motor interfaces, with a chain of several intermediate layers of neurons needing to fire to cause frontal neurons to fire caused by the environment. Among a range of other functions, the frontal lobe is involved in working memory tasks (see section 4). There is substantial evidence for CAs in all frontal areas. For example, the prefrontal cortex generates persistent activity that outlasts stimulus (Wang et al, 2006); single cell recordings of ferret medial prefrontal cortical neurons shows persistent activity. Comparisons to visual cortex indicate that the cellular features, synaptic properties, and connectivity properties of these prefrontal neurons favour persistent activity.

Primary motor cortical neurons (Brodmann area 4 BA4) have CAs. Electrode studies of owl monkey's primary motor cortical neurons (BA4), and dorsal premotor cortical neurons (BA6) were used to guide a robotic arm (Wessberg et al, 2000); these neurons persisted throughout the action. Recordings of neurons in motor areas of rats learning a reaction time task show elevated activity during the task, with most of the recorded neurons active (Laubach et al, 2000); analysis of firing can predict success or failure on particular instances of the task; and behaviour is learned. A human unable to move or sense his limbs for several years had electrodes implanted into his motor cortex. The monitored activity of the neural ensembles enabled control of computer devices




Fig. 2 Brodmann Areas, lateral surface of the brain (courtesy of Mark Dubin)

and via them external devices (Hochberg et al, 2006). The monitored neurons individually behaved quite differently at different times implying they participate in multiple CAs. Rhesus monkey motor cortical neurons are broadly tuned to the direction of arm movement with many neurons participating in a given action; when the action is delayed, the same neurons remain active (Georgopoulos et al, 1986). This shows BA4, which is close to the motor interface, has sound evidence for CAs, and these CAs need to persist to complete the appropriate action.

A learned response of frontal eye field (BA8) neurons of monkeys was associated with gaze target visual features (Bichot and Schall, 1999). These neurons fired persistently while searching for the associated target. Prefrontal neurons (BA9) of monkeys showed task specific firing during execution and delay portions of tasks (Assad et al, 2000). Rhesus monkey prefrontal cortical neurons stored action sequences in CAs (Averbeck et al, 2006b). Prefrontal cortical neurons (BA9 and BA10) also fire during the delay portion of a task with specific neurons correlated with specific visual cues (Quintana and Fuster, 1999). PET scans showed persistent elevated activity in premotor cortex (BA6), and frontal cortex (BA10) when the subject had learned to associate a visual stimulus with an auditory stimulus (McIntosh et al, 1998).

Orbitofrontal neurons (BA11 and BA12) of rats, measured by single unit recordings, exhibit persistent firing specific to reward or punishment (Gutier-

rez et al, 2006). Moreover, they fire in anticipation of the reward or punishment. Ventrolateral prefrontal cortical neurons of the macaque correspond to the inferior frontal gyrus in humans. These macaque neurons population code and respond to semantically specific macaque calls and do so persistently (Romanski et al, 2004). Macaque rostral inferior neurons (a homologue of human BA44) fired in a delay period before the macaque grasped specific objects, with neurons population coding specific actions (Gallese et al, 1996). Dorsolateral prefrontal cortical neurons (including BA46) and ventrolateral prefrontal cortical neurons (including BA45 and BA47) of monkeys responded in a cue task with population coding and persistently for abstract rules (Wallis et al, 2001).

Like the frontal lobe, the temporal lobe is also topologically distant from the sensory and motor interfaces. Though many areas contain CAs that persist, CAs in this lobe, in general, are driven more by environmental stimuli than in the frontal lobe. None the less, persistence after external stimulus has ceased is evident throughout.

There is strong evidence derived from single neuron recording that inferior temporal cortex (BA20) has CAs. One test of this area showed that neurons population code for specific images and specific features (Desimone et al, 1984). Face specific neurons were found in the fusiform gyrus (BA37), with neurons responding persistently, but the image remained present. In macaques performing a shape matching task, these neurons population code for specific shapes and spiking persists (Gochin et al, 1994). Similarly, shape specific neurons fire during delay in a paired association task (Naya et al, 1996).

Single cell recordings of middle temporal neurons (BA21) show that they collectively encode direction of motion of visual stimuli. Moreover, during a delay task, neurons have elevated activity showing persistence (Bisley et al, 2004). This area is also called V5 (a vision area) and it responds to visual stimuli, such as direction of motion, with population coding (Maunsell and Van Essen, 1983). In macaques, these neurons are broadly tuned to moving visual stimulus, and fire persistently and synchronously with neurons responding to different stimuli firing asynchronously (Kreiter and Singer, 1996). More recently, persistence beyond environmental stimulus has been noticed (Ranganath et al, 2005).

Responses from human neurons of the superior temporal gyrus (BA22) and nearby areas were explored before surgery for epilepsy. These neurons showed persistent activity for up to two seconds in response to auditory clicks (Howard et al, 2000). This is a relatively rare example of a single unit study of the human brain, and is strong evidence for human CAs.

Piriform cortical neurons (BA27) of rats, measured by single unit recording, fired persistently with population coding in response to particular odours (Rennaker et al, 2007). In the piriform cortex (BA27), Zelano et al (2009), using fMRI, found sustained activity associated with olfactory working memory. Similarly, entorhinal (BA28 and BA34) and hippocampal neurons of behaving rats fire persistently and synchronously (Chrobak and Buzaski, 1998), measured by single unit recording. Further studies showed that groups of neurons in this area fired with graded responses (Egorov et al, 2002).

Neurons from the parahippocampal gyrus (BA35 and BA36), hippocampal formation, rhinal cortex and inferior temporal cortex (BA20) of monkeys performing a delay task showed population coding and persistent firing during the delay period (Riches et al, 1991). Hippocampal place cells (BA35) can be used to predict the location a rat will move to in the next theta cycle (Dragoi and Buzsáki, 2006). Macaque temporal pole neurons (BA38) show population coding for a complex stimulus, including presentation of opposed halves of that stimulus, persisting for seconds after the stimulus (Nakamura et al, 1994).

Marmoset auditory cortical neurons (BA41 and BA42) population code for tones, and fire persistently (Bendor and Wang, 2008). Postsubicular rat neurons (BA48) population code for head direction. Moreover, persistent firing has been induced by direct current injection (Yoshida and Hasselmo, 2009).

3.1.2 Occipital and parietal lobes

The occipital and parietal lobes contain primary and higher sensory areas. While some researchers feel that even neurons in primary sensory cortices are parts of CAs (Harris, 2005), there is evidence of persistence in higher sensory areas, but less in primary sensory areas.

The occipital lobe consists of areas for early visual processing. There is evidence for population coding in primary (BA17) (e.g. Hubel and Wiesel, 1962), secondary (BA18) (Poggio and Fischer, 1977, e.g.) and association (BA19) visual cortices (e.g. Fischer et al, 1981), though there has been little evidence of persistence with single unit recording. However, using fMRI, BA17 shows increased activation when subjects recall visual stimulus (LeBihan et al, 1993), and functional connections during delay tasks in visual cortices (BA18 and BA19) (Gazzaley et al, 2004). Similarly, neurons in the primary visual cortex of blind subjects are activated during Braille reading (Sadato et al, 1996). Such studies suggest that there is input from higher areas to the lower visual cortical areas, and that this input is able to fire neurons in these areas.

In the parietal lobe, there is evidence for persistence in all areas. Somatosensory neurons (BA5 and BA7) of rhesus monkeys fired persistently in expectation of executing a task (Quintana and Fuster, 1999). Owl monkey posterior parietal neurons (BA7) were used to guide a robotic arm (Wessberg et al, 2000); these neurons persisted throughout the action. Parietal neurons (BA7) of rhesus monkeys fired while they considered and performed viewer and object centred tasks (Crowe et al, 2008); in all cases, state was population coded, and these neurons fired persistently. Single unit recordings of neurons from the supramarginal gyrus (BA40) of macaques show population coding and persistence associated with visual saccades (Nakamura et al, 1999).

fMRI studies show elevated activity in the angular gyrus (BA39) in response to linguistic tasks (Booth et al, 2004). Single unit recordings showed elevated firing rates in the parietal reach region during reach planning and this was population coded for direction (Scherberger et al, 2005). Similarly, fMRI showed elevated activity in the parietal operculum (BA43) in tacitly modulated visual tasks (Macaluso et al, 2000).

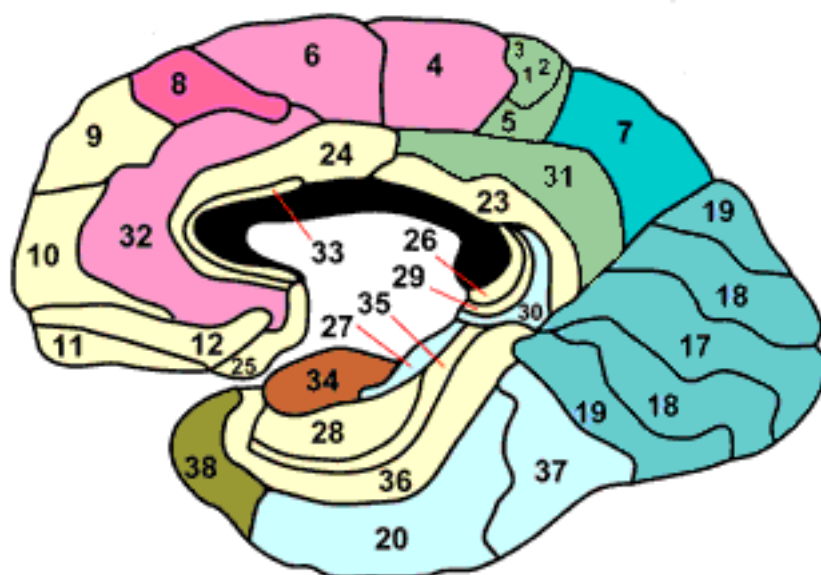


Fig. 3 Brodmann Areas, medial surface of the brain (courtesy of Mark Dubin)

Suites of rat primary somatosensory cortical neurons (BA1-3) respond to whisker stimulus with latencies increasing in deeper cortical areas (Moore and Nelson, 1998). While these neurons did not fire persistently after the stimulus, anterior parietal cortical neurons (BA1-3) of rhesus monkeys that showed elevated activity during tactile tasks showed elevated activity during delay tasks (Zhou and Fuster, 1996). These single cell recordings show a form of short-term haptic memory in early sensory processing.

The emerging view is that sensory cortex involved in relatively early processing are also used for short term storage (Pasternak and Greenlee, 2005). Thus CAs involve neurons in early sensory areas in parietal and occipital lobes.

3.1.3 Limbic and insular lobes

All Brodmann areas in the limbic lobe exhibit population coding and persistent neural firing. Suites of macaque anterior cingulate cortical neurons (BA24) responded persistently and correlated with cognitive phenomena such as search, order, and reward anticipation in a sequence learning task (Procyk et al, 2000). Anterior cingulate cortical neurons (BA33) population coded for stimuli during the delay period of a classical conditioning experiment (Kuo et al, 2009).

Rat anterior cingulate cortical neurons fire persistently during a delay task (Lapish et al, 2008). Monkey cingulate motor neurons (BA23 and BA6) exhibit persistent firing during the delay period of a delay task (Crutcher et al, 2004). Rabbit posterior cingulate cortical neurons (BA29) fired persistently during a delay period in a conditioned stimulus task (Talk et al, 2004). Rat retrosplenial cells (BA26, BA29 and BA30) fired in relation to one or more of the variables

location, direction, running speed, or movement. Some fired in anticipation, implying ignition before external stimulus (Cho and Sharp, 2001).

Macaque posterior cingulate cortical neurons (BA31) population code for direction of eye saccade, and firing persists beyond the movement (McCoy and Platt, 2005). Macaque pregenual cingulate cortical neurons (BA32) fired persistently with population coding for specific tastes (Rolls, 2008).

Using fMRI, Siegel et al (2006) noted varying levels of sustained reactivity to emotional stimuli in the subgenual cingulate cortex (BA25) and amygdala. Single cell recordings showed elevated firing rates in many neurons in BA25 when a macaque went to sleep (Rolls et al, 2003). Other neurons had elevated firing for novel visual stimulus, which decreased as the stimulus was repeated. As these elevated rates were less than 5 spikes per second, this is merely suggestive evidence for CAs in this area.

The insular lobe is small with only three Brodmann areas. Population coding is shown in all three, but persistence is less evident. Rat gustatory cortical neurons (BA13 and BA14) form distinct, population coded CAs that respond to particular stimulus by undergoing coupled changes in firing rates over several 100ms (Katz et al, 2002). Parainsular auditory cortical neurons (BA52) of squirrel monkeys responded to specific tones, but there was insufficient evidence to see if they persist for a significant time after the stimulus (Bieser and Muller-Preuss, 1996).

This paper's use of Brodmann areas to divide the cortex is intended to show that CAs are prevalent throughout the cortex, except perhaps BA17. Clearly, Brodmann areas can be subdivided, and other systems for partitioning the brain exist. Moreover, in most cases, homologues of other animals have been used to infer neural firing in the human brain. In some cases, e.g. premotor areas, these areas are quite similar, but in language areas for example, the function has changed. However, these invasive studies can be aligned with imaging studies such as fMRI, and other techniques.

Most of the evidence presented here of CAs in cortical areas is based on population coding and persistence, however, the evidence does not exhibit learning. Methodologically, it is more difficult to find learning in particular neurons than to record spikes. Single unit recording devices need to be placed before the learning is done, and then the learning has to affect those neurons being measured. There is evidence of learning (see Section 2.5), but this paper has not provided references to learning in each Brodmann area, assuming that the behaviour has been learned at some time.

The existence of CAs in all or almost all higher cortical areas aligns well with Hebb's idea of CAs as the basis of concepts. When a concept persists in the brain after the stimulus ceases, an active CA supports this STM. It is possible that some concepts are represented by some other neural mechanism than CAs, however the authors are unaware of any sound evidence of such a mechanism concept pair.

3.1.4 Subcortical areas

While it is unlikely that CAs exist in every area of the brain and the nervous system, there is evidence of CAs in some subcortical areas. However, as a rough guide, CAs have less prevalence as areas get further from the cortex.

There is evidence for CAs in noncortical forebrain areas. Spiny neurons in the striatum exhibit loosely synchronised up and down state dynamics that indicate they are parts of CAs that include cortical neurons (Stern et al, 1998). Neurons in the nucleus accumbens exhibited persistent states (O'Donnell and Grace, 1995). Neurons in the rat amygdala fired persistently during fear (Repa et al, 2001); moreover this response was learned over a series of several trials.

There is population coding in the thalamus. Neurons in the cat lateral geniculate nucleus respond reliably to a particular stimulus (Reinagel and Reid, 2000). Similarly, rat posterior medial neurons fired robustly in response to whisker stimulation with neurons responding to multiple whiskers and multiple neurons responding to particular whiskers (Nicolelis et al, 1993). While head direction neurons in the thalamus fire persistently (Stackman and Taube, 1998), there is little evidence of persistent firing in other parts of the thalamus.

The cerebellum is involved in eye blink conditioning. Classical conditioning experiments show that cerebellar neurons learn to respond to conditioned stimuli, and persist during a delay period while conditioning (Christian and Thompson, 2003). There is some evidence of CAs in the hindbrain. In this case, reverberatory activity is maintained without sensory or other external stimulus. However reverberation can be maintained via sensory stimulus, as is usually the case in primary visual cortex (BA17) and with thalamic head direction neurons. It is also possible the reverberatory activity is maintained by non-sensory input to the nervous system, as may be the case with some areas involved in emotion processing.

While this section has shown evidence that CAs are in many areas of the brain, it has not shown the extent of CA involvement in those areas. It is not clear that all neurons in, for example, the forebrain are part of CAs.

3.2 Topology of a Cell Assembly

Some aspects of the topology of a CA have extensive empirical evidence. CAs overlap with particular neurons participating in multiple CAs (e.g. Georgopoulos et al, 1986). CAs cross brain areas, so that some CAs are made of neurons in multiple brain areas (e.g. Pulvermuller, 1999).

A great deal of the theoretical structure of CAs has been discussed (e.g. Wickelgren, 1999). Central to this is the overlapping coding of CAs. The idea of overlapping coding is that some, perhaps most, neurons participate in multiple CAs. New concepts composed of pieces of old concepts will contain some new neurons, though will also contain some of the neurons from the base concepts.

Overlapping encoding relates to population coding (see Section 2.2), where a concept is composed of a large set of, but not the majority of, neurons. There

is extensive evidence for overlapping encoding (e.g. Tudusciuc and Nieder, 2007; Georgopoulos et al, 1986). The degree of overlap varies from CA to CA. For example, there is evidence that task specific CAs are largely orthogonal in frontal areas (Sigala et al, 2008), while there is extensive overlap in CAs in motor areas (Georgopoulos et al, 1986). Overlapping enables more CAs, but it also supports sharing of information between CAs. If two CAs share many neurons, they are probably closely related. An interesting mathematical model includes overlap in CAs for optimal decisions (Bogacz, 2007). A similar information theoretic work suggests that less overlap is better in sensory systems (Field, 1994). Note that the earlier use of atomic assemblies does not imply that a CA is orthogonal to other CAs, merely that it is not composed of other CAs.

A range of imaging and non-imaging data suggests that CAs for words are distributed in brain areas that are semantically specific (Pulvermuller, 1999). For instance, the CAs for nouns with strong visual associations include neurons in visual cortices. Similarly, concrete nouns can be predicted via fMRI from the brain areas active when the subject thinks of them (Just et al, 2010). Synchronous activity of neurons in behaving rats (Nicoletis et al, 1995) indicates that neural ensembles across several brain areas collaborate to integrate sensory information and determine action. However, it is not entirely clear when CAs cross brain areas. Moreover, the nature of CA processing and formation across laminar levels and brain areas is poorly understood (but see Douglas and Martin, 2004, for a proposal).

Orthogonality and CAs in multiple brain areas raise the question of how many neurons are in a CA. The answer to this question has less evidence, but single unit recordings show hundreds of neurons firing at elevated levels for a remembered concept (Histed et al, 2009). This is close to the lower range of 10^3 neurons in a CA that the authors suggested while introducing principle 1. This answer also raises the question: are there different types of CAs? As proposed in Section 1.1, smaller CAs are more likely to be atomic. That is, they cannot be decomposed into other CAs. Larger CAs are likely to be composed of other CAs, and can act to support composite structures such as cognitive maps.

It also seems likely that some neurons are more central to a particular CA than others. These central neurons may fire more frequently than less central ones, and may have closer links to other neurons in the CA, so membership of a neuron in a CA can be thought of as graded instead of binary.

Other topological issues such as neural type, involvement across laminar layers, and precise neural dynamics are poorly understood. However, it is likely that topological parameters will vary between CAs and within a CA during its life time.

3.3 Function of Cell Assemblies

From a cognitive perspective, CAs fulfil a wide range of cognitive roles. As Hebb stated, CAs are LTMs and STMs. Persistent firing beyond sensory stim-

ulus is a STM. The formation, via Hebbian learning, of a set of neurons, which can persistently fire, is an LTM. CAs vary in the degree that they are stimulated by the environment or from other parts of the brain. As commonly understood, CAs are ignited, in appropriate circumstances, by sensory stimulus, and categorise that stimulus. These prototypical CAs can also be used as STMs that can be ignited by internal processes such as working memory. In current day thinking, STM, LTM and categorisation are the primary functions of CAs.

Neurons in CAs can function without persistence, firing persistently or transiently based on sensory stimulus. By the principles introduced in Section 1.1, if a set of neurons can persist or occasionally persists without a sensory stimulus, it is a CA. None the less, the CA concept is more continuous than binary. Sets of neurons that are population coded and stimulus responsive, but rarely ignite due to internal processes are to a lesser degree CAs; for example neurons in early visual areas rarely persist without sensory stimulus, and thus are less prototypical examples of CAs. While prefrontal neurons associating a tone with a colour (Fuster et al, 2000) are members of a prototypical CA, as they persist without stimulus for seconds during delay tasks, are quickly learned and population coded. Beyond these factors, CAs vary by their ignition speed, duration of persistence, size, and learning properties.

As categorisers, CAs take a vast range of possible inputs and categorise them as the same thing; e.g. pictures of different dogs, their bark, the name, or even their smell can activate one's *dog* CA. An example from the literature involves inferior temporal neurons that respond to particular shapes in a wide range of similar rotations (Logothetis et al, 1995). A similar work shows that medial temporal neurons respond to particular objects from all angles and even different pictures (Quiroga et al, 2005). This goes beyond specific objects with specific neurons firing for specific tasks (Sakurai, 1998a). Note that categories do not need to be discrete as bump sensors have been used to model continuous categories (Wang, 2001). Also, categorisation may involve systems beyond the CA, gaining continued and directed input from senses, and contextual evidence and control from higher areas.

There is evidence that episodes are encoded by CAs (Dragoi and Buzsáki, 2006). For example, rat hippocampal place neurons (O'Neill et al, 2008) fire at an elevated rate during sleep after the associated places are visited, and places that were visited in sequence co-fire. Moreover, particular rat hippocampal neurons fired when the rat was at particular locations when running a maze, and in similar order during cognitive rehearsal of running the maze (Pastalkova et al, 2008). This showed internally generated CA sequences, that also predicted task performance errors. CAs involving mouse hippocampal place cells could even predict future paths (Dragoi and Tonegawa, 2011).

Top-down and bottom-up evidence (Engel et al, 2001) can be combined to ignite a CA. Anticipation can lead to priming in a CA that in turn leads to more rapid recognition or disambiguation when a sensory stimulus is provided. When ambiguous sensory information supports the ignition of different CAs, top-down information can support a particular CA, and thus it alone of

the options ignites. Synapses connect sensory areas to higher areas and vice-versa, so information flows between sensory and higher areas, supporting the combination of evidence (Lamme et al, 1998).

Beyond categorisation, STM and LTM, CAs support a range of other activities. They support working memory (see Section 4), associative memory, figure-ground separation, and imagery; single unit recordings of human neurons responded to both the presentation of pictures and the task of imagining the picture (Kreiman et al, 2000).

CAs support associative memory (Diester and Nider, 2007). Individual CAs are associated with other CAs via excitatory synapses from neurons in the first CA to neurons in the second, and via shared neurons (Sigala et al, 2008). Evidence comes from monkey cortical neurons, where monkeys were trained to associate pairs of pictures. Recordings of neurons in the anterior temporal cortex showed one set that responded to both pairs of stimuli, and a second that responded to one; the second type of neuron had elevated firing during the delay phase of picture presentation when the associated cue was presented (Sakai and Miyashita, 1991). There is also extensive simulation of CAs supporting associative memory (see Section 5.2.1). Similarly, association leads to priming. There is evidence that semantic priming via pictures for action words leads to an effect on error rates (Setola and Reilly, 2005). “This suggests the possibility of a broader account of priming effects in terms of CAs.” That is priming, and the closely related concept of associative memory, emerge naturally from a network of CAs.

It is well known that primates can make use of cardinality and this is supported by CAs. Monkeys were trained to associate particular pictures with particular cardinalities, providing a label. Particular neurons in the prefrontal cortex were recorded and responded more or less identically to the label and cardinal category (Diester and Nider, 2007). This type of association may provide a basis for symbol grounding in humans. Symbol grounding is particularly important as it shows how symbols gain their meaning (Taddeo and Floridi, 2005), and symbols are very important in modern society and in artificial intelligence. Indeed there is growing support for a neural basis to symbols in the larger cognitive science community (e.g. Gallese and Lakoff, 2005).

Hebb discusses figure ground separation as one of the benefits of CAs. The figure is activated by CA ignition and this causes it to psychologically pop out. Some evidence shows that edge detecting neurons in BA18 respond when the edge is part of the figure but not when they are part of the ground (Qiu and von der Heydt, 2005). This is consistent with the higher order information (e.g. from V3) boosting the activity of neurons in the figure at a lower level.

One very important aspect of CAs is that they are active. CAs that encode symbols are active symbols (Hofstadter, 1979; Kaplan et al, 1990). When a CA is active, it influences other CAs, by spreading activation or inhibition to them. This enables one CA to prime many others, inhibit many others, and be involved in other types of processing simultaneously. This active processing can also involve dynamic binding (see Section 5.2.2).

It appears that motion is driven by central pattern generators (CPGs) (Butt et al, 2002). CPGs consist of neurons, and they must be triggered. Neuron firing in the motor cortex also causes actions (Brecht et al, 2004). So, neurons in the motor cortex must interact with CPGs. For example, CPGs are involved in walking. It is not clear whether the neurons in the CPG are part of the CA; are those CPGs involved in walking part of the *walk* CA? However, on a continuum, those neurons are at most weakly part of the CA. The neurons in the motor cortex may or may not persist, and those that do not are weakly part of the CA. Those that do (Laubach et al, 2000) are clearly part of the CA. Moreover, CAs, like CPGs, persist, so CAs can be considered CPGs (Yuste et al, 2005).

CAs are involved in a wide range of cognitive processes, but it is not clear when and how they collaborate. There is some evidence about their collaboration in working memory tasks (see Section 4), and evidence from simulations (see Section 5).

4 Short Term and Working Memory

While it is simple to say CAs are involved in complex psychological systems, for example, associative memory, and to provide evidence that they are, building an accurate simulated neural model of associative memory is extremely difficult, and one can argue that associative memory cannot be truly understood without such a simulation. A good simulation provides explicit details of basic processes, and shows how more complex behaviour emerges. Progress is being made on a range of cognitive tasks, but the neural implementation of these tasks, like associative memory, is not well understood. This section will make a speculative description of one complex psychological system, the working memory system, and its task specific use.

One of the basic assumptions of Hebb's CA hypothesis is that an active CA is the neural basis of an STM. In the intervening 60 years, a related form of memory, working memory, has been studied extensively. Working memory combines short-term buffers and central executive processes (Jonides et al, 1993). Psychological evidence exists that the working memory and STM are separate systems with STM supporting working memory (Engel et al, 1999).

Central executive processing is largely based in the frontal cortex, and working memory leads to behaviour that is more complex than simple STM, while retaining access to a wide range of memories stored outside the frontal cortex (see Constantinidis and Procyk, 2004, for a review of region specific involvement). Executive processing is under voluntary control, and its mechanisms are an active area of research (Miyake et al, 2000).

An active CA is an STM as evident during sensing, where CAs are ignited from the environment with little input from frontal areas. However, frontal areas can also access these areas during sophisticated processes such as memory retrieval (Buckner and Wheeler, 2001). During memory retrieval, frontal areas may begin with, for example, a goal-directed attempt to remember. This, in

essence, query is sent to other areas of the brain in an attempt to retrieve the memory. For example, parts of the medial temporal lobe are involved in storing events. If the memory is retrieved, there is enhanced firing in the medial temporal lobe (Henson et al, 1999); that is the CA ignites.

There is not universal agreement on the psychological basis of working memory, but one popular model uses the central executive, the visuospatial sketchpad, and the phonological loop (Baddeley, 2003). The sketchpad and the loop act as limited capacity buffers, and the executive controls decisions.

There is growing evidence that working memory and LTM are closely related (see Burgess and Hitch, 2005, for a review of models). In this description, STM is maintained by neural firing (a traditional active CA) or short-term potentiation, and LTM by long-term potentiation. Different working memory subsystems can be used to perform different types of tasks such as retaining lists of words or lists of digits. Of course, STMs can be retained as LTMs, but the actual neural mechanism is not entirely clear, though it seems to involve the hippocampus (see Section 6).

The sketchpad, loop, and executive need to be implemented neurally. One computationally simulated model that could apply to the loop uses neurons with highly weighted synaptic connections to model a concept (a CA), inhibitory neurons to limit activity, and short-term synaptic depression for item order (Deco and Rolls, 2005). Once a sequence is presented, it will repeat.

There are working memory limitations, and while these limits are not entirely clear, one body of evidence supports one item as the focus of attention and a few others in STM (Nee and Jonides, 2008). Simulations of these limitations use winner take all effects (see Section 5.2), e.g. with a global inhibitory neural pool (Deco and Rolls, 2005).

Neurons involved in working memory in frontal areas are responsive to a broader range of inputs, while in more lateral areas they are more domain specific (Goldman-Rakic, 1996). For example, fMRI revealed that visual areas retained fine-tuned orientation features in a delayed orientation task (Harrison and Tong, 2009). Individual recording of macaque inferior temporal and prefrontal neurons during delayed matching to sample tasks shows how the two areas interact through working memory (Miller et al, 1996). The task is for the monkey to recognise an object, hold it in memory and then respond (for a reward) when a variant of the object is later presented. Neurons in both areas respond to visual stimulus though a larger percentage of the inferior temporal neurons were responsive. Only the prefrontal neurons fired at an elevated rate during the delay period. That is the CA for the remembered object persisted in the prefrontal area, but the linked CA in the inferior temporal area did not.

The working memory system is a re-entrant network (Goldman-Rakic, 1995) based in the prefrontal and hippocampal regions, that still communicates with other areas of the brain. During a working memory task, CAs in prefrontal areas are coactive with domain specific CAs in anterior areas (D'Esposito, 2007), and the two are dynamically bound. This evidence is largely based on imaging studies, but is consistent with a CA igniting when the memory be-

comes active. It resembles the AI theory, the society of minds, where different experts (brain areas) cooperate to achieve very complex tasks (Minsky, 1986).

There are a number of working memory tasks (see Durstewitz et al, 2000, for review). One task with a computationally simulated neural model is a variant of the Wisconsin Card Sorting Task. The neural model uses a store of a single item, but this item is replaced when a dynamic gating mechanism is signalled by error feedback (O'Reilly et al, 2002). A theory of working memory was proposed (Cowan, 1988) and has evolved. This theory explicitly uses CAs as the neural basis of STM and LTM. The role of CAs in working memory models is central and understanding of CAs continues to evolve.

5 Cell Assemblies and Models

Single unit recording is the best current method for understanding individual in vitro neural behaviour though only a relatively small number of neurons can be recorded at a given time. Lesioning, voltage sensitive dye techniques, transcranial magnetic stimulation, fMRI and other imaging techniques all give us a better picture of neural behaviour, but all are incomplete. Computational models fill the gaps that these techniques miss.

While there is extensive data to support the existence of CAs in the brain, and extensive data about the behaviour of individual neurons, groups of neurons, and psychological properties that emerge from behaving CAs, CAs themselves are still not well understood. Consequently, there are a wide range of theories as to how CAs behave. Some of this is closely tied to psychological data, and other theories are more speculative. In particular, it is very difficult to see complex human behaviour (e.g. language processing, or problem solving) at the neural level in the functioning brain. There is a long history of CA models, including computational models; for example, an early simulation showed that CAs emerged (Rochester et al, 1956) in response to environmental stimuli. While they persisted, this particular model could not show that one CA could activate another, but this model suffered from the lack of inhibition as this was not in the underlying theory at the time.

5.1 Basis of model

Models are based on axioms. One relatively simple way to model CAs is to assume that they are based on some relatively simple mathematical formulas. The more popular method is to base them on neural behaviour; of course, neural behaviour is very complex and also needs to be modelled.

5.1.1 Mathematical models

One set of models of CA behaviour is based on a relatively coarse set of assumptions, modelling CAs by a set of behaviours determined by mathematical

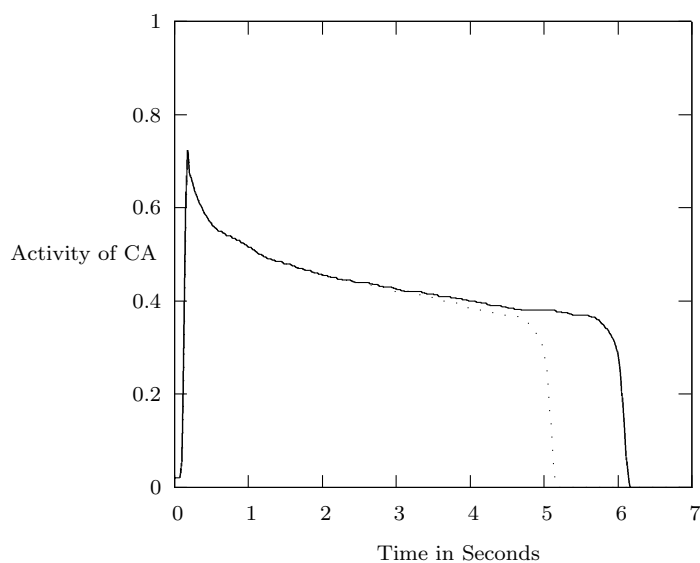


Fig. 4 Derived from (Kaplan et al, 1991). The CA is at rest for the first 100 ms. then externally stimulated for the second 100ms. This leads to the ignition of the CA with a high degree of activity, and the persistence beyond stimulation. Different CAs have a different length of persistence.

equations. The benefits of the mathematical approach include making the set of assumptions explicit, and simulations computationally inexpensive.

The TRACE equations are an example of this (Kaplan et al, 1991). This model takes advantage of neural parameters, but uses these to develop a model of a CA that has particular properties including activity, fatigue, short-term connection strength, long-term connection strength and external input. These equations are used to model learning new concepts and the duration of concepts in STM. Figure 4 shows a model ignition curve with neurons in a CA receiving some input, but then, via a cascade of firing, causing more neurons to fire. The CA persists long after the stimulus has ceased. Note that the lined CA persists longer than the dotted CA.

In another model, CAs are represented by variables including activity and connectivity that are used to drive behaviour (deVries, 2004). In this case, behaviour is binding objects to their locations.

Another mathematical model is based on graph theory. It considers neurons as nodes of a graph, and synapses as the arcs in that graph. Using a range of biologically realistic parameters for number of neurons, synaptic strength, and synapses per neuron, the number of neurons in a CA can be derived (Valiant, 2005). Moreover, this model supports associative memory (see Section 5.2.1).

5.1.2 Neural models

Any CA model based on neurons needs to define the neuron model. This section gives a very brief review of neural models (see Burkitt, 2006, for a more extensive review).

Simple or point neural models include integrate and fire neurons (Abbott, 1999; McCulloch and Pitts, 1943), leaky integrate and fire neurons (Amit, 1989), fatiguing leaky integrate and fire neurons (Huyck, 2007), and Boltzmann machines (Ackley et al, 1985). A particularly interesting point model gets a wide range of spiking behaviours from two parameters (Izhikevich, 2004). These are spiking models because the neurons simulate a spike to transfer information from one neuron to another. Models can also involve continuously valued output, e.g. the Spike Response Model (Gerstner and Kistler, 2002).

More precise simulation of neural behaviour is done by compartmental models (e.g. Hodgkin and Huxley, 1952), of which there are a wide range (see Brette et al, 2007, for a review). Compartmental models are expensive to simulate, but with fewer compartments there is less computation.

5.1.3 Topology

The neurons have to be connected, and biologically they are connected by synapses from one neuron to another. Topology involves the neurons, the connections, and properties of the connections. For simple neural models, the main property is synaptic weight, but this is relatively simplistic.

A great deal of simulation has taken advantage of unrealistic topologies where all neurons are connected to each other, and connections are bidirectionally weighted (e.g. Hopfield, 1982). This allows the system to move to a permanently stable state and the use of relatively simple statistical mechanics (Amit, 1989). This stable state, or attractor basin, is often considered a CA, and its persistence makes it consistent with the basic idea of a CA. Moreover, the CAs can be calculated by a type of Hebbian learning. Many simulations take advantage of this type of connectivity (e.g. Amit, 1989; Sougne, 2001). However, perhaps due to this lack of realism, the Hopfield model is falling out of favour as a model for processes in the real brain.

In the brain, neurons are not connected to every other neuron, but a single neuron has thousands or tens of thousands of synapses (Churchland and Sejnowski, 1999). There is some evidence for small world topologies (Bullmore and Sporns, 2009), and random topologies (Bertschinger and Natschlag, 2004; Valiant, 2005) have also been used in simulations; both have the benefit of existing mathematical expertise, and both are more realistic than full connectivity. Finally, more realistic topologies can explicitly copy known biological topology (Douglas and Martin, 1991).

5.1.4 Learning

Learning is important for CAs, and Hebbian learning rules are used widely in simulations. There are however a wide range of plausible Hebbian rules. A complete discussion of these rules is beyond the scope of this paper, as is a complete discussion of the underlying biology of learning (but see Section 6).

Typically learning is modelled as a permanent change in synaptic weight. However, learning can reasonably involve neural growth, neural death, synaptic growth, synaptic death, and transient synaptic weight change. Synapses that only rarely cause the post-synaptic neuron to fire can be pruned, and neurons that rarely cause their post-synaptic neurons to fire can be pruned. New neurons can be added, and their synapses can grow based on the success of nearby synapses. These are all forms of Hebbian learning.

A full-fledged Hebbian learning rule typically increases the synaptic weight when the two neurons co-fire, and decreases the weight when one fires and the other does not (see Gerstner and Kistler, 2002, for a review). These rules model the biological actions of long-term potentiation and depression.

In addition to permanent changes in synaptic weight, changes can be temporary. There is sound evidence for both short-term depression and potentiation (Zucker and Regehr, 2002; Hempel et al, 2000), but these two mechanisms are used less in current models.

Co-firing requires the neurons to fire at the same time, and biological evidence shows the timing involves the order of neural firing (Bi and Poo, 1998). If the presynaptic neuron fires first, and the postsynaptic neuron fires within 20ms there is an increase in strength; if they are reversed, there is a decrease.

A related problem is that for Hebbian learning to support the spread of a CA into new areas requires the neurons in the new area to fire. Background neural firing can support movement into new areas, but this requires the neural model to support neural firing without incoming activation.

There is also mathematical research on Hebbian learning. One mathematical treatment of learning shows that Hebbian learning can be used for principal and independent component analysis (Fyfe, 2005).

By making use of different learning mechanisms, a wide range of behaviours can be learned. Moreover, these learning mechanisms are based on pairs of neurons, and when combined with complex topologies, a vast range of learning becomes possible.

5.2 Simulations

Since Rochester et al (1956), there have been a vast range of CA simulations, and they have come with increasing frequency. For example, one system successfully modelled the learning of CAs (Hetherington and Shapiro, 1993). The neural model was continuous output, the network included excitatory and inhibitory neurons, there was sparse connectivity, and the learning rules capped synaptic weights. CAs persisted, were unique, and grouped similar stimuli.

Others have shown that neural models can implement CAs in isolation. One system uses CAs to categorise vowel sounds (Hoshino et al, 2002). Another uses a multi-level neural system to learn to categorise various line drawings of buildings (Knoblauch et al, 2007). Various layers in the system are modelled on human brain areas and their known processing. Recurrent connections allow a soft winner takes all system of CAs that interact to satisfy constraints at multiple levels.

CAs have also been applied to a range of model tasks. Perhaps the most common simulation is associative memory.

5.2.1 Models of associative memory

There is a review of associative memory models based on simulated CAs (Lansner, 2009), and it should be noted that CAs themselves are auto-associative memories; given one input, a similar pattern is retrieved. More typically, associative memory refers to hetero-associative memory with an input being associated with a different output. An early example of this (Fransen et al, 1992) uses relatively fine grained neural models to show pattern completion, and persistent behaviour. This acts as a proof that a CA model can implement associative memory.

An even earlier example uses a bistable neuron pool; the simulated system can learn concepts, and associations between concepts (Amari, 1977). In this case, and in most cases described in this section, the system learns the concept, i.e. a CA is formed, and associations between CAs are also formed.

Another model uses well connected neurons (Palm and Sommer, 1995) to store auto-associative memories (CAs) and associations between CAs. Patterns can be gradually learned via Hebbian learning. A related model (Wennekers and Palm, 2000) uses simulated spiking neurons to encode objects. During retrieval, multiple objects can be retrieved with neurons associated with particular concepts firing synchronously. If a neuron is associated with multiple objects, it is synchronous with both sets.

More recently, integrate and fire neurons were used to learn CAs, and then associate those with cognitive maps (Salihoglu et al, 2009), and an ambiguous stimulus could be disambiguated by the context of the current map. Learning took advantage of the recurrent multi-layer dynamics of the system. Similarly, fatiguing leaky integrate and fire neurons were used to learn and retrieve hierarchical categories, with super-class CAs associated with sub-class CAs by neural overlap (Huyck, 2007). Figure 5 shows a subset of the overlapping CAs; the super-category, *Mammal*, supports generalization if a novel type of *Mammal* is presented.

A graph theoretical model (Valiant, 2005) stores concepts and their associations. The model lacks a biologically plausible learning mechanism, but concepts and associations between concepts are stored as sets of neurons, and the model works with a range of biologically plausible parameters including numbers of neurons, neurons per CA, synaptic strength, and connections per neuron.

C	C ^M D		C	C	C ^M D
	D			D	C
	C	M _D	D	C ^M D	C
C	D	C	C	C ^M D	D
M		M	M		M
C	D	D	D	C	D

Fig. 5 Derived from (Huyck, 2007). Hierarchical categories stored as CAs with overlapping neurons. Each cell represents a neuron. Those with a *C* are part of the *Cat* CA, with *D* the *Dog*, and with *M* the super-category *Mammal*.

Inhibition is often used at a neural level to support competition between CAs. This can lead to winner take all effects (Lundqvist et al, 2006, and see Section 5.2). If an ignited CA has a large number of strong inhibitory connections to another CA, it will suppress the other CA. If both are mutually inhibitory, and inactive, an ambiguous stimulus will start to ignite both, and they will compete.

5.2.2 Decisions, binding and language

CAs can also be used in systems that perform a range of sophisticated processing tasks. Computationally, it has been shown that a large enough net with CAs is Turing complete³ (Byrne and Huyck, 2010). The real challenge is to build neuro-psychologically realistic simulations. While there are proposals to model an entire brain at the neural level on supercomputers (Markram, 2006), CA dynamics will still need to be understood before these models can successfully learn and process environmental input. Simulations using CAs include complex state transition mechanisms, variable binding, and learning.

A system of CAs can be used to store state as in a finite state automata. Imagine that there are several states each represented by a CA, if only one is active, the system is in that state. The neural system also has input, and if a new (non-state) CA is ignited by the environment, having categorised the environmental input, this input CA can act as input to the automata. Now input from two CAs, the current state and input, can ignite a third CA, the new state. The third CA can extinguish the first (via inhibitory synapses). This is a basic means of acting with CAs.

CAs have been used as the model of the basic categorising unit in networks that categorise visual input (Bogacz, 2007). These CAs emerge from continuous input continuous output neurons, and the system integrates information over time to make decisions. CAs have also been used as the model of the basic semantic unit in a system that models brain storming (Iyer et al, 2009).

³ If a language is Turing complete, anything that can be programmed can be programmed in it. So, to some degree, CAs are equivalent to Java.

A CA based model (Wichert, 2001) is able to use pictorial reasoning, simple rules, and goal directed behaviour using automata to implement an agent. This agent is able to learn spatial cognitive maps, and improve its own behaviour.

One criticism of current simulated neural systems is that they can not solve the binding problem (Jackendoff, 2002), so that there must be a CA for each concept including combined concepts. Without binding, a system would need a CA for the concepts *Red*, *Blue*, *Square*, *Red Square*, and *Blue Square* if it were to use them. Binding is important because it allows combinatorial expansion of states. For instance, if the system has ten noun CAs, and ten verb CAs, it can have 20 states. However, if it can bind these, it can have 100 states. Binding also allows rules with variables, which in themselves are Turing complete.

A relatively recent modification of neurocomputational theory supports a functional role for synchronous firing (see Section 2.4 and von der Malsburg, 1981). CAs remain central to this model, though precise timing of spikes becomes more important. Synchrony has been used to bind in several models. For example, spiking neurons are used to simulate binding by synchrony (Sougne, 2001). Though portions of this model have well connected topologies, binding via synchrony is a complex task. Another CA based model (Knoblauch et al, 2004) implements a regular language parser and binding via synchrony. A long standing model has used spiking neurons and global inhibition as a basis for feature binding (Terman and Wang, 1995).

CAs have been used with structured binding circuits as another method of enabling binding (van der Velde and de Kamps, 2006). A third method of binding is short-term potentiation (Huyck, 2009), which has been used in a CA based cognitive model of context-free natural language parsing. CAs are the grammar rules, words, and the semantic results of a parse.

Modelling learning is also important. Learning simulations have been done on a range of basic categories, for example visual categories (Knoblauch et al, 2007), for associations (see Section 5.2.1), and on learning weightings for rules (Belavkin and Huyck, 2010), but there are a range of more complex phenomena, like learning rules that use variable binding, that have not been explored. Learning new categories is closely related to the symbol grounding problem (Harnad, 1990), with learned CAs derived directly from environmental stimuli. For example, one system models the acquisition of spoken words using neurons that learn CAs (Gargnani et al, 2007).

Models can also be used to show that CAs can perform as expected. For example, a 3D neural topology is simulated to show that CAs can form in the striatum (Humphries et al, 2009). Another model simulates the striatal spiny neuron network showing that sequences of CAs emerge in response to a sequence of inputs (Ponzi and Wickens, 1977); this model is more closely linked to biological behaviour than most models.

There are a vast range of possibilities for neural models of cognition from non-neural connectionist models to highly accurate compartmental models. One set of principles, that is compatible with CAs (O'Reilly, 1999), permits

some exploration of the model space while still linking the model to a sensible degree of biological realism.

A great deal of modelling work has shown how a neural system can perform a task in principle. However, this work almost always has a relatively weak link to psychology much less neurophysiology. One challenge for future modelling work is to provide a more accurate neuropsychological explanation.

6 Discussion

While many researchers across many disciplines have worked on CAs, the complexity of the brain makes it difficult to gain a comprehensive understanding. This complexity includes the large number of neurons and synapses, the complexity and diversity of individual synapses and neurons, and a diverse range of effects over different time ranges. The CA provides conceptual support for understanding the behaviour of a group of neurons much smaller than the whole brain. How might the scientific community move forward from current knowledge to a richer understanding of CAs?

While there is extensive evidence that CAs persist (see Section 2.3), simulations that persist, and psychological theories that predict the strength and duration of a memory (Anderson and Lebiere, 1998), there is not a sound link between these. Assuming that the strength of a memory has a direct relation with the number of firing neurons in its CA, a good computational model would reflect that firing behaviour. The authors are not aware of such a model, with existing CA models either persisting forever, or stopping after a relatively brief time. Moreover, a good model would also be able to learn those CAs, and have the strength and duration emerge.

The issue of learning, or CA formation, is central to the study of CAs. Indeed, it has already been raised in the guise of data supporting the hypothesis (Section 2.5) and model learning rules (Section 5.1.4). The research on learning is vast, and a brief discussion is useful. While synaptic change may drive CA formation, it does occur in a large system of neural activity. This activity must support learning dynamics beyond the neuron (Abbott and Nelson, 2000), however this behaviour is not currently well understood.

CA formation is central to the standard model and its derivatives, with Hebb's learning rule as part of his original hypothesis. Biologically, it seems that the rule is instantiated by a form of long-term potentiation (Malenka and Nicoll, 1999) and long-term depression (Ito, 1989), and both lead to relatively long-term changes in the ability of one neuron to cause a particular other neuron to fire. Though there is extensive evidence for short-term plasticity (e.g. Zucker and Regehr, 2002), this is not typically included in CA models.

While it is relatively easy to look at the dynamics of synapses, it is very difficult to inspect the biological formation of a CA because it involves a large number of neurons. Typically, theory states that there is a bifurcation between a set of neurons that cannot persistently fire and, through a slight synaptic strength increase, a set of neurons that can persist. This theory has been

duplicated in simulations (Amit and Brunel, 1995), but leads to a problem of an abrupt appearance of CAs, and the inability of CAs to remember novel stimulus (Durstewitz et al, 2000). The use of short-term potentiation in models has reduced this problem, as the new CA is initially supported by short-term changes, which support longer-term changes via repeated neural co-firing.

This problem of rapid formation of CAs relates to the difference between semantic and episodic memories. Semantic memories tend to require the presentation of many instances to form an associated CA, while episodic memory forms in one presentation. CAs for episodic memories must be formed after one presentation, since events do not happen twice. This encoding may use a CA in the hippocampus combined with neurons to encode specific semantic elements (Eichenbaum, 2000). This CA may be reactivated in later stages (e.g. sleep) to support memory consolidation (Qin et al, 1997). Hippocampal damage could then lead to a loss of memory of those episodes (Paller, 1997).

Semantic memory provides another set of challenges, including questions involving the formation of CAs in the young animal; for example, when do CAs for particular concepts form? How do CAs for a concept learned early (e.g. animal) evolve into sets of concepts (e.g. cat and dog)? These questions are particularly important as it is very likely that the development of later CAs depends on those early CAs, and computational models would be more robust if they were based on relatively biologically and psychologically accurate developmental CA models. These questions are difficult to answer biologically because they involve relatively long-term study of a developing brain.

The authors are unaware of any significant evidence of the neural evolution of categorisation in infants using single cell recording. There is of course long standing evidence of early cell specialisation in response to stimulus (Hubel and Wiesel, 1962); there is just not evidence of these or later neurons forming CAs. None the less, existing data does not contradict Hebb's hypothesis. Improving EEG and other non-invasive techniques (Reynolds and Richards, 2009) may help uncover the mechanisms of early CA formation.

The dynamics of CA formation are thus, at least potentially, quite complex. For example, the spontaneous reoccurrence of neural firing patterns, neural avalanches, have been proposed as a potential mechanism involved in CA formation (Plenz and Thiagarajan, 2007). As this obeys power law dynamics in size, every neuron can be linked with every other neuron, allowing CAs to involve any set of neurons in the brain.

There are further problems of CA learning with existing CAs. CAs may recruit new neurons and lose neurons, and since membership of a CA is not binary, this loss or gain of a given neuron is a matter of degree. Similarly, CAs may fractionate into separate CAs. There is also research in circuit formation making use of, for instance, laminar architecture (Raizada and Grossberg, 2003), and brain areas (Knoblauch et al, 2007). Learning is an important key to understanding CAs.

In the authors' view, the CA is a graded concept. We suspect, and evidence provided in this paper supports the idea that, prototypical CAs are the basis of typical concepts and their associated symbols. As such, prototypical CAs

are the neural basis of basic symbols used in, for instance, natural language processing. As CAs are a graded category, less typical CAs, but things still reasonably called CAs, may be the basis of a wide range of cognitive phenomena ranging from cognitive maps to complex motor skills.

There are many questions about how CAs relate to a wide range of processes (e.g. path finding, language, and theory formation), and higher order structures. These are difficult questions, but as scientific understanding of these mechanisms is weak, even suggestive biological and psychological evidence, and plausible psychological and computational models would be useful.

7 Conclusion

The standard model is that CAs are the neural basis of concepts. The CA is a small set of interconnected neurons; that can persist without external stimulus; is learned; and is supported by synchronous firing behaviour. The model has some alternatives.

CAs are in a large range of brain areas. CAs are involved in varying degrees in a large range of cognitive tasks; they are categorisers and act as STMs and LTMs, but they are also intimately involved in priming and associative memory; they are important, but less central, in complex tasks such as language and working memory.

The full CA hypothesis states that CAs are the basis of all concepts. This paper has not provided support for this claim. Though the authors have not found a contradiction, it is plausible that some other mechanism is used for some concepts; however, no mechanism currently suggests itself. While CAs may or may not be the neural correlate of all concepts, there is evidence that they are the neural correlate of many. For these concepts, CAs provide a crucial bridge between categorisation, LTM and STM.

CAs are not the sole answer to understanding the brain. For example, higher order structures, and systems play crucial roles in human behaviour. None the less, CAs are a key link between biology and psychology, providing an intermediate level of analysis between neurons and the brain. CAs at this intermediate level are also the basis of symbolic concepts and thus symbols. So, CAs provide the bridge between subsymbolic and symbolic reasoning. However, our current models of CAs, both computational and theoretical, are relatively crude and underspecified, and the links to neuropsychological behaviour are far from complete.

Evidence for CAs can be evaluated based on neural and psychological behaviour. Tying psychological states and processes to neural behaviour is difficult but can be done with single unit recording, and with less certainty via other techniques. While this paper has provided evidence that supports the general CA model, it has not fully shown the structure of a single functioning CA much less the structure of all CAs.

Similarly, models of CAs should account for both neural and psychological data. Better neural models are obviously more reliable, and reasonable

topologies are important; input to simulations is better when it follows known biological paths, for example via retinal centre surround detectors. Also, the models are better when they adhere to the four theoretical principles of CAs.

The prototypical CA, as defined by these four principles, is a flexible yet solid concept that can be used across disciplines. There is sufficient evidence for this concept that it can be used now for discussion and model development. Since understanding of neuroscience and psychology is incomplete and will develop, the concept of CA needs to remain flexible to accommodate the evolving understanding of neuropsychology.

Improving our understanding of CAs will lead to a better understanding of animal cognition. CAs provide a very important mechanism that is central to cognitive behaviour supporting categorisation, STM and LTM, and providing support for other processes. Computational systems that make use of CAs to perform in complex environments and learn about those environments can make use of the knowledge of neuropsychology.

The authors hope that this review will distribute existing information on CAs, that it will stimulate further discussion and research on the prevalence, function, and behaviour of CAs, and that this will lead to significant advances in the understanding of cognitive neuropsychology.

Acknowledgements:

This work was supported by EPSRC grant EP/D059720. Thanks to Ray Adams, Eric Chown, Dan Diaper, Kailash Nadh and Pieter DeVries for comments on this paper, and Mark Dubin for the use of his Brodmann figures.

References

- Abbott L (1999) Lapicque's introduction of the integrate-and-fire model neuron (1907). *Brain Research* 50:303–304
- Abbott L, Nelson S (2000) Synaptic plasticity: taming the beast. *Nature Neuroscience* 3:1178–1183
- Abeles M, Bergman H, Margalit E, Vaadia E (1993) Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology* 70(4):1629–1638
- Ackley D, Hinton G, Sejnowski T (1985) A learning algorithm for boltzmann machines. *Cognitive Science* 9:147–169
- Amari S (1977) Neural theory of association and concept-formation. *Biological Cybernetics* 26:175–185
- Amit D (1989) *Modelling Brain Function: The world of attractor neural networks*. Cambridge University Press
- Amit D, Brunel N (1995) Learning internal representations in an attractor neural network with analogue neurons. *Network: Computation in Neural Systems* 6:359–388
- Anderson J, Lebiere C (1998) *The Atomic Components of Thought*. Lawrence Erlbaum

- Assad W, Rainer G, Miller E (2000) Task-specific neural activity in the primate prefrontal cortex. *Journal of Neurophysiology* 84:451–459
- Auseng P, Grismayr B, Freunberger R, Klimesch W (2010) Control mechanisms in working memory: A possible function of EEG theta oscillations. *Neuroscience and Biobehavioural Reviews* 34:1015–1022
- Averbeck B, Latham P, Pouget A (2006a) Neural correlations, population coding and computation. *Nature Reviews Neuroscience* 7:358–366
- Averbeck B, Sohn J, Lee D (2006b) Activity in prefrontal cortex during dynamic selection of action sequences. *Nature Neuroscience* 9:2:276–282
- Baddeley A (2003) Working memory: looking back and looking forward. *Nature Reviews Neuroscience* 4:829–839
- Baeg E, Kim Y, Kim J, Ghim J, Kim J, Jung M (2007) Learning-induced enduring changes in functional connectivity among prefrontal cortical neurons. *Journal of Neuroscience* 27:4:909–918
- Barlow H (1972) Single units and sensation: A neuron doctrine for perceptual psychology? *Perception* 1:371–394
- Belavkin R, Huyck C (2010) Conflict resolution and learning probability matching in a neural cell-assembly architecture. *Cognitive Systems Research* 12:93–101
- Bendor D, Wang X (2008) Neural response properties of primary, rostral, and rostraltemporal core fields in the auditory cortex of marmoset monkeys. *Journal of Neurophysiology* 100:888–906.
- Bennett B, Callaway J, Wilson C (2000) Intrinsic membrane properties underlying spontaneous tonic firing in neostriatal cholinergic interneurons. *Journal of Neuroscience* 20(22):8493–8503.
- Bertschinger N, Natschlagler T (2004) Real-time computation at the edge of chaos in recurrent neural networks. *Neural Computation* 16(7):1413–1436
- Beurle R (1956) Properties of a mass of cells capable of regenerating pulses. *Transactions of the Royal Society London B* 240:55–94
- Bevan M, Wilson C (1999) Mechanisms underlying spontaneous oscillation and rhythmic firing in rat subthalamic neurons. *Journal of Neuroscience* 19:7617–7628.
- Bi G, Poo M (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience* 18:24:10,464–10,472
- Bichot N, Schall J (1999) Effects of similarity and history on neural mechanisms of visual selection. *Nature Neuroscience* 2:6:549–554
- Bieser A, Muller-Preuss P (1996) Auditory responsive cortex in the squirrel monkey: neural responses to amplitude-modulated sounds. *Experimental Brain Research* 108:273–284
- Bisley J, Zaksas D, Droll J, Pasternak T (2004) Activity of neurons in cortical area MT during a memory for motion task. *Journal of Neurophysiology* 91:286–300
- Bogacz R (2007) Optimal decision network with distributed representation. *Neural Networks* 20:564–576

- Booth J, Burman D, Meyer J, Gitelman D, Parrish T, Marsel Mesulam M (2004) Development of brain mechanisms for processing orthographic and phonologic representations. *Journal of Cognitive Neuroscience* 16:7:1234–1249
- Braitenberg V (1978) Cell assemblies in the cerebral cortex. In: Heim R, Palm G (eds) *Lecture Notes in Biomathematics* 21, *Theoretical Approaches to Complex Systems*, Berlin: Springer, pp 171–188
- Braitenberg V (1989) Some arguments for a theory of cell assemblies in the cerebral cortex. In: Nadel, Cooper, Culicover, Harnish (eds) *Neural Connections, Mental Computation*, MIT Press
- Brecht M, Schneider M, Sakmann B, Margrie T (2004) Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex. *Nature* 427:704–710
- Bressler S (1995) Large-scale cortical networks and cognition. *Brain Research Reviews* 20:288–304
- Brette R, Rudolph M, Carnevale T, Hines M, Beeman D, Bower J, Diesmann M, Morrison A, Goodman P, Harris F, Zirpe M, Natschalager T, Pecevski D, Ermentrout B, Djurfeldt M, Lansner A, Rochel O, Vieville T, Muller E, Dafison A, ElBoustani S, Destexhe A (2007) Simulation of networks of spiking neurons: A review of tools and strategies. *Journal of Computational Neuroscience* 23:349–398
- Bruno R, Sakmann B (2006) Cortex is driven by weak by synchronously active thalamocortical synapses. *Science* 312:1622–1627
- Buckner R, Wheeler M (2001) The cognitive neuroscience of remembering. *Nature Reviews Neuroscience* 2:624–634
- Bullmore E, Sporns O (2009) Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience* 10:186–198
- Buonomano D, Merzenich M (1998) Cortical plasticity: from synapses to maps. *Annual Review of Neuroscience* 21:149–186
- Burgess N, Hitch G (2005) Computational models of working memory: putting long-term memory into context. *Trends in Cognitive Sciences* 9:11:535–541
- Burkitt A (2006) A review of the integrate-and-fire neuron model: I. homogeneous synaptic input. *Biological Cybernetics* 95(1):1–19
- Butt S, Harris-Warrick R, Kiehn O (2002) Firing properties of identified interneuron populations in the mammalian hindlimb central pattern generator. *Journal of Neuroscience* 22:22:9961–9971
- Buzsáki G, Draguhn A (2004) Neuronal oscillations in cortical networks. *Science* 304:1926–1929
- Byrne E, Huyck C (2010) Processing with cell assemblies. *Neurocomputing* 74:76–83
- Cho J, Sharp P (2001) Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behavioral Neuroscience* 115:1:3–25
- Christian K, Thompson R (2003) Neural substrates of eyeblink conditioning: acquisition and retention. *Learning and Memory* 10:427–455

- Chrobak J, Buzaski G (1998) Gamma oscillations in the entorhinal cortex of the freely behaving rat. *Journal of Neuroscience* 18(1):388–398
- Churchland P, Sejnowski T (1999) *The Computational Brain*. MIT Press
- Constantinidis C, Procyk E (2004) The primate working memory networks. *Cognitive, Affective, and Behavioral Neuroscience* 4(4):444–465
- Cossart R, Aronov D, Yuste R (2003) Attractor dynamics of network up states in the neocortex. *Nature* 423:283–288
- Cowan N (1988) Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information-processing system. *Psychological Bulletin* 1-4:2:163–191
- Crowe D, Averbach B, Chafee M (2008) Neural ensemble decoding reveals a correlate of viewer-to-object-centered spatial transformation in monkey parietal cortex. *Journal of Neuroscience* 28(20):5218–5228
- Crutcher M, Russo G, Ye S, Backus D (2004) Target-, limb-, and context-dependent neural activity in cingulate and supplementary motor areas of the monkey. *Experimental Brain Research* 158:278–288
- Curtis C, D’Esposito M (2003) Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Science* 7:9:415–423
- Deco G, Rolls E (2005) Sequential memory: a putative neural and synaptic dynamical mechanism. *Journal of Cognitive Neuroscience* 17:2:294–307
- Desimone R, Albright T, Gross C, Bruce C (1984) Stimulus-selective properties of inferior temporal neurons in the macaque. *Journal of Neuroscience* 4:8:2051–2062
- D’Esposito M (2007) From cognitive to neural models of working memory. *Philosophical Transactions of the Royal Society* 362:761–772
- deVries P (2004) Effects of binning in the identification of objects. *Psychological Research* 69:41–66
- Diester I, Nieder A (2007) Semantic associations between signs and numerical categories in prefrontal cortex. *PLOS Biology* 5:11:2684–2695
- Douglas R, Martin K (1991) A functional microcircuit for cat visual cortex. *Journal of Physiology* 440:735–769
- Douglas R, Martin K (2004) Neuronal circuits of the neocortex. *Annual Review of Neuroscience* 27:419–451
- Dragoi G, Buzaski G (2006) Temporal encoding of place sequences in hippocampal cell assemblies. *Neuron* 50:145–157
- Dragoi G, Tonegawa S (2011) Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* 469:397–401
- Durstewitz D, Seamans J, Sejnowski T (2000) Neurocomputational models of working memory. *Nature neuroscience supplement* 3:1184–1191
- Egorov A, Hamam B, Franssen E, Hasselmo M, Alonso A (2002) Graded persistent activity in entorhinal cortex neurons. *Nature* 420:173–178
- Eichenbaum H (2000) A cortical-hippocampal system for declarative memory. *Nature Reviews Neuroscience* 1:41–50
- Eliasmith C, Thagard P (2001) Integrating structure and meaning: a distributed model of analogical mapping. *Cognitive Science* 25:245–286

- Engel A, Fries P, Singer W (2001) Dynamic predictions: oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience* 2:704–716
- Engel R, Tuholski S, Laughlin J, Conway A (1999) Working memory, short-term memory, and general fluid intelligence: a latent-variable approach. *Journal of Experimental Psychology* 12:3:309–331
- Field D (1994) What is the goal of sensory coding. *Neural Computation* 6:559–601
- Fischer B, Boch R, Bach M (1981) Stimulus versus eye movements: comparison of neural activity in the striate and prelunate visual cortex (a17 and a19) of trained rhesus monkey. *Experimental Brain Research* 43:69–77
- Fransen E, Lansner A, Liljenstrom H (1992) A model of cortical associative memory based on hebbian cell assemblies. In: Niklasson L, Boden M (eds) *Connectionism in a Broad Perspective*, Ellis Horwood Springer-Verlag, pp 165–171
- Freedman D, Assad J (2006) Experience-dependent representation of visual categories in parietal cortex. *Nature* 433:85–88
- Freiwald W, Kreiter A, Singer W (2001) Synchronization and assembly formation in the visual cortex. In: Nicolelis M (ed) *Progress in Brain Research*, vol 130, Elsevier
- Fujii H, Hiroyuke I, Aihara K, Ichinose N, Tsukada M (1996) Dynamical cell assembly hypothesis - theoretical possibility of spatio-temporal coding in the cortex. *Neural Networks* 9:8:1303–1350
- Funahashi S (2001) Neuronal mechanisms of executive control by the prefrontal cortex. *Neuroscience Research* 39:147–165
- Fusi S (2008) A quiescent working memory. *Science* 319:1495–1496
- Fuster J, Alexander G (1971) Neuron activity related to short-term memory. *Science* 173:652–654
- Fuster J, Bodner M, Kroger J (2000) Cross-modal and cross-temporal association in neurons of frontal cortex. *Nature* 405:347–351
- Fyfe C (2005) *Hebbian Learning and Negative Feedback Networks*. Springer
- Gallese V, Lakoff G (2005) The brain's concepts: the role of the sensory motor system in conceptual knowledge. *Cognitive Neuropsychology* 22(3/4):455–479
- Gallese V, Fadiga L, Fogassi L, Rizzolatti G (1996) Action recognition in the premotor cortex. *Brain* 119:593–609
- Gargnani M, Wennekers T, Pulvermuller F (2007) A neuronal model of the language cortex. *Neurocomputing* 70:1914–1919
- Gazzaley A, Rissman J, D'Esposito M (2004) Functional connectivity during working memory maintenance. *Cognitive, Affective, & Behavioral Neuroscience* 4(4):580–599
- Georgopoulos A, Schwartz A, Kettner R (1986) Neuronal populations coding of movement direction. *Science* 233:1416–1419
- Gerstner W, Kistler W (2002) Mathematical formulations of hebbian learning. *Biological Cybernetics* 87:404–415
- Gochin P, Colombo M, Dorfman G, Gerstein G, Gross C (1994) Neural ensemble coding in inferior temporal cortex. *Journal of NeuroPhysiology*

- 71:6:2325–2337
- Golari G, Ghahremani D, Whitfield-Gabrieli S, Reiss A, Eberhardt J, Gabrieli J, Grill-Spector K (2007) Differential development of high-level visual cortex correlates with category-specific recognition memory. *Nature neuroscience* 10:4:512–522
- Goldman-Rakic P (1995) Cellular basis of working memory. *Neuron* 14:477–485
- Goldman-Rakic P (1996) Regional and cellular fractionation of working memory. *PNAS* 93:13,473–13,480
- Gray C, Singer W (1989) Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *PNAS* 86:1698–1702
- Grill-Spector K, Golarai G, Gabrieli J (2008) Developmental neuroimaging of the human ventral visual cortex. *Trends in Cognitive Science* 12:152–162
- Grinvald A, Arieli A, Tsodyks M, Kenet T (2003) Neuronal assemblies: Single cortical neurons are obedient members of a huge orchestra. *Biopolymers* 68:422–436
- Gutierrez R, Carmena J, Nicolelis M, Simon S (2006) Orbitofrontal ensemble activity monitors licking and distinguishes among natural rewards. *Journal of Neurophysiology* 95:119–133
- Harnad S (1990) The symbol grounding problem. *Physica D* 42:335–346
- Harris K (2005) Neural signatures of cell assembly organization. *Nature Reviews Neuroscience* 6:399–407
- Harris K, Csicsvari J, Hirase H, Dragoi G, Buzsaki G (2003) Organization of cell assemblies in the hippocampus. *Nature* 424:552–556
- Harrison S, Tong F (2009) Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458:632–635
- Hebb DO (1949) *The Organization of Behavior: A Neuropsychological Theory*. J. Wiley & Sons
- Hempel C, Hartman K, Wang X, Turrigiano G, Nelson S (2000) Multiple forms of short-term plasticity at excitatory in rat medial prefrontal cortex. *Journal of Neurophysiology* 83:3031–3041
- Henson R, Rugg M, Shallice T, Josephs O, Dolan R (1999) Recollection and familiarity in recognition in memory: an event-related functional magnetic resonance imaging study. *Journal of Neuroscience* 19(10):3962–3972
- Hetherington P, Shapiro M (1993) Simulating hebb cell assemblies: the necessity for partitioned dendritic trees and a post-net-pre ltd rule. *Network: Computation in Neural Systems* 4:135–153
- Histed M, Pasupathy A, Miller E (2009) Learning substrates in the primate prefrontal cortex and striatum: sustained activity related to successful actions. *Neuron* 63(2):244–253
- Hochberg L, Serruya M, Friehs G, Mukand J, Saleh M, Caplan A, Branner A, Chen D, Penn R, Dooghue J (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442:13:164–171
- Hodgkin A, Huxley A (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology* 117:500–544

- Hofstadter D (1979) *Godel, Escher and Bach: an Eternal Golden Braid*. Basic Books
- Hopfield J (1982) Neural nets and physical systems with emergent collective computational abilities. *Proc of the Nat Academy of Sciences USA* 79:2554–2558
- Hoshino O, Miyamoto M, Zheng M, Kuroiwa K (2002) A neural network model for encoding and perception of vowel sounds. *Neurocomputing* 44–46:435–442
- Howard M, Volkov I, Mirsky R, Garell P, Noh M, Granner M, Damasio H, Steinschneider M, Reale R, Hind J, Brugge J (2000) Auditory cortex on the human posterior superior temporal gyrus. *Journal of Comparative Neurology* 416:79–92
- Hubel D, Wiesel T (1962) Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *Journal of Physiology* 160:106–154
- Humphries M, Wood R, Gurney K (2009) Dopamine-modulated dynamic cell assemblies generated by the GABAergic striatal microcircuit. *Neural Networks* 22:1174–1188
- Huyck C (2007) Creating hierarchical categories using cell assemblies. *Connection Science* 19:1:1–24
- Huyck C (2009) A psycholinguistic model of natural language parsing implemented in simulated neurons. *Cognitive Neurodynamics* 3(4):316–330
- Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, Yuste R (2004) Synfire chains and cortical songs: Temporal modules of cortical activity. *Science* 304:23:559–564
- Isomura Y, Harukuni R, Takekawa T, Aizawa H, Fukai T (2009) Microcircuitry coordination of cortical motor information in self-initiation of voluntary movements. *Nature Neuroscience* 12:12:1586–1593
- Ito M (1989) Long-term depression. *Annual Review of Neuroscience* 12:85–102
- Iyer L, Doboli S, Minai A, Brown V, Levine D, Paulus P (2009) Neural dynamics of idea generation and the effects of priming. *Neural Networks* 22:674–686
- Izhikevich E (2004) Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks* 15:5:1063–1070
- Jackendoff R (2002) *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford University Press
- James W (1892) *Psychology: The Briefer Course*. University of Notre Dame Press
- Jezek K, Henriksen E, Treves A, Moser E, Moser M (2011) Theta-paced flickering between place-cell maps in the hippocampus. *Nature* 478:246–249
- Jonides J, Smith E, Koeppe R, Awh E, Minoshima S, Mintun M (1993) Spatial working memory in humans as revealed by pet. *Nature* 363:623–625
- Just M, Chherkassky V, Aryal S, Mitchell T (2010) A neurosemantic theory of concrete noun representation based on the underlying brain code. *PLoS ONE* 5:1:e8622
- Kaplan S, Weaver M, French R (1990) Active symbols and internal models: Towards a cognitive connectionism. *AI and Society* 4:51–71

- Kaplan S, Sontag M, Chown E (1991) Tracing recurrent activity in cognitive elements(trace): A model of temporal dynamics in a cell assembly. *Connection Science* 3:179–206
- Katz D, Simon S, Nicolelis M (2002) Taste-specific neuronal ensembles in the gustatory cortex of awake rats. *Journal of Neuroscience* 22:5:1850–1857
- Keri S (2003) The cognitive neuroscience of category learning. *Brain Research Reviews* 43:85–109
- Klausberger T, Magill P, Marton L, Roberts J, Cobden P, Buzaski G, Somogyi P (2003) Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature* 421:844–848
- Kleinsmith L, Kaplan S (1963) Paired-associate learning as a function of arousal and interpolated interval. *Journal of Experimental Psychology* 65:2:190–193
- Knoblauch A, Markert H, Palm G (2004) An associative model of cortical language and action processing. In: *Proceedings of the Ninth Neural Computation and Psychology Workshop*
- Knoblauch A, Kupper R, Gewaltig M, Korner U, Korner E (2007) A cell assembly based model for the cortical microcircuitry. *Neurocomputing* 70:1838–1842
- Krahe R, Gabbiani F (2004) Burst firing in sensory systems. *Nature Reviews Neuroscience* 5:13–24
- Kreiman G, Koch C, Fried I (2000) Imagery neurons in the human brain. *Nature* 408:357–361
- Kreiter A, Singer W (1996) Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *Journal of Neuroscience* 16(7):2381–2396
- Kuo C, Chiou R, Liang K, Yen C (2009) Differential involvement of the anterior cingulate and primary sensorimotor cortices in sensory and affective functions of pain. *Journal of Neurophysiology* 101:1201–1210
- Lamme V, Super H, Spekreijse H (1998) Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Biology* 8:525–535
- Lansner A (2009) Associative memory models: from the cell-assembly theory to biophysically detailed cortex simulations. *Trends in Neuroscience* 32:3:178–186
- Lapish C, Durstewitz D, Chandler L, Seamans J (2008) Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex. *PNAS* 105:33:11,963–11,968
- Larimer P, Strowbridge B (2009) Representing information in cell assemblies: persistent activity mediated by semilunar granule cells. *Nature Neuroscience* 13:213–222
- Laubach M, Wessberg J, Nicolelis M (2000) Cortical ensemble activity increasingly predicts behaviour outcomes during learning of a motor task. *Nature* 405:567–570
- LeBihan D, Turner R, Zeffiro T, Cuenod C, Jezzard P, Bonnerot V (1993) Activation of human primary visual cortex during visual recall: A magnetic resonance imaging study. *PNAS* 90:11,802–11,805

- Levy N, Horn D, Meilijson I, Ruppin E (2001) Distributed synchrony in a cell assembly of spiking neurons. *Neural Networks* 14:815–824
- Liebenthal E, Uhlmann O, Camhi J (1994) Critical parameters of the pike trains in a cell assembly: coding of turn direction by giant interneurons of the cockroach. *Journal of Comparative Physiology A* 174:281–296
- Logothetis N, Pauls J, Poggio T (1995) Shape representation in the inferior temporal cortex of monkeys. *Current Biology* 5:5:552–563
- Lundqvist M, Rehn M, Djurfeldt M, Lansner A (2006) Attractor dynamics in a modular neural network model of neocortex. *Network: Computation in Neural Systems* 17:253–276
- Macaluso E, Frith C, Driver J (2000) Modulation of human visual cortex by crossmodal spatial attention. *Science* 289:1206–1208
- Malenka R, Nicoll R (1999) Long-term potentiation - a decade of progress? *Science* 285:1870–1874
- Maquet P (2001) The role of sleep in learning and memory. *Science* 10:1048–1052
- Markram H (2006) The blue brain project. *Nature Reviews Neuroscience* 7:153–160
- Martini F (2001) *Fundamentals of Anatomy and Physiology*. Prentice Hall, New Jersey
- Maunsell J, Van Essen D (1983) Functional properties of neurons in the middle temporal visual area of the macaque monkey. i. selectivity for stimulus direction, speed and orientation. *Journal of Neurophysiology* 49:5:1127–1147
- McCoy A, Platt M (2005) Risk-sensitive neurons in the macaque posterior cingulate cortex. *Nature Neuroscience* 8:9:1220–1227
- McCulloch W, Pitts W (1943) A logical calculus of ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics* 5:115–133
- McIntosh A, Cabeza R, Lobaugh N (1998) Analysis of neural interactions explains the activation of occipital cortex by an auditory stimulus. *Journal of American Physiological Society* 90:2790–2796
- McNaughton B, Battaglia F, Jensen O, Moser E, Moser M (2006) Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience* 7:663–678
- Miller E, Erickson C, Desimone R (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience* 16:16:5154–5167
- Milner P (1957) The cell assembly: Mark II. *Psychological Review* 64:4:242–252
- Minsky M (1986) *The Society of Mind*. Simon and Schuster, New York
- Miyake A, Friedman N, Emerson M, Witzki A, Howeter A, Wager T (2000) The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: a latent variable analysis. *Cognitive Psychology* 41:49–100
- Mongillo G, Barak O, Tsodyks M (2008) Synaptic theory of working memory. *Science* 319:1543–1546

- Moore C, Nelson S (1998) Spatio-temporal subthreshold receptive fields in the vibrissa representation of rat primary somatosensory cortex. *Journal of Neurophysiology* 80:2882–2892
- Morrison J, Hof P (1997) Life and death of neurons in the aging brain. *Science* 278:412–419
- Nakamura K, Matsumoto K, Mikami A, Kubota K (1994) Visual response properties of single neurons in the temporal pole of behaving monkeys. *Journal of Neurophysiology* 71:3:1206–1221
- Nakamura K, Chung H, Graziano M, Gross C (1999) Dynamic representation of eye position in the parieto-occipital sulcus. *Journal of Neurophysiology* 81:2374–2385
- Naya Y, Sakai K, Miyashita Y (1996) Activity of primate inferotemporal neurons related to a sought target in pair-association task. *PNAS* 93:2664–2669
- Nee D, Jonides J (2008) Neural correlates of access to short-term memory. *PNAS* 105:37:14,228–14,233
- Nicolelis M, Lin R, Woodward D, Chapin J (1993) Dynamic and distributed properties of many-neuron ensembles in the ventral posterior medial thalamus of awake rats. *PNAS* 90:2212–2216
- Nicolelis M, Baccala L, Lin R, Chapin J (1995) Sensorimotor encoding by synchronous neural ensemble activity at multiple levels of the somatosensory system. *Science* 268:1353–1358
- Nishimura M, Scherf S, Behrmann M (2009) Development of object recognition in humans. *F1000 Biology Reports* 1:56
- Nitz D, Cowen S (2008) Crossing borders: sleep reactivation as a window on cell assembly formation. *Nature Neuroscience* 11:126–128
- O'Donnell P, Grace A (1995) Synaptic interactions among excitatory afferents to nucleus accumbens neurons: hippocampal gating of prefrontal cortical input. *Journal of Neuroscience* 15:3622–3639
- Ohl F, Scheich H, Freeman W (2001) Change in pattern of ongoing cortical activity with auditory category learning. *Nature* 412:733–736
- Olufsen M, Whittington M, Camperi M, Kopell N (2003) New roles for the gamma rhythm: Population tuning and preprocessing for the beta rhythm. *Computational Neuroscience* 14:33–54
- O'Neill J, Senior T, Allen K, Huxter J, Csicsvari J (2008) Reactivation of experience-dependent cell assembly patterns in hippocampus. *Nature Neuroscience* 11:209–215
- O'Reilly R (1999) Six principles for biologically-based computational models of cortical cognition. *Trends in Cognitive Sciences* 2:455–462
- O'Reilly R, Noelle D, Braver T, Cohen J (2002) Prefrontal cortex and dynamic categorization tasks: representational organization and neuromodulatory control. *Cerebral Cortex* 12:62–101
- Paller K (1997) Consolidating dispersed neocortical memories: the missing link in amnesia. *Memory* 5(1/2):73–88
- Palm G, Sommer F (1995) Associative data storage and retrieval in neural networks. In: Domany E, van Hemmen J, Schulten K (eds) *Models of Neural Networks III*, Springer

- Pastalkova E, Itskov V, Amarsingham A, Buzsáki G (2008) Internally generated cell assembly sequences in the rat hippocampus. *Science* 321:1322–1327
- Pasternak T, Greenlee M (2005) Working memory in primate sensory systems. *Nature Reviews Neuroscience* 6:97–107
- Pasupathy A, Connor C (2002) Population coding of shape in area v4. *Nature Neuroscience* 5:12:1332–1338
- Plate T (1995) Holographic reduced representations. *IEEE Transactions on Neural Networks* 6:3:623–641
- Plenz D, Thiagarajan T (2007) The organizing principles of neuronal avalanches: cell assemblies in the cortex? *Trends in Neurosciences* 30:3:101–110
- Poggio G, Fischer B (1977) Binocular interaction and depth sensitivity in striate and prestriate cortex of behaving rhesus monkeys. *Journal of Neuroscience* 40:6:1392–1405
- Ponzi A, Wickens J (1977) Input dependent cell assembly dynamics in a model of the striatal medium spiny neuron network. *Journal of Neuroscience* 40:6:1392–1405
- Procyk E, Tanaka Y, Joseph J (2000) Anterior cingulate activity during routine and non-routine sequential behaviours in macaques. *Nature Neuroscience* 3:502–508
- Pulvermuller F (1999) Words in the brain's language. *Behavioral and Brain Sciences* 22:253–336
- Purushothaman G, Bradley D (2005) Neural population code for fine perceptual decision in area MT. *Nature Neuroscience* 8:1:99–106
- Qin Y, McNaughton B, Skaggs W, Barnes C (1997) Memory reprocessing in corticocortical and hippocampocortical neuronal ensembles. *Philosophical Transactions of the Royal Society B* 352:1525–1533
- Qiu F, von der Heydt R (2005) Figure and ground in the visual cortex: V2 combines stereoscopic cues with gestalt rules. *Neuron* 47:155–166
- Quintana J, Fuster J (1999) From perception to action: Temporal integrative functions of prefrontal and parietal neurons. *Cerebral Cortex* 9:213–221
- Quiroga R, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435:1102–1107
- Raizada R, Grossberg S (2003) Towards a theory of the laminar architecture of cerebral cortex: computational clues from the visual system. *Cerebral Cortex* 13(1):100–113
- Ranganath C, Cohen M, Brozinsky C (2005) Working memory maintenance contributes to long-term memory formation: Neural and behavioral evidence. *Journal of Cognitive Neuroscience* 17:994–1010
- Reinagel P, Reid R (2000) Temporal coding of visual information in the thalamus. *Journal of Neuroscience* 20(14):5392–5400
- Rennaker R, Chen C, Ruyle A, Sloan A, Wilson D (2007) Spatial and temporal distribution of odorant-evoked activity in the piriform cortex. *Journal of Neuroscience* 27(7):1534–1542
- Repa J, Muller J, Apergis J, Desrochers T, Zhou Y, LeDoux J (2001) Two different lateral amygdala cell populations contribute to the initiation and

- storage of memory. *Nature Neuroscience* 4(7):724–731
- Reynolds G, Richards J (2009) Cortical source localization of infant cognition. *Developmental Neuropsychology* 34(3):312–329
- Riches I, Wilson F, Brown M (1991) The effects of visual stimulation and memory on neurons of the hippocampal formation and the neighboring parahippocampal gyrus and inferior temporal cortex of the primate. *Journal of Neuroscience* 11(6):1763–1779
- Riesen A (1947) The development of visual perception in man and chimpanzee. *Science* 106:107–108
- Rochester N, Holland J, Haibt L, Duda W (1956) Tests on a cell assembly theory of the action of the brain using a large digital computer. *RE Transaction on Information Theory* IT-2:80–93
- Roelfsma P, Engel A, Konig P, Singer W (1997) Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature* 385:157–161
- Rolls E (2008) Functions of the orbitofrontal and pregenual cingulate cortex in taste, olfaction, appetite and emotion. *Acta Physiologica Hungarica* 95(2):131–164
- Rolls E, Inoue K, Browning A (2003) Activity of primate subgenual cingulate cortex neurons is related to sleep. *Journal of Neurophysiology* 90:134–142
- Romanski L, Averbeck B, Diltz M (2004) Neural representation of vocalizations in the primate ventrolateral prefrontal cortex. *Journal of Neurophysiology* 93:734–747
- Sadato N, Pascual-Leone A, Grafman J, Ibanez V, Deiber M, Dold G, Hallett M (1996) Activation of primary visual cortex by braille reading in blind subjects. *Nature* 380:526–528
- Sakai K, Miyashita Y (1991) Neural organization for the long-term memory of paired associates. *Nature* 354:152–155
- Sakurai Y (1998a) Cell-assembly coding in several memory processes. *Neurobiology of learning and memory* 70:212–225
- Sakurai Y (1998b) The search for cell assemblies in the working brain. *Behavioural Brain Research* 91:1–13
- Sakurai Y, Tuakahashi S, Inoue M (2004) Stimulus duration in working memory is represented by neuronal activity in the monkey prefrontal cortex. *European Journal of Neuroscience* 20:1069–1080
- Salihoglu U, Bersini H, Yamaguchi Y, Molter C (2009) Online unsupervised formation of cell assemblies for the encoding of multiple cognitive maps. *Neural Networks* 22:687–696
- Scherberger H, Jarvis M, Anderson R (2005) Cortical local field potential encodes movement intentions in the posterior parietal cortex. *Neuron* 46:347–354
- Schoenbaum G (1998) Cell assemblies and the ghost in the machine. In: Eichenbaum H, Davis J (eds) *Neuronal Ensembles: Strategies for Recording and Decoding*, Wiley, pp 81–116
- Seamans J, Nogueira L, Lavin L (2003) Synaptic basis of persistent activity in prefrontal cortex in vivo and in organotypic cultures. *Cerebral Cortex*

- 13:1242–1250
- Setola P, Reilly R (2005) Words in the brain's language: An experimental investigation. *Brain and Language* 94:251–259
- Shoham D, Glaser D, Arieli A, Kenet T, Wijnbergen C, Toledo Y, Hildesheim R, Grinvald A (1999) Imaging cortical dynamics at high spatial and temporal resolution with novel blue voltage-sensitive dyes. *Neuron* 24:791–802
- Siegel G, Carter C, Thase M (2006) Use of fmri to predict recovery from unipolar depression with cognitive behavior therapy. *American Journal of Psychiatry* 163:735–738
- Sigala N, Gabbiani F, Logothetis N (2002) Visual categorization and object representation in monkeys and humans. *Journal of Cognitive Neuroscience* 14:2:187–198
- Sigala N, Kusunoki M, Nimmo-Smith I, Gaffan D, Duncan J (2008) Hierarchical coding for sequential task events in the monkey prefrontal cortex. *PNAS Neuroscience* 105:33:11,969–11,974
- Silvanto J, Muggleton N (2008) New light through old windows: Moving beyond the "virtual lesion" approach to transcranial magnetic stimulation. *NeuroImage* 39:549–552
- Singer W, Engel A, Kreiter A, Munk M, Neuenschwander S, Roelfsema P (1997) Neuronal assemblies: necessity, signature and detectability. *Trends in Cognitive Sciences* 1:7:252–261
- Smith K (2010) Settling the great glia debate. *Nature* 468:160–162
- Sougné J (2001) Binding and multiple instantiation in a distributed network of spiking neurons. *Connection Science* 13:99–126
- Stackman R, Taube J (1998) Firing properties of rat lateral mammillary single units: head direction, head pitch, and angular head velocity. *Journal of Neuroscience* 18(21):9020–9037
- Steriade M, Nunez A, Amzica F (1993) A novel slow (≈ 1 Hz) oscillation of neocortical neurons in vivo: Depolarizing and hyperpolarizing components. *Journal of Neuroscience* 13(8):3252–3265
- Stern E, Jaeger D, Wilson C (1998) Membrane potential synchrony of simultaneously recorded striatal spiny neurons in vivo. *Nature* 394:475–478
- Sutherland G, McNaughton B (2000) Memory trace reactivation in hippocampal and neocortical neuronal ensembles. *Current Opinion in Neurobiology* 10:180–186
- Taddeo M, Floridi L (2005) Solving the symbol grounding problem: a critical review of fifteen years of research. *Journal of Experimental and Theoretical Artificial Intelligence* 17:4:419–445
- Talk A, Kang E, Gabriel M (2004) Independent generation of theta rhythm in the hippocampus and posterior cingulate cortex. *Brain Research* 1015:15–24
- Tark K, Curtis C (2009) Persistent neural activity in the human frontal cortex when maintaining space that is off the map. *Nature Neuroscience* 12:1463–1468
- Terman D, Wang D (1995) Global competition and local cooperation in a network of neural oscillators. *Physica D* 81:148–176

- Tudusciuc O, Nieder A (2007) Neuronal population coding of continuous and discrete quantity in the primate posterior parietal cortex. *PNAS* 104:36:14,513–14,518
- Usher M, Donnelly N (1998) Visual synchrony affects binding and segmentation in perception. *Nature* 394:179–182
- Valiant L (2005) Memorization and association on a realistic neural model. *Neural Computation* 17:527–555
- van der Velde F, de Kamps M (2006) Neural blackboard architectures of combinatorial structures in cognition. *Behavioral and Brain Sciences* 29:1–72
- von der Malsburg C (1981) The correlation theory of brain function. Tech. rep., Dept. of Neurobiology, Max-Planck-Institute for Biophysical Chemistry
- Wallace D, Kerr J (2010) Chasing the cell assembly. *Current Opinion in Neurobiology* 20:296–305
- Wallis J, Anderson K, Miller E (2001) Single neurons in prefrontal cortex encode abstract rules. *Nature* 411:953–956
- Wang X (2001) Synaptic reverberation underlying mnemonic persistent activity. *Trends in Neurosciences* 24:8:455–463
- Wang Y, Markram H, Goodman P, Berger T, Ma J, Goldman-Rakic P (2006) Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nature Neuroscience* 9:534–542
- Webb S, Long J, Nelson C (2005) A longitudinal investigation of visual event-related potentials in the first year of life. *Developmental Science* 8:6:605–616
- Wennekers T, Palm G (2000) Cell assemblies, associative memory and temporal structure in brain signals. In: Miller R (ed) *Time and the Brain: Conceptual Advances in Brain Research* (Vol. 2), Harwood Academic Publishers, pp 251–274
- Wessberg J, Stambaugh C, Kralik J, Beck P, Laubach M, Chapin J, Kim J, Biggs S, Srinivasan M, Nicolelis M (2000) Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 408:361–365
- Wichert A (2001) Pictorial reasoning with cell assemblies. *Connection Science* 13:1:1–42
- Wickelgren W (1999) Webs, cell assemblies, and chunking in neural nets. *Canadian Journal of Experimental Psychology* 53:1:118–131
- Ylinen A, Bragin A, Nadasdy Z, Jando G, Szabo I, Sik A, Buzsaki G (1995) Sharp wave-associated high-frequency oscillation (200hz) in the intact hippocampus: network and intracellular mechanisms. *Journal of Neuroscience* 15:1:30–46
- Yoshida M, Hasselmo M (2009) Persistent firing supported by an intrinsic cellular mechanism in a component of the head direction system. *Journal of Neuroscience* 29:15:4945–4952
- Yuste R, MacLean J, Smith J, Lansner A (2005) The cortex as a central pattern generator. *Nature Reviews Neuroscience* 6:477–483
- Zelano C, Montag J, Khan R, Sobel N (2009) A specialized odor memory buffer in primary olfactory cortex. *PLOS One* 4 4:3:e4965
- Zhou Y, Fuster J (1996) Mnemonic neuronal activity in somatosensory cortex. *PNAS* 93:10,533–10,537

Zucker R, Regehr W (2002) Short-term synaptic plasticity. *Annual Review of Physiology* 64:355–405