Charles R. Legéndy*

Synaptic and extrasynaptic traces of long-term memory: the ID molecule theory

Abstract: It is generally assumed at the time of this writing that memories are stored in the form of synaptic weights. However, it is now also clear that the synapses are not permanent; in fact, synaptic patterns undergo significant change in a matter of hours. This means that to implement the long survival of distant memories (for several decades in humans), the brain must possess a molecular backup mechanism in some form, complete with provisions for the storage and retrieval of information. It is found below that the memory-supporting molecules need not contain a detailed description of mental entities, as had been envisioned in the 'memory molecule papers' from 50 years ago, they only need to contain unique identifiers of various entities, and that this can be achieved using relatively small molecules, using a random code ('ID molecules'). In this paper, the logistics of information flow are followed through the steps of storage and retrieval, and the conclusion reached is that the ID molecules, by carrying a sufficient amount of information (entropy), can effectively control the recreation of complex multineuronal patterns. In illustrations, it is described how ID molecules can be made to revive a selected cell assembly by waking up its synapses and how they cause a selected cell assembly to ignite by sending slow inward currents into its cells. The arrangement involves producing multiple copies of the ID molecules and distributing them at strategic locations at selected sets of synapses, then reaching them through small noncoding RNA molecules. This requires the quick creation of entropy-rich messengers and matching receptors, and it suggests that these are created from each other by small-scale transcription and reverse transcription.

Keywords: cell assembly; high-entropy messenger; ignition; long-term memory; RNA; synapse.

*Corresponding author: Charles R. Legéndy, Department of Psychology, Columbia University, 406 Schermerhorn Hall, 1190 Amsterdam Ave., New York, NY 10027, USA, e-mail: crl2122@columbia.edu

Introduction

Traditionally, memory had been a concept of old-time psychology, the science of the 'psyche.' It was an abstract notion divorced from mechanistic reality, and the attempt to explain it in terms of the known hardware of the nervous system (Cajal, 1894; McCulloch and Pitts, 1943; Konorski, 1948; Hebb, 1949; von Neumann, 1952) has marked a qualitative change in the approach to understanding the brain.

From very early on, the explanations have centered on the synapses and were built around the proposition that the contribution of the individual synapse to the electrical potential inside a neuron, formally expressible in terms of a synaptic weight, is capable of plastic change.

Formal exploration of the molecular mechanisms underlying synaptic change was begun in the 1960s by Kandel and coworkers in the marine mollusk *Aplysia* (Frazier et al., 1967; Kandel, 2001). After locating the neurons and synapses involved in a simple conditioned reflex, the gill withdrawal reflex (Pinsker et al., 1970), the Kandel group was able to extend the *Aplysia* work to learning and succeeded, over a number of years, in mapping out the detailed synaptic events of learning and memory at the molecular level. Significantly, the results obtained on *Aplysia* have been found to carry over to higher animals remarkably well (Kandel et al., 2014).

Meanwhile, the study of synapse-level molecular mechanisms has been extended to higher animals and has given rise to the large body of literature under the heading of long-term potentiation (LTP) and long-term depression (LTD) (Lømo, 1966; Bliss and Gardner-Medwin, 1973; Bliss and Lømo, 1973; Dudek and Bear, 1992; Bliss and Collingridge, 1993; Collingridge et al., 2010).

The logistics of information flow underlying LTP/LTD, central to the present paper, has been addressed only tangentially, notably in the synaptic tagging hypothesis (Krug et al., 1984; Frey and Morris, 1997; Sanes and Lichtman, 1999; Martin and Kosik, 2002; Sajikumar and Frey, 2004; Reymann and Frey, 2007; Sajikumar et al., 2007; Redondo

and Morris, 2011), to which I will return in a moment. (In the following, for the sake of brevity, I will simply write to 'LTP' instead of 'LTP/LTD,' with the understanding that LTD and LTP actually go together in most of the contexts mentioned.)

A separate line of inquiry, also directed at the participation of molecules in learning and memory, has linked performance in behavioral tasks to the application of protein synthesis inhibitors (see Davis and Squire, 1984; Nader et al., 2000; Gold, 2008). In an extension of the same effort, a long list of small polypeptide molecules, the neuropeptides, has been compiled and sequenced over the past years (see Fricker, 2012; Borbély et al., 2013).

An ambitious but unsuccessful series of papers appeared in the 1960s and 1970s, the so-called 'memory molecule papers,' whose line of inquiry has by now been largely discredited and forgotten. Morange (2006) reviewed this piece of the literature, carefully separating the discredited papers (McConnell, 1962; Babich et al., 1965) from the ones which obtained legitimate results and merely overinterpreted them (Hydén and Egyházi, 1962; Flexner et al., 1963; Ungar et al., 1972).

The present hypothesis revisits the 'memory molecule' concept in the light of knowledge acquired in the years since the old papers and recasts it in less ambitious terms. In the present model, the molecules do not contain any structural details of mental entities; instead, they carry unique identifiers for them. The identifiers, which take the form of molecular 'code words,' are randomly generated and assigned to mental entities, then coordinated with macroscopic information processing through cell assemblies. Accordingly, they are not transferable from one individual to another – which explains why the paradigm of the old papers could not have worked.

The recent discovery of microDNA molecules (Shibata et al., 2012; Dillon et al., 2015) offers an insight into the way in which random nucleotide sequences, usable as templates for the code words, may become available within the cells (see below).

Mental entities, in the rudimentary form considered here, will be represented by cell assemblies, noting that the cell assembly concept can be extended to accommodate complex syntactic structures between cell assemblies (Legéndy, 2009). In this way, the problem of accessing mental entities reduces to the problem of accessing their cell assemblies.

The cell assembly model to be used in the present paper is based on Hebb's (1949) hypothesis combined with the ignition concept (Rapoport, 1952; Legéndy, 1967, 2009;

Palm, 1981b, 1982; Wickelgren, 1999). The general framework of the cell assembly hypothesis relies on several lines of evidence (Gerstein et al., 1989; Abeles, 1991; Harris et al., 2003; Ikegaya et al., 2004; Buzsáki, 2010) which, in addition to formal data, include the inferences that can be made from dendritic spike data and cross-correlation data (see 'The size of cell assemblies').

One property of cell assemblies useful in supporting the communication between macroscopic events and the molecules is, as will be seen, that the many degrees of freedom in the choice of cell assembly membership, under proper conditions, make a cell assembly recognizable on the basis of small subsets of its cells, such as the cells synapsing on a piece of dendrite.

Three types of molecules and molecular systems are described below. Two of them are in the nature of unique identifiers, the 'ID molecules' in the title of this paper, and the third one is a molecular complex implementing the communication between them and the cell assemblies.

The ID molecule hypothesis

Critique of the 'classical' view of memory

The existing literature on LTP, in its approach to information storage, is consistent with the 'classical' view of memory that states, broadly speaking, that memory is governed by a gradual evolution of synaptic strengths, under Hebbian and other rules, and that the responses of neurons to their synaptic input is determined by whatever happens to be the latest configuration of the synaptic weights. Put in another way, this view of memory states that the memory trace is made up of the synaptic weights and nothing else (for a review, see Martin et al., 2000).

The present paper presents a challenge to this classical view and raises a need for an extrasynaptic contribution to memory, one that is intrinsically multisynaptic, with no single synapse indispensable, and that includes a mechanism for enabling the network to store whole patterns of synaptic weights, as units, and then restoring them on demand later. The group-oriented handling of synaptic patterns can, among other things, vastly increase memory capacity.

The need for modifying the classical concept has become apparent in view of recent data, along with the reevaluation of old data, which all point to some problems that cannot be addressed within the old framework.

The two most serious ones are described in the next sections.

Synapses are not permanent

Synapses do not last forever. New synapses can arise in an adult brain, while other synapses disappear in comparable numbers (Lendvai et al., 2000; Trachtenberg et al., 2002; Stettler et al., 2006). A quantitative value has been attached to the synaptic turnover rate by Stettler et al. (2006), who examined the appearance and disappearance of axonal boutons in the intact visual cortex in monkeys, in a context where no experimentally induced learning or trauma could be held responsible, and found the turnover rate to be 7% per week – which would give the average synapse a lifetime of a little over 3 months.

Dramatic changes in overall synaptic spine counts have also been found during certain changes of the global environment. In an electron microscope study, Kirov et al. (1999) compared spine counts between perfusion-fixed hippocampal material and slices in which the neurons were kept alive but were deprived of their usual input. They found that the slices contained 40%–50% more synaptic spines and determined that the added spines appeared in the first few hours after slicing. The extra spines were verified, through serial-section electron microscopy, to carry full-fledged synapses with all the usual presynaptic and postsynaptic elements.

Spine counts can also undergo dramatic change without laboratory intervention. Female rats in the 24 h between proestrus and estrus show a 30% decrease in the spine counts in CA1 pyramidal cells, then the spine counts return to their earlier levels during the rest of the cycle (Woolley et al., 1990).

It has further been found that, under stress, axons in the cortex can sprout new branches and let old ones disappear in substantial numbers, taking all their synapses with them (Yamahachi et al., 2009), indicating that even the configuration of synapse-carrying fibers in the brain is not stable.

The mutually opposing effects of LTP and LTD further add to the eventual disappearance of the memory maintained in the form of synaptic strengths. Successive events of LTP and LTD, occurring in diverse and unrelated contexts, counteract and overwrite each other and will, as time goes by, tend to obliterate old patterns of synaptic weights, covering them with layers of new ones. Once again, we are led to the conclusion that the pattern of synaptic strengths cannot be relied upon to preserve, for instance, childhood memories.

The issue of brain capacity

A second problem with the classical view of memory is just as serious as the first, even though it does not yield to similarly clear-cut experimental demonstration; it arises from theoretical considerations of storage capacity.

It is sometimes tacitly assumed, in justifying the classical view, that the vast number of neurons and synapses in the brain ensures a correspondingly large memory capacity because of all the possible permutations that can be formed from all the synapses. However, this is a fallacy. In the higher animals, single neurons and single synapses are useless on their own, as the neurons must act in coordinated groups when exerting their effect on the network. The number of independently adjustable variables is determined by the group structure and is therefore drastically reduced. The viable forms of group structure in the brain have been explored in the cell assembly literature.

In the wake of Hebb's (1949) publication of the cell assembly hypothesis, a number of computations and computer simulations have been published, exploring ways in which the Hebbian conceptual framework could be translated into concrete circuitry (reviewed by Scott, 1975, 1977). The most elaborate one of the brain capacity calculations (Legéndy, 1967) was aimed at determining the number of cell assemblies able to fit into a randomly connected uniform network before interference renders their information irretrievable. The calculations note that cell assemblies must be large enough to be ignitable and interconnectable, and the large size of cell assemblies significantly reduces the number of independently adjustable parameters in the network. (The randomly connected uniform network, as has been learned since then, is a good approximation to the hippocampal CA3 network; see Buzsáki, 2006; Wittner et al., 2007).

In the 1967 calculations, the objective was to determine an upper limit to the number of cell assemblies possible under the purely synapse-based ('classical') view, which meant that all parameters not known at the time were selected so as to maximize the estimated capacity. This in fact led to a fairly large number, of the order of 10^9 cell assemblies, and favorable comments by Scott (1975, 1977) and Hebb himself (1976) encouraged the idea, for a while, that a classical memory model might lead to an acceptably large storage capacity.

However, in hindsight, the assumptions leading to the large capacity were unrealistic. Cell assemblies are likely to be several orders of magnitude larger than the size allowed in the 1967 paper (see 'The size of cell assemblies' section below), which means that correspondingly

fewer of them will fit into the model network because the classical memory model does not permit any cell assemblies to be made temporarily unavailable so as to reduce overlap and make room for more cell assemblies. Accordingly, the classical model (Legéndy, 1967), when updated to account for currently available data, falls short of the originally calculated capacity of 10^9 cell assemblies by several orders of magnitude.

To this day, as far as I know, no purely synapse-based memory model exists with an acceptable storage capacity.

The result is, then, that even if synapses were to last forever, a memory model relying on them alone could not hold up because it could not explain the large capacity of actual brains.

The linkage between LTP and surprising events

One challenge in understanding memory is the difficulty of modeling the information flow between neurons. To appreciate the problem, it is enough to point out that in the rat hippocampus, a typical pyramidal cell has about 30,000 excitatory synapses (Megías et al., 2001); in the mouse cerebral cortex, the corresponding number is 8000 (Braitenberg and Schüz, 1998).

The continual shower of spikes arriving to all these synapses differs in an essential way from the input streams to gates in a digital computer, which are precisely known and can be combined to provide knowable output signals. The input streams to neurons come from sources largely unknown. In the higher animals, the specificity of cell adhesion molecules and other determinants of connectivity are only enough to identify the broader biological neuron types, not the individual neurons. Accordingly, the genetic blueprint (unlike the computer blueprint) provides no basis for assigning a functional specification of each neuron's output.

The standard approach to dealing with data incompletely understood is the use of statistics. The simplest statistics-based concept available for use in the nervous system, and one often used in the literature, ignores individual neuronal spikes and confines the discussion to the average spike rates of neurons (see for example Caianiello, 1961; Wilson and Cowan, 1972; Bienenstock et al., 1982). However, the spike-rate-based events and changes predicted in neuronal networks are relatively slow, and their functional relevance to learning is at best obscure.

When quick events and quick changes need to be described, and a straightforward interpretation is sought

in terms of mental entities and known data, it is preferable to invoke another statistical notion, the *surprise* concept (Legéndy, 1970, 1975, 2009; Palm, 1981a, 2012; von der Malsburg, 1999), tailored to the demands of quick neuron-level decision making. (The term 'surprise,' as used in the context of the neuronal networks, has been defined to quantify the degree to which a pattern of firing is unlikely to occur by accident, in other words, unlikely to occur under the baseline spike statistics – for instance under Poisson distribution of firing in the case of *Poisson surprise*; Legéndy and Salcman, 1985).

It so happens that the neuronal events inducing LTP closely match the published definition of surprising events. The interplay between N-methyl-D-aspartate (NMDA) receptors, alpha-amino-3-hydroxy-5-methyl-4-isoxazole-propionic acid (AMPA) receptors, and Mg^{2+} ions results in responsiveness to spike coincidence and repetition (Larkum et al., 2009; Major et al., 2013; Palmer et al., 2014). Since a large class of surprising events is distinguished by nothing more than spike coincidence and repetition, these receptor systems make for highly effective surprise detectors.

In general, it can be said that the role of surprising events in the brain, originally postulated on purely theoretical grounds, is powerfully demonstrated by the success of LTP in producing plastic change in a wide range of contexts.

Experimental induction of LTP usually involves dramatic application of surprising input to the cells of interest, such as 100 or more repetitions of intense electrical stimulation. Probably the least extreme of the experiments successfully inducing LTP have been those of Remy and Spruston (2007), with five volleys of spikes delivered 10 ms apart to a set of 25–40 synapses (on hippocampal CA1 pyramidal cells, via stimulation of the Schaffer collaterals). The latter stimulation initiated dendritic spikes most of the time, and at the times when it did, LTP was found to be induced.

Under natural conditions, the events inducing LTP are expected to be far less dramatic than the experimentally produced events seen in the literature, but still enough to permit the crucial inference that, when viewed from the perspective of a neuronal locality, they appear unlikely to be accidental.

The linkage between surprising events and underlying causes

By the usual syllogism of statistics, when a highly surprising firing event occurs, even once, its occurrence

lends support to the assumption that the occurrence has reasons outside the mechanisms determining the baseline statistics. In statistics-based laboratory research, the equivalent of a surprising 'event' can sometimes be an entire research project from beginning to end, and its 'surprise' reflects the probability that the whole of the data is a fluke, its results all accidental.

In the case of neuronal spike trains, the surprising event is generally brief, for instance brief enough to induce LTP. The underlying cause of the event is not knowable at the locality of the occurrences; however, the local descriptive details of the surprising event, such as the set of synapses affected by its volleys, are capable of being recorded.

Not all surprising events are expected to induce LTP; an important class of them relies on earlier LTP. The surprise in these may be called 'prediction surprise.' It does not require repetition of a volley but can arise from a single occurrence of a volley, which must, however, be preceded by a predictive event that 'wakes up' the same set of synapses as would be visited by the volley (see Figure 3), thus setting up the synapses for a volley that may subsequently arrive.

An important distinction between entropy (Shannon, 1948) and surprise gains prominence at the stage of recording. The surprise content indicates whether the pattern is distinct from the surrounding noise; the entropy content indicates whether the pattern is rich enough in descriptive detail to be recognized later if recorded. For instance, tetanic stimulation of a single synapse is highly surprising but remains entropy-poor. (The classic example of such entropy-poor synaptic excitation is the mossy fiber input to the hippocampal CA3 pyramidal cells – which, significantly, is also an example of NMDA-receptor-independent LTP induction; Harris and Cotman, 1986; Langdon et al., 1995).

When a surprising event is sufficiently entropy-rich to be recognized later, it becomes meaningful to speak of 'recurrences' of the event. It can then be said that when the event recurs, the recurrence is likely to signify the presence of the same underlying cause as it did earlier, and that if a neuron is equipped to respond to the recurrences, its responses (with high probability) signify the same underlying cause as they did at recording.

This, then, is the rationale behind Figure 1 and the assertion, implicit in the drawings, that the indicated spike patterns signify describable objects. The patterns are understood to be LTP-inducing patterns, past or ongoing, and the accompanying icons embody the assumption that such patterns can correspond to underlying causes.
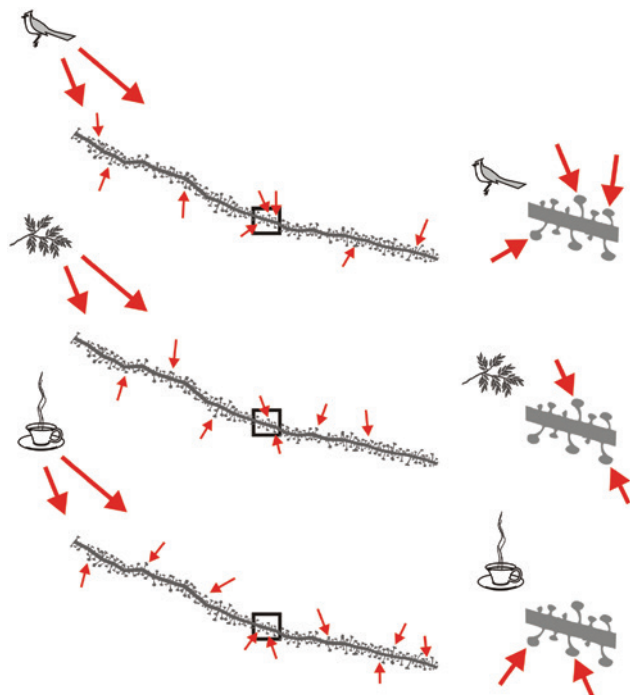
Figure 1: Meaningful input patterns incident on a segment of dendrite.

Synaptic patterns of input volleys, issued by area-wide signal fields (igniting cell assemblies), can be designated by symbols representing the underlying causes. A segment of a dendrite, shown as having about 160 viable synapses on it, is reproduced in three copies, each receiving a volley of spikes (arrows) on a subset of its synapses from presynaptic elements (not shown). The icons of 'bird,' 'leaves,' and 'coffee' represent the underlying causes responsible for the synaptic input patterns. A short segment of the dendrite (framed) is magnified at the right of each drawing and is reproduced in some of the subsequent figures. Note: here and in the figures below, I chose icons representing visually clear-cut items, such as birds and the like, since those are the easiest to illustrate in a drawing. The intent, however, is to symbolize mental entities in the most general sense of the word, including locations, smells, sounds, goals, sequences, actions, and the rest, together with parts of them as well as parts of parts, and, as will be appreciated, most of these are not suitable for being shown in drawings.

Ignition of cell assemblies: the source of the 'surprising' events

Cell assemblies are static objects; they are merely collections of neurons with strengthened synaptic contacts between them. What enables them to be heard is that they can be made to ignite (Rapoport, 1952; Legéndy, 1967, 2009; Palm, 1981b, 1982; Wickelgren, 1999). During ignitions, the strengthened contacts between cell assembly members facilitate rapid spread of firing through the assembly, which continues until all neurons of assembly have fired.

To neurons in the rest of the brain, the ignitions can, under proper conditions, be made highly surprising. Clearly, the ignitions are not 'surprising' to a human observer who knows the mechanism of their generation, but it is noted that the definition of surprise is formulated from the vantage point of an observer unaware of events outside the observable baseline spike statistics – and, more to the point, from the vantage point of the neurons themselves.

Far from being spurious events, the ignitions can be purposefully induced, as described in the sections below. Any selected cell assembly, even one that has been dormant for a long time, can be revived and made ignitable through targeted strengthening of its connecting synapses and any ignitable cell assembly can be made to ignite through targeted current injection into its neurons.

Each time a cell assembly ignites, it gives off a set of spikes that, helped by the completion property inherent in ignitions, carry a unique signature of the assembly. The uniqueness, paired with the surprise of the coactivity, makes for a powerful mechanism of broadcasting information; it enables detailed information to reach the neurons despite the constant background noise throughout the network.

The time structure of an ignition determines whether the neurons it reaches receive volleys or barrages of spikes. When the ignition concept was first introduced and later applied to cell assemblies (Rapoport, 1952; Legéndy, 1967), it was described as a quick changeover from the resting activity of cells to prolonged reverberation among cell assembly members or, what is similar in its effect, to attractor behavior (Hopfield, 1982; Amit, 1989). Either of these amounts to a sudden upward shift in the spike rates by all cell assembly members, acting on other neurons as a barrage of spikes, lasting until silenced.

Another form of ignition (Legéndy, 2009, p. 18) sends out one brief volley each time it occurs and can be made to recur a few times to increase the surprise. (For simplicity, I will refer to the output events from both forms of ignition as 'volleys.') In the latter kind of ignition, assumed here to be common in the cortex and hippocampus, the spread of firing through the cell assembly lasts only until all members have fired once and (ideally) only once. Reverberation is suppressed by a relatively long refractory period, and a few surprise-enhancing repetitions can be induced via controlled instability (Legéndy, 2009, p. 35).

Ignitions of either kind can impart substantial surprise to the neurons they reach, and since they both share the completion property, their volleys carry the signatures of the igniting cell assemblies.

The problem of saturation in purely synapse-based memory

In Figure 1, each set of arrows has been drawn to cover roughly 5% of the synapses (Losonczy and Magee, 2006); in the context of interest, the marked synapses are the ones reinforced when the corresponding pattern induces LTP. With the thresholds appropriately set, and no other synapses reinforced, a later volley precisely reproducing the LTP-inducing volley will bring the dendrite beyond threshold, whereas other random volleys over the same number of synapses will not. In this sense, the reinforcement of synapses can 'imprint' a volley on the dendrite, making it selectively responsive to the volley.

However, there is clearly a limit to the number of patterns that can be imprinted on dendrites in this way. If all three of the patterns shown, or three comparable ones, were imprinted on the dendrite of Figure 1, the percentage of synapses reinforced would rise to about 14% (accounting for overlap); the enhanced-weight synapses would still remain fairly 'sparsely' distributed (Palm et al., 2014), and selective responsiveness would still remain moderately reliable. The estimated number of patterns that can be imprinted and later retrieved becomes somewhat greater when it is noted that different events can end up imprinted on different dendritic segments and on different neurons, but all the same, as the number of imprinted patterns is increased, the dendrites gradually become 'saturated,' until all synapses, or almost all, are reinforced and selective responsiveness disappears.

In practice, repeated events of LTD steadily reduce the synaptic weights and, together with events of LTP, achieve subtractive normalization (Legéndy, 1978; Bienenstock et al., 1982; Oja, 1982; Miller and MacKay, 1994; Song et al., 2000), but while this lowers the synaptic weights, it will not undo the effects of saturation. Over time, all synapses will have been reinforced and depressed multiple times, making the synaptic weights more or less uniformly random, and ruin selective response to all but the most recent imprints. In the most recent imprints, the synaptic weights are raised after the rest have already been lowered, and for a while, the enhanced synapses remain *sparsely* distributed; they stand out above the rest well enough to support selective response.

When the distribution is not sparse, overlap between cell assemblies raises the chance of parasitic spreading of ignitions from one cell assembly to the others, and possibly to the whole network. Up to a point, probabilities can guard against the spread of firing (Legéndy, 1967), but during times of intense function (Buzsáki, 2005), when many ignition sequences are simultaneously active, the

likelihood of devastating chain-reaction-like spread of firing to the whole network increases, and a wave of inhibition must repeatedly step in to bring the spread of firing under control. It is possible that such defensive inhibition plays a role during the inhibitory phase of gamma and theta waves (Singer, 1993; Buzsáki, 2002).

Retrieval of information remains possible, despite the waves of inhibition, as long as the parasitic spread of firing is slow enough that the ignitions can achieve completion before inhibition shuts them off. Recognitions will only fail when inhibition is forced to overwhelm all firing before the ignitions can achieve completion.

The picture described in the last few paragraphs, as will be appreciated, is of the 'classical' picture of memory, with all its weaknesses. The molecule-assisted model of memory, outlined in the remainder of this paper, leaves much of the classical picture in place, including recognition via enhanced synapse sets, saturation, and the depression of weights via LTD, but it adds a 'backup' mechanism that can, on demand, raise the weights of selected sets of synapses, making them similar to recently enhanced synapses and enabling them to support selective response.

Next, let me digress for a moment and describe a concept, borrowed from coding theory, which will play a central role in the model.

The issue of unique labels: Shannon's random code ensemble

Mental entities within the brain are accessible at will; they are not constrained to any predetermined order in which they can be called up. In the implementation of memory retrieval, this means that the calling agent must have a way to single out the mental entity to be retrieved, for instance through a unique identifier of some form.

The uniqueness of such identifiers is crucial to their function. Unique labels, such as serial numbers, are present in computerized databases, where their uniqueness can be guaranteed simply by making sure that no two items get the same label.

The problem is that in the brain the same thing is not possible. In the brain, new items, such as new mental entities and new components of mental entities, are added to the system from varied and mutually independent sources all the time, and since these sources often do not communicate with each other, it is not feasible to design protocols to guard against repetitions in the assignment of unique identifiers. This is where the idea of a random code offers a solution.

It is said that no two people have the same fingerprint, but in fact the uniqueness of fingerprints is only probabilistic. It is not rigorously arranged in a systematic way; the fingerprints are randomly generated during early development of the skin. Their usefulness comes from the fact that the random variation has so many degrees of freedom that there are many more distinguishable ridge patterns than there are people.

Calculating the probability of accidental coincidences between randomly generated patterns is easiest in a formalized paradigm where the patterns are simply random sequences of binary digits; for instance one may envision a set of n-bit binary numbers ('code words'), each of them generated by flipping a coin n times. The arrangement, in its general form, is known as 'Shannon's random code ensemble' (Shannon, 1957; Mézard and Montanari, 2009, p. 107).

One can estimate the probability of accidental coincidences between random binary code words by using the '2n rule,' which states that, if an ensemble has N code words, each made up of n randomly generated bits, then by choosing n to be at least twice the minimum number of bits needed to create N distinct words, the probability that any two of them will coincide can be kept below 0.5. Further, each additional binary digit doubles the safety factor. For instance, if an ensemble has 256 code words, then (since $256=2^8$) one needs 16 random digits per code word to ensure a 0.5 probability that no two code words are equal, and if instead the strings have 26 bits, the probability that any two are equal becomes 1/1024.

(The '2n rule' can be proven as follows:

Let N=number of code words; n=number of random digits in each code word; p=the probability that some two code words are the same. With these:

$$p = \frac{N(N\text{-}1)}{2} 2^{-n} \cong 2^{-n-1} N^2 = 2^{-n-1+2\log_2 N}$$

Then if, for instance, the requirement is that the $p<\frac{1}{2}$, the latter result can be written as $2^{-n-1+2\log_2 N}<2^{-1}$ or, alternatively, as $n>2 \log_2 N$, which is what the '2n rule' states. As seen from the previous forms, each added binary digit in the code words halves p or doubles the safety factor.)

Random code ensembles are often used in communication networks because their code words are easy to generate, and when sufficient redundancy is introduced, they have powerful error correcting capability (Hamming, 1950).

As will be appreciated, molecules made up of sequences of amino acids or nucleic acids are well suited for implementing a system of random coding – and they do not need to be very long to permit a vast number of variations.

The nature of macromolecular contribution: the ID molecule 'code words' are probably RNA and DNA sequences

The objective of the present paper, briefly stated, is to outline a method of putting cell assemblies into molecular storage, by way of recording their synaptic patterns before those patterns are buried under layers of subsequent synaptic changes, then to outline a method of reconstituting the cell assemblies from the molecules at a later time.

The challenge is that the things to be put into storage and retrieved later are *sets* of synapses and sets of neurons, to be treated together, as units, which means that the underlying machinery needs to be able to specify in detail which synapses are to be restored when reconnecting a cell assembly, and which neurons must receive current injection when igniting it. Since cell assemblies can contain very many neurons, this is no mean feat.

The group-oriented restoration protocols are necessary if the system is to utilize the combinatorial diversity of synapse choices and neuron choices and unleash the full memory capacity achievable by neuron grouping.

A multisynaptic pattern arriving to a piece of dendrite can carry a considerable amount of information (entropy). In Figure 1, for instance, the pattern of active synapses in the 'bird' drawing singles out 1 out of about 10^{13} possible combinations (the number of ways in which eight synapses can be selected from among 160). But, as noted above, if recognition relied solely on threshold summation, only a few dozen of these could be utilized. To utilize the large number of combinatorially distinct synaptic combinations, the intracellular communication must be able to distinguish a correspondingly large number of possibilities; its information-carrying molecules must have correspondingly large information capacity.

This means that, to do justice to the rich synaptic input, the intracellular signal transduction pathways need to include a class of messengers, which may be called 'high-entropy messengers,' with much greater information capacity than for instance calcium ions or cyclic AMP (which in this context would be called 'low-entropy messengers'). Similarly, the pathways must include matching receptors, which, again, may be called 'high-entropy receptors,' and these, too, need to single out the code word they are built for, from among millions of others, with precision and reliability.

In efforts to identify such messengers and receptors in the laboratory, it is of interest that these are not expected to be proteins; they are expected to be DNA/RNA chains. The brain acquires new mental entities at a steady rate, and since a new random code word must be found for each new entity, then messengers synthesized for the code words and receptors for the messengers (see below), the system must have a quick and reliable way to synthesize a sequence to match any sequence arriving to it. While proteins are capable of key-and-lock matching, the rules of protein folding are notoriously capricious, and the creation of high-entropy matches to given chain segments tends to be imprecise and time consuming.

It is noted that the immune system, which is tasked with high-entropy matching, does not have a way to create antibodies to given antigens on the spot and at a moment's notice. It can only create them by means of a lengthy selection process from among a large pool of premanufactured molecules (Edelman, 1972; Market and Papavasiliou, 2003).

However, unlike the proteins, the nucleic acids are well suited to serve as code words. It is known that RNA and DNA can be transcribed into each other with speed and precision, which means that they can support quick creation of ligands for given receptors and receptors for given ligands.

The implied prediction is that both transcription and reverse transcription (of short sequences) are common elements of everyday brain function. In the rest of this paper, I will go on the assumption that this is indeed so and that, in particular, the stationary memory traces (which must be highly stable) use DNA chains as their code words and the movable messengers between them (which do not need to be quite as stable) use RNA chains ('noncoding messenger RNAs' or 'ncmRNAs').

This still leaves the question as to the way the random nucleotide sequences for the code words are generated. Here, the encouraging recent discovery of microDNA molecules (Shibata et al., 2012; Dillon et al., 2015) may offer an answer.

The microDNAs are small (mostly 100–400 bp) extrachromosomal DNA molecules, present in healthy cells and in all the cell classes examined, including brain cells. They are circular in the form isolated in the experiments and are found to arise as the products of DNA breaks and mismatch repair after replication slippage. The microDNAs are irregular in their sequential pattern, and the inherently error-prone way in which they arise ensures great variability among them; it makes them well suited for a role as templates for Shannon-style random code words.

Three molecular systems

In walking through the requirements of a system of molecular data storage, it becomes clear that at least three

components are needed, each of them along the lines of a molecular complex. Two of these are built around ID molecules and serve as unique identifiers, or code words; I will refer to them as 'mental entity ID' (MEID) and 'neuron ID' (NID). The third one provides the interface between them and the macroscopic network; I will call it the 'event recognition complex' (ERC).

The next few sections are organized around a few illustrations showing how the relatively simple 'ID molecule' concept can be used to achieve precisely controllable mass action; then the sections under the heading of 'Putting it all together: the steps of storing a new item of memory' list the steps of setting up the supporting details, presupposed in the illustrations, with emphasis on the 'informational logistics' of the process, in other words on the details of arranging that the necessary information is always present wherever it is needed.

MEID molecules: reviving a dormant cell assembly by 'waking up' its synapses

The nucleotide sequence around which the MEID is built (the 'code word' part) identifies a mental entity, or equivalently its cell assembly. The code word can also be considered a local 'event ID,' to emphasize that local recognition of cell assemblies is achieved through recognition of their ignitions, which manifest themselves at neuronal localities as surprising *events* of multineuronal firing.

Since a mature brain contains very many mental entities (noting that whole items and details of items, of all levels and all modalities, are counted separately), and since ideally all of these should be addressable, the MEID (Figure 2) needs to contain a correspondingly long random code word (maybe 30–40 bases).

It is not immediately obvious that a small molecular system like a MEID is able to 'control' the storing of memories and convey enough detail to specify the precise selection of millions of neurons and synapses. The 'trick' is that the act of storing a piece of memory amounts to manufacturing multiple copies of its MEID receptors and placing them close to the access points through which the selected neurons and the synapses can be reached. The synapses are reached from dendrites and the neurons from astrocytes (see Figures 3 and 4, respectively). In that way, when a mental entity needs to be converted to a macroscopic excitation pattern, a swarm of its MEID messengers is sent out, identifying the mental entity in question, and these messengers do the selective long-distance work of finding their matching receptors (see 'Putting it all together: the steps of storing a new item of memory'), then the rest of
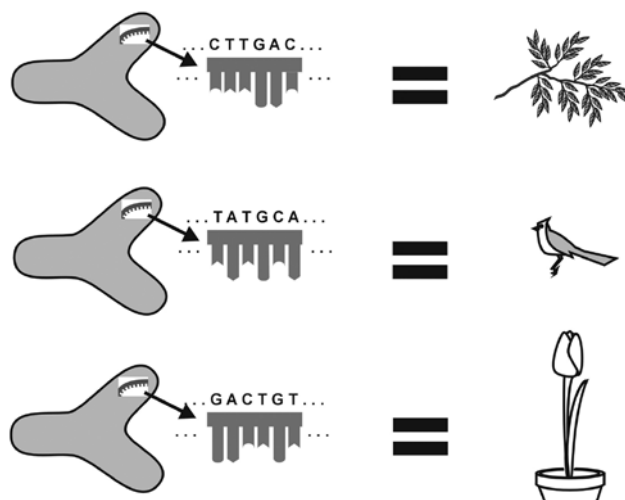
Figure 2: Principle of the 'code word' part of a MEID system. Different MEID molecules differ in the random sequence contained in them. The amorphous designs on the left symbolize the fixed MEID molecules, containing all their supporting hardware, with an inset indicating a part of their variable segment, the 'code word,' with the sequence shown as a piece of RNA arbitrarily chosen for illustration. The icons on the right stand for the associated mental entities (not necessarily tangible objects, as noted in Figure 1).

the signal transduction (the waking up of synapses or destabilizing of neurons) can proceed via messengers of lower entropy.

Figure 3 illustrates how a selected swarm of MEID messengers, seeking out a set of strategically placed MEID receptors matching them, can 'wake up' the desired set of synapses.

Because cell assemblies only differ from disconnected sets of cells by having the synapses between their members reinforced, the collective wake-up of synapses, which makes these synapses strong again, like freshly reinforced synapses, has the effect of waking up a dormant cell assembly.

In practice, the waking up and reestablishment of synapses are expected to be initiated somewhat in advance of their actual use, through prestored environmental clues (see the 'Extensions of ERC function: sequences, syntax, recursive processing' section).

The act of waking up a set of synapses greatly contributes to the *surprise* brought by a subsequent volley if it arrives to the same synapses (it enables 'prediction surprise'); the result is that, after a set of synapses have been woken up, a single volley coming to them is sufficient, when conditions are right, to fire their postsynaptic neuron. The effectiveness of prediction surprise in eliciting response is essential in quick processing of sequential information.
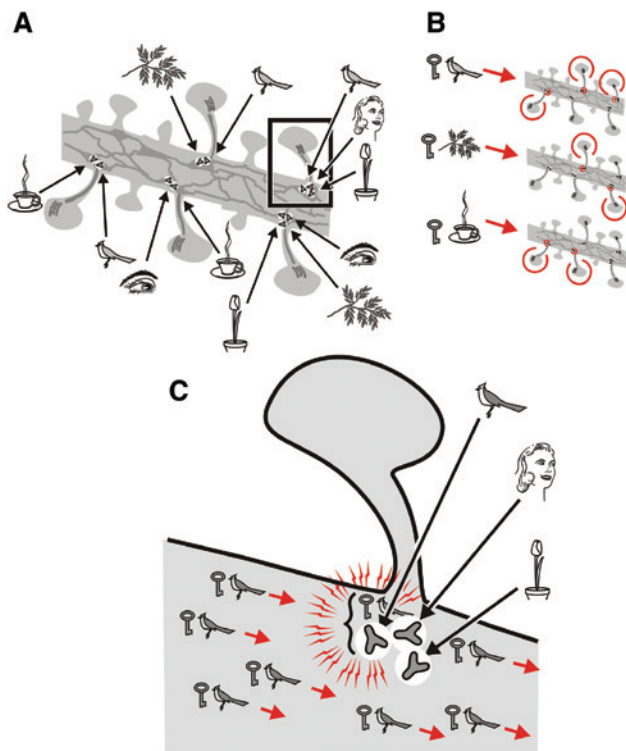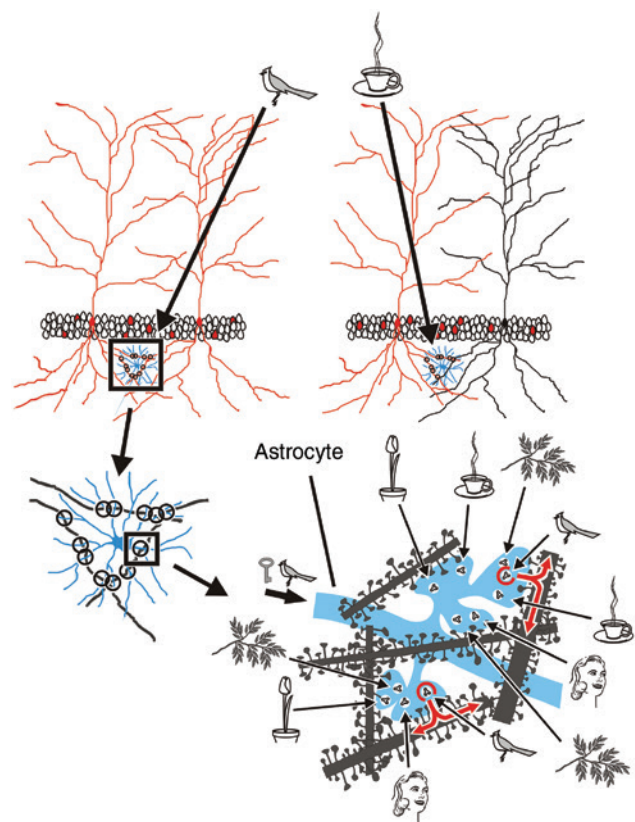
Figure 3: MEID systems in dendrites: targeted wake-up of synapse sets.

Waking up the synapses of a cell assembly by letting messengers of its MEID single them out, via matching receptors prestored near the synapses. A piece of dendrite (reproduced from Figure 1) is shown, including a number of MEID receptors, represented as stylized 'Y' shapes (see Figure 2). Messengers are marked with 'keys.' (A) A sample set of MEID receptors for a number of entities, placed near synapses and marked with descriptive icons. (B) The effect of three different swarms of MEID messengers, each waking up a different set of synapses (circled), as determined by the placement of their receptors (which is the same as in Figure 1). (C) The participants in the top drawing of (B) magnified. A number of 'bird' messengers moves past; one of them combines with its matching MEID receptor (curly bracket) and activates local signal transduction (sparks) to wake up the nearby synapse. The MEID molecules take up less space than shown in the stylized drawings since their code words only need to be 10–13 nm long (30–40 bases) and the supporting proteins can be shared between a number of receptors.

MEID molecules: igniting a selected cell assembly by directing slow inward currents into its cells

If neurons act in cooperative groups, it follows that initiating any action, mental or physical, requires a 'prime mover' in some form (Legéndy, 2009, chapter 7), which must have the ability to initiate coordinated firing by a preselected group of neurons, scattered among other neurons and not including the neurons outside the group. The 'trick' of the MEID messengers, mentioned in the last section, makes this seemingly impossible task feasible.

Figure 4: MEID systems in astrocytes: targeted ignition of a cell assembly.

Slow inward currents (SICs) can be sent into the cells of a cell assembly by using MEID receptors placed inside an astrocyte near where it touches the dendrites from the assembly. Top: two cell assemblies (see icons), highlighted by red coloring of the cell bodies and dendrites, together with an astrocyte (blue), which touches most of the pyramidal cells with cell bodies lying within the width of the figure. Bottom: MEID receptors are scattered inside an astrocytic process (right), as in Figure 3A. 'Bird' messengers (left, marked with 'key') are moving through the astrocytic process and cause SICs (red arrows) to enter the dendrites near their matching MEID receptors (circled in red). The neuron-glia apposition is stylized in the drawing; for the realistic 3D arrangement, see Witcher et al. (2007).

When electrical current is slowly pumped into the interior of a neuron, it will quickly spread to all branches and the cell body, gradually lowering the firing threshold and making the neuron increasingly sensitive to synaptic input, until the neuron 'spontaneously' fires, in the sense that the random background firing within the network is enough to fire it.

And when a cell assembly whose interconnecting synapses have recently been woken up (Figure 3) receives current injection individually targeted to each of its neurons, the firing from any spontaneously discharging neurons immediately spreads to the other neurons of the assembly (because they have also been made

hypersensitive by the current injection), until they have all fired. The cells briefly recover after each action potential, but the continued inward current sensitizes them again, and the process repeats. As a result, the group emits randomly repeated volleys, or a continued barrage, depending on the parameters, until the current stops.

As it happens, the slow inward currents (SICs) phenomenon, which has been well documented (Araque et al., 1998; Fellin et al., 2004; for reviews, see Fellin, 2009; Halassa and Haydon, 2010) and shown to be controlled by astrocytes, has the ability to pump electrical current into neurons apposed to astrocytes, for brief periods (roughly half a minute), lowering the thresholds of the affected neurons and making them hypersensitive.

In the experimental situations, the SICs were variously induced by astrocyte stimulation, neuron stimulation, or by pharmacological means, each time resulting in Ca^{2+} elevation within the astrocytes, which, in each case, caused them to release glutamate and activate a set of extrasynaptic NMDA receptors, setting in motion the inward currents. For the purposes of the present paper, it is significant that the SICs are targetable (Kozlov et al., 2006; Halassa et al., 2007) and often synchronized (Fellin et al., 2004), both of which are necessary if the SICs are to be instruments of targeted cell assembly ignition.

The spatial arrangement of astrocytes in relation to neurons is such that one astrocyte is in contact with some hundreds of neurons, or more, due to close apposition of its processes to dendrites, engulfing many of the synapses (Witcher et al., 2007). This makes astrocytes geometrically well suited to send currents into neuron groups large enough to extrapolate to cell assemblies. The branching of astrocytes has been shown to divide up space into essentially nonoverlapping 'domains' (Bushong et al., 2002; Oberheim et al., 2008), and whichever astrocyte arborizes over a domain touches essentially every dendritic process passing through the domain.

The concept of targeting long-distance high-entropy signaling to multiple preselected locations, with the help of strategically placed MEID receptors, has been described in the last section in the context of reviving a set of synapses from inside a dendrite (Figure 3). The same principle can be adapted to the targeted injection of current into a set of neurons from inside an astrocyte, as is shown in Figure 4.

The required exocytosis is expected to occur at the *punctum adherens* junctions, usually near synapses (Spacek and Harris, 1998), and the SIC-inducing signal transduction is initiated, as in the dendritic variant of the concept, when the messengers identifying the chosen cell assembly combine with receptors. The receptors

are deposited at their places, near the cell assembly members, at an earlier time when the cell assembly is active.

Recovering lost synapses: the NID molecules

The method described in Figure 3 for recovering dormant synapses presupposes that the synapses involved, while dormant, are anatomically still present. In reality, as time passes, synapses are expected to come and go, and if the findings of Stettler et al. (2006) can be generalized, the average synapse only lasts a few months. The MEID receptor planted near a synapse, for use in Figure 3, may, a few years later, be sitting in the middle of nowhere inside a dendritic shaft (as in Figure 5B below).

However, the same general principle used in recovering dormant synapses, by utilizing distributed MEID receptors, is still applicable, in a modified form, for recovering those synapses that are gone, at least where the presynaptic axon has a branch near the former synapse.

The modification requires a second kind of ID molecule, one whose code word serves as an ID tag for individual neurons. In describing multisynaptic input volleys, the neuron ID (NID) can also be considered as being a more robust version of a 'synapse ID,' specifying what is more permanent about the synapse, which is its presynaptic neuron.

The NID code word is expected to be (randomly) generated in the neuronal cell body and remain constant for the lifetime of the neuron. The NID is more selective than the (mammalian) cell adhesion molecule, as it must make a distinction between neurons within the same biological neuron pool. At the same time, it is less selective than the MEID code word because in any one situation, an NID only needs to distinguish between the few hundred or few thousand neurons, the ones accessible within a given neighborhood. Accordingly the NID messengers are expected to be smaller (and more mobile) than the MEID messengers.

The reason the NIDs are indispensable to the recovery of lost synapses is that the recovery requires growth, targeted to processes from the formerly presynaptic neuron. When an axonal process from that neuron is found, the growing filopodium must have, available to it, an identifier of its target neuron. In practice, the NID must be stored together with the MEID receptor that initiates the growth, making the information available when the advancing tip encounters an axonal segment.

This means that, to prepare the information needed for reconnection through growth, each MEID receptor,
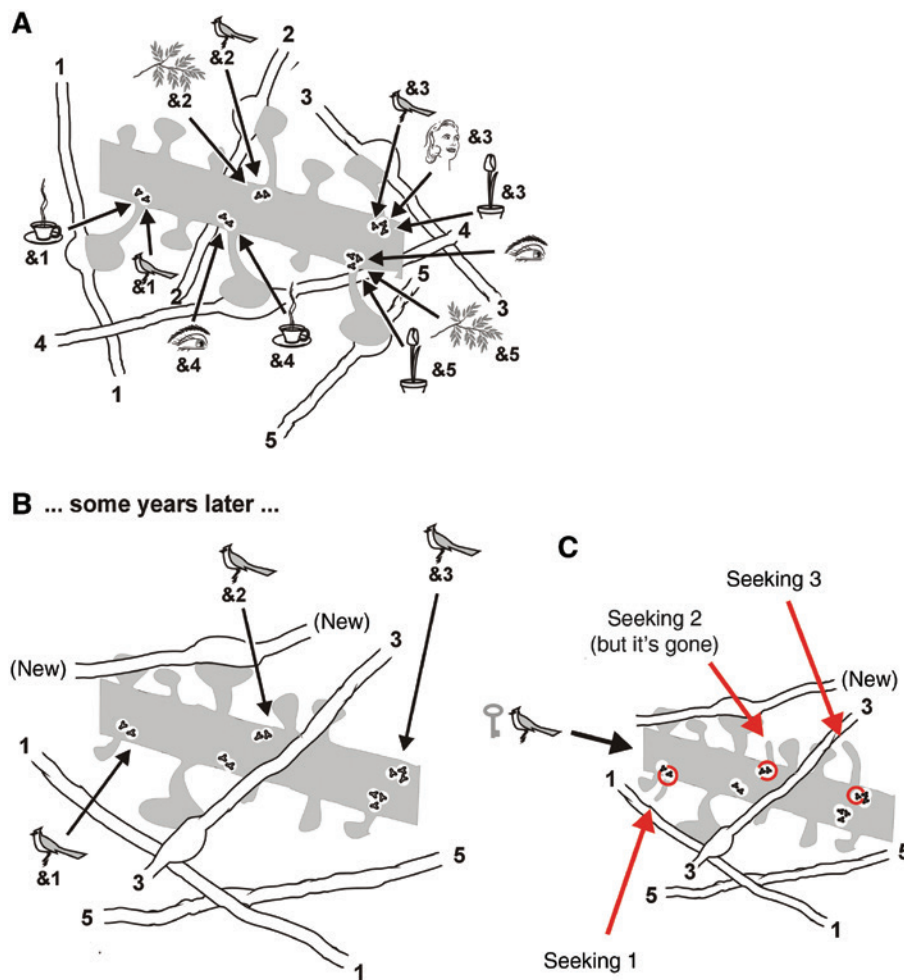
Figure 5: Reviving lost synapses, guided by paired MEID and NID molecules.
Recovery of synapses that have disappeared is possible only if the identities of the presynaptic neurons are preserved earlier, when the synapses are still present, and their NID code words are attached to the MEID receptors responsible for wake-up. Pairing of ID molecules is shown by linking the MEID icons and NID numbers with '&'. (A) The dendrite segment of Figure 3A, redrawn to include the presynaptic axons with their 'names' (1, 2, …, 5). (B) The same segment, as it may appear after some years. The whole synaptic pattern is shown as changed, along with the pattern of nearby axon branches; the only remaining clue to the former synaptic arrangement is the presence of MEID/NID pairs in the surviving dendrite. (Only the 'bird' icons are shown, to reduce clutter). (C) In the recall step, a swarm of 'bird' messengers enters the locality, as in Figure 3, but since the original synapses are gone, the arriving messengers must instead initiate filopodium growth, seeking out the axons of former presynaptic cells. Some axons are no longer within reach, as shown in the example of axon 2.

when originally assigned to a synapse, must be securely attached to an NID code word designating the presynaptic neuron (see below). With the synaptic information thus preserved, the problem of reattachment reduces to biological implementation of a way to let the NID code word guide fiber growth and find a suitable axon branch (Figure 5C).

Super-redundant random ensembles and the hippocampus

The last sections described how a suitable protocol can 'translate' a system of ID molecules into cell assemblies; it

still remains to be demonstrated that cell assemblies can also be translated into ID molecules. The challenge here is that cell assemblies are widespread and their identity is inscribed in their collective membership. The solution proposed here (see the next section) utilizes the idea mentioned above, that the multisynaptic volleys imparted to pieces of dendrites (or to astrocytes) from cell assemblies, even though confined to relatively small localities in the neuropil, can carry enough information, if the parameters are right, to permit recognition of the cell assembly to a fair degree of certainty.

One way to look at the concept is by considering a special case of Shannon's random code ensemble, which

may be called the 'super-redundant random ensemble,' where the code words are so large that small portions of them can carry enough information to constitute a full and reliable ensemble. For instance, as was seen above, an ensemble of 256 code words, each of them made up of 40 random binary bits, can achieve fairly good isolation of the code words. But let us say that instead of 40 bits, we use 1000 bits for each of the 256 code words. Then by taking the first 40 bits (or any fixed subset of 40 bits) out of the 1000, the 40-bit subsets will also make a highly reliable ensemble of 256 words. (For a formal discussion, see Mézard and Montanari, 2009.)

The recognition of cell assemblies based on volleys arriving to dendrite segments makes use of the same concept. If the cell assemblies in an area are all such that they extend over the whole area and take up for instance 5% of the neurons there (Losonczy and Magee, 2006), and if the collection of these large cell assemblies is regarded as an ensemble of code words for mental entities, then it can be said that volleys arriving to a piece of dendrite, like the one in Figure 1, are large enough to serve as signatures of the underlying mental entities.

But there is one catch: sampling can only identify the cell assemblies, no matter where the samples are taken, if the assemblies are truly random and homogeneous. This implies, among other things, that the synapses from each cell assembly must be spread out fairly uniformly throughout the region. In this regard, it is to be noted that the connectivity of the hippocampal CA3 pyramidal cells appears to be uniformly random (Buzsáki, 2006; Wittner et al., 2007). Further, it can be argued that the hippocampal dentate gyrus, by virtue of its peculiar connectivity, interacts with CA3 to perform a randomizing function, analogous to stirring a liquid when mixing it (Legéndy, in preparation).

The possibility of correctly recognizing cell assemblies in many places, by localized means, makes it meaningful to assign different code words to the same cell assembly at different localities and establish their equivalence in a separate and possibly later step (see Figure 7 and the section 'The astrocyte builds ERCs for the coincident signals' below). In this way, even without direct communication between the localities, their contents can be reliably coordinated.

From the standpoint of the neocortex, it is significant that the homogenizing effect of the hippocampus is passed down to the cortex as the cell assemblies themselves are passed down (noting that the natural tendency of cascaded processes is to move toward greater randomness). The result is that different cortical columns can independently recognize the same global ignition. The hippocampus, in turn, ultimately receives its input from the cortex and attaches the same CA3 cell assemblies to the differently labeled cortical assemblies, erasing their differences.

Two-way linkage between molecules and mental entities: the event recognition complex (ERC)

The recognition of cell assemblies based on their volleys, as received on a dendrite, requires a piece of molecular hardware inside the dendrite, the ERC, able to gather inputs from the synapses involved.

In order to utilize the combinatorial diversity of synapse choices, the ERC must do more than just count the active synapses (which the cell membrane is already doing); it must be equipped to utilize the synapse identities – in the form of the identities of the presynaptic neurons. This means that the synaptic information from the active synapses can be brought by the same NID messengers that are stored at the synapses for the case they are lost and need to be recovered.

The ERC, in turn, must be equipped with built-in receptors for each of the NIDs it is designed to expect and, when enough of the receptors have combined with their matching ligands, to activate its transcription machinery and synthesize multiple copies of its output molecule.

Every ERC is designed to recognize one cell assembly, and the output molecules it emits in response are MEID messengers carrying the code word identifying that assembly.

The version of ERC discussed up until this point, and illustrated in Figure 6, is the rudimentary version where all input receptors are for NIDs emitted at synapses. Further elaborations on the ERC concept, where ERCs can also receive MEIDs emitted from other ERCs and where input channels can interact, are also envisioned (Figure 7 and the next section); but even in these cases, the input ultimately comes, directly or indirectly, from NIDs.

When the NIDs, destined for ERCs, leave their respective synapses, they arise in response to presynaptic input, but not every presynaptic spike is expected to generate an NID. Since much of the firing arriving to synapses is noise, it is desirable to switch on NID synthesis selectively, when and where the NIDs sent into the cellular interior are expected to be pieces of meaningful volleys. Since meaningful input is best recognized by the fact that it is *surprising*, this means that NID production is expected to be switched on either at synapses where LTP is being
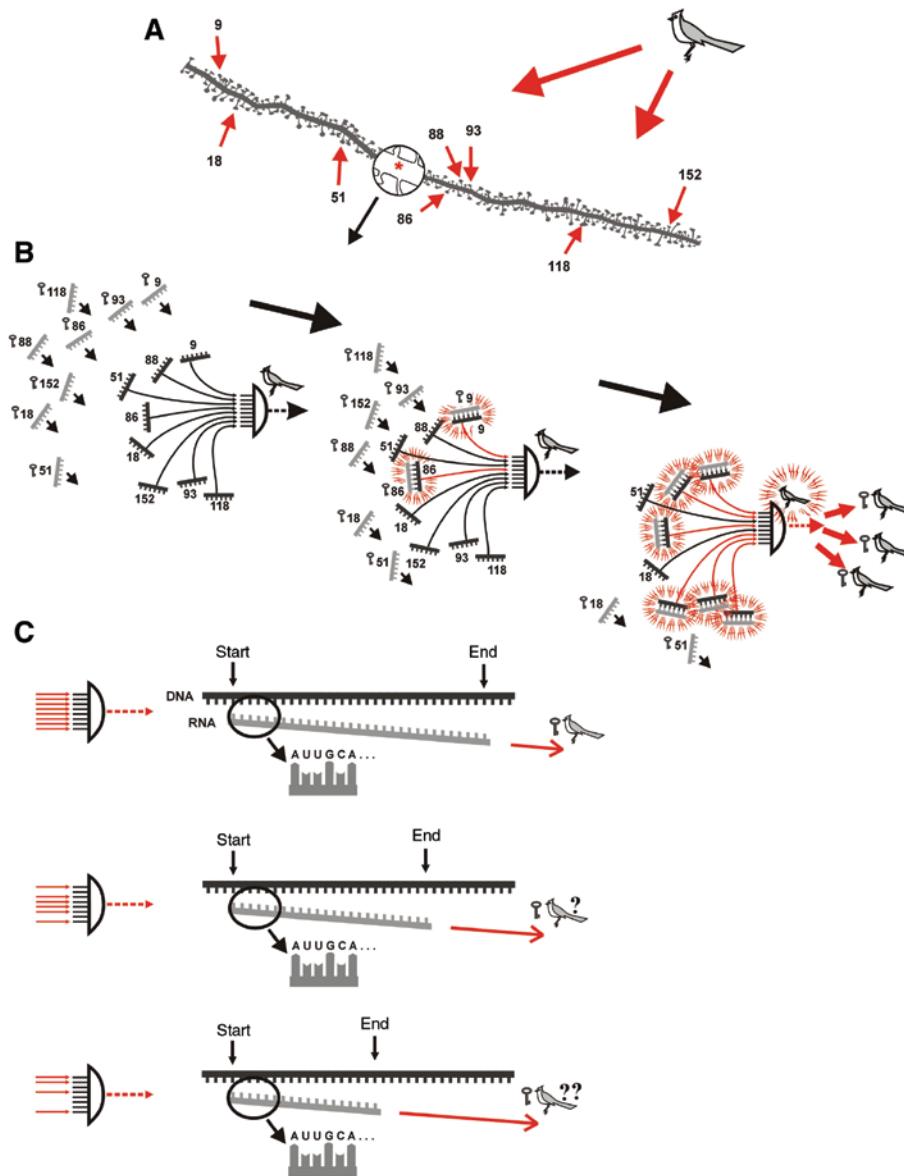
Figure 6: An ERC recognizing a cell assembly by the synapses it excites.
A multisynaptic pattern brought by synapse-specific messengers causes an ERC to turn on the synthesis of MEID messengers. (A) A dendrite receives the 'bird' input pattern (copied from Figure 1), with the 'names' of the presynaptic neurons shown as numbers. An ERC is shown inside a magnified portion (red asterisk). (B) NID messengers from synapses (shown with their numbers) approach the receptors and, when finding enough of them, start the synthesis of MEID messengers. (C) The uncertainty of recognition (due to incomplete input) is incorporated in the output through controlled truncation of the MEID messengers emitted. Note: The ERC shown in this illustration only carries receptors for NIDs at its input end; in general, ERC inputs can also carry MEID receptors (Figure 7).

induced or at ones that had undergone LTP at some time in the past and have recently been 'woken up,' as in Figure 3 (enabling 'prediction surprise').

The MEIDs emitted by ERCs, when they recognize their preprogrammed inputs, are the same ones that show up in Figure 3, where they wake up a cell assembly, or in Figure 4, where they ignite a cell assembly. It is noted that the cell assembly being woken up or being ignited in these situations is not the one causing the ERC to emit its MEIDs; rather, they are elements required in subsequent action.

One crucial issue concerning ERCs, touched upon in Figure 6C, is that the input patterns arriving to the receptors are usually incomplete.

The operation of a pattern-recognizing device is, by its nature, a probabilistic affair; because, in practice,
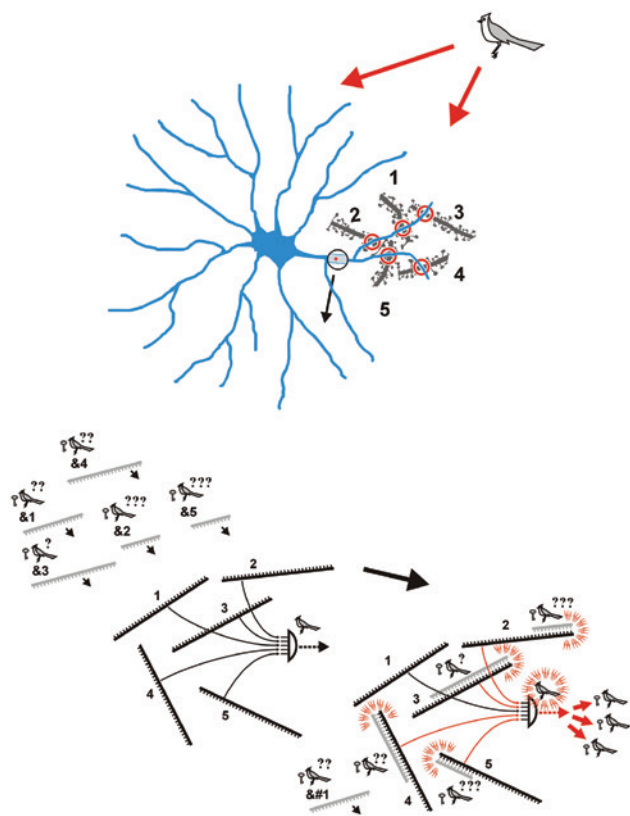
Figure 7: An aERC combining uncertain inputs and improving the certainty.
Combining the evidence from a number of MEIDs, all reporting the same cell assembly with varying degrees of certainty, to improve the resulting reliability. Top: an astrocyte is shown with a few dendrite segments touched by its processes, each dendrite from a different neuron on which the same cell assembly induces LTP. The coincident LTPs, signaled to the astrocyte and passed to it via NIDs identifying the neurons (1, 2, …, 5), give rise to a tentative aERC, with receptors for the NIDs (not shown). Bottom (prelude to events shown): MEIDs produced by ERCs in the neurons are sent over to the astrocyte, there attached to NIDs designating their neuron of origin (see '&'), and establish MEID receptors with the corresponding NID receptors of the aERC. Based on the coincidence creating the aERC, the aERC operates on the premise that all the MEIDs for which it has created receptors designate the same cell assembly as the aERCs itself ('bird'). In later uses of the aERC (shown), the degree to which the MEIDs are truncated, as in Figure 6C (see also question marks on the 'birds'), permits the aERC to estimate its confidence in their combined recognition.

pattern recognition (for instance speech recognition; see for instance White et al., 2007) is never 100% certain. In the case of the schematic system illustrated in Figure 6, the uncertainty arises from the fact that only rarely will all the NIDs expected from a cell assembly arrive to the ERC when they should; therefore, most recognitions will be at best tentative. At the same time, it can be said that any method of increasing the number of synapses contributing

to an ERC will act to increase the entropy underlying the recognition and improve the certainty.

One way in which an ERC can effectively increase the pool of contributing synapses is to monitor some of the farther-away synapses indirectly, through other ERCs responding to the same cell assembly. This means that, unlike the ERC shown in Figure 6, which only uses receptors for NIDs, coming directly from the synapses, an ERC can also carry receptors for items of 'indirect input' in the form of MEIDs, arising as outputs from other ERCs and representing whole groups of additional synapses that respond to the same cell assembly.

It is noted that the MEIDs from other ERCs are only useful in such a role if the ERC utilizing them has access to the level of certainty carried by them, hence the uncertainty indication, through truncation, shown in Figure 6C.

Then the next rank of ERCs (Figure 7) needs to contain apparatus for using the truncation and 'computing' its own certainty of recognition. The method used (not shown in Figure 7) may simply amount to making use of the logarithmic definition of entropy and effectively adding the truncated code word lengths (after appropriate correction) to construct the output code word length.

Extensions of ERC function: sequences, syntax, recursive processing

In the last section, indirect input was only used for giving the ERCs, in effect, access to more synapses, and accordingly, in that description, an ERC and the ERCs generating its inputs all dealt with the same macroscopic mental entity. But the concept of indirect input is more general; the MEIDs at the input terminals need not designate the same mental entity as the MEID produced at the output end. A few examples are mentioned below, without elaboration, as they go beyond the scope of this paper.

One example of ERCs in which the input and output entities are different is the class of ERCs created for sequential events, ones whose components are linked by reproducible relative timing, as for instance in sequential movements, or bird songs, or musical pieces. It is noted that sequence processing adds complexity to the ERC, as it raises the need for sensitivity to relational information.

The task of 'waking up' a selected set of synapses prior to their expected use (Figure 3) and the task of igniting a cell assembly when it is needed (Figure 4) rely on the

prior recognition of contextual cues; accordingly, these tasks offer further uses for indirect input.

A more general extension, adding to the conceptual complexity, is equipping the ERCs to process syntactic relations between events of ligand binding. The cell assembly hypothesis is able to support syntactic relations between ignitions, as has been demonstrated in the context of visual shape perception (Legéndy, 2009, chapters 8 and 9). In general, it can be argued that the brain supports syntactic structures at all levels, not only in the speech areas and not only in humans (Legéndy, 1970, 2009; Buzsáki, 2010), which suggests that molecular memories, as implemented through ERCs and ID molecules, should also be syntactically structured.

By extrapolation from the earlier work, it can be expected that the inclusion of syntax between ERC input terminals calls for 'cross-wires' between the receptors (referring to Figure 6B), and calls also for multicomponent receptors. This is partly because the structure must allow for the presence of definable links connecting some receptors, with an identifiable 'bridgehead' at each end of each link (Legéndy, 2009, chapter 15) and partly because additional ID code words are needed to distinguish objects from actions and from other broad groups of elements, so as to assign them syntactic roles.

The prominent role of episodic memory (Scoville and Milner, 1957; Vargha-Khadem et al., 1997; Tulving and Markowitsch, 1998) suggests that 'episodes' are among the classes of mental entities (see Figure 1, note) that can be identified in MEIDs; and that, accordingly, some MEIDs are expected to describe 'episode IDs.' Since episodes are, in essence, brief stories, extending over the full range and complexity of mental imagery, it is probably safe to say that any ERCs encoding them would need to support syntactic relations between their receptors.

It may be said, schematically speaking, that ERCs permitting indirect inputs 'define' mental entities in terms of other mental entities, which define those in terms of further mental entities, in a recursive network of definitions within definitions, without any hierarchical structure and without any inherent limitation on the number of stages it can accommodate – a potentially powerful construct.

Putting it all together: the steps of storing a new item of memory

Memory storage begins with the neurons, not the molecules, because only the neurons have the ability to reach far-apart locations in the brain quickly. In the present simple model, where mental entities do not have syntactic structure or time structure, the mental entities to be stored are simply made up of single large cell assemblies.

A new cell assembly announces itself by inducing LTP

A new mental entity enters the picture as a new cell assembly. (The present paper bypasses the issues of converting items of the world to cell assemblies and takes the position that the rest of the brain can achieve that somehow.) As discussed above, cell assemblies, through their ability to ignite, cause coincidence and repetition of synaptic events and lead to LTP induction.

The LTP, at least its superficial manifestation of enhanced postsynaptic response to presynaptic stimulation, is only temporary; the synaptic changes seldom last more than a few weeks (Abraham, 2003). The next few subsections make a case that the superficially detectable synaptic change fits into a much wider process whereby a permanent memory trace is installed.

The LTP enables NID synthesis and places tags at its synapses

A secret to the large capacity of molecular storage is the combinatorial diversity of synapse choices, and the capability to utilize this diversity hinges on assigning identity markers to the individual synapses. Without identity markers, a device gathering the synaptic inputs could at best count up the number of contributing synapses. To enable the marking of individual synapses, it is necessary to set up the machinery for synthesizing the NID messengers at the synapses when needed.

In the desired eventual system, where events of multisynaptic coincidence are made detectable at ERCs, the synaptic impulses must trigger NID synthesis. But the machinery for NID synthesis cannot remain enabled at all times; if it were, the continual rainfall of firing, arriving to thousands of synapses, would flood the interior of the dendrites and frustrate any internal correlation detection.

This means that the NID-synthesizing machinery must be enabled only at selected times and at selected synapses, in particular at the times and at the synapses where meaningful (surprising) input is expected.

This means that during LTP induction, transcription is expected to be switched on as soon as the surprise is detected. Then, if the repetitions continue a little longer, the rest of the volleys, which are tacked on to the same surprising event, can be converted to NIDs.

Another situation where transcription can be switched on arises at a later time, when ERCs are already in place; it is when a set of synapses is 'woken up' by means of MEIDs (as in Figures 3 and 5). If the same synapses subsequently receive a volley, the compound event, made up of the targeted wake-up followed by a volley arriving to the synapses woken up, will be highly surprising (with 'prediction surprise').

Returning to the events preceding the formation of a new ERC during LTP induction, an additional step needed is to put place holders at the synapses being changed – the 'synaptic tags' of Frey and Morris (1997). These are to act as guide posts to later MEID messengers by which they can find the synapses later, in case the subsequently released NIDs happen to give rise to a new ERC and the ERC responds (see below).

An ERC building site detects converging NIDs and records them

After LTP induction has gotten underway, and volleys arriving to the dendrite have begun to set off swarms of NIDs inside the dendrite, it becomes possible to record the volleys by creating an ERC for them.

It is not predictable where an ERC is going to form, just as there is no way to predict where the coincident events of firing would appear most prominently. In order to maintain the freedom to select favorable locations, more potential 'building sites' for ERCs must exist than would end up being built upon, and each site must be equipped with the enzymatic apparatus needed to make a quick imprint from arriving NIDs, whenever it is detected that unusually many of them arrive at a site. It is noted that the imprinting made from the NID messengers (which are RNAs) amounts to small-scale reverse transcription, and accordingly, the resulting receptor code word arises as complementary DNA (cDNA).

It can be expected that the imprinting apparatus has a way to require repetition; this means that it is equipped to make quick temporary imprints and allow them to disappear unless confirmed by subsequent arrivals of the same NIDs. If the number of confirmed receptors is large enough to form an acceptably unique combination, it can be said that the first part of the ERC construction, the building of a viable input end, is complete.

The ERC building site recruits a random template for its output

The next step is to assign a unique output code word to the newly assembled input combination. Ideally, the

input combination is complex enough, therefore unique enough, that a successful matching of the receptors indicates, with probability close to 1.0, the presence of the same macroscopic circumstances that originally gave rise to the ERC. The random code word of the output MEIDs must also be sufficiently unique, meaning that it must contain a sufficiently long random sequence that its entropy can do justice to the input.

Designing a generator of random sequences is not trivial, and in a biological context, there is, up until now, only one known example of randomly generated macromolecular sequences that are expressly meant to achieve variety: in the case of the immunoglobulins (Edelman, 1972; Market and Papavasiliou, 2003) – which go a long way toward achieving sufficient variety to present an antibody for every possible invading protein.

Significantly, the recent discovery of microDNAs (Shibata et al., 2012; Dillon et al., 2015) adds what may turn out to be one more example. The microDNAs contain a substantial irregular segment with the required random property and large entropy capacity. As mentioned above, it is assumed in this paper that these molecules are in fact enlisted in the design of the ERCs, as templates for the random sequences to be used in output MEIDs, meaning that when an ERC is ready to be launched, a usable portion from a nearby microDNA can be 'harvested' for use in its output code word.

This stage, the stage when the output apparatus is being built, is the time to assemble the machinery for truncating the output MEID, by an amount reflecting the uncertainty of the recognition (in accordance with the number of receptors missed at the input end; see Figure 6C). It is also the time to assemble the machinery for computing the improved certainty of the output when it refers to the same entity as do a number of input MEIDs (see Figure 7).

ERCs are not expected to be hard-wired organelles; for instance in Figure 6B, the 'wires' drawn between the receptors and the 'gate' should probably be envisioned as being either electrochemical fields or systems of low-entropy messengers. The reason is that the receptors themselves are expected to be in constant flux. Synapses on neurons continually appear and disappear, and the ERC, in order to continue tracking the same mental entity that originally formed it, must be able to acquire new receptors and drop old ones. This means that it cannot be bound by a rigid physical architecture.

MEID receptors are created for the ERC output at the tagged synapses

At this point, it will be appreciated that the ERC merely emits a unique but uninterpretable response to a unique

but unknowable input – which in this form is useless. However, it becomes useful in the moment the output molecules can be made to 'wake up' (Figure 3) a set of synapses that describe a meaningful synaptic pattern.

This requires setting up the situation that forms the basis of Figure 3, where the first step is for the nascent ERC to emit a swarm of its output MEIDs in the hopes that they will land at the right synapses. What enables them to land there is that the synapses are marked, as part of the 'synaptic tagging' step (Frey and Morris, 1997), which occurred right after the synaptic stimulation, when the immediate effects of the stimulation have not yet disappeared.

The 'tagging' amounts to setting up the enzymatic apparatus able to 'capture' the freely floating MEID messengers coming from ERCs (we may note the reference to 'capture' in Redondo and Morris, 2011), and having captured them, to reverse-transcribe them to cDNAs, which are to serve as the code words for the receptors being set up at the synapses (in the 'Y-shaped' elements in Figure 3). The receptors are then attached to the signal transduction pathway, eventually waking up the synapses. Thereafter, whenever the same MEID messengers combine with the receptors, they will wake up the same synapses.

MEID receptors at the synapses are paired up with NIDs

The next step sets up the situation forming the basis of Figure 5, where the MEID receptors are attached to NID code words (recently active at the same synapses, when sending out the same NID code words to contribute to the new ERC); this will allow them to provide the information necessary to recreate the synapse if it is lost at a later time and its recreation requires growth.

As noted in the text accompanying Figure 5, after the filopodial growth has acquired its intended target, it must be able to recognize the target before commencing synapse formation and the specificity of cell adhesion molecules is (in higher animals) insufficient for identifying the right axon. The identifying information must be contained in a high-entropy code word (an NID), which in turn must be accessible from the MEID receptor that initiates the reconnection.

Many neurons record a cell assembly and initiate signals to an astrocyte

Since the surprising events inducing LTP tend to arise from macroscopic causes resulting in cell assembly ignition, it is likely that when one neuron undergoes LTP from ignitions

of a cell assembly, so do many others. For instance in Figure 1, in each sample event, the neurons sending coordinated spikes to the dendrite are only a small fraction of the neurons actually participating in the same macroscopic event. It follows that if the signal transduction pathways of LTP include signaling from neurons to astrocytes (Fellin, 2009), an astrocyte receiving such signals from one neuron will also receive them from many others.

It is now known that neurons that undergo LTP do in fact send signals to astrocytes apposed to them. When neurons in CA1 are subjected to intense synaptic bombardment through stimulation of the Schaffer collaterals, stimulation known to induce LTP, events of SICs are found to occur in some of the neurons, and the SICs can be shown to be governed by the astrocytes (Fellin et al., 2004).

The astrocyte builds ERCs for the coincident signals

From the standpoint of the astrocyte, the coincident signals from a collection of neurons, all responding to the same cell assembly, make a detectable surprising event, and the geometry of astrocyte-neuron relations (Witcher et al., 2007) positions the astrocyte well for detecting it.

The signals from the neurons convey the presence of LTP in the signaling neurons, and each such indication conveys evidence of the fact that the sending neuron itself has detected multineuronal activity. A separate indicator of multineuronal cooperation available to astrocytes is the presence of waves of K^+ ions, released by widespread neuronal firing and known to be processed within the astrocytes (Kuffler et al., 1966; Kofuji and Newman, 2004).

If the astrocyte is equipped to build an ERC in response to the coincident neuronal signaling, the resulting astrocytic ERC (aERC) is in a favorable position to develop highly reliable indication of the presence of the underlying object, more reliable than the ERCs in the participating neurons, because it can make use of multiple independent measurements of the same thing (Figure 7).

Taking advantage of this possibility involves a two-step process because the surprise detection step cannot make use of the MEID code words from the ERCs (as different ERCs assign different MEID code words to the same event); it can only rely on detecting the coincidence between the single-neuron signals, communicated to their location by NIDs passed to the astrocyte from the neurons (along the same lines as described in the above sections 'The LTP enables NID synthesis and places tags at its synapses' and 'An ERC building site detects converging NIDs and records them').

Once the tentative aERC is equipped with a set of neuronal NID receptors, it is ready for the second round of the determination, which confirms and strengthens the conclusions of the first, as it involves pairing up the NID receptors with the MEID receptors generated by ERCs in the corresponding neurons. This round is somewhat slower than the first because it involves movement of the MEID messengers from the neuronal ERCs to the astrocyte and then transporting them inside the astrocyte to the aERC. Since the MEID messengers are larger than the NID messengers, they move more slowly.

After the imported MEIDs have been added to the aERC, it becomes justifiable to combine their individual levels of certainty in order to obtain greater certainty, as shown in Figure 7 (a step that implicitly relies on noting that the different-looking MEIDs all designate the same thing).

The astrocyte places tags on the interfaces to each of the neurons

The role of aERCs just described, as tools of improving the certainty of recognition, is only one of their roles. The other one is in the trick, illustrated in Figure 4, which achieves addressable ignition of cell assemblies through SICs.

The steps needed in setting up the situation required in Figure 4 mirror the steps described in connection with Figure 3, where the addressable wake-up of synapses is described, at least in regard to the logistics of information flow.

The first of the steps is the astrocytic version of the synaptic tagging; in the present case, it amounts to placing tags, inside the astrocyte, at the locations of apposition to the targeted neurons. These are labeled by NID code words deposited earlier, at the time the correlated events of LTP signaling initiated the aERC creation. The astrocytic tags must be placed at the locations where the SIC-initiating message is later to be transmitted to the neuron, which are near where the same neurons earlier sent their LTP signals (and where consequently their NIDs are deposited).

As in the synaptic version of tagging, the tags are expected to include the molecular apparatus designed to capture the MEID messengers, which are subsequently expected to be sent out from the aERCs (see next section) and permit them to be reverse-transcribed into the cDNAs, which are to serve as receptor code words.

The aERCs send out the aMEIDs, which enable SIC initiation

The concept that astrocytes can cause selected cell assemblies to ignite, through targeted initiation of SICs

(Figure 4), is predicated on the assumption that astrocytes have a unified way to address all exit points toward a cell assembly, all with the same MEID code word. Since the different ERCs use different code words for the same cell assembly (and one neuron may even contain a number of ERCs), but the aERC translates them to a common MEID code word (Figure 7), the MEID messengers emitted from the aERC are the perfect choice for supplying the unified code word.

Accordingly, in the final phase of their installation, the aERCs must send out a batch of messengers from their freshly minted aMEIDs (the unified code word) into the cell, in the hopes that they will encounter the recently placed tags (at the appropriate places of neuronal apposition), and be reverse-transcribed there, into receptors. Subsequently, the receptors can initiate assembly of the signal transduction pathway sending SICs into the neurons across the gap from them, when receiving the same aMEIDs in the future.

Discussion

The size of cell assemblies

The concept that cell assemblies can be recognized simply by the firing they send to a few hundred synapses on a piece of dendrite and that the dendrite can be anywhere in the area relies on the assumption that the cell assemblies themselves are extremely large.

The cell assembly hypothesis relies on various lines of evidence and reasoning (Gerstein et al., 1989; Abeles, 1991; Harris et al., 2003; Ikegaya et al., 2004; Buzsáki, 2010). Specifically, the statement of interest here, the hypothesis that there exist some classes of very large cell assemblies, relies on inferences that can be made from dendritic spike data and cross-correlation data.

Dendritic spikes are now a well-documented phenomenon; they have been seen both in slices and *in vivo* (Svoboda et al., 1997; Kamondi et al., 1998; Gasparini et al., 2004; Losonczy and Magee, 2006; Spruston, 2008; Major et al., 2013; Smith et al., 2013; Palmer et al., 2014). This is of interest in the present context because each time a dendritic spike is generated, the conditions for its generation must, by necessity, be satisfied, and these include coordinated assaults of spikes on a segment of dendrite, assaults intense enough that they are unlikely to arise as part of random spike fluctuations.

The quantitative conditions of dendritic spike generation have been established by recently available glutamate

uncaging data, which tell us that to initiate a dendritic spike requires time-concentrated volleys arriving to some 5% of the synapses on the spike-initiating dendrite segments; for instance to about 20 synapses within about 6 ms (Losonczy and Magee, 2006; Polsky et al., 2009; Major et al., 2013). It is clear that only very large neuron groups can be expected to contact small segments of a dendrite at a sufficiently large number of synapses.

If one estimates the rate at which random fluctuations would give rise to volleys meeting the above requirements, one finds the rate to be once in eons. However, in fact, spontaneous dendritic spikes (ones not arising in response to experimental stimulation) occur in awake animals frequently enough to be studied (Kamondi et al., 1998), which implies that coordinated firing of very large groups of neurons must occur at a corresponding rate. This agrees with the conclusion stated by Kamondi et al. (1998).

Cross-correlation data also support the large size of ignitable neuron groups. The probability of finding neuron pairs with strong positive cross-correlation is often quite high, especially when the pairs are selected to have similar response properties (Ts'o et al., 1986; Hampson et al., 1996) – and this is true even in neuron pools, where only 1%–2% of all neuron pairs are synaptically connected (Deuchars and Thomson, 1996; Hampson et al., 1996). Since in the latter cases direct excitation can add very few counts, the substantial numbers of extra counts observed in the central bins of cross-correlograms are more likely to come from ignitions of cell assemblies in which both of the recorded neurons are members. If the cell assemblies were small, the probability of positive cross-correlation would be small.

Dendrite segments as decision-making elements

It will be noted that Figure 1 deals with segments of dendrites rather than whole neurons, and so do the other drawings. Part of the motivation is that a segment of dendrite is considerably smaller and simpler than the whole neuron, with less propagation delay and fewer spike trains to deal with. A functionally more important point is that a dendrite segment can initiate dendritic spikes, which often proceed to the cell body and generate action potentials; further, it can induce LTP (Remy and Spruston, 2007) and therefore initiate synaptic change.

In other words, in the context of interest here, a segment of dendrite can act as a more or less autonomous decision-making and memory-storing element (Branco and Häusser, 2010; Major et al., 2013); it offers an

attractive platform on which to discuss the microstructure of learning and memory.

In LTP-driven change, which singles out sets of synapses to the exclusion of all other synapses mixed in with them, it is tempting to assume that cellular membrane properties will force adjacent synapses to act together and will prevent them from being interspersed with unaffected synapses; but in fact this is not so.

By use of two-photon glutamate uncaging, combined with two-photon Ca^{2+} imaging, Losonczy and Magee (2006) verified that the synaptic inputs initiating dendritic Na^+ spikes do not need to be adjacent but can be widely scattered, and the same thing has been shown for dendritic NMDA spikes (Major et al., 2013). Furthermore, the linkage between dendritic spikes and LTP (Remy and Spruston, 2007) tells us that the synapses affected in LTP also do not need to be adjacent.

The problem of glial motility: the 'tripartite synapse' arrangement

In the discussion of implementing current injection into cell assemblies, via SICs (see Figure 4 and the section 'The aERCs send out the aMEIDs, which enable SIC initiation'), one issue has been bypassed, having to do with the motility of neuronal and glial processes. In placing the current-injecting apparatus inside the controlling glia cell, it is essential to place it close to an interface between the glia cell and the targeted neuron. But the small processes of dendrites and astrocytes move around so much that after some time it is quite easy for the apparatus to land near the wrong neuron.

As was noted earlier, neurons continually grow new synaptic spines and retract others, a problem that must be solved by any protocol for targeted wake-up of synapses (see Figure 5). A similar issue arises, but more severely, in the case of glia cells because of the continual and vigorous change in the shape of the small processes engulfing the dendrites and synaptic spines (Witcher et al., 2007).

Here, too, the only solution I can think of is that the MEID deposited at each intercellular interface (which are aMEIDs in the glial case) must be paired up with a molecule identifying the neuron to be targeted by the SICs (via NID). In that way, when changes of cellular shape move the aMEIDs away from their intended neurons, the NIDs will stay with them and preserve the information needed for finding the neuron.

The motility of the small processes tends to be much more lively in the glial cells than in the neurons. In the case of neurons, the time of noticeable change is measured

in hours or days (Woolley et al., 1990; Kirov et al., 1999; Stettler et al., 2006); in the case of astrocytes, it is measured in minutes (Haber et al., 2006). There is evidence, from time-lapse confocal imaging (Haber et al., 2006), that astrocytic processes a small distance away from a dendritic spine, but not directly apposed to it, can grow several microns in a few minutes, until they are in direct apposition to it, and the pictures indicate that corresponding neuronal elements are much less motile. The slower rate of synaptic changes has been directly verified, also through time-lapse confocal imaging, by Stettler et al. (2006).

All this means that in the case of the astrocytes, it is especially important to make provisions for the continual growth and disappearance of small cellular processes because the continual shifting keeps changing the intracellular locations where pieces of apparatus must be assembled, if they are to be within communication range of their target neurons.

It may be that the well-known 'tripartite synapse' arrangement (Araque et al., 1999) is in fact nature's way of addressing this problem, at least in part. Since synaptic spines are more stable than the small glial processes, the glial apparatus requiring direct apposition is kept close to the synapses, so as to 'lean against them,' as it were.

It is known that at any given time, only about half of the synapses are able to carry a tripartite arrangement; for instance in CA1, about 57% of the spines are found to have perisynaptic astrocytic processes (Ventura and Harris, 1999). Where coverage is needed but is not present, it is expected that the astrocytes have a way to respond to neuronal signals telling them to come closer.

Because the induction of SICs can commence only after tight apposition of the intended neurons is accomplished, the aMEID receptors, as the MEID receptors in Figure 5, must be equipped with means to initiate growth.

References

Abeles, M. (1991). Corticonics: Neural Circuits of the Cerebral Cortex (Cambridge, UK: Cambridge University Press).

Abraham, W.C. (2003). How long will long-term potentiation last? Phil. Trans. Roy. Soc. Lond. B *358*, 735–744.

Amit, D.J. (1989). Modeling Brain Function: The World of Attractor Neural Networks (Cambridge, UK: Cambridge University Press).

Araque, A., Parpura, V., Sanzgiri, R.P., and Haydon, P.G. (1998). Glutamate-dependent astrocyte modulation of synaptic transmission between cultured hippocampal neurons. Eur. J. Neurosci. *10*, 2129–2142.

Araque, A., Parpura, V., Sanzgiri, R.P., and Haydon, P.G. (1999). Tripartite synapses: glia, the unacknowledged partner. Trends Neurosci. *22*, 208–215.

Babich, F.R., Jacobson, A.L., Bubash, S., and Jacobson, A. (1965). Transfer of a response to naive rats by injection of ribonucleic acid extracted from trained rats. Science *144*, 656–657.

Bienenstock, E.L., Cooper, L., and Munro, P. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in the visual cortex. J. Neurosci. *2*, 32–48.

Bliss, T.V.P. and Collingridge, G.L. (1993). A synaptic model of memory: long-term potentiation in the hippocampus. Nature *361*, 31–39.

Bliss, T.V.P. and Gardner-Medwin, A.R. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path. J. Physiol. (Lond.). *232*, 357–374.

Bliss, T.V.P. and Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. J. Physiol. (Lond.) *232*, 331–356.

Borbély, É., Scheich, B., and Helyes, Z. (2013). Neuropeptides in learning and memory. Neuropeptides *47*, 439–450.

Braitenberg, V. and Schüz, A. (1998). Cortex: Statistics and Geometry of Neuronal Connectivity, 2nd ed. (Berlin, Heidelberg, New York: Springer).

Branco, T. and Häusser, M. (2010). The single dendritic branch as a fundamental functional unit in the nervous system. Curr. Opin. Neurobiol. *20*, 494–502.

Bushong, E.A., Martone, M.E., Jones, Y.Z., and Ellisman, M.H. (2002). Protoplasmic astrocytes in CA1 stratum radiatum occupy separate anatomical domains. J. Neurosci. *22*, 183–192.

Buzsáki, G. (2002). Theta oscillations in the hippocampus. Neuron *33*, 325–340.

Buzsáki, G. (2005). Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory. Hippocampus *15*, 827–840.

Buzsáki, G. (2006). Rhythms of the Brain (Oxford, UK: Oxford University Press).

Buzsáki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. Neuron *68*, 362–385.

Caianiello, E.R. (1961). Outline of a theory of thought-processes and thinking machines. J. Theoret. Biol. *2*, 204–235.

Cajal, S.R. (1894). La fine structure des centres nerveux. Proc. Roy. Soc. Lond. *55*, 444–468.

Collingridge, G.L., Peineau, S., Howland, J.G., and Wang, Y.T. (2010). Long-term depression in the CNS. Nat. Rev. Neurosci. *11*, 459–473.

Davis, H.P. and Squire, L.R. (1984). Protein synthesis and memory: a review. Psych. Bull. *96*, 518–559.

Deuchars, J. and Thomson, A.M. (1996). CA1 pyramid-pyramid connections in rat hippocampus *in vitro*: dual intracellular recordings with biocytin filling. Neuroscience *74*, 1009–1018.

Dillon, L.W., Kumar, P., Shibata, Y., Wang, Y.-H., Willcox, S., Griffith, J.D., Pommier, Y., Takeda, S., and Dutta, A. (2015). Production of extrachromosomal microDNAs is linked to mismatch repair pathways and transcriptional activity. Cell Rep. *11*, 1749–1759.

Dudek, S.M. and Bear, M.F. (1992). Homosynaptic long-term depression in area CA1 of hippocampus and effects of *N*-methyl-D-aspartate receptor blockade. Proc. Natl. Acad. Sci. USA *89*, 4363–4367.

Edelman, G.M. (1972). Antibody structure and molecular immunology. Nobel Lecture.

Fellin, T. (2009). Communication between neurons and astrocytes: relevance to the modulation of synaptic and network activity. J. Neurochem. *108*, 533–544.

Fellin, T., Pascual, O., Gobbo, S., Pozzan, T., Haydon, P.G., and Carmignoto, G. (2004). Neuronal synchrony mediated by astrocytic glutamate through activation of extrasynaptic NMDA receptors. Neuron *43*, 729–743.

Flexner, J.B., Flexner, L.B., and Stellar, E. (1963). Memory in mice is affected by intracerebral puromycin. Science *141*, 57–59.

Frazier, W.T., Kandel, E.R., Kupfermann, I., Waziri, R., and Coggeshall, R.E. (1967). Morphological and functional properties of identified neurons in the abdominal ganglion of *Aplysia californica*. J. Neurophysiol. *30*, 1288–1351.

Frey, U. and Morris, R.G. (1997). Synaptic tagging and long-term potentiation. Nature *385*, 533–536.

Fricker, L.D. (2012). Neuropeptides and Other Bioactive Peptides: From Discovery to Function (San Rafael, CA, USA: Morgan and Claypool).

Gasparini, S., Migliore, M., and Magee, J.C. (2004). On the initiation and propagation of dendritic spikes in CA1 pyramidal neurons. J. Neurosci. *24*, 11046–11056.

Gerstein, G.L., Bedenbaugh, P., and Aertsen, A.M.H.J. (1989). Neuronal assemblies. IEEE Trans. Biomed. Eng. *36*, 4–14.

Gold, P.E. (2008). Protein synthesis inhibition and memory: formation vs amnesia. Neurobiol. Learn. Mem. *89*, 201–211.

Haber, M., Zhou, L., and Murai, K.K. (2006). Cooperative astrocyte and dendritic spine dynamics at hippocampal excitatory synapses. J. Neurosci. *26*, 8881–8891.

Halassa, M.M. and Haydon, P.G. (2010). Integrated brain circuits: astrocytic networks modulate neuronal activity and behavior. Annu. Rev. Physiol. *72*, 335–355.

Halassa, M.M., Fellin, T., Takano, H., Dong, J.H., and Haydon, P.G. (2007). Synaptic islands defined by the territory of a single astrocyte. J. Neurosci. *27*, 6473–6477.

Hamming, R.W. (1950). Error detecting and error correcting codes. Bell Syst. Tech. J. *29*, 147–160.

Hampson, R.E., Byrd, D.R., Konstantopoulos, J.K., Bunn, T., and Deadwyler, S.A. (1996). Hippocampal place fields: relationship between degree of field overlap and cross-correlations within ensembles of hippocampal neurons. Hippocampus *6*, 281–293.

Harris, E.W. and Cotman, C.W. (1986). Long-term potentiation of guinea pig mossy fiber responses is not blocked by *N*-methyl D-aspartate antagonists. Neurosci. Lett. *70*, 132–137.

Harris, K.D., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsáki, G. (2003). Organization of cell assemblies in the hippocampus. Nature *424*, 552–556.

Hebb, D.O. (1949). The Organization of Behavior (New York: Wiley).

Hebb, D.O. (1976). Physiological learning theory. J. Abnormal. Child Psychol. *4*, 309–314.

Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. Proc. Natl. Acad. Sci. USA *79*, 2554–2558.

Hydén, H. and Egyházi, E. (1962). Nuclear RNA changes during a learning experiment in rats. Proc. Natl. Acad. Sci. USA *48*, 1366–1372.

Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D., and Yuste, R. (2004). Synfire chains and cortical songs: temporal modules of cortical activity. Science *304*, 559–564.

Kamondi, A., Acsády, L., and Buzsáki, G. (1998). Dendritic spikes are enhanced by cooperative network activity in the intact hippocampus. J. Neurosci. *18*, 3919–3928.

Kandel, E.R. (2001). The molecular biology of memory storage: a dialogue between genes and synapses. Science *294*, 1030–1038.

Kandel, E.R., Dudai, Y., and Mayford, M.R. (2014). The molecular systems biology of memory. Cell *157*, 163–186.

Kirov, S.A., Sorra, K.E., and Harris, K.M. (1999). Slices have more synapses than perfusion-fixed hippocampus from both young and mature rats. J. Neurosci. *19*, 2876–2886.

Kofuji, P. and Newman, E.A. (2004). Potassium buffering in the central nervous system. Neuroscience *129*, 1045–1056.

Konorski, J. (1948). Conditioned Reflexes and Neuronal Organization (Cambridge, UK: Cambridge University Press).

Kozlov, A.S., Angulo, M.C., and Charpak, S. (2006). Target cell-specific modulation of neuronal activity by astrocytes. Proc. Natl. Acad. Sci. USA *103*, 10058–10063.

Krug, M., Lössner, B., and Ott, T. (1984). Anisomycin blocks the late phase of long-term potentiation in the dentate gyrus of freely moving rats. Brain Res. Bull. *13*, 39–42.

Kuffler, S.W., Nicholls, J.G., and Orkand, R.K. (1966). Physiological properties of glial cells in the central nervous system of amphibia. J. Neurophysiol. *29*, 768–787.

Langdon, R.B., Johnson, J.W., and Barrionuevo, G. (1995). Posttetanic potentiation and presynaptically induced long-term potentiation at the mossy fiber synapse in rat hippocampus. J. Neurobiol. *26*, 370–385.

Larkum, M.E., Nevian, T., Sandler, M., Polsky, A., and Schiller, J. (2009). Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle. Science *325*, 756–760.

Legéndy, C.R. (1967). On the scheme by which the human brain stores information. Math. Biosci. *1*, 555–597.

Legéndy, C.R. (1970). The brain and its information trapping device. Progress in Cybernetics, Vol. 1. J. Rose, ed. (New York: Gordon and Breach), pp. 309–338.

Legéndy, C.R. (1975). Three principles of brain function and structure. Int. J. Neurosci. *6*, 237–254.

Legéndy, C.R. (1978). Cortical columns and the tendency of neighboring neurons to act similarly. Brain Res. *158*, 89–105.

Legéndy, C.R. (2009). Circuits in the Brain: A Model of Shape Processing in the Primary Visual Cortex (Berlin, Heidelberg, New York: Springer).

Legéndy, C.R. and Salcman, M. (1985). Bursts and recurrences of bursts in the spike trains of spontaneously active striate cortex neurons. J. Neurophysiol. *53*, 926–939.

Lendvai, B., Stern, E.A., Chen, B., and Svoboda, K. (2000). Experience-dependent plasticity of dendritic spines in the developing rat barrel cortex *in vivo*. Nature *404*, 876–881.

Lømo, T. (1966). Frequency potentiation of excitatory synaptic activity in the dentate area of the hippocampal formation. Acta Physiol. Scand. *68*, 128.

Losonczy, A. and Magee, J.C. (2006). Integrative properties of radial oblique dendrites in hippocampal CA1 pyramidal neurons. Neuron *50*, 291–307.

Major, G., Larkum, M.E., and Schiller, J. (2013). Active properties of neocortical pyramidal neuron dendrites. Annu. Rev. Neurosci. *36*, 1–24.

Market, E. and Papavasiliou, F.N. (2003). V(D)J recombination and the evolution of the adaptive immune system. PLoS Biol. *1*, 24–27.

Martin, K.C. and Kosik, K.S. (2002). Synaptic tagging – who is it? Nat. Rev. Neurosci. *3*, 813–820.

Martin, S.J., Grimwood, P.D., and Morris, R.G.M. (2000). Synaptic plasticity and memory: an evaluation of the hypothesis. Annu. Rev. Neurosci. *23*, 649–711.

McConnell, J.V. (1962). Memory transfer through cannibalism in planarians. J. Neuropsychiat. *3*, 42–48.

McCulloch, W.S. and Pitts, W. (1943). A logical calculus of ideas immanent in nervous activity. Bull. Math. Biophys. *5*, 115–133.

Megías, M., Emri, Z.S., Freund, T.F., and Gulyás, A.I. (2001). Total number and distribution of inhibitory and excitatory synapses on hippocampal CA1 pyramidal cells. Neuroscience *102*, 527–540.

Mézard, M. and Montanari, A. (2009). Information, Physics, and Computation (Oxford, UK: Oxford University Press).

Miller, K.D. and MacKay, D.J.C. (1994). The role of constraints in Hebbian learning. Neural Comput. *6*, 100–126.

Morange, M. (2006). What history tells us VI. The transfer of behaviours by macromolecules. J. Biosci. *31*, 323–327.

Nader, K., Schafe, G.E., and Le Doux, J.E. (2000). Fear memories require protein synthesis in the amygdale for reconsolidation after retrieval. Nature *406*, 722–726.

Oberheim, N.A., Tian, G.F., Han, X., Peng, W., Takano, T., Ransom, B., and Nedergaard, M. (2008). Loss of astrocytic domain organization in the epileptic brain. J. Neurosci. *28*, 3264–3276.

Oja, E. (1982). Simplified neuron model as a principal component analyzer. J. Math. Biol. *15*, 267–273.

Palm, G. (1981a). Evidence, information and surprise. Biol. Cybernetics *42*, 57–68.

Palm, G. (1981b). Towards a theory of cell assemblies. Biol. Cybernetics *39*, 181–194.

Palm, G. (1982). Neural assemblies (Berlin, Heidelberg, New York: Springer-Verlag).

Palm, G. (2012). Novelty, information and surprise (Berlin, Heidelberg, New York: Springer-Verlag).

Palm, G., Knoblauch, A., Hauser, F., and Schüz, A. (2014). Cell assemblies in the cerebral cortex. Biol. Cybernetics *108*, 559–572.

Palmer, L.M., Shai, A.S., Reeve, J.E., Anderson, H.L., Paulsen, O., and Larkum, M.E. (2014). NMDA spikes enhance action potential generation during sensory input. Nat. Neurosci. *17*, 383–390.

Pinsker, H., Kupfermann, I., Castellucci, V., and Kandel, E.R. (1970). Habituation and dishabituation of the gill-withdrawal reflex in *Aplysia*. Science *167*, 1740–1742.

Polsky, A., Mel, B., and Schiller, J. (2009). Encoding and decoding bursts by NMDA spikes in basal dendrites of layer 5 pyramidal neurons. J. Neurosci. *29*, 11891–11903.

Rapoport, A. (1952). "Ignition" phenomena in random nets. Bull. Math. Biophys. *14*, 35–44.

Redondo, R.L. and Morris, R.G. (2011). Making memories last: the synaptic tagging and capture hypothesis. Nat. Rev. Neurosci. *12*, 17–30.

Remy, S. and Spruston, N. (2007). Dendritic spikes induce single-burst long-term potentiation. Proc. Natl. Acad. Sci. USA *104*, 17192–17197.

Reymann, K.J. and Frey, J.U. (2007). The late maintenance of hippocampal LTP: requirements, phases, 'synaptic tagging', 'late-associativity' and implications. Neuropharmacology *52*, 24–40.

Sajikumar, S. and Frey, J.U. (2004). Late-associativity, synaptic tagging, and the role of dopamine during LTP and LTD. Neurobiol. Learn. Mem. *82*, 12–25.

Sajikumar, S., Navakkode, S., and Frey, J.U. (2007). Identification of compartment- and process-specific molecules required for "synaptic tagging" during long-term potentiation and long-term depression in hippocampal CA1. J. Neurosci *27*, 5068–5080.

Sanes, J.R. and Lichtman, J.W. (1999). Can molecules explain long-term potentiation? Nat. Neurosci. *2*, 597–604.

Scott, A.C. (1975). Neurodynamics (a critical survey). MRC Technical Summary Report No. 1548 (Mathematics Research Center, University of Wisconsin-Madison), October 1975. Note: This is a preprint of the article "Scott (1977)," cited by Hebb (1976).

Scott, A.C. (1977). Neurodynamics: a critical survey. J. Math. Psychol. *15*, 1–45.

Scoville, W.B. and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. J. Neurol. Neurosurg. Psychiatry *20*, 11–21.

Shannon, C.E. (1948). A mathematical theory of communication. Bell Syst. Tech. J. *27*, 379–423, 623–656.

Shannon, C.E. (1957). Certain results in coding theory for noisy channels. Inform. Contr. *1*, 6–25.

Shibata, Y., Kumar, P., Layer, R., Willcox, S., Gagan, J.R., Griffith, J.D., and Dutta, A. (2012). Extrachromosomal microDNAs and chromosomal microdeletions in normal tissues. Science *336*, 82–86.

Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. Annu. Rev. Physiol. *55*, 349–374.

Smith, S.L., Smith, I.T., Branco, T., and Häusser, M. (2013). Dendritic spikes enhance stimulus selectivity in cortical neurons *in vivo*. Nature *503*, 115–120.

Song, S., Miller, K.D., and Abbott, L.F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. Nat. Neurosci. *3*, 919–926.

Spacek, J. and Harris, K.M. (1998). Three-dimensional organization of cell adhesion junctions at synapses and dendritic spines in area CA1 in the rat hippocampus. J. Comp. Neurol. *393*, 58–68.

Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. Nat. Neurosci. *9*, 206–221.

Stettler, D.D., Yamahachi, H., Li, W., Denk, W., and Gilbert, C.D. (2006). Axons and synaptic boutons are highly dynamic in adult visual cortex. Neuron *49*, 877–887.

Svoboda, K., Den, W., Kleinfeld, D., and Tank, D.W. (1997). *In vivo* dendritic calcium dynamics in neocortical pyramidal neurons. Nature *385*, 161–165.

Trachtenberg, J.T., Chen, B.E., Knott, G.W., Feng, G., Sanes, J.R., Welker, E., and Svoboda, K. (2002). Long-term *in vivo* imaging of experience-dependent synaptic plasticity in adult cortex. Nature *420*, 788–794.

Ts'o, D.Y., Gilbert, C.D., and Wiesel, T.N. (1986). Relationships between horizontal interactions and functional architecture it cat striate cortex as revealed by cross-correlation analysis. J. Neurosci. *6*, 1160–1170.

Tulving, E. and Markowitsch, H.J. (1998). Episodic and declarative memory: role of the hippocampus. Hippocampus *8*, 198–204.

Ungar, G., Desiderio, D.M., and Parr, W. (1972). Isolation, identification, and synthesis of a specific-behaviour-inducing brain peptide. Nature *238*, 198–202.

Vargha-Khadem, F., Gadian, D.G., Watkins, K.E., Connelly, A., Van Paesschen, W., and Mishkin, M. (1997). Differential effects of early hippocampal pathology on episodic and semantic memory. Science *277*, 376–380.

Ventura, R. and Harris, K.M. (1999). Three-dimensional relationships between hippocampal synapses and astrocytes. J. Neurosci. *19*, 6897–6906.

von der Malsburg, C. (1999). The what and why of binding: the modeler's perspective. Neuron *24*, 95–104.

von Neumann, J. (1952). Probabilistic logics and the synthesis of reliable organisms from unreliable components (based on a series of lectures delivered at the California Institute of Technology, January 4–15, 1952). Automata Studies. C.E. Shannon and J. McCarthy, eds. (Princeton, NJ, USA: Princeton University Press, 1956), pp 43–98.

White, C., Droppo, J., Acero, A., and Odell, J. (2007). Maximum entropy confidence estimation for speech recognition. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), *4*, pp. IV-809–IV-812.

Wickelgren, W.A. (1999). Webs, cell assemblies, and chunking in neural nets: introduction. Can. J. Exp. Psych. *53*, 118–131.

Wilson, H.R. and Cowan, J.D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. Biophys. J. *12*, 1–24.

Witcher, M.R., Kirov, S.A., and Harris, K.M. (2007). Plasticity of perisynaptic astroglia during synaptogenesis in the mature rat hippocampus. Glia *55*, 13–23.

Wittner, L., Henze, D.A., Záborszky, L., and Buzsáki, G. (2007). Three-dimensional reconstruction of the axon arbor of a CA3 pyramidal cell recorded and filled *in vivo*. Brain Struct. Funct. *212*, 75–83.

Woolley, C.S., Gould, E., Frankfurt, M., and McEwen, B.S. (1990). Naturally occurring fluctuation in dendritic spine density on adult hippocampal pyramidal neurons. J. Neurosci. *10*, 4035–4039.

Yamahachi, H., Marik, S.A., McManus, J.N.J., Denk, W., and Gilbert, C.D. (2009). Rapid axonal sprouting and pruning accompany functional reorganization in primary visual cortex. Neuron *64*, 719–729.