



Incremental change in the set of coactive cortical assemblies enables mental continuity

Jared Edward Reser, PhD

University of Southern California, 16380 Meadow Ridge Road, Encino, CA 91436–3607, United States



HIGHLIGHTS

- The term state-spanning coactivity (SSC) is introduced to refer to shared content between successive brain states.
- SSC is made possible by sustained firing, binding, and recurrent processing.
- Sustained activity ensures that consecutive topographic maps are interrelated.
- SSC also permits continuous, algorithmic, and progressive alterations to mental imagery.
- Incremental change in SSC (icSSC) allows the continuity necessary for updating, modeling, and systemization.

ARTICLE INFO

Article history:

Received 13 January 2016

Received in revised form 9 September 2016

Accepted 17 September 2016

Available online 19 September 2016

Keywords:

Association cortex

Coactivation

Cortical assembly

Neural assemblies

Primary sensory cortex

Short-term memory

Systems neuroscience

Working memory

ABSTRACT

This opinion article explores how sustained neural firing in association areas allows high-order mental representations to be coactivated over multiple perception-action cycles, permitting sequential mental states to share overlapping content and thus be recursively interrelated. The term “state-spanning coactivity” (SSC) is introduced to refer to neural nodes that remain coactive as a group over a given period of time. SSC ensures that contextual groupings of goal or motor-relevant representations will demonstrate continuous activity over a delay period. It also allows potentially related representations to accumulate and coactivate despite delays between their initial appearances. The nodes that demonstrate SSC are a subset of the active representations from the previous state, and can act as referents to which newly introduced representations of succeeding states relate. Coactive nodes pool their spreading activity, converging on and activating new nodes, adding these to the remaining nodes from the previous state. Thus, the overall distribution of coactive nodes in cortical networks evolves gradually during contextual updating. The term “incremental change in state-spanning coactivity” (icSSC) is introduced to refer to this gradual evolution. Because a number of associated representations can be sustained continuously, each brain state is embedded recursively in the previous state, amounting to an iterative process that can implement learned algorithms to progress toward a complex result. The longer representations are sustained, the more successive mental states can share related content, exhibit progressive qualities, implement complex algorithms, and carry thematic or narrative continuity. Included is a discussion of the implications that SSC and icSSC may have for understanding working memory, defining consciousness, and constructing AI architectures.

© 2016 The Author. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The present article will delineate a simplistic but previously overlooked nonlinear dynamic pattern of brain activity. Two hypothetical constructs are introduced to describe this pattern. The first construct is state-spanning coactivity (SSC), which occurs when cortical nodes exhibit sustained coactivity during the span of short-term memory. The gradual evolution of SSC exhibits a distinctive spatiotemporal

pattern of turnover as it plays out in real time. The second construct introduced here, incremental change in state-spanning coactivity (icSSC), refers to this pattern of turnover. icSSC conveys that the set of nodes that are simultaneously coactive changes incrementally as newly activated nodes are added and others are deactivated while a distinct subset remains in SSC. Spreading activity from the nodes in SSC select: 1) inactive neural nodes for activation, 2) active nodes for deactivation, and 3) active nodes for sustained activation. Because a distinct subset of nodes is always conserved from one brain state to the next, each state is embedded recursively in the previous state, amounting to an iterative process that has the potential to progress algorithmically toward a complex

E-mail address: jared@jaredreser.com.

result. The general intention of the present article is to propose a qualitative model delineating the theoretical functions of SSC and icSSC from the perspective of cognitive neuroscience.

The term SSC can be used either to denote a property or to designate a set of neurons. *icSSC* denotes a property or process (Table 1). Both are related to the construct of working memory, which is defined as a system responsible for the transient holding and processing of attended information. The fundamental assumption made by this article is that the content of working memory can be said to be in SSC; and as working memory progresses over time, the content can be said to exhibit *icSSC*. This assumption is applied not only to working memory as the same could be said of attention, consciousness or short-term memory. *icSSC* can be taken to be the underlying neural substrate of mental continuity. As proposed here, mental continuity is a process where a gradually changing collection of mental representations held in attention/working memory emerges from the *icSSC* of neural nodes. The thematic and narrative quality created by this continuity during internally generated thought may be largely congruent with key facets of conscious experience. In the course of exploring how neural continuity creates mental continuity, this article will attempt to integrate current theoretical approaches while remaining consistent with prevailing knowledge.

Animals are information-processing agents. They receive unprocessed data through sensory receptors, expose it to a massively parallel network of nodes and channels, and allow the interaction between the activity and the existing network to determine behavior. Even small invertebrates with elementary nervous systems exhibit ongoing, internally generated neural activity that temporarily biases the network weights. Because it involves mechanisms that include sustained firing, this continuous endogenous processing constitutes a fleeting form of SSC, even in animals like the nematode and fruit fly. In vertebrates, however, SSC involves the coactivation of high-level representations from long-term memory within a single, massively interconnected representational network (telencephalon). Each such representation is a record of the distribution of past neural activity corresponding to a recognizable stimulus or motor pattern. An instantaneous attentional state is composed of a novel combination of these template-like representations which together create contextual, cognitive content. The mammalian neocortex can hold a number of such mnemonic representations coactive for hundreds of milliseconds, using them to make predictions by allowing them to spread their activation energy together, throughout the thalamocortical network. This activation energy converges on the inactive representations in long-term memory that are the most closely connected with the current group of active representations, making them active and pulling them into SSC. Thus, new representations join the representations that recruited them, are incorporated into the set of coactive parameters in SSC and used in subsequent searches.

When the activity of certain nodes can be sustained for several seconds at a time, as in primate association cortex, the complexity of search in such a system increases. Highly sustained activity allows prioritized representations to act as search parameters for multiple perception-action cycles. This permits more dynamic *icSSC*, whereby goal-relevant representations can be held constant as others are allowed to change.

Table 1
Definition of key terms.

| | |
|--|---|
| Instantaneous coactivity | The coactivity of a set of cortical nodes in a single instant or state. |
| State-spanning coactivity (SSC) | Sustained coactivity exhibited by a set of two or more cortical nodes that spans two or more consecutive brain states. |
| Incremental change in state-spanning coactivity (<i>icSSC</i>) | The process in which a set of three or more neural nodes exhibiting SSC undergoes a shift in group membership, where at least two nodes remain in SSC and at least one is deactivated and replaced by a new node. |
| Mental continuity | The recursive interrelatedness of consecutive mental states made possible by <i>icSSC</i> . |

The *icSSC* taking place in association areas allows task-pertinent representations to be maintained over multiple cycles, in order to direct complex sequences of interrelated mental states. The individual states in a sequence of such states can be considered interrelated because they share representational content. The associations linking these sequences are saved to memory, impacting future searches and ultimately permitting semantic knowledge, planning, and systemizing.

2. Sustained firing, attentional updating, and memory decay

Mammals regularly encounter scenarios involving sets of stimuli that may remain present (or relevant) throughout the experience. In order to systemize such a scenario, it may be necessary to maintain mental representations of the pertinent contextual stimuli during the experience, and even afterward. Mammalian brains are well-equipped to do exactly this. The glutamatergic pyramidal neurons in the prefrontal cortex (PFC), parietal cortex, and other association cortices, are specialized for sustained firing, allowing them to generate action potentials at elevated rates for several seconds at a time [35]. In contrast, neurons in other brain areas, including cortical sensory areas, often remain persistently active for periods of mere milliseconds unless sustained input from either the environment or association areas makes their continued activity possible [35]. A neuron may exhibit tonic sustained firing due to temporary changes in the strength of certain synapses (short-term synaptic modification [80]), its intrinsic biophysical properties, extrinsic circuit properties (reverberatory circuits), or dopaminergic innervation [25]. Prolonged activity of neurons in association areas is largely thought to allow the maintenance of specific features, patterns and goals [8].

Goldman-Rakic [37,38] first suggested that the phenomenon of sustained firing in the PFC is responsible for the information maintenance capabilities of the temporary storage buffers of working memory. Goldman-Rakic [39] also proposed that the PFC is parceled into several specialized regions, each of which is responsible for detecting, representing and sustaining a different extraction of multimodal information. Since then, the PFC, along with a number of association areas, has been divided into increasingly smaller modules, each with unique receptive/projective fields and functional properties including faculties such as short-term spatial memory, short-term semantic memory, response switching, error detection, reward anticipation, impulse suppression, and many others. Working memory, executive processing and cognitive control are now widely thought to stem from the active maintenance of patterns of activity in the PFC, especially the dorsolateral PFC, that correspond to goal-relevant features and patterns [33,34]. The temporary persistence of these patterns ensures that they continue to transmit their effects on network weights as long as they remain active, biasing ongoing processing, and affecting the interpretation of stimuli that occur during their episode of continual firing [57]. This persistence ensures that context from the recent past is taken into account during action selection.

During any experience, some neural nodes exhibit more prolonged sustained firing than others. I will assume that in general the most enduringly active nodes correspond to what attention is most focused on, or the underlying theme that remains most constant as other contextual features change. From subjective introspection we know that when we envision a scenario in our mind's eye, we often notice it transform into a related but distinctly different scenario [46]. These two scenarios are related because our brain is capable of *icSSC*. In other words, the distribution of active neurons in the brain transfigures incrementally from one configuration to another, instead of changing all at once. If it were not for the phenomenon of *icSSC*, instantaneous information processing states would be time-locked and isolated (as in most serial and parallel computing architectures), rather than continuous with the states before and after them.

These observations point to the notion that every cortical state is composed of a subset of elements from the previous state, and also

composed of increasingly smaller subsets of elements of states directly before that. In fact, when comparing successive cortical states, the shorter the time difference between two states (on the order of seconds to fractions of milliseconds), the more similar in composition the two states will be. For instance, over the span of 10 milliseconds, a relatively large proportion of nodes will exhibit uninterrupted coactivity; however, over 10 s, this proportion will be much smaller. Here, we will be concerned with neural nodes exhibiting SSC at two distinct levels: A) short-term memory/priming, i.e., elements of long-term memory activated above baseline (for seconds to minutes); and B) the focus of attention/immediate memory, i.e., a small, perhaps more active subset of A (for milliseconds to a few seconds). Items in SSC within the focus of attention likely demonstrate active neural binding whereas items in SSC within short-term memory may not.

Mental continuity and iSSC require a densely interconnected representational system such as a neural network that is capable of holding two or more representations (each specifying discrete and separate informational content) active over the course of two or more points in time (Fig. 1). The sustained activity of a single representation over time does not provide any context or associative/relational content, and so should not be taken to be sufficient for mental continuity. More than one representation is needed. Although its limits are presently being debated, the human neocortex is clearly capable of holding numerous representations active over numerous points in time.

In Fig. 1 above, representations B, C, D, and E are active during t_1 , and C, D, E and F are active during t_2 . Thus representations C, D, and E demonstrate SSC because they exhibit continuous and uninterrupted activity from t_1 through t_2 . The brain state at t_1 and the brain state at t_2 share C, D, and E in common and therefore can be expected to share other commonalities such as: similar information processing operations, similar memory search parameters, similar mental imagery, similar cognitive and declarative aspects, and similar experiential and phenomenal characteristics. The active nodes that have demonstrated SSC over any specific time interval can be thought of as constituting a unit with emergent functional properties. Together, these nodes impose sustained information processing demands on the lower-order sensory and motor

areas within the reach of their long-range connections. The longer the activity in these higher-order neurons is sustained, the longer they remain engaged in hierarchy-spanning, recurrent (and reentrant) broadcasting throughout the cortex and subcortex.

Compared to those of other mammals, human association areas contain more neurons, more intrinsic and extrinsic connections, and a higher capacity for sustained firing [33,34]. These characteristics presumably permit us to retain more information, for a longer time before it decays. This likely allows humans to better retain elements from recent thoughts, and allows the computational results of previous processes to more thoroughly inform subsequent ones. This once influenced the present author to assume that somehow thoughts are “longer” in humans than they are in other animals; however, if thought has an architectural geometry marked by length, then mustn't it also have starting and stopping points? If persistent activity of individual representations in SSC is staggered and overlapping, then there cannot be objective stopping or starting points of thought. Instead, thought itself must be composed of the startings and stoppings of huge numbers of individual elements that could be depicted graphically in the form of a continuous, stream-like distribution (Fig. 2). Therefore, it is not that human thoughts are somehow longer than in other animals; rather, human thought is composed of larger sets of representations that are capable of remaining coactivated longer [70,71].

The reallocation of processing resources in Fig. 2 is similar to the behavior of treads on a military tank. Individual treads are continually placed on the ground temporarily, and the treads that have sat on the ground for the longest are withdrawn in series. The total set of treads touching the ground in one moment partially overlaps with the total set in the next. Our mental set of active representations may cycle in an analogous, although more flexible and stochastic manner. A more precise analogy and schematic will be introduced in Section 5.

PFC neurons are likely tuned throughout life to best determine what aspects of the present environment should be maintained in SSC (or released from maintenance) given the current scenario and its preceding circumstances. When confronted with a complex configuration of stimuli, the PFC may select the representations that it “predicts” should be




Fig. 1. Schematic depiction of iSSC in the focus of attention. Each bracket represents the active time span of a representation. The x-axis represents time, and the y-axis demarcates the cortical area where the representation is active. Red brackets denote representations that have exhibited uninterrupted activity from the point when they became active, whereas blue brackets denote representations that have not been sustained. In time sequence 1, representations B, C, D and E have remained active until t_1 . In time sequence 2, B has deactivated, C, D and E have remained active, and F is newly active. The figure depicts a system with SSC because more than one representation (C, D, and E) has been maintained over more than one point in time (t_1 and t_2). The figure depicts iSSC because the removal of B and addition of F were incremental and did not disrupt the SSC of C, D and E. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 2. A simplified schematic representation of icSSC depicting it as a gradually shifting, stream-like distribution. This figure extends Fig. 1 over five time intervals revealing a repeating pattern: remnants of the preceding state are consistently carried over to the next state. If this distributional plot were modeling neurons rather than representations, there might be millions of units per time period rather than four; however, this fundamental reoccurring pattern would remain.

temporarily maintained for their processing utility in the immediate future. This selection process is likely determined by the incoming stimulus configuration itself, prior probability as encoded in the network, and the network-biasing representations already in SSC. Initially during development, the process of selecting neurons for persistent activity may be random and heavily influenced by innate connectivity. The expertise of the PFC is probably garnered slowly, over developmental time, after connections between groups of neurons exhibiting sustained firing are strengthened for their role in mediating task proficiency and reward achievement. The selection process for SSC is perhaps best exemplified by the ability to identify and maintain strategically important representations from a forthcoming scenario. A sentence (spoken or written) is a suitable example. A sentence will be comprehended if: 1) the relevant representations are identified and enter into SSC; 2) all the necessary representations are sustained throughout the duration of the sentence; 3) the network has enough experience with this particular combination of representations to build the appropriate imagery, depicting them in the way they were intended. Most people have had the experience where either the wrong representations were anchored upon, or the right representations could not be maintained for long enough, and the sentence had to be repeated or reread.

The quantity of SSC can be thought of as directly proportional to the number of sustained nodes and the average length of time of their activity [69–71]. It should also be possible, in theory, to quantify icSSC by determining the proportion of previously active neural nodes that have remained active over a given time period. One way to do this would be to determine how long it takes half of the currently firing association neurons to sufficiently reduce their firing. Employing the idea of a “half-life” may be a useful concept even though the “decaying quantity” may not exhibit constant exponential decay, and despite the fact that current scanning and recording methods could not produce the necessary data without significant extrapolation. If the average rate of decay was properly operationally defined and could be measured, then cognitive neuroscientists would be able to discuss the “icSSC half-life” associated with individuals or even species. Would it be informative if it were found that Wistar rats have an average icSSC half-life of, say, one second, macaque monkeys twice this and humans twice that? Even in a single individual, this number is likely to vary depending on the task at hand, level of arousal, motivational state, brain oscillation factors, and brain regions assayed. Moreover, short-term memory/priming would have a much longer half-life than the focus of attention. An SSC/icSSC

profile featuring numerous such assays could be computed for an individual based on various standardized criteria. If characterized correctly and averaged meaningfully, these numbers could prove to be consistent and reliable psychometric markers. Tononi [83] developed a method for calculating a measure of “integrated information” within a single, static brain state. The concept of icSSC could be used to expand on this measure in order to calculate the integration of information between two brain states, or across multiple brain states.

It is not always the case that the majority of representations are conserved from one thought to the next. When they become a lower priority, nearly all items in the focus of attention can be displaced at the same time. This readily happens when we are exposed to a new, salient, perhaps emotionally laden stimulus. Whenever a person loses their train of thought, and forgets what they were just thinking, SSC in the focus of attention (though not necessarily in short-term memory) is interrupted. SSC “jumps,” reallocating attentional resources, and reorienting to the new stimulus configuration and its accompanying set of features. Such a jump would constitute a disruption of, or fluctuation in, mental continuity. The degree of fluctuation in continuity varies depending on the proportion of neural activity that is abruptly deactivated (Fig. 3). Because icSSC is the change in SSC, as attention shifts, SSC decreases, and icSSC increases.

In the most intelligent mammals, late motor output and early sensory activity are heavily influenced by several seconds of sustained input from association areas. In mammals with smaller association areas, capable of less SSC, motor and sensory output are informed by a much briefer window of continuous activity. High SSC likely allows “behavioral continuity” where sequential behaviors can be complexly interrelated and mutually informed. This can be contrasted with the more isolated and impulsive behaviors seen in individuals with injuries to the PFC (i.e., field-dependent behavior in which the patient’s behavior is dictated by incidental cues and distractions). In fact, the temporal extent of SSC may be a major facet of the “general factor” of intelligence. SSC may be related to, and a primary determinant of, attention span, behavioral flexibility, working memory capacity, short-term memory capacity, reasoning ability, and general fluid intelligence. Furthermore, significant individual differences in SSC may exist in humans where deficits in this capacity may map onto a variety of clinical syndromes such as schizophrenia, mental retardation, cognitive aging, chronic stress, various forms of intoxication, and prefrontal injury. Nevertheless, why did SSC and icSSC evolve, what purposes do they serve, and how do




Fig. 3. This figure expands on Fig. 2 by comparing four possible incremental changes in SSC. In the first transition 75% of representational continuity is maintained between time periods 1 and 2. The other transitions depict 50%, 25%, and 0% maintenance of continuity respectively. According to the definition of mental continuity, neither the graphic marked “25% continuity” nor the one marked “0% continuity” depict mental continuity, because they do not feature the maintenance of more than one representation.

they relate to dopaminergic functions? Mammals most likely evolved the capacity to sustain certain representations so that hypothetical groupings of representations could be modeled and systemized.

2.1. Dopamine, task-relevance, and contextual modeling

The dopamine (DA) system exerts complex actions within several interrelated systems of the mammalian brain: motor function, motivation, reward attention, learning, and delayed response [33,34]. DA sent from the ventral tegmental area (VTA) modulates the activity and timing of sustained neural firing in the PFC, association cortices, and elsewhere [76]. DA neurotransmission in the PFC is thought to underlie the ability to internally represent, maintain, and update contextual information [11]. Given this role, the mesocortical DA system, and the systems on which it acts are likely to be heavily involved in mental continuity. Widespread projections from prefrontal areas are thought to utilize DA to regulate the transmission of incoming stimuli from sensory areas, suppressing irrelevant stimuli and enhancing relevant ones [73]. DA may serve a “gating” function in the PFC, regulating the selection of task-relevant information and non-routine actions in the face of interference from (potentially stronger) task-irrelevant information and routine impulses [9,57]. Relevant actions and perceptions must be actively sustained in SSC so that information related to the context surrounding behavioral goals can bias behavior and perception over temporally extended periods.

DA/PFC system activity fluctuates with environmental demand [32]. The system is engaged by prediction error, when reward or punishment contingencies change. Novel appetitive and aversive events have been shown to increase dopamine release in the VTA, causing sustained firing of PFC neurons [75]. Thus, emotional areas like the nucleus accumbens and the amygdala help to structure the progression of thought by determining what elements remain in SSC the longest. Even in mice, DA levels often remain high for a duration far longer than that of the experience itself. Seamans and Robbins [75] elaborated a functional

explanation. They have stated that the DA system is phasically activated in response to novel rewards and punishments because it is adaptive for the animal to anchor upon and further process unpredicted opportunities and dangers even after the experience itself has ended.

How does simply holding a suite of related representations in SSC allow the animal to make useful associative inferences? The representations in memory for the component parts of the experience coactivate, spreading their activation energy in an attempt to converge on and activate historically associated representations that are not found in the experience itself [70]. Because memory traces for important features remain active and primed, they can probably be used repeatedly as specifications that guide the generation of apposite mental imagery in other processing modules, such as the phonological loop and visuospatial sketchpad [70,71].

It is probable that a succession of lower-order topographic images or maps created in sensory processing modules depict and explore hypothetical, causal relationships between the higher-order, top-down specifications held in SSC. Consecutive topographic images may model the scenario because some of the contextual elements are held constant while others are allowed to change. The fact that newly activated search terms are combined with search terms from the previous cycle makes this process demonstrate qualities of progressive iteration. This allows the important search terms to “work” actively with newly introduced terms in a systematic effort to interrogate the nature of the situation being simulated.

The present article intends to further explore how icSSC enables mental modeling through the mediation of the relationship between the high-order representations held in working memory and the low-order topographic content of sensory modules. First, however, we should make an effort to identify the neurobiological units of SSC. In the present paradigm, two neural units are proposed: assemblies (previously referred to here as *nodes*) and ensembles (previously referred to as *representations*). Both of these demonstrate SSC at different levels of granularity.

3. Pattern recognition, cell assemblies, and neural ensembles

The present model is consistent with connectionism and parallel distributed processing in that it conceptualizes cortical architecture as a plurality of densely interconnected networks composed of decentralized, semi-hierarchically organized, pattern-recognizing nodes that have multiple inputs and outputs [41,47]. Like many biologically plausible neural network models, it envisions these nodes as microscopic, modular neural units in the cerebral cortex and assumes that each individual unit represents an elementary feature or stable “microrepresentation” of long-term memory [43,55,87]. Like other models (e.g., [20,59]), this model views cognition as a system responsible for using active representations from long-term memory to guide goal-directed processing [66]. This is, in some ways, consistent with Fuster’s [35] concept of cognits—distributed items of knowledge defined by patterns of connections between neuron populations associated through experience. Fuster emphasized that his cognits are hierarchically organized, link noncontiguous neurons, overlap, and interconnect profusely. We will assume that to build a higher-order mental representation, the cortex must combine a subset of the numerous, fragmentary lower-order units at its disposal into an improvised composite. In this model, all high-order mental states are constrained in that they can be built only from combinations of preexisting lower-order units.

The fundamental unit of cognition may comprise a number of similarly tuned neurons that are synaptically bound to create a functionally discrete assembly [51]. These have been called “cognitive building blocks,” “lego-like” elements, and “cognitive atoms” [65]. Because the neurons of such an assembly share highly similar receptive fields, the assembly responds to a particular “preferred” aspect of the environment, and can be said to have a unique although primitive “window on the world.” An assembly is maximally activated when the simultaneous firing of a sufficient subset of its inputs converges upon it [31]. In general, when an assembly becomes active, the subsymbolic pattern that it represents has been recognized and it communicates this recognition with the assemblies that it is connected to through its outputs. Thus, the assemblies, like the neurons that compose them, recognize conjunctions and function as “coincidence detectors” or “pattern recognition nodes.” More intelligent mammals likely: 1) have a larger number of available assemblies in the cortex to select from, 2) coactivate a larger number of assemblies simultaneously, and 3) have the capacity to prolong the activation of goal-relevant assemblies for extended periods.

Minicolumns of cells in the cortex are strong candidates for these hypothetical assemblies because they are somewhat spatially distinct, contain neurons with qualitatively similar receptive fields, span each of the cortical layers, and communicate with each other as well as with subcortical structures. Neurons are inelegant candidates for these building blocks because, despite the fact that each neuron has a distinct and singular receptive field, their functional properties vary widely depending on cell type and the layer in which they are found. Hypercolumns are also not good candidates, as they can be divided into subunits with receptive fields that are qualitatively different or mutually inhibitory [44].

It is unlikely that individual cell assemblies represent consciously perceptible constructs. We will refer to a coalition of assemblies that acts as an engram for a whole consciously perceptible construct as an “ensemble.” Ensembles are congruent with the “representations” discussed earlier. Ensembles in association areas encode invariant, high-order representations such as objects, people, places, rules and concepts as well as the propositions that define the nature of the relationships between these. An ensemble is composed of cortical assemblies that have become highly connected in a Hebbian manner owing to approximately simultaneous activity in the past. Ensembles encode statistical, combinatorial information about the most frequently concomitant features of a reoccurring stimulus configuration. Ensembles in frontal cortex will be taken to encode information involving actions,

from discrete muscle-driven movements in premotor cortex to intended behaviors in motor association cortex. In this formulation, ensembles exist within association areas but not within primary sensory or motor areas, because a single construct can take countless neural instantiations in primary areas as neurons here do not capture invariance.

It is unclear exactly what the neural organization of an ensemble is, but perhaps they can take many forms owing to the vast number of different types of representations that can be held in the numerous specialized modules of the brain. I expect that ensembles are not static but instead are constantly transmuted as additional information is injected into them during the top-down, bottom-up and horizontal reciprocations between functional regions. For instance, as a person thinks about a concept, only a particular subset of assemblies that are ordinarily statistically associated with this concept will be active. In a sense, an ensemble is continually fleshing out but never completely encompassing the concept (or Platonic ideal) that it embodies. Such a higher-order map could be referred to as a “macroensemble.” This conceptualization of an ensemble is highly theoretical, but it will allow us to continue in our systematization.

A single assembly probably does not form a crucial part of an ensemble to which it belongs, and surely must contribute to multiple ensembles. An ensemble may be composed of other ensembles, although an assembly cannot be composed of other assemblies. When an assembly’s neurons are targeted by a bombardment of excitatory post-synaptic potentials, this will lead to the “ignition” of the assembly: cells within the input layers become excited enough to activate the assembly’s pyramidal projection neurons, causing it to fire out rapidly to the cells of other assemblies in the cortex. Ensembles may ignite autoassociatively after a sufficient number of their assemblies have been targeted. When a new ensemble is converged upon, it brings a new representation into SSC. In the present model, all active assemblies contribute to processing in a highly interconnected cortical workspace. This processing involves the currently active assemblies pooling their activation energy in order to converge on the most closely associated, task-appropriate ensembles.

3.1. Spreading activation, the cortical workspace, and polyassociativity

The manner in which assemblies and ensembles are selected for activity in this model is consistent with spreading activation theory. According to this theory, associative networks can be searched by labeling a set of source nodes, which then spread their activation energy to closely associated nodes [17]. Here, the propagation of activation energy is passed among assemblies and follows weighted links. The links are found in the axons and dendrites, and the weights are found in the synapses. Several alternate paths through these links can reach the same end assembly. When enough of these links excite the same assembly, it is activated. Unlike subcortical areas, information processing in the cortex is not compartmentalized into individual nuclei that are relatively isolated from processing occurring elsewhere. Rather, the spread of activity in the cortex is thought to involve many-to-one (convergence) and one-to-many (divergence) interactions within a massively interconnected network of assemblies [55]. Because hippocampal-dependent processes impose an additional degree of complexity on the selection of cortical assemblies in SSC, they will be ignored here in the interest of simplification.

Cortical assemblies work cooperatively by spreading the activation energy (both excitatory and inhibitory) necessary to converge upon the next set of ensembles that will be coactivated with the remaining ensembles from the previous cycle. This pooling of activity during the coactivation of multiple assemblies, and the manner that new ensembles are selected for entry into the cortical workspace, will be referred to as “polyassociativity.” This term may be apt because the next representational addition to thought is converged upon by multiple associated specifiers from the previous state. One ensemble does not activate the next in linear sequence. Rather, several assemblies/ensembles are

coactivated together polyassociatively, and they pool their activation energy to determine which assemblies/ensembles will be activated next in a nonlinear yet mechanistic manner. Table 2 demonstrates how active, subsymbolic representations work together to select symbolic representations, resulting in a self-perpetuating “stream” or “train” of thought that demonstrates icSSC.

The brain employs a blind heuristic, summoning up the long-term memory fragments that are most probabilistically related to the currently active set of fragments. Thus, polyassociativity is an automatic and unconscious form of pattern completion that could ultimately be responsible for, not only the execution of automatic associations in an unconscious processing stream, but also sudden insight and even “deliberate” declarative associations. From a subjective perspective, the contents of our next thought are chosen for us based on how the currently active assemblies interact with the existing neuro-/nomological network. The current state will “fall” deterministically into the next most probable or stable state according to the prevailing neurophysiological conditions. Many authors favor the idea of “selfish,” “rivalrous” assemblies that compete with each other for activation energy. However, it may be equally appropriate to view them as “patient,” “democratic” assemblies that wait until enough of their peers signal them before they become active. Furthermore, they are “cooperative” in the sense that they work together to spread activation energy.

Fig. 4 depicts working memory as the interaction between icSSC and polyassociativity. Working memory utilizes polyassociativity to select the representations that will be maintained in SSC in the service of memory recall and behavioral control. In this figure, the letters represent ensembles held in association areas and are largely congruent with the items or chunks that can be held in a limited capacity in working memory. Four ensembles are used after Cowan's four chunks (1995). In time sequence 1 of Fig. 4, the assemblies that comprise ensembles B, C, D and E are each individually associated with a very large number of ensembles, but as a group they are most closely associated with ensemble F. This is analogous to the event where together the words “radioactive,” “nylon,” “arthropod” and “responsibility” may together prompt the recollection of “Spiderman,” yet none of the words alone would be sufficient.

In Fig. 4, between time periods 1 and 2, C, D and E exhibit SSC, whereas, between time periods 1 and 3, only C and E exhibit SSC. C and E are active over all three time periods, meaning that these representations are being used as search function parameters for multiple cycles, and are likely the subject of attention. When we are motivated to think about the relationship between two concepts, the ensembles for these concepts may remain active as other ensembles filter in and out. Alternatively, we can imagine a scenario where B, C, D, and E from step one of Fig. 3 were immediately replaced by F, G, H, and I. Such a processing system may still be using previous states to determine subsequent states; however, because no activity is sustained over time, there would be no continuity in such a system. This is generally how standard, von Neumann computing systems operate, using rule-based symbolic processing to execute serial instructions. Even in contemporary artificial neural networks, recursion marked by icSSC and polyassociativity are not harnessed for information processing.

Table 2
Determinants of polyassociativity.

| |
|---|
| Gradual additions to and subtractions from a pool of simultaneously coactivated ensembles occur as: |
| 1. Assemblies that continue to receive sufficient activation energy from the network are maintained; |
| 2. Assemblies that receive sufficiently reduced activation energy, or are inhibited, are deactivated; |
| 3. New assemblies, which are tuned to receive sufficient activation energy from the currently coactive assemblies, are converged upon, and incorporated into the remaining pool of active assemblies from the previous cycle; |
| 4. The ensembles that most closely correspond to the constellation of presently coactive assemblies are recalled autoassociatively. |

In Fig. 4, ensembles C and E have fired together over three individual time intervals, and thus will show a propensity to wire together, increasing their propensity to fire together in the future. It is more likely, however, that they will recruit one another if the other contextual ensembles are also active. Thus, after the events depicted in Fig. 4, C and E are more likely to fire together in the future, especially if accompanied by B, D, F and G. This may allow mammals to associate two concepts but only within the context of the other concepts that fired with them historically. The association between C and E will become increasingly less specific, as they fire together with more ensembles not included in Fig. 4. The associations between cortical ensembles that are made during a brief episodic event are only slightly strengthened during encoding and the long-term potentiation (LTP) that accompanies it. Coincidental or rare associations between the ensembles of a single experience are probably mostly lost from non-hippocampal-dependent cortical memory. However, reoccurring associations are heavily encoded, persist as semantic knowledge, and may come to underpin beliefs and habitual behaviors. Statistical codependencies between ensembles start early in cortical development when an infant coactivates ensembles for perceptual stimuli with ensembles for: 1) other temporally proximate stimuli, 2) innate responses to these stimuli, and 3) the environmental outcomes of these responses.

Polyassociativity may be able to inform our understanding of the formation of some types of implicit memory. When a group of ensembles becomes coactive and converges on a new representation, the ensembles have a propensity to wire together. This strengthening of the group of ensembles is likely to allow a subset of them to activate the ensemble that previously required them all. In Fig. 4, B, C, D, E and F might become so strongly bound that only three of the four might be sufficient to activate F henceforth. Thus, some ensembles need not be present in SSC, and can become implicit in this process. We are probably often unaware of the statistical regularities in our environment that underlie why some coactivates can be omitted without consequence. If B, C, D, and E were to converge on F, strengthened associations among these representations may make it possible, for instance, for C, D, E and F to converge on B. In essence, associative cognition may stem from the ability of discrete neural ensembles to sensitize (and habituate) to one another during group coactivation. Coactive ensembles likely fuse or differentiate their respective maps depending on whether their active connections are excitatory or inhibitory. We often have two completely unrelated thoughts in quick succession, and surely elements from both will remain coactive in SSC. Because of this, the brain must have neural mechanisms to ensure that ensembles which appear close together in time, and are exposed to one another, do not wire together unless they have some logical connection that was formed within the focus of attention. The brain may use a combination of dopaminergic processes, binding processes, and/or molecular tagging processes, to accomplish this.

Polyassociativity allows the introduction of a hypothetical phenomenon referred to here as a “novel convergence event.” A novel convergence event occurs when ensembles B, C, D and E have each individually and repeatedly contributed to the firing of node F in the past, although they have never fired all together as a group to activate F, until now. This amounts to a pseudo-Hebbian form of convergence that may be prevalent in the brain. For instance, we may have never seen the trivia prompt “The name of a planet, an element, and a Roman god,” but each of the clues may contribute independently to unconscious neural convergence onto the ensemble representing the construct of “mercury.” It is the present author's belief that most thought is driven by heuristic approximations. In other words, the cortex constantly spreads activation energy from novel combinations of active ensembles that have never been coactive before, and attempts to converge upon the statistically most relevant association without certain or exact precedence, resulting in a solution that is subject to recall errors and is not guaranteed to be optimal. Optimality could be approached if a specific group of ensembles (say, C and E) have been thoroughly




Fig. 4. Schematic depicting “polyassociativity” and illustrating the ways in which high-level representations (chunks, ensembles) are maintained, displaced, and newly activated. 1) Representation A has already been deactivated and B, C, D and E are coactivated, mirroring the pattern of activity shown in Fig. 1. When coactivated, these representations pool and spread their activation energy, resulting in the convergence of activity onto a new representation, F. Once F becomes active, it immediately becomes a coactivate, restarting the cycle. 2) B has been deactivated, while C, D, E, and F are coactivated and G is newly activated. 3) D but not C, has been deactivated. Thus, what is deactivated is not necessarily what entered first but what receives the least converging activity. C, E, F, and G coactivate and converge on H.

associated with many others, and if a type of expertise with these concepts has developed due to either extensive operant conditioning, mental modeling, or a combination of both.

There are at least two components that could exhibit natural variation in such a system: 1) the temporary storage capacity, and 2) the temporal duration of storage. Fig. 4 features four representations for the former, and four time periods per representation for the latter. The storage capacity probably remains relatively constant over time, whereas as attention shifts, the storage duration may shrink to zero in order to accommodate a new set of activates. If the number of items that can be held in the temporary storage capacity were increased to 5, and A, B, C, D, and E coactivated simultaneously, this could change the network dynamics such that node F might no longer be the node upon which the system converged. This could also be the case if the capacity were to shrink to 3 items rather than 4. Altering capacity or duration may lead to discrepant associations and thus different advantages and deficits. For example, a decrease in the storage capacity or duration of SSC might reduce working memory span but increase creative and spontaneous thought, facilitate reaction time, and increase attentiveness to the environment. It may additionally result in constraint relaxation and chunk decomposition. On the other hand, an increase in storage capacity and duration may increase working memory span and specificity of memory but necessitate a higher degree of prior network training to produce adaptive behavior.

3.2. Associations between nonsimultaneous but causally related stimuli

Due to their sustained activity, neurons in the PFC can span a wide delay time or input lag between associated occurrences [35,89] and thereby allow elements of prior events to become coactive with elements of subsequent events. Sustained activity allows neurons that would otherwise never fire together to both fire and wire together, and also allows features that never co-occur in the environment to be present together in topographic imagery. Thus, it may be reasonable to assume that SSC underlies the brain's ability to make internally derived associations between representations that never occur simultaneously in the environment. The longer sustained firing in association

cortex lasts, the better the animal will be at capturing information about causally linked stimuli that present apart in time. The higher the SSC, the longer the delay can be. The same regularity may happen persistently in the environment, where a stimulus is followed several seconds later by another stimulus, concern, or opportunity; however, if the animal lacks sufficient SSC, this statistical regularity will not be captured by the neocortical system because the ensembles for them will never be exposed to each other.

Few if any mammals have evolved a human-like capacity for sustained firing in PFC neurons, and thus the mental lives of most mammals likely involve associations made between temporally proximate stimuli and concepts. This may suggest that in most ecological niches it is not helpful to create memories for relationships between stimuli that occur in delayed succession and instead it is better to focus on analyzing stimuli that present in quick succession [68,72]. There may therefore be two strategies, on opposite ends of a continuum, for holding recent information active: immediate and delayed succession strategies. The delayed succession strategy, involving high SSC and low icSSC, is optimal for environmental scenarios that are prolonged over time, where temporally distant cues may retain contextual relevance. This strategy is likely associated with certain ecological or life-history conditions such as low extrinsic mortality, intergenerational resource flows, meme transference, and the K-selection strategy in general.

How can the brain trust that an association between two concepts that are removed in time and never co-occur simultaneously in the environment is valid? Each of the contents of SSC contribute to the selection of the next addition to SSC, and this may help to ensure that the contents held in SSC at any moment are veridically concordant rather than incongruous. This is because the system is narrowly constrained to only combining ensembles that have been highly associated in the past. If this is true, it suggests that at an early age the first associations are between stimuli that are nearly simultaneous, but that these can create foundational knowledge upon which to base reliable inferences about associations between stimuli that are removed from each other by a delay in time.

Because the frontal lobes of infants are underdeveloped, their brains probably exhibit far less continuity between brain states. Very young

children can trust the connections that their early sensory areas have made concerning the spatiotemporal associations between near simultaneous features because these events show high order and regularity. This may be why sensory areas myelinate so early in life. Perhaps association areas are programmed genetically not to finish myelinating until early adulthood because it is a time-intensive process to form and test higher-order hypotheses about relationships between constructs that are more distributed through time.

The processes of iSCC and polyassociativity may work within the brain at a number of neural levels. They probably influence network dynamics at the level of assemblies, ensembles, dedicated modules and even cross-module interactions. The remainder of this article, however, will focus on the role of iSCC and polyassociativity in representational updating at the level of the global workspace [4,5,24]. At this point, it is necessary to consider mental imagery and the hierarchical reciprocations that construct it. The following sections frame the thinking process as a succession of interrelated maps or images guided by a continually updating buffer of higher-order representations that demonstrates iSCC.

4. Cortical hierarchy, topographic mappings, and mental imagery

Perceptual sensory processing is thought to be accomplished using hierarchically arranged, nonlinear transformations [16]. The cortical hierarchy from sensory to association cortex arises because simple patterns converge upon second-order patterns, which in turn converge upon third-order patterns and so on. This amounts to a hierarchy of increasingly complex representations stretching from subsymbolic neurons and assemblies in sensory areas to analogically symbolic ensembles in association areas. Many pathways in the brain, such as the ventral visual pathway, appear to use a “structurally descriptive” architecture with “compositional syntax” in which neurons or neural populations that encode low-level, nonaccidental features are allowed to converge onto those that encode more abstract, higher-order, generic, template-like features [28]. For instance, a neuron in the retinotopically arranged primary visual area may receive inputs from a row of contiguous upstream neurons, and interpret these adjacent “dots” as a structurally inferred line. Downstream areas of extrastriate visual cortex are tuned to recognize when these lines and curves come together to create more complex visual features, eventually amounting to the recognition of structures such as objects (inferior temporal cortex), scenes, faces (fusiform face area), and the like [81]. It seems that evolution tuned cortical assemblies higher in this hierarchy to fire for longer spans of time, predisposing mammals to retain high-order generics rather than precategorical specifics.

Early sensory areas create topographic mappings from patterns recognized in the external environment, but are also thought to be able to combine top-down inputs from higher association cortex to build internally-derived imagery [54,61]. For instance, top-down modulation by the dorsolateral PFC has been observed to be responsible for retrieval of viewed images in visual cortex [42], and retrieval of associative memories from inferior temporal cortex [82]. In his convergence-divergence framework (1989), which may be largely compatible with the present model, Damasio proposed that during recall, association cortices send back divergent projections about the records needed to reconstruct original perceptual maps in early sensory cortices [23]. The predictive feedback from the template-like association ensembles (analogous to Damasio's CDRs) may make incomplete or noisy perceptions in early sensory cortices more complete by retroactivating generic features (analogous to Damasio's CDZs) normally associated with the stimulus and completing the expected pattern. A major difference between this formulation and Damasio's is that not one, but multiple ensembles (CDRs), exhibiting iSCC, spread divergent activity in a polyassociative manner back toward sensory cortex to guide the construction of topographic maps.

The PFC and other associative areas do not direct processing in V1 or the lateral geniculate nucleus (LGN) of the thalamus (the earliest of visual processing areas), but can profoundly influence the activity in extrastriate visual areas via re-entrant projections, which in turn influence V1 [56]. Internally-derived sensory imagery, such as that seen in the “mind's eye” probably appears topographically organized because it is created by the same lower-order networks responsible for perceiving external stimuli. Thus it may be safe to assume that when we visualize and imagine, we construct and manipulate maps in early perceptual networks [22]. The available population of assemblies in sensory cortex may serve as an active canvas for either the environment (via feedforward connections) or expectation and imagination (via feedback connections) to paint on. During perception, the bottom-up activity may be driving and the top-down may be modulatory; however, during imagination the top-down activity may be driving and the bottom-up may be modulatory.

The manner in which sensory areas integrate inputs when they construct images is informed by reality, as they have been tuned directly by real environmental inputs, unlike associative areas, which are tuned indirectly by reality owing to the intervening effects of motivation, temporal delay, and inference. The sensory areas do what association areas cannot do on their own. They take various features and rapidly and unconsciously integrate them into a map that conforms to strict, environmentally imposed constraints. These assumptions are consistent with the “consolidation hypothesis,” which states that memory is stored in the same areas that allow active, real-time perception and function [60]. They are also consistent with the finding that remembering or imagining a particular sensory construct largely activates the same neural networks that are involved in perceiving the construct in the environment [21].

Unlike some popular models, in the present framework, activity in association areas does constitute imagery. Like sensory areas, their neurons have become correlated with the occurrence of spatiotemporal structure from the environment. Prefrontal areas such as the dorsolateral PFC likely contain ensembles that are reflections of experiences, albeit abstracted postcategorical ones that are disoriented from modality-specific topographic coordinates. Thus, purporting that association areas do not hold true imagery, as some neuroscientists do, is akin to saying that imagery is held in the “dots” of primary visual cortex but not in the “contours” of secondary visual cortex. The topology of PFC imagery is not unimodal (e.g., retinotopic) but rather must correspond to the interspersed placement of lower-order projections from various modalities, arranged during the evolution of the frontal lobe. This suggests that the spatial format of associative imagery can only be understood through neurocartographic investigation of the unique connective geometries found in higher-order areas. Neurons near the top of the hierarchy in areas such as the frontal pole, which seem to code for very densely multimodal abstractions, appear to be calibrated to sustain their firing for longer than other more posterior neurons. Describing the exact relationship between mental states and brain states will ultimately involve “functional modeling,” i.e., scrutinizing the receptive fields of these anterior neurons, better determining which facets of our mental lives span several seconds of time, and then matching these together.

The discrepant firing properties of sensory and association cortex make their interaction dynamic. An important attribute of this cooperative system is that imagery in lower-order sensory areas can be quickly wiped clean to accommodate a new mapping, whereas activity in association areas is more enduring. Early sensory cortex is thought to have its own very-short-term memory called “sensory memory.” Sensory memory has been shown to hold more information than does working memory, but it does so very transiently [2,3,79]. Crick and Koch [21] advocated that a helpful way to consider the reciprocal activity between sensory and association areas is to imagine that the front of the brain is “looking at” the sensory systems in the back of the brain. This would be similar to watching a television program that can depict ideas and conceptualizations. This retinotopic “screen” may be able to perform

important functions such as keeping track of, or binding, particular objects with their respective traits [84]. However, the TV analogy does not greatly reduce the problem at hand because it depicts association cortex as another kind of conscious entity that “watches” the TV. The next section will attempt to show how icSSC allows the interaction between association and sensory cortex to amount to more than the relationship between a homunculus and a “Cartesian Theater.”

4.1. Top-down to bottom-up transformations, and progressive imagery modification

Object recognition, associative recall, planning, decision making and other important cognitive processes involve two-way signal activity among various neural maps that stretch transversely through the cortex from early sensory areas to late association areas [50]. Bottom-up sensory areas deliver fleeting sensory information, whereas top-down association areas deliver lasting perceptual expectations in the form of templates or prototypes. During recognition and recall, these two systems may determine category belongingness by constantly comparing their respective outputs [12]. These exchanges involve feedforward and feedback (recurrent) connections in the corticocortical and thalamocortical systems that bind topographic information from lower-order sensory maps with dispositional information from higher-order maps. These bindings form synchronous constellations of activity, often presenting in the gamma range, which can remain stable for tens or hundreds of milliseconds [21].

icSSC influences this reciprocating crosstalk. These reciprocations may create progressive sequences of related thoughts, specifically because the topographic maps generated by lower-order sensory areas are guided by the enduring representations that are held active in association areas [69–71]. Because they are drawing from a register with sustained contents, sequential images formed in sensory areas have similar subject matter and thus can be expected to be symbolically or semiotically related to one another. It is probable that the higher the SSC, the more related two sequential sensory images can be on average.

Also, association areas maintain representations from not only the most recent topographic mapping but also from several recent topographic maps produced in sensory areas (Fig. 5).

SSC and recurrent processing make it possible for current states to spill over into subsequent states, creating the context for them in a recursive fashion. This creates a cyclical, nested flow of information processing marked by icSSC (Fig. 6). In a sense, each new topographic map is embedded in the previous one. This form of hierarchical crosstalk could be termed “progressive imagery modification.” The same general process may take place in language areas and be responsible for the internal monologue. This process could similarly drive the progression of complex behavior in that the motor-relevant representations that pass in and out of SSC impose iterative constraints on motor commands, and thereby regulate the animal’s preparatory set, action sequencing, and planning behavior.

Feedback activation from top-down association areas passes specifications to early sensory cortex for use in imagery building. Disparate chunks of information are integrated into a plausible map and transiently bound together. This suggests that one can immediately visualize the relationship between two abstract representations only if one already has implicit information in the visual cortex about how to integrate them into a single image. This integrative process may be very rapid and use the structurally descriptive perceptual hierarchy in reverse to go from abstractions to specifics. In the opposite direction, the topographic bindings of the map are “disintegrated” as feedforward activation from bottom-up sensory areas converges upon relevant assemblies in association cortex. After this “fast feedforward sweep,” the PFC and association areas could be protagonized as “saying” the following: “We have identified the important features from the last bottom-up image and combined them with the other features that we’ve been holding online from previous images. Let’s engage in another round of imagery generation, this time with more emphasis on the salient elements from the last image. It should be informative to see how the visual system combines this updated set of higher-order features into another, composite, topographical map. Perhaps this next

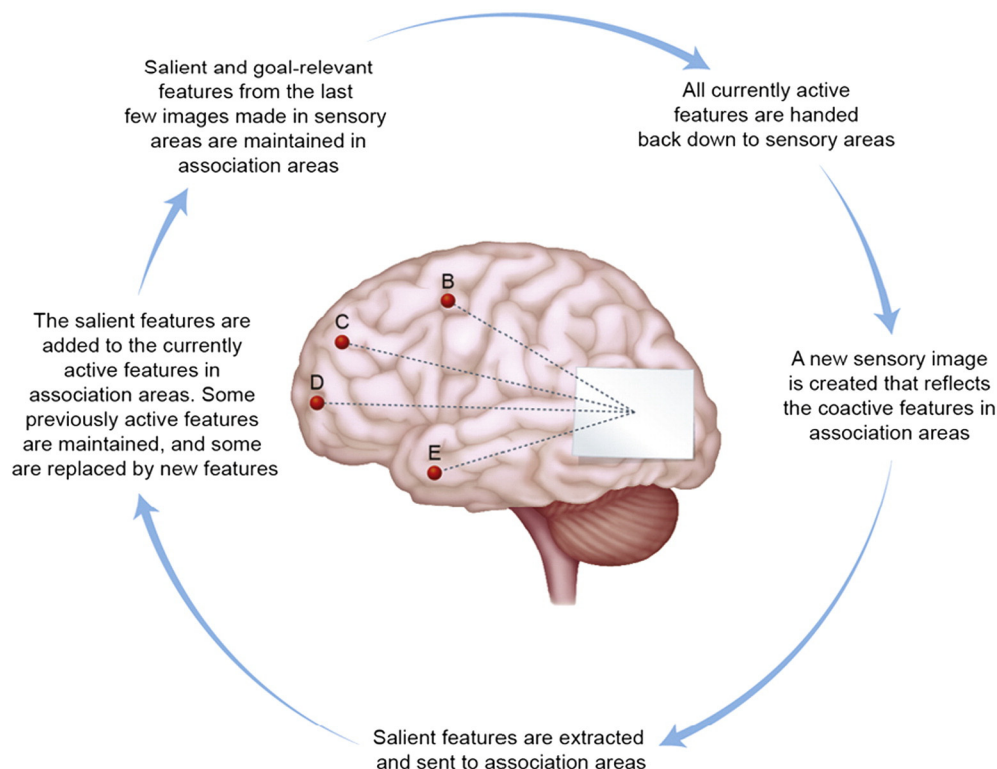


Fig. 5. Depiction of the reciprocal transformations of information between lower-order sensory maps and higher-order association area ensembles during internally generated thought. Sensory areas can create only one topographic map at a time, whereas association areas hold the salient or goal-relevant features of several sequential maps at the same time.




Fig. 6. Depiction of icSSC enabling progressive imagery modification. 1) Representations B, C, D, and E, which are held active in association areas, all spread their activation energy to early visual cortex, where a composite image is built that is based on prior experience with these representations. 2) Features involved in the topographic imagery from time sequence 1 converge on the PFC neurons responsible for F. B drops out of activation, and C, D, E and F remain active and diverge back onto visual cortex. 3) The same process leads to G being activated and D being deactivated, mirroring the pattern of activity shown in Fig. 4.

map will introduce informative representational material for subsequent cycles.”

The rapid, rhythmic oscillations of information between sensory and association areas allow them to learn from each other and integrate their knowledge like two people in conversation. The fact that they have fundamentally different perspectives makes the “conversation” dynamic and informative for both of them due to the lack of redundancy. The crosstalk is similar to that between two specialists in related areas, interrogating each other about the nature of their common interests. One of these specialists speaks figuratively and keeps a list of previous topics on hand; the other is a literalist and approaches each turn in the conversation with an unbiased, almost amnesic approach.

Sensory areas may use the specifications handed down from association areas to generate appropriate imagery, but they probably often elaborate on what they have been given with closely associated but unforeseeable embellishments. This would be expected given that the creation of these maps involves polyassociativity and constitutes “novel convergence events.” The unspecified, extemporized features built into early imagery probably provide much new content for the stream of thought. For example, if higher-order nodes come to hold features supporting the representations for “pink,” “rabbit,” and “drum,” then the subsequent mappings in lower-order visual nodes may activate representations for a well-known battery advertisement, and the auditory nodes may activate the representation for the word “Energizer Bunny.” Thus, the concept of a battery might be introduced into the thought process in an unconscious, polyassociative manner, despite the fact that none of the previous concepts alone had any close associations with batteries. The mental imagery that is generated may

constitute only a slight modification to the previous imagery or may be a paradigm shift away from it; however, unless the focus of attention shifts dramatically, it is likely that many of the high-order representations will remain in SSC.

In reality, association areas have much more to converse with than simply a single retinotopic map as depicted in Fig. 6. In fact, they feed their specifications to and receive specialized input from dozens of known topographic mapping areas [48]. Relative to Baddeley’s model of working memory [7,8], this relationship is congruent with that of the central executive and its “slave systems” such as the visuospatial sketchpad and the phonological (articulatory) loop. Areas of various modalities are constantly responding to incoming activity in an attempt to pull up the most context-appropriate map in their repertoire. Further compounding the complexity, the sensory modules that build these maps take specifications not only from association areas but also from other sensory modules [50]. Moreover, motor and premotor modules give specifications to and receive specifications from this common workspace while building their musculotopic imagery for movement. Thus, association areas may utilize icSSC to direct progressive sequences of mental imagery in a large number of topographic sensory and motor modules (Fig. 7). Distinct regions of association cortex can be thought of as specialized modules that regulate this workspace in distinct fashions. This frames the brain as a system of numerous interacting modules that are each generating a topological mapping that is their most internally consistent, composite interpretation of the other module’s mappings. Because the postcategorical and multimodal modules have neurons that fire for sustained periods, they are better positioned to direct this activity through time.

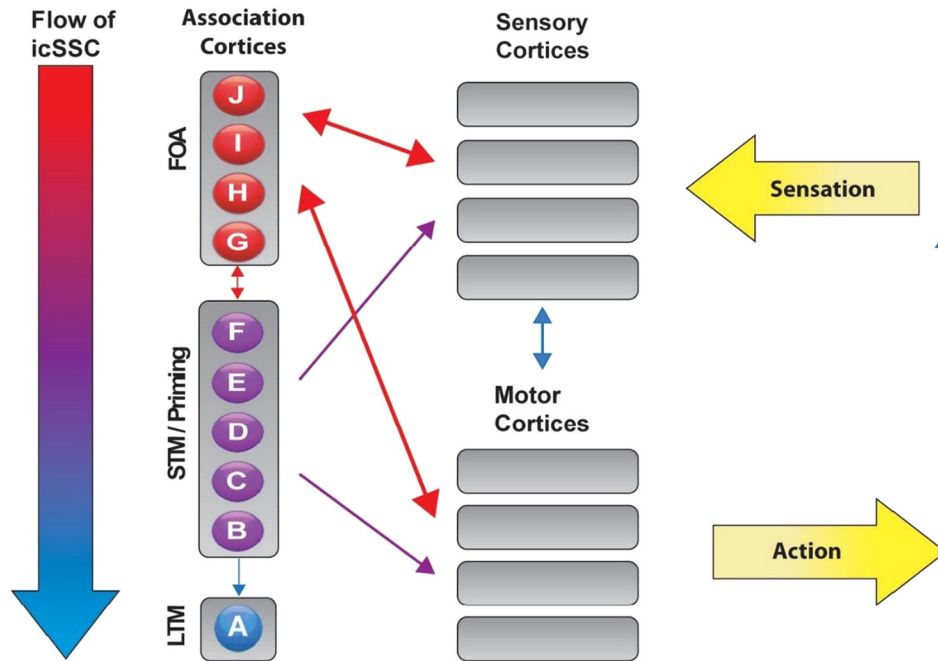


Fig. 7. Incoming information from motor and sensory cortices enters the focus of attention where it can then influence other sensory and motor cortices. The arrow on the far left indicates that as time passes, representations in the focus of attention decay toward a less active form of short-term memory from which point they can 1) reenter the focus of attention, 2) spread their instructions to sensory and motor cortex, or 3) deactivate and return to inert long-term memory.

5. Capacity limits, chunks, and an octopus analogy

A very simple and inflexible caricature of icSSC was suggested by the tank tread analogy and Fig. 2. The dynamics are better captured by an analogy involving an octopus grabbing and releasing footholds as it pulls itself from place to place. This is an activity known as “sea floor walking.” The analogy illustrates that the thought process involves the simultaneous coactivation of several representations at a time (multiple footholds held by an octopus) as well as the deactivation of previously active representations (the releasing of footholds), and the activation of previously inactive representations (the placement of an arm on a new foothold). This analogy may be valuable because it depicts a system, that even a child can understand, where specific representations are conserved through time as others are actively repositioned. This shares attributes with Edelman and Tononi’s [29] conceptualization of a “functional cluster” or “dynamic core,” where reentrant neuronal interactions yield a succession of differentiated yet unitary metastable states, and with Varela’s conception of dynamic links mediated by synchrony over multiple frequency bands [85].

In the present analogy, each octopus arm corresponds to an active ensemble, and the suction cups on each arm can be taken to correspond with the assemblies that comprise the ensemble, while the grains of sand under each suction cup represent active cortical neurons. Fig. 8 offers a hypothetical example of icSSC that, though contrived, is more realistic than the emblematic example of Fig. 2.

The “octopus” has a limited and fixed number of arms, so it must release footholds to free up tentacles to grab new footholds. Similarly, our working memory has a limited capacity and is constantly forced to reallocate its bioelectrical and hemodynamic resources [18]. Coincidentally, the number of representational chunks (psychologically perceptible units of perception and meaning) that can be held in working memory—7 plus or minus 2 by some estimates [58]—neatly coincides with the number of tentacles on an octopus.

I do not mean to suggest that the sustained firing involved in the focus of attention and the working memory system is fully responsible for mental continuity. Cortical priming, early LTP, the slow action of neuromodulators, emotional priming, and the hippocampus must each make unique contributions to continuity as well. Take priming

for instance. Cortical nodes recently used in the focus of attention maintain above-baseline activity, and contribute to non-hippocampal dependent short-term memory [13,88]. These primed nodes also must contribute to imagery and polyassociativity because of the way they continue to spread their activation energy within the cortical network.

Episodic pattern completion made possible by the hippocampus [53] must make similar, independent contributions to the mental continuity that we experience. The conceptual recurrences that the hippocampus makes possible may lend a continuity to consciousness in the same way that a motif (a distinctive and recurring form) gives coherency to a literary, artistic, or musical work. Sustained firing and short-term synaptic plasticity may allow a form of “short-term continuity” on the order of milliseconds to seconds, whereas priming, early LTP, and emotions may allow an “intermediate-term continuity” on the order of seconds to hours, while late LTP and the episodic memory of the hippocampus may allow “long-term continuity” on the order of minutes to years.

6. Unconscious processes demonstrating icSSC result in executive function

In cognitive psychology, dual process theory suggests that implicit and explicit processes use two fundamentally different cognitive processing architectures: system 1 and system 2 [63]. System 1 is implicit, automatic, and unconscious. It operates quickly, effortlessly and has a large capacity [74]. System 2 is explicit, controlled, and conscious. It operates with effort and has a smaller capacity [30]. These two systems are thought to be subserved by distinct neural substrates [49], although they may share the same general neuroanatomical foundation. It is possible that system 2 is actually system 1 exhibiting icSSC [70]. That is, what is referred to as system 2 may simply be the processing architecture of system 1 (top-down to bottom-up reciprocations within the thalamocortical hierarchy) implementing the same suite of representations for a prolonged period. Again, system 1 is making automatic, intuitive, flash judgments; however, because of SSC, these rapid associations are able to support each other in a progressive and additive manner.

Rather than arriving at a judgment about something based on a single impulse, system 2 may involve the sequential arrangement of several impulses that cooperatively, lead to a higher-order judgment. For




Fig. 8. A hypothetical distribution of ensembles held in the focus of attention of SSC. Unlike Fig. 2, the order of entry does not rigidly determine the order of exit. Here, there are 7 representational constructs (rather than 4) that can be maintained. Representations C and F demonstrate continuity over all 5 time intervals. Representation G enters during time 1, is deactivated during time 2, then reenters during time 5. Representations that remained active until time 5 are shown in red. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

example, multiplying two 2-digit numbers may be a mental calculation that involves a series of system 1 impulses coordinated to implement a learned pattern of algorithmic steps. All learned mental operations probably have distinct, predefined algorithmic sequences of topological mappings that must be completed in sequence to achieve the solution. Each brain state would correspond to a different step in the algorithm, and because each was trained to do so, its activity would recruit the next step.

System 2 cognition may be involved in a mental task whenever ensembles related to the task exhibit sustained firing and are utilized in progressive imagery modification, culminating in sensory imagery or motor output that could not be informed by any of the intermediate steps alone, or that is capable of solving a problem too difficult for any system 1 process itself. For example, early processes may provide premises or propositional stances that can be used algorithmically (e.g., syllogistically) to induce or justify a conclusion in subsequent processes. Perhaps reciprocating activity between the SSC updating system and the imagery generation systems produces sequences of interrelated mental images that build upon and provide scaffolding for one another to form logical deductions and abductive expectations.

This preparatory modeling process may employ learned rules and schemas to actively test relationships between higher-order representations. Associations that are perceived as, or deemed to be, valid may then be saved to memory. This has the effect of training the initial impulses in advance, influencing future behavior without necessitating that the scenario be modeled again. Researchers in psychology suggest that system 2 processes are initiated by motivation, surprise, and curiosity [30]. Of course these are the same three factors involved in recruiting the mesocortical dopamine system necessary for the sustained firing responsible for mental continuity. So, is there a fundamental equivalency between conscious processes and unconscious ones? Perhaps consciousness and metacognitive processing can be reduced to preattentive processing in which certain ensembles exhibit SSC over successive instances of imagery generation. This may afford the executive system the capacity to handle novel situations outside the domain of automatic processes, prepotent responses, and set behaviors.

7. Implications for artificial intelligence

The present model may have implications for structuring artificial intelligence (AI) architectures. There are existing computing architectures

with limited forms of recursion in which the current state is a function of the previous state and active data are entered into a limited capacity buffer to inform other processes. However, there are no AI systems in which this buffer is multimodal and positioned at the top of a hierarchical system so that it can inform and interact with topographic maps in a recursive and polyassociative manner. The model described here could be used to inform the construction of a multimodal, neural network architecture that harnesses icSSC for use in pattern recognition, analytics, prediction, adaptive control, decision making, and response to query.

The program would be composed of several highly interconnected multilayer neural networks of pattern-recognizing nodes organized into a hierarchical architecture similar to that of the mammalian neocortex. The multiple interfacing neural networks would be arranged biomimetically, i.e., connected and tuned to simulate sensory and association cortices (Fig. 9). Like neural assemblies, the nodes would exhibit a continuous gradient from low-order nodes that code for sensory features (and serve in feature extraction) to high-order nodes that are designed to capture and encode temporally or spatially extended relationships between such features. In each sensory network, nodes would operate both “competitively” and “cooperatively” to build “self-organizing” topographic maps. Nodes higher in the hierarchy would be increasingly multimodal, and have larger capacity for sustained activity. These nodes would identify and sustain pertinent, high-level features over elapsing time. Prioritized features could be used to guide the construction of topographic maps as well as the construction of natural language and robotic movement. Properly implemented, the architecture should feature the aspects of both Kohonen and Hopfield networks, backpropagation/bidirectionality, and Hebbian learning as well as a combination of principal-components learning and competitive learning. Replicating icSSC using deep learning algorithms could make it possible for large multimodal neural networks to recognize invariant patterns whose components are distributed across time. This capacity may then allow them to find high-order structure in unlabeled data, integrate abstract knowledge, and perform logical inferences in uncertain environments.

If the sustained activity of higher-order nodes was programmed to persist for extended intervals, the system could exhibit a superhuman capacity for continuity (Fig. 10). This could amplify the ability of the network to make associations between causally linked but temporally distant representations, and allow its processing and actions to be

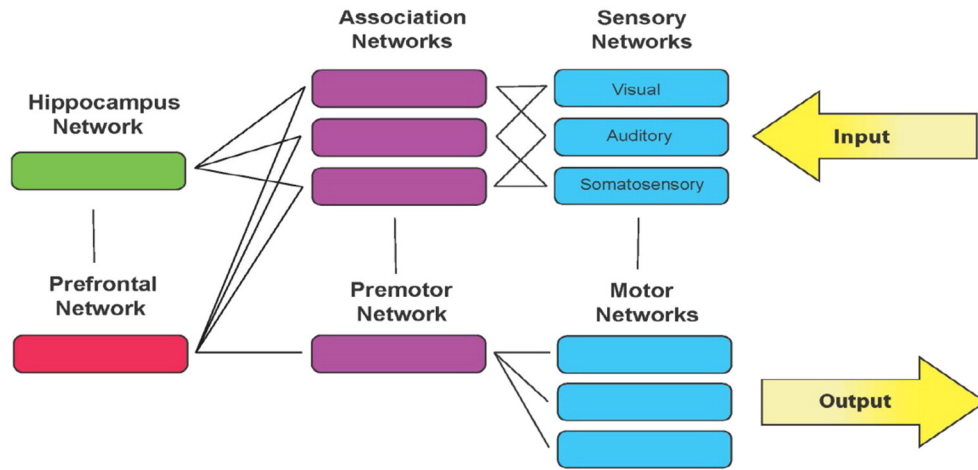


Fig. 9. A plausible biomimetic arrangement of interfacing neural networks. Black lines indicate massive connectivity.

informed by more temporally distant concerns. It could also allow the imagery that is created to be informed by a larger number of parameters, and ensure that important features were not omitted simply because their activity could not be sustained due to biological

limitations. To operate meaningfully, and reduce its propensity for recognizing “false patterns,” such an ultraintelligent system would require extensive supervised training and unsupervised machine learning.

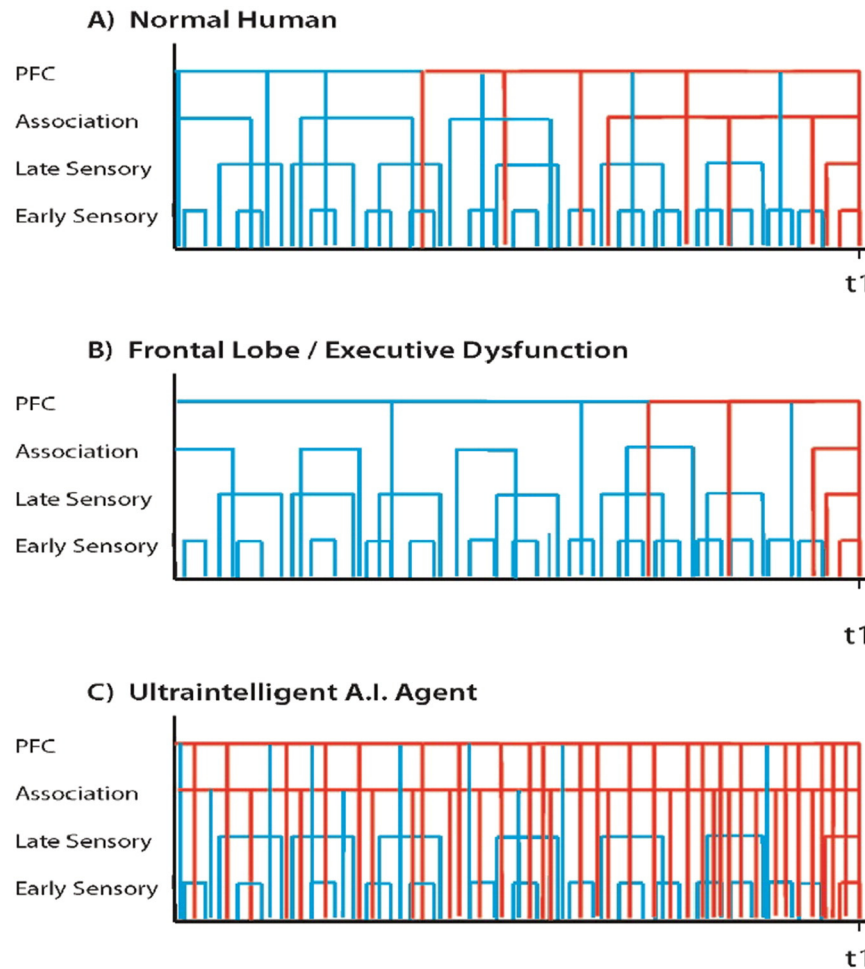


Fig. 10. Three examples of sustained firing in a neural network. The figure compares the number of nodes that have remained active until the present time (t_1), in a normal human, in a human with frontal lobe dysfunction, and in a hypothetical AI agent. The AI agent can maintain a larger number of nodes over a longer time span, ensuring that its perceptions and actions in t_1 will be informed by a larger amount of recent information.

8. Conclusions

The model presented here is qualitative and exploratory, contains unverified assumptions, and makes untested claims. A more complete discussion could focus on better integration of existing knowledge from clinical neuropsychology, the sustained firing behavior of cortical columns, the functional neuroanatomy of working memory, and the cognitive neuroscience of attention. In the absence of pertinent data and related literature this article attempts to introduce new perspectives using plausible, convergent reasoning. The model is intended to inspire more detailed hypotheses that can be tested experimentally, and against already-existing experimental data. Further research should search for the distinct neural signature of icSSC (see Figs. 1, 2 and 8) within the brain. This search for the icSSC signal should involve extensive time series analyses of simultaneous recordings, on submillisecond time scales, from small groups of neurons, in many places in the brain. It is not clear how the neurobiological substrate of mental continuity should be best examined, but it is reassuring that there are diverse disciplines which use advanced methodologies to track, analyze and predict the evolution of gradual changes in complex, dynamic systems, including population genetics, geology, material science, and meteorology.

Crick and Koch [21] have advocated that neuroscientists should concentrate on finding the “neural correlates of consciousness,” defined as the minimal set of brain mechanisms and events sufficient for some specific phenomenal state. icSSC as described here may be a strong candidate for a neural correlate of consciousness. The reasoning here suggests a consideration of consciousness as an elaborative process, where the longer a set of representations is maintained within SSC, the more conscious one becomes of that subset. icSSC may be necessary to form higher-order, or transitive, representations (i.e., a representation of a representation), and thus may be necessary for active self-awareness. In other words, self-consciousness may necessitate that common notions about the autobiographical self remain in SSC, where they help to direct the progression of images and impulses. Qualia (i.e., instances of subjective conscious experience) may gain their experiential qualities when the features of a quale are maintained in SSC, volleyed between association and sensory areas, and used to uncover apposite memories and images. This processes could be conceptualized as “phenomenal continuity.”

The present line of thinking may cause us to ask: 1) whether icSSC could be used to quantify the extent of consciousness in animals or in coma/vegetative state, 2) whether icSSC is necessary for human-level intelligence in machines, and 3) what form of phenomenal experience or sentience an entity could possess if it exhibited no icSSC. Even though mental continuity and icSSC may resemble consciousness in some respects, they are certainly not identical to it, and consciousness must involve a number of additional phenomena and conditions.

Suppose that ... we take half a dozen glances at a strange and curious flower. We have not as many complex presentations which we might symbolise as F1, F2, F3. But rather, at first, only the general outline is noted, next the disposition of petals, stamen, etc., then the attachment of the anthers, form of the ovary, and so on ... It is because the earlier apprehensions persist that the later are an advance upon them and an addition to them.

John Locke [52]

The fact that “earlier apprehensions” persist makes them immediately available to contribute to the formation of later, more complex apprehensions. Newly attended stimuli create gradual, progressive modifications to the elements held in SSC as the series advances, and interrelated perceptions accumulate. A new impression in this series may be irrelevant, or instrumental, or anything in between. Sometimes the final construct in such a compiled series produces the effect necessary

to reframe all of the preceding stimuli, creating an end product that was unforeseeable before its entry.

Swiftly the head mass becomes an enchanted loom where millions of flashing shuttles weave a dissolving pattern, always a meaningful pattern though never an abiding one; a shifting harmony of subpatterns.

Charles Sherrington [78]

I like to think that Sherrington's dissolving pattern may directly correspond to the gradually transforming distribution of temporarily sustained representations in icSSC. This purely biological process that simultaneously exhibits both deterioration and accretion underlies our ability to transition between closely related thoughts and may constitute the physical “fabric” of consciousness. Whether a series of actions, words, or musical notes, each is delivered separately in its own discrete span of time, but our brain's capacity for SSC allows us to knit these temporally isolated stimuli together into performances, speeches, and melodies. Merging such a group of occurrences as if they occurred simultaneously permits the use of the entire series at once to search for related memories. This attribute, permits mammals to weave together threads of isolated, objective occurrences into a continuous tapestry of subjective experience.

References

- [2] R.C. Atkinson, R.M. Shiffrin, Chapter: human memory: a proposed system and its control processes, in: K.W. Spence, J.T. Spence (Eds.), *The Psychology of Learning and Motivation* (Volume 2), Academic Press, New York 1968, pp. 89–195.
- [3] R.C. Atkinson, R.M. Shiffrin, Storage and retrieval processes in long-term memory, *Psychol. Rev.* 76 (2) (1969) 179–193.
- [4] B.J. Baars, *In the Theater of Consciousness*, Oxford University Press, New York, NY, 1997.
- [5] B.J. Baars, The conscious access hypothesis: origins and recent evidence, *Trends Cogn. Sci.* 6 (1) (2002) 47–52.
- [7] A.D. Baddeley, The episodic buffer: a new component of working memory? *Trends Cogn. Sci.* 4 (2000) 417–423.
- [8] A.D. Baddeley, *Working Memory, Thought and Action*, Oxford University Press, Oxford, 2007.
- [9] A.D. Baddeley, G.J. Hitch, Developments in the concept of working memory, *Neuropsychology* 8 (4) (1994) 485–493.
- [11] T.S. Braver, J.D. Cohen, On the control of control: the role of dopamine in regulating prefrontal function and working memory, in: S. Monsell, J. Driver (Eds.), *Attention and Performance XVIII: Control of Cognitive Processes 2000*, pp. 713–737.
- [12] G.A. Carpenter, S. Grossberg, Adaptive resonance theory, in: M.A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks*, Second Edition MIT Press, Cambridge, MA 2003, pp. 87–90.
- [13] L.S. Cermak, N. Talbot, K. Chandler, L.R. Wolbarst, The perceptual priming phenomenon in amnesia, *Neuropsychologia* 23 (5) (1985) 615–622.
- [16] G. Cohen, Hierarchical models in cognition: do they have psychological reality? *Eur. J. Cogn. Psychol.* 12 (1) (2000) 1–36.
- [17] A.M. Collins, E.F. Loftus, A spreading-activation theory of semantic processing, *Psychol. Rev.* 82 (6) (1975) 407–428.
- [18] A.R.A. Conway, Individual differences in working memory capacity: more evidence for a general capacity theory, *Memory* 4 (6) (1996) 577–590.
- [20] N. Cowan, *Working Memory Capacity*, Psychology Press, New York, NY, 2005.
- [21] F. Crick, C. Koch, A framework for consciousness, *Nat. Neurosci.* 6 (2) (2003) 119–126.
- [22] A.R. Damasio, Time-locked multiregional retroactivation: a systems level proposal for the neural substrates of recall and recognition, *Cognition* 33 (1989) 25–62.
- [23] A.R. Damasio, *Self Comes to Mind: Constructing the Conscious Brain*, Pantheon Books, New York NY, 2010.
- [24] S. Dehaene, L. Naccache, Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework, *Cognition* 79 (1–2) (2001) 1–37.
- [25] D. Durstewitz, J.K. Seamans, The computational role of dopamine D1 receptors in working memory, *Neural Netw.* 15 (2002) 561–572.
- [28] S. Edelman, Computational theories of object recognition, *Trends Cogn. Sci.* 1 (8) (1997) 296–304.
- [29] F. Edelman, G. Tononi, *A Universe of Consciousness: How Matter Becomes Imagination*, Basic Books, 2000.
- [30] J. Evans, In two minds: dual-process accounts of reasoning, *Trends Cogn. Sci.* 7 (10) (2003).
- [31] H. Fuji, H. Ito, K. Aihara, N. Ichinose, M. Tsukada, Dynamical cell assembly hypothesis – theoretical possibility of spatio-temporal coding in the cortex, *Neural Netw.* 9 (8) (1998) 1303–1350.
- [32] J.M. Fuster, G.E. Alexander, Neuron activity related to short-term memory, *Science* 173 (3997) (1971) 652–654.
- [33] J.M. Fuster, Frontal lobe and cognitive development, *J. Neurocytol.* 31 (3–5) (2002) 373–385.

- [34] J.M. Fuster, Physiology of executive functions: the perception-action cycle, in: D.T. Stuss, R.T. Knight (Eds.), *Principles of Frontal Lobe Function*, Oxford University Press, Oxford 2002, pp. 96–108.
- [35] J.M. Fuster, Cortex and memory: emergence of a new paradigm, *J. Cogn. Neurosci.* 21 (11) (2009) 2047–2072.
- [37] P.S. Goldman-Rakic, Circuitry of the prefrontal cortex and the regulation of behavior by representational memory, in: V.B. Mountcastle, F. Plum, S.R. Geiger (Eds.), *Handbook of Neurobiology*, American Physiological Society, Bethesda 1987, pp. 373–417.
- [38] P.S. Goldman-Rakic, Cellular and circuit basis of working memory in prefrontal cortex of nonhuman primates, in: H.B.M. Uylings, C.G.V. Eden, J.P.C. DeBruin, M.A. Corner, M.G.P. Feenstra (Eds.), *Progress in Brain Research*, 85, Elsevier Science Publications 1990, pp. 325–336.
- [39] P.S. Goldman-Rakic, Cellular basis of working memory, *Neuron* 14 (1995) 477–485.
- [41] K.N. Gurney, Reverse engineering the vertebrate brain: methodological principles for a biologically grounded programme of cognitive modeling, *Cogn. Comput.* 1 (1) (2009) 29–41.
- [42] I. Hasegawa, T. Fukushima, T. Ihara, Y. Miyashita, Callosal window between prefrontal cortices: cognitive interaction to retrieve long-term memory, *Science* 281 (1998) 814–818.
- [43] D. Hebb, *The Organization of Behavior*, Wiley, New York, 1949.
- [44] J.C. Horton, D.L. Adams, The cortical column: a structure without a function, *Philos. Trans. R. Soc. Lond. Ser. B Biol. Sci.* 360 (1456) (2005) 837–862.
- [46] W. James, *The Writings of William James: A Comprehensive Edition*, University of Chicago Press, Chicago, 1978.
- [47] P.N. Johnson-Laird, *Computer and the Mind: An Introduction to Cognitive Science*, Harvard University Press, 1998.
- [48] J.H. Kaas, Topographic maps are fundamental to sensory processing, *Brain Res. Bull.* 44 (2) (1997) 107–112.
- [49] D. Kahneman, *Thinking, fast and Slow*, 1st ed. Farrar, Straus and Giroux, New York, 2011.
- [50] W. Klimesch, R. Freunberger, P. Sauseng, Oscillatory mechanisms of process binding in memory, *Neurosci. Biobehav. Rev.* 34 (7) (2010) 1002–1014.
- [51] A. Lansner, Associative memory models: from the cell-assembly theory to biophysically detailed cortex simulations, *Trends Neurosci.* 32 (3) (2009) 179–186.
- [52] J. Locke, *Essay Concerning Human Understanding*, 1689 ii. x., 5.
- [53] B.L. McNaughton, Associative pattern completion in hippocampal circuits: new evidence and new questions, *Brain Res. Rev.* 16 (1991) 193–220.
- [54] E. Mellet, L. Petit, B. Mazoyer, M. Denis, N. Tzourio, Reopening the mental imagery debate: lessons from functional anatomy, *Neuroimage* 8 (2) (1998) 129–139.
- [55] K. Meyer, A. Damasio, Convergence and divergence in a neural architecture for recognition and memory, *Trends Neurosci.* 32 (7) (2009) 376–382.
- [56] K. Meyer, Primary sensory cortices, top-down projections and conscious experience, *Prog. Neurobiol.* 94 (4) (2011) 408–417.
- [57] E.K. Miller, J.D. Cohen, An integrative theory of prefrontal cortex function, *Annu. Rev. Neurosci.* 24 (2001) 167–202.
- [58] G. Miller, The magical number seven, plus or minus two: some limits on our capacity for processing information, *Psychol. Rev.* 63 (1956) 81–97.
- [59] M. Moscovitch, Memory and working-with-memory: a component process model based on modules and central systems, *J. Cogn. Neurosci.* 4 (3) (1992) 257–267.
- [60] M. Moscovitch, J.M. Chein, D. Talmi, M. Cohn, Learning and memory, in: B.J. Baars, N.M. Gage (Eds.), *Cognition, Brain, and Consciousness: Introduction to Cognitive Neuroscience*, Academic Press, London, UK 2007, p. 234.
- [61] Y. Miyashita, Cognitive memory: cellular and network machineries and their top-down control, *Science* 306 (2005) 435–440.
- [63] A. Paivio, *Mind and its Evolution: A Dual Coding Theoretical Approach*, Lawrence Erlbaum Associates, Mahwah, NJ, 2007.
- [65] R. Perin, T.K. Berger, H. Markram, A synaptic organizing principle for cortical neural groups, *PNAS* 108 (13) (2010) 5419–5424.
- [66] B.R. Postle, Activated long-term memory? The bases of representation in working memory, in: N. Osaka, R.H. Logie, M. D'Esposito (Eds.), *The Cognitive Neuroscience of Working Memory*, Oxford University Press, Oxford, UK, 2007.
- [68] J. Reser, Schizophrenia and phenotypic plasticity: schizophrenia may represent a predictive, adaptive response to severe environmental adversity that allows both bioenergetic thrift and a defensive behavioral strategy, *Med. Hypotheses* 69 (2) (2007) 383–394.
- [69] J.E. Reser, *What Determines Belief: The Philosophy, Psychology and Neuroscience of Belief Formation and Change*, Verlag Dr. Muller, Saarbrücken, Germany, 2011.
- [70] J.E. Reser, *Assessing the psychological correlates of belief strength: Contributing factors and role in behavior*, Doctoral Dissertation, Retrieved from University of Southern California, 2012 Usctheses-m2627.
- [71] J.E. Reser, *The Neurological Process Responsible for Mental Continuity: Reciprocating Transformations Between a Working Memory Updating Function and an Imagery Generation System*, Association for the Scientific Study of Consciousness Conference, San Diego CA, 12–5th July 2013.
- [72] J.E. Reser, Chronic stress, cortical plasticity and neuroecology, *Behav. Process.* 129 (2016) 105–115.
- [73] M. Sarter, B. Givens, J.P. Bruno, The cognitive neuroscience of sustained attention: where top-down meets bottom-up, *Brain Res. Rev.* 35 (2) (2001) 146–160.
- [74] W. Schneider, R.M. Shiffrin, Controlled and automatic human information processing: 1. Detection, search, and attention, *Psychol. Rev.* 84 (1977) 1–66.
- [75] J.K. Seamans, T.W. Robbins, Dopamine modulation of the prefrontal cortex and cognitive function, *The Dopamine Receptors* 2010, pp. 373–398.
- [76] J.K. Seamans, C.R. Yang, The principal features and mechanisms of dopamine modulation in the prefrontal cortex, *Prog. Neurobiol.* 74 (1) (2004) 1–58.
- [78] C.S. Sherrington, *Man on his Nature*, Cambridge University Press, 1942 178.
- [79] G. Sperling, The information available in brief visual representations, *Psychol. Monogr.* 74 (1960) (1960) 1–29.
- [80] M.G. Stokes, Activity-silent working memory in prefrontal cortex: a dynamic coding framework, *Trends Cogn. Sci.* 19 (7) (2015) 394–405.
- [81] K. Tanaka, Mechanisms of visual object recognition: monkey and human studies, *Curr. Opin. Neurobiol.* 7 (4) (1997) 523–529.
- [82] H. Tomita, M. Ohbayashi, K. Nakahara, I. Hasegawa, Y. Miyashita, Top-down signal from prefrontal cortex in executive control of memory retrieval, *Nature* 401 (1999) 699–703.
- [83] G. Tononi, An information integration theory of consciousness, *BMC Neurosci.* 5 (2004) 42.
- [84] A. Treisman, Solutions to the binding problem: progress through controversy and convergence, *Neuron* 24 (1999) 105–110.
- [85] F. Varela, J.P. Lachaux, E. Rodriguez, J. Martinerie, The brain web: phase synchronization and large-scale integration, *Nat. Rev. Neurosci.* 2 (2001) 229–239.
- [87] C. von der Malsburg, The what and why of binding: the modeler's perspective, *Neuron* 24 (1999) 95–104.
- [88] G.S. Wig, S.T. Grafton, K.E. Demos, W.M. Kelley, Reductions in neural activity underlie behavioral components of repetition priming, *Nat. Neurosci.* 8 (9) (2005) 1228–1233.
- [89] T.P. Zanto, M.T. Rubens, A. Thangavel, A. Gazzaley, Causal role of the prefrontal cortex in top-down modulation of visual processing and working memory, *Nat. Neurosci.* 14 (2011) 656–661.