Graded Attractors: Configuring Context-Dependent Workspaces for Ideation

Ali A. Minai[1]**, Laxmi R. Iyer**[2]**, Divyachapan Padur**[1]**, Simona Doboli**[3] **and Vincent R. Brown**[4]
[1]Department of Electrical & Computer Engineering, [2]Department of Computer Science
University of Cincinnati, Cincinnati, OH 45221-0030
ali.minai@uc.edu, laxmi.r.iyer@gmail.com, divyachapan@gmail.com

[3]Department of Computer Science, [4]Department of Psychology
Hofstra University, Hempstead, NY11549
simona.doboli@hofstra.edu, vincent.r.brown@hofstra.edu

Abstract

Thought is an essential aspect of mental function, but remains very poorly understood. In this paper, we take the view that thought is a *response process* – the emergent and dynamic configuration of structured response, i.e., *ideas*, by composing response elements, i.e., *concepts*, from a repertoire under the influence of afferent information, internal modulation and evaluative feedback. We hypothesize that the process of generating ideas occurs at two levels: 1) The identification of a context-specific subset – or workspace – of concepts from the larger repertoire; and 2) The configuration of plausible/useful ideas within this workspace. Workspace configuration is mediated by a dynamic selector network (DSN), which is an internal attention/working memory system. Each unit of the DSN selectively gates a subset of concepts, so that any pattern of activity in the DSN defines a workspace. The configuration of efficient and flexible workspaces is mediated by dynamical structures termed *graded attractors* – attractors where the set of active units can be varied in systematic order by inhibitory modulation. A graded attractor in the DSN can project a selective bias – a "searchlight" – onto the concept repertoire to define a specific workspace, and inhibitory modulation can be used to vary the breadth of this workspace. As it experiences various contexts, the cognitive system can configure a set of graded attractors, each covering a *domain* of similar contexts.

In this paper, we focus on a mechanism for configuring context-specific graded attractors, and evaluate its performance over a set of contexts with varying degrees of similarity. In particular, we look at whether contexts are clustered appropriately into a minimal number of workspaces based on the similarity of the responses they require. While the focus in this paper is on semantic workspaces, the model is broadly applicable to other cognitive response functions such as motor control or memory recall.

A cognitive system's primary function is to continually construct responses to afferent stimuli. These responses can include percepts, recognitions, memories, thoughts, ideas, actions, etc. We postulate that all these responses are constructed through structured configuration of *response elements* from an appropriate repertoire. For example, there is considerable evidence that complex movements are constructed as combinations of motor synergies – coordinated spatiotemporal activation patterns of motor neurons (d'Avella, Saltiel, and Bizzi 2003). Similarly, ideas are combinations of concepts (Boden 1995; Iyer et al. 2009), which are themselves combinations of features (Warrington and Shallice 1984; McRae, de Sa, and Seidenberg 1997). The search for appropriate responses is, therefore, the search for the right combination of response elements, organized into a structure by some self-organization process (e.g., sensemaking for episodic memories). Given the vast amount of information encoded in the brain, this search can only succeed if it is focused on a *relevant* subset of response elements, and combines them only in potentially productive ways; random and/or exhaustive search is not an option. We call this relevant subset of response elements organized in a context-appropriate fashion a *workspace* – a term also used by others (Dehaene and Naccache 2001; Baars and Franklin 2003) in a somewhat different (but related) sense.

Once the workspace is identified, there must also be a dynamic search process that generates structured responses within it, as well as higher-level processes to modify the search appropriately if needed. It is obvious that different contexts can potentially require distinct workspaces, but it is inefficient to maintain a separate workspace for every context. A successful cognitive system must be able to infer a *minimal* or *optimal* set of distinct workspaces based on experience across many contexts – merging contexts requiring similar responses into a single workspace, while creating distinct workspaces for contexts requiring qualitatively different responses. For example, the workspace for assembling a sandwich in one's kitchen may be the same as that for assembling one in a cafeteria, but distinct from one used to decide what to wear to work. The configuration of such context-dependent workspaces is the focus of this paper.

While the above framework should apply to all mental processes, we focus in this paper on the generation of ideas, which are seen as combinations of concepts. We do not consider the structural organization of such conceptual combinations, which is a complex issue in its own right (van der Velde and de Kamps 2006).

Thought, Ideation and Creativity

In this paper, we begin with the reasonable assumption that thought – like perception, action and emotion – is an *emergent property* of the material organism, and especially of organisms with well-developed cortical structures. In this view, what we recognize as "thought" – including creative thought – arises from the natural dynamics of the nervous system embedded within the body, and is part of a *continuum* with perception and action, using the same structures, processes and informational strategies as those subserving these other mental processes. For example, just as actions arise from the combination of motor synergies, thoughts may be seen as combinations of "semantic synergies", i.e., concepts.

This perspective is consistent with recent views of mental function as a whole (Edelman and Tononi 2000; Fuster 2003), and especially with the "agents of mind" model proposed by Houk (Houk 2005). Interestingly, our perspective also implies that, since thought is not a unique mental process, it is also not a uniquely human attribute. Rather, it simply encompasses the internally grounded processes that became possible as the brains of organisms grew larger and more complex over the course of evolution.

Directed, context-specific thinking is usually seen as the process of *ideation*, i.e., generating relevant ideas (which may be trivial or profound), and has generally been studied through behavioral experiments on brainstorming (Paulus and Dzindolet 1993; Coskun et al. 2000; Paulus and Brown 2003; Dugosh and Paulus 2005; Nijstad and Stroebe 2006). Only recently have there been attempts to determine the neural correlates of creativity and insight (Bowden et al. 2005; Heilman, Nadeau, and Beversdorf 2003). Behavioral studies by Bowden, Jung-Beeman and colleagues suggest that insight involves collaborative processing by left and right hemispheres (Bowden et al. 2005), and fMRI data from the same group shows that specific patterns of activity in the right hemisphere (RH) are correlated with insight. A useful body of data for understanding at least some types of ideation comes from the literature on semantic cognition (Martin 2007), which has clarified how and where the brain represents different types of information. In particular, this work has shown that nouns and verbs are represented in modal (Martin et al. 1996; Martin 2007) and amodal (Patterson, Nestor, and Rogers 2007) ways, and that words related to humans, animals, tools, etc., activate distinct regions of the left hemisphere (LH) (Damasio et al. 1996; Martin 2007), as well as regions in the RH (Damasio et al. 1996; Bowden et al. 2005), which may play a critical role in linking disparate concepts to generate creative ideas (Bowden et al. 2005; Schilling 2005; Duch 2007). Another important source of information on ideation is the extensive literature on cognitive control (Miller 2000), elucidating the mechanisms of attentional switching (Graybiel 1995; Houk 2005), working memory (Goldman-Rakic 1995), and reward (Schultz 2000). These mechanisms are all critical in the generation and recognition of novel emergent arrangements of existing knowledge.

A Computational Model for Ideation

We have recently proposed a computational model for ideation, which is described in more detail elsewhere (Iyer et al. 2009). Here we just give a qualitative description to motivate the study of workspace formation. The core semantic system is modeled as a two-level neural network representing *concepts* and *features*, where each concept is represented both modally and amodally. The amodal representation for a concept involves the activation of a specific *concept unit* (corresponding to a cell assembly) in a neural system termed the *concept network* (CN). Each concept also corresponds to a pattern of activity across *feature units* in another neural system called the *feature layer* (FL). This is the concept's modal representation, which allows similarities and differences between concepts to be encoded. Concept units are connected with each other through excitatory connections whose strength reflects their joint utility in previously experienced ideas. The CN as a whole is subject to global inhibition, which makes the activity of concept units competitive K-of-N. An activated unit can only remain active for a finite duration, after which it enters a refractory period of finite duration. This reflects the natural neural mechanisms of resource depletion through activity, and the requirement to replenish them.

Any K units of the CN active simultaneously represent a set of concepts, i.e., a potential idea. However, this pattern of activity can only persist in the face of competition if these active units are mutually strongly linked, i.e., they "make sense" together based on previous experience. Such a persistent activity pattern can be seen as a metastable attractor – precluded from true stability because of the limit on the activity duration of concept units. Such a pattern of activity is recognized as an idea. Non-persistent patterns are regarded as cognitive noise. Thus, the dynamics of the CN is an itinerant trajectory that pauses at metastable attractors representing ideas before moving on. This is an example of *winnerless competition* that has been proposed as a mechanism for cognitive dynamics (Rabinovich et al. 2001).

The actual pool of concept units available at a time is determined by a *context input* representing the current task or problem situation, and is mediated by a highly flexible biasing system called the *dynamic selection network* (DSN). This is the system primarily responsible for configuring the workspace, and is the main focus of this paper. It is described in more detail below.

Context-Dependent Workspace Model

We assume that the concepts in the CN are grouped into overlapping *categories* through a subspace clustering process driven by experience. The DSN consists of a large number of *selector units* organized into *modules* with dense internal connectivity. Each selector unit, k, is tuned to a specific category, Q^k, and sends a *gating signal* to all the concept units in its category. This set of concept units is designated as the *member set* for the selector unit, and is denoted by Γ^k. The activation of unit k, denoted by $z_k^s = 1$, causes all the concepts in its category to become available in the workspace. The workspace at any time is thus defined as $W \equiv \bigcup_{k|z_k^s=1} \Gamma^k$.

Numerous selector units in the DSN are tuned to each category, but are assumed not be repeated within the same module. Since the number of units in each module is taken to be significantly smaller than the total number of categories, each module provides a subset of categories that can be seen as defining a partial workspace, or a "building block" for the overall workspace. We hypothesize that the rapid construction of efficient workspaces involves the selective activation of modules whose component categories, taken as a whole, define an appropriate set of concepts in the current context.

Context-dependent workspaces must be both efficient and flexible. When needed, they must initially include only the most relevant concepts so that good ideas can be found quickly. However, if this does not happen, the workspace should expand by including more concepts, albeit in a way that still takes the context into account. This means that, upon being stimulated by a context, only a small set of DSN units should initially be activated, in turn selecting a small set of highly relevant concept units. Then, if needed, more DSN units should be added to the active set, *gradually* expanding the set of available concept units, as illustrated in Figure 1. We introduce a dynamical entity termed a *graded attractor* to mediate this function.

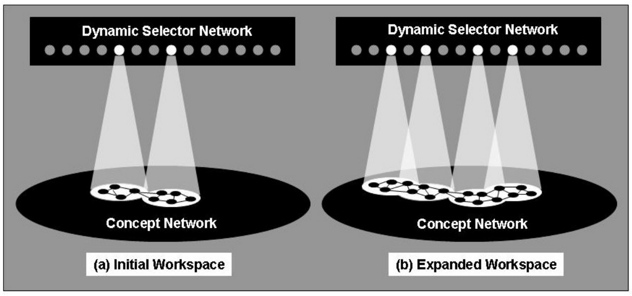

Figure 1: Selection of context-dependent workspace by DSN units. (a): A small workspace selected initially by a few active DSN units (shown in white). (b): An expanded workspace obtained when more DSN units are activated.

Graded Attractors

A graded attractor is defined over a population of neural units, and like a normal attractor, comprises an activity pattern that persists because of preferentially strong mutual excitation among its active units. However, the number of active units in a graded attractor can be varied by changing the level of global inhibition in the system. When the inhibition is highest, only a small *core set* of units is active after convergence. Then, as inhibition is lowered, more units become active until, at the lowest level of inhibition, the pattern has the maximum number of active units, comprising the *base set* for the attractor. This is illustrated in Figure 2. Crucially, this pattern of widening activity is *context-specific*, ensuring that the most potentially relevant categories (and concepts) are activated earlier in the search process. The graded attractor in the DSN can thus be seen as embodying a context-specific ordering of categories (and, thus, concepts) to be included in the workspace.

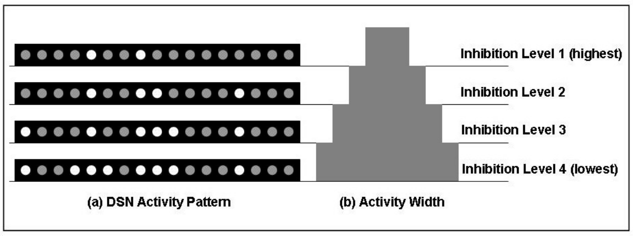

Figure 2: Graded attractor. (a) The activity pattern in the DSN (white units are active, gray inactive); the number of active units in the attractor increase with decreasing inhibition level. (b): A figurative view of the graded attractor indicating the activity width at each inhibition level.

The use of attractors rather than activity patterns generated by the feed-forward projection of the context is also crucial because it ensures that the workspace is stable against noise, and provides generalization across contexts requiring similar workspaces. After experience with a large number of contexts, the system should construct an efficient set of graded attractors, each defining the workspace for a broad but distinct contextual domain – a "frame of mind", so to speak. It is often noticed that, during the performance of a task, the prefrontal and parietal cortical regions seem to have units tuned to all relevant contingencies, stimuli and decisions. This has led to the idea of task-specific *functional networks* (Varela et al. 2001; Engel, Fries, and Singer 2001; Bressler and Tognoli 2006) that emerge dynamically in a given context and provide the scaffolding on which detailed responses can be configured. Graded attractors represent a mechanism for the formation of such functional networks.

System Implementation

The key issue of interest in this paper is the efficient configuration of suitable graded attractors based on data. Our approach exploits modularity, which is known to be a fundamental enabling feature for complexity in biological systems (Callebaut and Rasskin-Gutman (eds.) 2005). As with all modularity-based approaches, the method trades a little optimality for much greater efficiency of configuration.

We simulate a system with N selector units organized into M modules. The CN is not simulated explicitly, but it is assumed that the concepts (response elements) available are organized into N_q categories with possibly overlapping memberships, where $N \gg N_q$. Since each selector unit is tuned to one of these categories, every module represents a combination of a few unique categories. The set of all categories is denoted as $\mathcal{Q}$.

The system is configured using a *training set* based on N_x contexts. Each context, χ^k, is represented as a m_x-bit vector, X^k, where each bit represents an abstract feature (see

(Iyer et al. 2009) for examples of concrete context features). It is also associated with a *characteristic category distribution* (CCD), $\mu^k = [\mu_1^k\ \mu_2^k\ ...\ \mu_{N_q}^k]$ over the set $\mathcal{Q}$, where μ_j^k is the probability that an idea in context χ^k includes at least one concept from category Q^j. The categories for which $\mu_j^q > 0$ comprise the *support set* of context χ^k, and is divided further into three parts: 1) Primary categories, with $\mu_j^k > 0.8$; 2) Secondary categories, with $0.8 \geq \mu_j^k > 0.3$; and 3) Tertiary categories, with $0.3 \geq \mu_j^k > 0$.

To see whether contexts with similar CCDs (i.e., requiring ideas constructed from similar categories) cluster together to form unified workspaces, the contexts are clustered into several *groups* by construction. Contexts within a group share two primary categories and two secondary categories, while one primary category for each is a secondary category for the other. Contexts within a group also have similar feature representations. In the current simulations, we use 11 contexts, grouped into one group of size 5, another of size 3, and three contexts dissimilar to all others. Thus, the goal is for the system to infer 5 workspaces.

For training, 100 N_q-dimensional binary vectors are generated as *exemplars* for each context, where each bit represents a unique category, indicating that the idea includes a concept from that category. Thus, the probability of bit j being 1 in an exemplar for context χ^k is μ_j^k, i.e., the exemplars conform with the CCD for their contexts. Each exemplar can be seen as a category-level representation of an underlying idea.

Network Configuration

The DSN comprises $N = 1137$ selector units organized into 195 modules of 4 to 6 units each. Units within each module have strong excitatory connections, while connectivity across modules is relatively sparse. The category assignments for selector units are made such that some modules group the high probability categories for the given contexts, ensuring that these modules would be tuned on strongly by exemplars from those contexts. The remaining modules are constructed as random combinations of these. Thus, the network represents a broad repertoire of category combinations across all its modules, including some tuned to the given contexts.

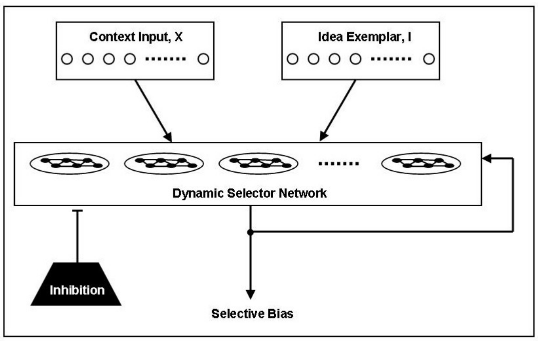

Figure 3: Architecture of the graded attractor system.

The architecture of the simulated system is shown in Figure 3. Each DSN unit, i receives connections from three sources: 1) A connection with fixed weight $w_{ij}^{sc} = 1$ from idea input I_j; 2) Modifiable associative connections, w_{ij}^{sx} from the context bits, X_j; and 3) Initially weak but modifiable real-valued connections, w_{ij}^{ss}, from other DSN units, j. The dynamics of selector unit i is given by:

$$u_i^s(t) = (1 - \alpha^s)u_i^s(t-1) + \alpha^s y_i^s(t); \quad 0 < \alpha^s < 1 \quad (1)$$

where α^s is an inertial parameter and y_i^s is the current input to the unit from all sources. As in the CN, selector units fire competitively, with the $K(t)$units with the highest $u^s(t) > 0$ activated at time t.

During training, the 1100 exemplars from all contexts are presented in random order to the network, with the corresponding context input clamped on. Each exemplar is presented for several consecutive steps termed a *presentation cycle*. At the beginning of the presentation cycle, the DSN activity is determined purely by the exemplar input. The system is then allowed to relax through recurrent dynamics, still without any influence from the context input. Only the K_s most excited units are allowed to be active at a time, where K_s represents the desired size of the base set for the graded attractor. The activity pattern obtained at the end of the presentation cycle is embedded auto-associatively in the DSN, and hetero-associated with the context input. The key element in this process is that the learning rates for both associations are modulated nonlinearly by the overall activity of the module to which the postsynaptic unit belongs, i.e., $\eta(t) = \eta_{fixed}S[\sum_k z_k^s(t)]$, where k indexes all units in the module for unit i, z_k^s is the activation of unit k, η_{fixed} is a fixed number between 0 and 1, and $S[.]$ is a highly nonlinear sigmoid function. This soft-competitive learning process ensures that highly active modules learn much more than those with low total activity. Essentially, each module acts as a multi-unit representational element (Lin, Osan, and Tsien 2006) in a competitive learning network, and stores the resulting activity pattern as a graded attractor with the most activated modules forming its core. Each attractor becomes associated with one or more contexts, and is activated from the core outward (depending on the inhibition level) when a similar context is presented.

Results

After training, the network was tested by probing it with each of the 11 training contexts as well as several novel contexts similar to the original context groups. The resulting dynamics was evaluated to see: 1) whether each context recalled the correct graded attractor; 2) how close its pattern of widening activity was to the optimal pattern with respect to the CCDs for the contexts it was associated with; and 3) how much of the CCD for the context was covered. The evaluation was done as follows. First, each context, χ^k, was presented to the system at progressively decreasing inhibition levels starting with the highest, and the system was allowed to relax to convergence in each case. The set of categories covered by activated units before relaxation and after convergence was noted. This gave a sequence of *covered category sets*, $s^k = \{s_1^k,\ s_2^k,\ ...,\ s_r^k\}$, where each s_j^k is a set of categories, and its cardinality is denoted by $|s_j^k|$. Only sets where the covered categories changed were stored, since changing the inhibition level did not always result in a change in active categories. For each s_j^k, *coverage* was calculated as, $\beta_j^k = \sum_{i|Q^i \in s_j^k} \mu_i^k$, and a *quality* value obtained as $\zeta_l^a = \beta_j^k / \omega_j^k$, where ω_j^k is the highest possible sum that could be obtained from the CCD of context χ^k using $|s_j^k|$ categories. Two metrics were calculated from these quantities:

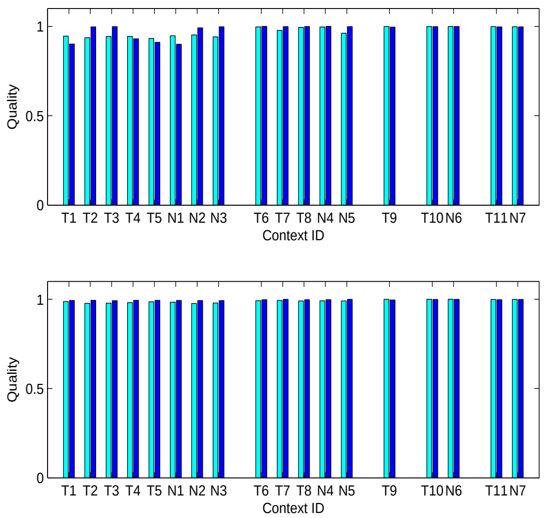

Figure 4: Average attractor quality for 11 learned and 7 novel contexts. The lighter bars show the initial value, and the dark bars after relaxation. The Tx labels indicate learned contexts and Nx labels novel ones. Context groups are separated by spaces Top: Context-level metric; Bottom: Group-level metric.

1. **Average Attractor Quality**, A^k, for context χ^k is defined by the mean of β_j^k over all j. This was calculated for category sequences obtain both before and after relaxation.
2. **Total Coverage**, B^k, for context χ^k is defined as $B^k = \beta_r^k / \nu^k$, where $\nu^k = \sum_{i=1}^{N_q} \mu_i^k$. This was evaluated only for the category sequences obtained after relaxation to convergence.

Since the clustering of multiple contexts onto one graded attractor necessarily makes the attractor suboptimal for each individual context, we evaluated two versions of each metric. The first – termed the *context-level metrics* – were as described above, while the second – called the *group-level metrics* – used a single combined CCD for each context group. This combined CCD was obtained by taking the maximal value of each CCD element across the group.

Figures 4 and 5 show the results for both versions of both metrics. The system does very well for all contexts on both quality and total coverage. As expected, the context-level

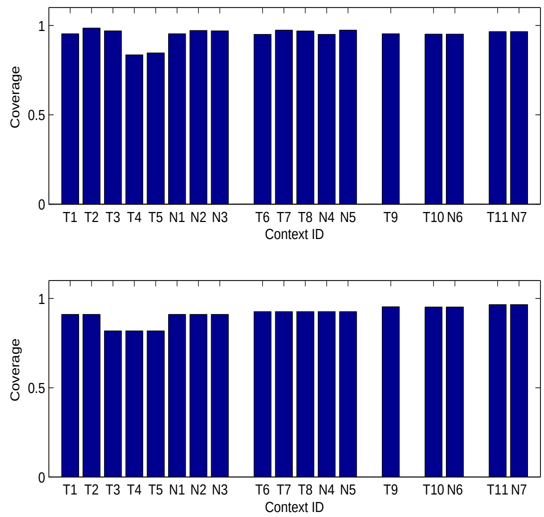

Figure 5: Post-relaxation total coverage for 11 learned and 7 novel contexts. All labels and conventions are as in Figure 4 Top: Context-level metric; Bottom: Group-level metric.

metrics for the multi-context workspaces are not as good as those for singleton contexts. However, when group-level metrics are used, the attractors are all close to optimal. It is interesting to note that auto-associative dynamics makes only a slight difference to attractor quality, i.e., the pattern of activity recalled initially by a context is already very good. However, the recurrent activity is crucial for stabilizing the attractor against noise. The figures also show that novel contexts are appropriately mapped to the workspaces for similar contexts. As pointed out earlier, this can be supplemented with mechanism for switching attractors if necessary, as described in (Iyer et al. 2009).

Conclusion

In this paper, we have presented a modular recurrent network model that can learn highly flexible dynamical objects called graded attractors. These can be used to instantiate context-depend workspaces or functional networks in many cognitive tasks, including ideation. Our simulations show that a simple unsupervised learning mechanism can configure multiple graded attractors based on experience, and that these attractors can be recalled in useful form when needed, both by familiar and novel contexts. This model is still in its early stages of development, but we believe that graded attractors will prove to be useful in modeling response selection and working memory functions across many cognitive tasks.

Acknowledgment

This work was supported by collaborative National Science Foundation Human and Social Dynamics Program grants to Ali Minai (BCS-0728413), Simona Doboli (BCS-0729470) and Paul Paulus (BCS-0729305), which includes support from the Deputy Director of National Intelligence. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Helpful suggestions by Paul Paulus, Dan Levine, Shane Mueller, Jim Houk and Steve Bressler are gratefully acknowledged.

References

Baars, B., and Franklin, S. 2003. How conscious experience and working memory interact. *Trends in Cognitive Sciences* 7:166–172.

Boden, M. 1995. Creativity and unpredictability. *Stanford Humanities Review* 4(2).

Bowden, E.; Jung-Beeman, M.; Fleck, J.; and Kounios, J. 2005. New approaches to demystifying insight. *Trends in Cognitive Sciences* 9:322–328.

Bressler, S., and Tognoli, E. 2006. Operational principles of neurocognitive networks. *International Journal of Psychophysiology* 60:139–148.

Callebaut, W., and Rasskin-Gutman (eds.), D. 2005. *Modularity: Understanding the Development and Evolution of Natural Complex Systems*. MIT Press, Cambridge, MA.

Coskun, H.; Paulus; P.B., Brown, V.; and Sherwood, J. 2000. Cognitive stimulation and problem presentation in idea generation groups. *Group Dynamics: Theory, Research, and Practice* 4:307–329.

Damasio, H.; Grabowski, T.; Tranel, D.; Hichwa, R.; and Damasio, A. 1996. A neural basis for lexical retrieval. *Nature* 380:499–505.

d'Avella, A.; Saltiel, P.; and Bizzi, E. 2003. Combinations of muscle synergies in the construction of natural motor behavior. *Nature Neuroscience* 6:300–308.

Dehaene, S., and Naccache, L. 2001. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* 79:1–37.

Duch, W. 2007. Intuition, insight, imagination and creativity. *IEEE Comput. Intell.* 40–52.

Dugosh, K., and Paulus, P. 2005. Cognitive and social comparison processes in brainstorming. *Journal of Experimental Social Psychology* 41:313–320.

Edelman, G., and Tononi, G. 2000. *A Universe of Consciousness: How Matter Becomes Imagination*. Basic Books.

Engel, A.; Fries, P.; and Singer, W. 2001. Dynamic predictions: oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience* 2:704–716.

Fuster, J. 2003. *Cortex and Mind: Unifying Cognition*. Oxford University Press.

Goldman-Rakic, P. 1995. Cellular basis of working memory. *Neuron* 14:477–485.

Graybiel, A. 1995. Building action repertoires: memory and learning functions of the basal ganglia. *Current Opinion in Neurobiology* 5:733–741.

Heilman, K.; Nadeau, S.; and Beversdorf, D. 2003. Creative innovation: possible brain mechanisms. *Neurocase* 9:369–379.

Houk, J. 2005. Agents of the mind. *Biol. Cybern.* 92:427–437.

Iyer, L.; Doboli, S.; Minai, A.; Brown, V.; Levine, D.; and Paulus, P. 2009. Neural dynamics of idea generation and the effects of priming. *Neural Networks* 22:674–686.

Lin, L.; Osan, R.; and Tsien, J. 2006. Organizing principles of real-time memory encoding: neural clique assemblies and universal neural codes. *Trends in Neurosciences* 29:57.

Martin, A.; Wiggs, C.; Ungerleider, L.; and Haxby, J. 1996. Neural correlates of category-specific knowledge. *Nature* 379:649–652.

Martin, A. 2007. The representation of object concepts in the brain. *Annual Review of Psychology* 58:25–45.

McRae, K.; de Sa, V.; and Seidenberg, M. 1997. On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General* 126:99–130.

Miller, E. 2000. The prefrontal cortex and cognitive control. *Nature Reviews: Neuroscience* 1:59–65.

Nijstad, B., and Stroebe, W. 2006. How the group affects the mind: A cognitive model of idea generation in groups. *Personality and Social Psychology Review* 3:186–213.

Patterson, K.; Nestor, P.; and Rogers, T. 2007. Where do you know what you know? the representation of semantic knowledge in the human brain. *Nature Rev. Neurosci.* 8:976–987.

Paulus, P., and Brown, V. 2003. Enhancing ideational creativity in groups: Lessons from research on brainstorming. In Paulus, P., and Nijstad, B., eds., *Group Creativity*. New York: Oxford University Press. 110–136.

Paulus, P., and Dzindolet, M. 1993. Social influence processes in group brainstorming. *Journal of Personality and Social Psychology* 64:575–586.

Rabinovich, M.; Volkovskii, A.; Lecanda, P.; Huerta, R.; Abarbanel, H.; and Laurent, G. 2001. Dynamical encoding by networks of competing neuron groups: winnerless competition. *Physical Review Letters* 87:068102–1.

Schilling, M. 2005. A small-world network model of cognitive insight. *Creativity Res. J.* 17:131–154.

Schultz, W. 2000. Multiple reward signals in the brain. *Nature Reviews Neuroscience* 1:199–207.

van der Velde, F., and de Kamps, M. 2006. Neural blackboard architectures of combinatorial structures in cognition. *Behavioral and Brain Sciences* 29:37–108.

Varela, F.; Lachaux, J.-P.; Rodriguez, E.; and Martinerie, J. 2001. The brainweb: phase synchronization and large-scale integration. *Nature Reviews Neuroscience* 2:229–239.

Warrington, E., and Shallice, T. 1984. Category specific semantic impairments. *Brain* 107:829–854.