

Modularity and Self-Organized Functional Architectures in the Brain

Laxmi Iyer and Ali A. Minai
University of Cincinnati
iyerlr@email.uc.edu
Ali.Minai@uc.edu

Simona Doboli and Vincent R. Brown
Hofstra University
Simona.Doboli@hofstra.edu
Vincent.R.Brown@hofstra.edu

It is generally believed that cognition involves the self-organization of coherent dynamic functional networks across several brain regions in response to incoming stimulus and internal modulation. These context-dependent networks arise continually from the spatiotemporally multi-scale structural substrate of the brain configured by evolution, development and previous experience, persisting for 100-200 ms and generating responses such as imagery, recall and motor action. In the current paper, we show that a system of interacting modular attractor networks can use a selective mechanism for assembling functional networks from the modular substrate. We use the approach to develop a model of idea-generation in the brain. Ideas are modeled as combinations of concepts organized in a recurrent network that reflects previous associations between them. The dynamics of this network, resulting in the transient co-activation of concept groups, is seen as a search through the space of ideas, and attractor dynamics is used to “shape” this search. The process is required to encompass both rapid retrieval of old ideas in familiar contexts and efficient search for novel ones in unfamiliar situations (or during brainstorming). The inclusion of an adaptive modulatory mechanism allows the network to balance the competing requirements of exploiting previous learning and exploring new possibilities as needed in different contexts.

1 Introduction

A consensus is gradually developing that cognition involves the continual self-organization and dissipation of functional networks across several brain regions – especially the neocortex – in response to incoming stimulus and internal modulation [25, 19, 18, 21, 23, 42, 3, 37]. Each functional network emerges from the brain’s substrate in response to contextual information, persists while the context applies, and then dissolves back into the substrate to allow a new network – or networks – to emerge.

As pointed out by Doyle and colleagues [8, 7], useful systems are typically *heterogeneous* and *specific* rather than homogeneous and generic. The configuration of such specific heterogeneity usually requires optimization, but that is not feasible in real-time for a complex system like the brain. Instead, such systems must work through self-organization arising naturally from the structure and dynamics of the system. However, rapid self-organization of functional networks is only possible in the brain if it provides the structure and mechanisms that facilitate the process – and it does! The cortex is organized into modules called *cortical columns* that group together into larger modules termed *hypercolumns* [28]. Such modularity is a fundamental enabling mechanism for self-organized complexity in living systems [38], and provides exactly the sort of flexibility that is needed for efficient reconfiguration of functional networks.

We postulate that five factors combine to produce the emergence of effective, flexible, robust and reliable functional networks in the brain. These are:

1. *A modular substrate with sufficient diversity:* The underlying network, which is configured over the multiple time-scales of evolution, development and experiential learning, provides modules with a wide variety of functional micro-behaviors.
2. *A dynamic selective process to bind functional structures:* This process selectively combines an appropriate set of modules so that the correct macro-functionality emerges from the interaction of their micro-behaviors [4].
3. *A dynamic modulatory process to control scope and switching:* This process modulates the excitability of neural units to determine which ones participate in the current functional network, thus controlling the *effective breadth* of these networks, and the transition between networks.
4. *An evaluative feedback process:* This provides a reinforcement signal back to the system so that it can appropriate functional networks can be configured and triggered.
5. *A repertoire of learning processes:* These include: a) Self-organization of micro-behaviors in modules to provide a good behavioral basis; b) Reinforcement-driven Hebbian learning to associate contexts with appropriate functional networks; c) Reinforcement-driven Hebbian learning to

configure the interactions among modules so that useful functional networks become embedded in the substrate through self-organization.

In order to study these principles in a concrete, albeit much simplified, framework, we consider a model for idea generation in a neural system.

A major motivation for the development of the model presented here is our goal to develop a detailed, neurally plausible model of the process of creative idea generation commonly called ‘brainstorming’. Brainstorming refers to idea generation under specific guidelines designed to promote quantity and creativity with minimal censorship and criticism [32]. Although these guidelines were designed for use in groups, individuals can obviously engage in creative idea generation as well. In fact the vast majority of laboratory research on brainstorming finds that an equal number of solitary brainstormers outperform interactive groups by almost a 2 to 1 margin when quantity of ideas is counted [13, 29, 35]. Nonetheless being exposed to the ideas of others can be stimulating and, under conditions designed to reduce social inhibition, groups can match or exceed the performance of an equal number of solitary brainstormers, thus closing or eliminating the ‘group productivity gap’ [33, 36]. Theoretically groups provide the stimulation necessary to get individuals “out of a rut” by activating less accessible categories of ideas and activating atypical sequences of ideas that can be fuel for novel conceptual combinations [6, 30]. There is accumulating evidence for the stimulating effects of the exposure to others’ ideas [10, 16, 26] (although see Nijstad et al. [31] for evidence that exposure to other ideas can also have interfering effects on brainstorming). Recent models based on brainstorming as activation, search, and recall of ideas in associative memory can account for a number of these empirical results and have proven fruitful in generating testable hypothesis about a number of important cognitive processes involved in brainstorming, including attention, working memory, memory accessibility, and convergent vs. divergent thinking [6, 34, 30]. One major limitation of these models is the inability to account for the important process of conceptual combination in generating creative responses. One goal of the model presented here is to provide a neurocomputational mechanism for the generation of novel conceptual combinations. In addition, with recent work on the neuroscience of creativity [22, 24, 15], there is the need for more detailed neurally-inspired models of the brainstorming process. Prototypes of the model presented here based on the attractor network architecture have shown promise in accounting for some of the basic empirical brainstorming results in both individuals and groups [5, 14].

2 Problem Formulation

We begin by postulating that *ideas* are combinations of *concepts*, which are the basic representational units in our model. Novel ideas are conceptual combinations whose elements have not been combined in the past, while groups that have been formed previously represent familiar ideas. A cognitive system must be able to retrieve familiar ideas and to generate novel ones.

Ideas always arise – and make sense – within a *context*, e.g., a task situation. Each context – even if it is not completely familiar – tends to elicit a set of ideas by association and suppresses others, presumably making the search for useful ideas more efficient. Thus, a context, Φ , can be seen as a semantic biasing mechanism that preferentially *implicitly* un.masks a subset, I^Φ , of ideas from the space of all ideas, allowing a *search process* to explore this subset and “discover” good ideas. Since ideas are combinations of concepts, we model this process as the selective activation of a *context-dependent concept network (CCN)* whose states represent ideas, and whose dynamics embodies the search process. The CCN is a functional network.

Ideas can be seen more broadly as internal responses of the cognitive system, and idea-generation is essentially no different from the generation of motor responses, memory recall or mental imagery (indeed, mental images can be regarded as ideas too!) It is well known that mental responses fall into two broad categories: Automatic and effortful [12]. The former (automatic) are faster, stereotypical and fluent, while the latter (effortful) appear to entail some type of constructive process or search. It has been proposed that effortful responses require the involvement of working memory [2] or global workspace [1], which functions essentially as a temporary “hidden layer” facilitating direct linkage between stimulus and response. Once the direct linkage is consolidated, the mediating process is cut out, leaving behind an automatic response for the future [12]. Alternatively, working memory may remain involved in complex, context-dependent tasks — again as a biasing hidden layer [27, 17, 11]. While these abstract formulations clearly capture important elements of the phenomenology of response generation, they represent an implicitly teleological view, where automatization of response is seen as an end towards which the mechanism of learning works. In contrast, we seek a neural system where the emergence of novel responses and facilitation of familiar ones *both* arise naturally from the intrinsic dynamics of the system. Embedded in an environment, such a system continually creates responses to the information flowing through it — perpetually adjusting its internal constraints to facilitate especially useful responses. Behavior, in this view, is not a goal or a purpose of the system, but a property.

In the case of idea generation, we develop a simple neural model showing how modular organization, selectivity, modulation, evaluative feedback and reinforcement-driven adaptation combine to produce an efficient search process through the system’s natural dynamics.

3 The Idea Generation Process

Search is useful only if it is efficient. Exhaustive search, while guaranteed to succeed, is typically not an option. Thus, our primary focus is on how the search for ideas is *shaped* and *guided* by the system’s dynamics so that it is as broad as necessary and as narrow as possible. We only consider this criterion heuristically, and will refine it further in future studies.

Efficiency in search requires effective use of information. There are three

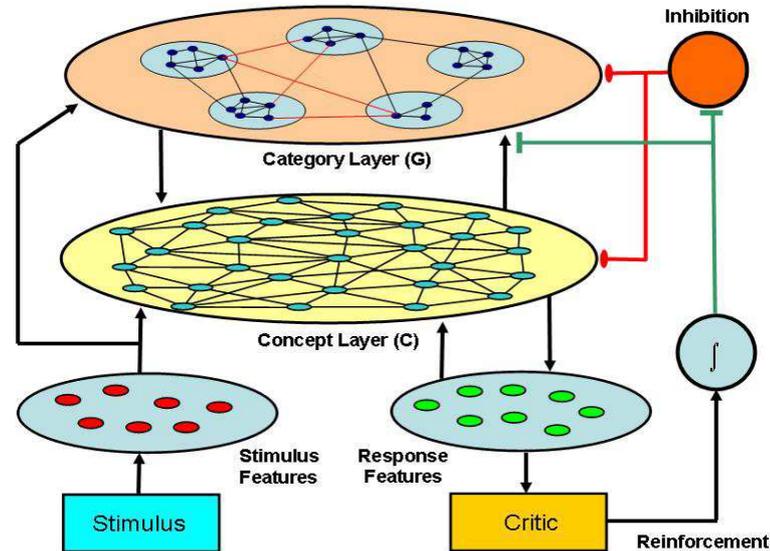


Figure 1: Architecture of the idea generation model.

major sources of information available to guide the search for ideas in our model, and form the basis of its utility. These are:

1. *Current task or domain context:* The search should focus on concepts that, in combination, are likely to be relevant to the context at hand. A neural system does this by associating context representations with specific patterns of module activation, forming functional networks.
2. *Previous experience:* The flow of the search should be guided, and its extent determined, by the experience of the system in inferring the regularities of its operating environment. The only way for a neural system to do this is by embedding such experiential knowledge in its structure and dynamical parameters, so that the emergent flow of the system's activity is appropriately constrained to productive regions and trajectories in the search space.
3. *The progress of the search so far:* As the search proceeds, it should be guided continually by the incoming evaluative feedback towards more productive regions. This is necessary because novel ideas often require the formation of functional networks other than those triggered initially by the context. This can only happen if feedback can overcome the initial bias and create new structures in a systematic way – again, guided by the experiential knowledge implicit in the system's structure.

The primary hypothesis behind our model can be stated in two parts: 1)

The interaction between afferent context/stimulus information, network structure and modulation create a dynamic energy landscape in concept space; and 2) The itinerant flow of activity over this landscape represents the search for ideas, which are metastable attractors. Thus, the mechanisms for rapidly generating productive energy landscapes in concept space is the fundamental focus of the research.

3.1 Convergent vs. Divergent Thinking

An important issue in idea generation is the style of thinking. In the *divergent thinking* mode (or exploration), there is some possibility that concepts or categories that are not strongly connected to each other normally or not strongly associated with the same context, will still become co-active, leading to a large number of relatively random combinations, but occasionally to useful novel associations. These useful novel associations are less likely to occur in the *convergent thinking* mode (exploitation) where the search of the conceptual space is more restricted to existing ideas.

We hypothesize that these differences can be understood through the dynamics of activity on the dynamic energy landscape described above, and arise from the relative flexibility of this landscape. A system with insufficient modulation and short-term learning is only able to produce stereotypical energy landscapes with relatively high barriers between attractors. It, therefore, tends to get trapped in suboptimal regions of search space and leads to convergent thinking. In contrast, a more flexible system can adapt the energy landscape to create new attractors through the recombination of old ones, thus evincing divergent thinking.

The model we develop embodies this view of idea generator, implementing it in a connectionist framework.

4 Model Description

The model we propose is shown in Figure 1. It comprises the following components:

- A *stimulus* or *input layer*, I , providing a set of n_I *stimulus feature units* (SFUs) encoding afferent stimuli, including context information.
- A *response layer*, R , comprising a fixed pool of n_R *response feature units* (RFUs) denoting response-relevant characteristics of concepts.
- A *concept layer*, C , comprising a pool of n_C *concept units* (CUs). Each concept unit has connections (both excitatory and inhibitory) with a subset of RFUs, thus implicitly defining the response semantics of the concept. The concept layer also has fast recurrent excitatory connections among concept units, which tend to stabilize specific patterns of coactive units, and slower

recurrent self-inhibition that *stochastically* limits how long a unit can remain continuously active. These competing tendencies generate a flow of activity patterns — similar to itinerant dynamics over a set of metastable attractors [43] — consistent with the hypothesis that mental constructs are represented as slightly persistent patterns of coherent cortical activity [20, 41]. The patterns of coactive concepts temporarily stabilized by fast connections represent ideas. Each idea activates the response features corresponding to its constituent concepts, which is the effective response of the system. As a whole, the concept layer can be seen as comprising a distributed *semantic network*.

The pattern of recurrent connectivity within the concept layer reflects the utility of associations between concepts based on previous experience. Thus, if CUs i and j have been coactive in several “good” ideas, they have strong positive reciprocal connections, while the connection may be weak, non-existent or even negative if this is not the case. Thus, given the noisy activation and refractoriness of concept units, the dynamics of the concept layer tends to stochastically reactivate attractors that are previously seen good ideas or their mixtures.

- A *category layer*, G , comprising a set of N_S *concept group units* (CGUs) organized into n_Q *cliques*, $\{Q^k\}$, $k = 1, \dots, n_Q$. These cliques corresponding (roughly) to utilitarian categories. CGUs within the same clique are (relatively) densely interconnected by excitatory connections, while pairs of cliques are symmetrically connected by relatively sparse excitatory or inhibitory connections *in a specific pattern* reflecting the system’s experience of whether joint activity by clique pairs is useful or otherwise. Each CGU has reciprocal connections with a subset of concept units, thus defining a *conceptual group* (CG) which can be seen roughly as “basis functions” for ideas. Several cliques may include the same CG — each associated with its own CGU, and the concepts within a conceptual group may be quite different in terms of their features. The presence of two CGs in the same clique indicate that they have been useful together *in the same context* at various times, implying that good ideas can be elicited in similar contexts simply by activating CGs in this clique. This utilitarian clustering typically means that cliques are also semantically distinct from each other, so that every clique does not include concepts with all possible features. Thus, while the concepts and CGs associated with a clique, Q^k , are quite heterogeneous, they cover only a subset, R^k , of response features, and only a subset in the space of possible CGs.

It should be noted that the modular structure imposed by layer G on layer C is consistent with but *not* identical to the structure implicit in the recurrent connectivity of C . Concepts may be strongly connected even if they do not share a clique if they have been part of good ideas in the past. Over time, learning should adjust clique memberships and conceptual groups to remove this “mismatch”, but the current model does

not address that issue yet.

- A critic, Ψ , that takes input from layer R and compares the activity pattern of this layer with an internally stored *criterion pattern* of features, generating a graded scalar *response evaluation*. This response evaluation is used as feedback by the system to modulate its parameters, and to control learning. The critic is intended to be a *phenomenological* model of internal and external evaluative mechanisms (e.g., a dopamine signal triggered by a reinforcement signal [40, 39]).

Afferent input from the stimulus layer drives all three layers through modifiable connections. The category layer is also subject to adaptive inhibitory modulation, which controls the total amount of activity allowed in the layer. Competitive inhibition in the concept layer also constrains the number of simultaneously active concept units.

5 System Functionality

As discussed earlier, the system’s response is generated through a transiently stable (metastable) spatiotemporal pattern of activity spanning the whole system — a reverberatory (or resonant) [9] pattern involving the category, concept and response layers. However, it is specifically the activity of the concept layer that represents the system’s internal response at the level of ideas. This response is projected into the common semantic space spanned by the response features, where it can be evaluated. These features can thus be seen as “internal actuators”, or as verbalization components. Given a context (and possibly a stimulus stream), the goal is to generate responses that meet the functional criteria known to the critic as efficiently, rapidly, and copiously as possible.

Processing starts when the system is stimulated by a *context stimulus*, Φ^j through layer S . This results in the activation of one or more cliques in G based on the association between the stimulus pattern and CGUs, and creates a stable activation pattern in layer G that persists even after the context stimulus is removed. This G -layer activity projects a selective bias onto Layer C , creating an implicit energy landscape in concept space with ideas as attractors. However, since only a small number of CUs can be co-active and individual CUs can remain active only for limited durations, the energy landscape keeps changing as ideas emerge, persist for a brief time, and dissolve. This “sticky” flow of ideas in C is the search process shaped by the interaction of bias from G and the recurrent connectivity in C .

As ideas are produced, they generate a stream of evaluations, $\phi(t)$, through the critic. This is integrated by a low-pass filter

$$\bar{\phi}(t) = \alpha\phi(t) + (1 - \alpha)\bar{\phi}(t - 1); \quad 0 < \alpha < 1$$

and modulates the inhibition level on layer G . As $\bar{\phi}(t)$ becomes lower (because no good responses have been found), inhibition on G is also lowered,

eventually leading to the activation of more cliques and their CGs. This changes the bias on Layer C and, therefore, the energy landscape for the flow of ideas. The order in which new cliques are activated as inhibition is lowered is robustly dependent on two factors:

1. The original context stimulus, which determined the original clique activity pattern.
2. The inter-clique connectivity reflecting the expected utility of co-activating certain cliques (or CGUs within cliques).

This broadening of the search for ideas is useful, but is not likely to be very efficient without a concurrent process for narrowing the search so it can focus in more productive regions of idea space. This is achieved using feedback from the concept layer to the category layer. As the search is widened and some good ideas start appearing, this results in a gradual increase in $\bar{\phi}$, which then tries to *decrease* the activity in G . Also, each time a good idea is rewarded sufficiently, two processes are triggered:

1. *Short-term performance-dependent reorganization of G -layer connectivity:* The reward signal causes the connections from the concept layer to the category layer to be transiently, but strongly, potentiated (shown by the gate in Figure 1). Thus, the concept units active at the time of the reward (which presumably comprised the good idea) send an unusually strong excitation to the CGUs active at the same time, which respond to this jolt by firing strongly (e.g., bursting). This, in turn, leads to short-term synaptic potentiation of connections between bursting units [18], and a short-term depression of connections between bursting and non-bursting units. The net result is that Layer G units projecting to concepts in good ideas become relatively more strongly connected to each other, as do the cliques in which they reside, *thus changing the G -layer connectivity to reflect recent experience.*
2. *Permanent performance-based change in association between concepts:* The reward signal causes the concept units active at the time of the reward to permanently strengthen their excitatory connectivity.

The effect of (2) is that good ideas become more embedded in the concept layer, which makes them likelier to arise in the future. The effect of (1) is more pertinent to the current search. The change in G -layer connectivity means that the cliques likely to survive the process of increasing inhibition due to greater reward are not necessarily those originally activated by the context, but those responsible for the recent good ideas. If more good ideas are forthcoming from this restricted set of cliques, the system will go on restricting the search further. If the process becomes too restrictive (and good ideas cease), inhibition will drop again and the search will widen. However, the pattern of widening will be affected by the recent experience of the system because the original inter-clique

connectivity has been modified (transiently) by this experience. Eventually, the system will reach a stage where a set of cliques dense in good ideas is activated, so that the increased reinforcement sustains this set even at maximum inhibition. Basically, the system has arrived at the best set of cliques – and thus the best pool of ideas – through a process of “intelligent annealing” that takes into account the current context, the system’s past experience, and the results of the system’s current search. This discovered set of cliques and/or concepts/ideas can be associated with the original context, and slightly change the inter-clique connectivity to reflect the new discovery. More significant changes such as realignment of clique membership are left for the future.

6 Results

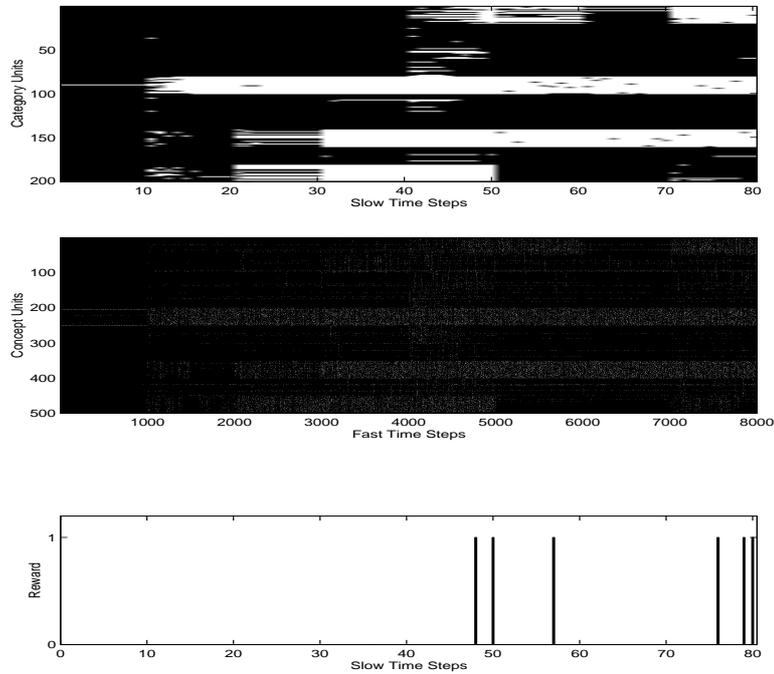


Figure 2: Simulation of the search process.

A somewhat simplified version of the model described above has been implemented, and search processes simulated using this model. The main simplification in the simulations is that each concept is typically associated with CGs from only one or two cliques. The cliques themselves are connected but non-overlapping (i.e., they do not share CGs), and both category and concept layer units are arranged such that units associated with the same clique are plotted together in figures. The dynamics of the simulated network incorporates two

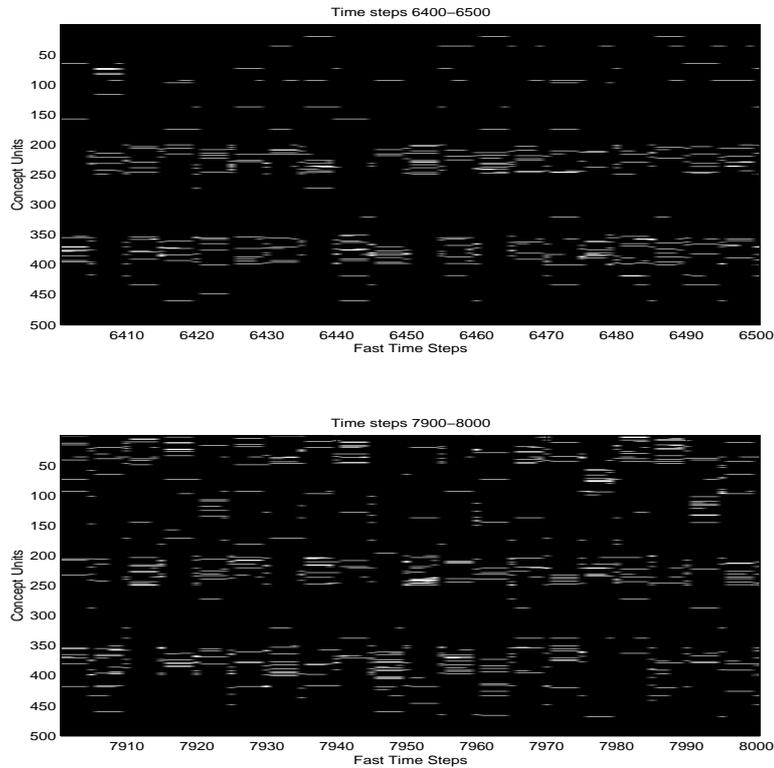


Figure 3: Simulation of the search process.

time-scales: A slow scale over which the CG activations are updated, and a fast scale used to update the concept units. Thus, each cycle of G layer update embeds 100 steps of the fast cycle. Each concept unit, however, can only be active for a random duration between 4 and 8 fast steps, and must then be refractory for a similar random duration. This allows for the context-dependent itinerant dynamics to emerge in concept space.

Figure 2 shows an example run of the simulator. The G layer has 10 cliques of 20 units each, while the C layer has 500 concept units. The system emulates a situation where good ideas arise preferentially from concepts associated with cliques 3, 6 and 10 *together*. The search starts by activating clique 6, and gradually discovers the right combination of cliques through the process described above. At that point, rewards become much more dense in time, indicating success.

Figure 3 shows a close-up view of the concept layer dynamics at different stages of the search process. The itinerant dynamics across idea space can be seen clearly.

7 Conclusions

The model described in this paper embodies a general, complex systems approach to idea generation and cognitive response – albeit in a simplified way. Our results show that the model is capable of performing an intelligent search in idea space through its inherent dynamics, and can represent the convergent and divergent modes of thinking.

A major motivations for the current model is to provide a better understanding of the cognitive, computational and neural processes underlying “brainstorming”, i.e., idea generation by individuals interacting in a group according to specific guidelines. Two mechanisms whereby groups may lead to improved idea generation are the priming of low-accessibility categories and the facilitation of novel conceptual combinations. Both of these mechanisms are inherent in the framework described here, and our model will be used to explore these issues in a systematic fashion.

Acknowledgements: The authors wish to thank Paul Paulus and Daniel Levine for useful ideas and insights. This work was partially supported by the National Science Foundation through award nos. BCS-0728413/0729470/0729305 to Ali Minai, Vincent Brown and Paul Paulus. Partial support also came from the Deputy Director of National Intelligence for Analysis.

Bibliography

- [1] BAARS, B.J., and S. FRANKLIN, “How conscious experience and working memory interact”, *Trends in Cognitive Sciences* **7** (2003), 166–172.
- [2] BADDELEY, A., *Working Memory*, Oxford University Press, New York (1986).
- [3] BRESSLER, S.L., and E. TOGNOLI, “Operational principles of neurocognitive networks”, *International Journal of Psychophysiology* **60** (2006), 139–148.
- [4] BROOKS, R.A., *Cambrian Intelligence: The Early History of the New AI*, MIT Press (1999).
- [5] BROWN, V., and S. DOBOLI, “A neural network simulation of interactive group brainstorming”, *Eighteenth Annual Convention of the Association for Psychological Science*.
- [6] BROWN, V., M. TUMEO, T. LAREY, and P. PAULUS, “Modeling cognitive interactions during group brainstorming”, *Small Group Research* **29** (1998), 495–526.
- [7] CARLSON, J.M., and J. DOYLE, “Highly optimized tolerance: Robustness and design in complex systems”, *Physical Review Letters* **84** (2000), 2529–2532.

- [8] CARLSON, J.M., and J. DOYLE, “Complexity and robustness”, *Proceedings of the National Academy of Sciences USA* **99 Suppl. 1** (2002), 2539–2545.
- [9] CARPENTER, G.A., and S. GROSSBERG, “A massively parallel architecture for a self-organizing neural pattern recognition machine”, *Computer Vision, Graphics, and Image Processing* **37** (1987), 54–115.
- [10] COSKUN, H., PAULUS, V. P.B., BROWN, and J.J. SHERWOOD, “Cognitive stimulation and problem presentation in idea generation groups”, *Group Dynamics: Theory, Research, and Practice* **4** (2000), 307–329.
- [11] DECO, G, and E.T. ROLLS, “Attention, short-term memory, and action selection: A unifying theory”, *Progress in Neurobiology* **76** (2005), 236–256.
- [12] DEHAENE, S., M. KERSZBERG, and J.-P. CHANGEUX, “A neuronal model of a global workspace in effortful cognitive tasks”, *Proceedings of the National Academy of Sciences, USA* **95** (1998), 14529–14534.
- [13] DIEHL, M., and W. STROEBE, “Productivity loss in brainstorming groups: Toward the solution of a riddle”, *Journal of Personality and Social Psychology* **53** (1987), 497–509.
- [14] DOBOLI, S., A.A. MINAI, and V.R. BROWN, “Adaptive dynamic modularity in a connectionist model of context-dependent idea generation”, *Proceedings of the International Joint Conference on Neural Networks*, (2007).
- [15] DUCH, W., “Intuition, insight, imagination and creativity”, *IEEE Computational Intelligence Magazine* (August 2007), 40–52.
- [16] DUGOSH, K.L., and P.B. PAULUS, “Cognitive and social comparison processes in brainstorming”, *Journal of Experimental Social Psychology* **41** (2005), 313–320.
- [17] DUNCAN, J., “An adaptive coding model of neural function in prefrontal cortex”, *Nature Reviews: Neuroscience* **2** (2001), 820–829.
- [18] EDELMAN, G.M., and G. TONONI, *A Universe of Consciousness: How Matter Becomes Imagination*, Basic Books (2000).
- [19] FREEMAN, W.J., *How Brains Make Up Their Minds*, Columbia University Press (2000).
- [20] FREEMAN, W.J., “Cortical aperiodic shutter enabling phase transitions at theta rates”, *Proceedings of the International Joint Conference on Neural Networks*, (2007).
- [21] FUSTER, J.M., *Cortex and Mind: Unifying Cognition*, Oxford University Press (2003).

- [22] HEILMAN, K.M., S.E. NADEAU, and D.O. BEVERSDORF, “Creative innovation: Possible brain mechanisms”, *Neurocase* **9** (2003), 369–379.
- [23] IZHIKEVICH, E.M., *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*, MIT Press (2007).
- [24] JUNG-BEEMAN, M., E.M. BOWDEN, J. HEBERMAN, J.L. FRYMAIRE, S. ARAMBEL-LIU, R. GREENBLATT, P.J. REBER, and J. KOUNIOS, “Neural activity when people solve verbal problems with insight”, *PLoS Biology* **2** (2004), 500–510.
- [25] KELSO, J.A.S., *Dynamic Patterns: The Self-Organization of Brain and Behavior*, Cambridge, MA: MIT Press (1995).
- [26] LEGGETT, K.L., “The effectiveness of categorical priming in brainstorming” (1997).
- [27] MILLER, E.K., “The prefrontal cortex and cognitive control”, *Nature Reviews: Neuroscience* **1** (2000), 59–65.
- [28] MOUNTCASTLE, V.B., “The columnar organization of the neocortex”, *Brain* **120** (1997), 701–722.
- [29] MULLEN, B., C. JOHNSON, and E. SALAS, “Productivity loss in brainstorming groups: A meta-analytic integration”, *Basic and Applied Social Psychology* **12** (1991), 3–23.
- [30] NIJSTAD, B.A., and W. STROEBE, “How the group affects the mind: A cognitive model of idea generation in groups”, *Personality and Social Psychology Review* **3** (2006), 186–213.
- [31] NIJSTAD, B.A., W. STROEBE, and H.F.M. LODEWIJKX, “Cognitive stimulation and interference in groups: Exposure effects in an idea generating task”, *Journal of Experimental Social Psychology* **38** (2002), 535–544.
- [32] OSBORN, A.F., *Applied Imagination*, New York: Scribner’s (1957).
- [33] PAULUS, P., and V. BROWN, “Dennis, a.r. and williams, m.l.”, *Group Creativity*, (P. PAULUS AND B. NIJSTAD eds.). Oxford University Press New York (2003), pp. 160–178.
- [34] PAULUS, P., and V. BROWN, “Enhancing ideational creativity in groups: Lessons from research on brainstorming”, *Group Creativity*, (P. PAULUS AND B. NIJSTAD eds.). Oxford University Press New York (2003), pp. 110–136.
- [35] PAULUS, P.B., and M.T. DZINDOLET, “Social influence processes in group brainstorming”, *Journal of Personality and Social Psychology* **64** (1993), 575–586.

- [36] PAULUS, P.B., and H.C. YANG, “Idea generation in groups: A basis for creativity in organizations”, *Organizational Behavior and Human Decision Processes* **82** (2000), 76–87.
- [37] PERLOVSKY, L.I., and R. (eds.) KOZMA, *Neurodynamics of Cognition and Consciousness*, Springer (2007).
- [38] SCHLOSSER, G., and G.P. WAGNER (EDS.), *Modularity in Development and Evolution*, University of Chicago Press, Chicago, IL (2004).
- [39] SCHULZ, W., “Multiple reward signals in the brain”, *Nature Reviews Neuroscience* **1** (2000), 199–207.
- [40] SCHULZ, W., P. DAYAN, and P.R. MONTAGUE, “A neural substrate of prediction and reward”, *Science* **275** (1997), 1593–1599.
- [41] SPIVEY, M., *The Continuity of Mind*, Oxford University Press (2007).
- [42] TONONI, G., “An information integration theory of consciousness”, *BMC Neuroscience* **5** (2004), 42.
- [43] TSUDA, I., “Towards an interpretation of dynamic neural activity in terms of chaotic dynamical systems”, *Behavioral and Brain Sciences* **24** (2001), 793–847.