

Effects of Relevant and Irrelevant Primes on Idea Generation: A Computational Model

Laxmi R. Iyer, Ali A. Minai, Vincent R. Brown, Paul B. Paulus and Simona Doboli

Abstract—Brainstorming is the process of generating ideas in a specific task or problem context. We have previously presented a connectionist framework to study the dynamics of idea generation in individuals. In this paper, we develop this model further, and apply it to studying qualitatively the effects of priming on the process of ideation. Motivated by experimental data from a previous study, we explore the differential effects of relevant and irrelevant primes on productivity of idea generation in specific problem/task contexts. Simulations using our model suggest that even irrelevant primes can provide a modest productivity boost in contexts that are familiar or are similar to familiar contexts, but no benefit when the context is unfamiliar. We propose possible explanations for these results and make predictions for future experiments.

I. INTRODUCTION

Behavioral experiments [1], [2], [3], [4], [5] show that the idea generation process is affected by both social and cognitive processes. Brainstorming models based on associative memory have been developed before [6], [3], [5], and while they account for results observed in behavioral experiments, they provide limited insight into the underlying neural processes. Our group has recently proposed a connectionist model for the dynamics of idea generation [7], [8], [9], [10], [11], building on experimental results from research on the neurobiology of semantic cognition. After being refined through application to simple non-semantic problems [11], the model is now being extended to simulate and explain actual data from behavioral experiments. In this paper, we use the model to investigate experimental data on the effects of cuing with relevant and irrelevant primes during the ideation process.

II. THE IDEA GENERATION PROCESS

Semantic information in the brain is represented at several levels, ranging from combinations of sensorimotor features [12], [13], through amodal concepts [14], [15], to semantic categories [16], [17]. Considerable evidence now supports

Laxmi Iyer is with the Department of Computer Science, University of Cincinnati, Cincinnati OH 45221, Email: iyerlr@email.uc.edu, Ali Minai is with the Department of Electrical and Computer Engineering, University of Cincinnati, Cincinnati OH 45221, Email: Ali.Minai@uc.edu, Vincent Brown is with the Department of Psychology, Hofstra University, Hempstead, NY11549. Email: Vincent.Brown@hofstra.edu, Paul Paulus is with the College of Science, University of Texas Arlington, Arlington TX 76019, Email: paulus@uta.edu, and Simona Doboli is with the Department of Computer Science, Hofstra University, Hempstead, NY11549, Email: Simona.Doboli@hofstra.edu.

Acknowledgement: This project is sponsored by collaborative National Science Foundation Human and Social Dynamics program grants to Simona Doboli (BCS-0729470), Ali Minai (BCS-0728413) and Paul Paulus (BCS-0729305), including funds from the Deputy Director of National Intelligence for Analysis.

the idea that semantic processing involves several cortical functional networks that process and integrate information at all these levels [18], [19], [20], [13]. The areas involved include the left temporal lobe, the prefrontal cortex, the anterior cingulate cortex, the orbitofrontal cortex and parts of the occipital cortex. Regions of the right hemisphere temporal and parietal cortices are also involved [21], and may provide a crucial non-linguistic component needed for the intuitive generation of novel ideas [21], [22], [23]. The flow of information in these cortical networks is controlled by switching processes in the basal ganglia [24], [25], and is modulated by dopaminergic [26], [27] and noradrenergic [28] reflecting judgments of value.

We postulate that concepts comprise the key elements of the semantic space, and that ideas are combinations of concepts that arise through the dynamics of the functional networks linking them with each other and with underlying sensorimotor features, and clustering them into partial categorical groups [9], [10], [11]. The features allow all concepts to be embedded in a single consistent semantic space, while the categorical clustering provides a coarse-grained repertoire of basis functions from which useful context-specific semantic spaces can be constructed [29].

The process of idea generation is seen as an itinerant flow [30] through the space of concepts, where groups of concepts are co-activated by the system's natural dynamics, persist for a time if they are strongly linked together by previous learning, and then dissolve away. Other co-activated groups do not persist but dissolve immediately because their component concepts are not strongly associated, and are not recognized as ideas – except under situations when novel ideas are explicitly being sought. The ideas generated by this itinerant flow are evaluated by an internal critic based on context-specific information, and are reinforced if they meet a quality metric. This process, in combination with the system's inherent connectivity, guides the flow towards more productive regions of the semantic space. Any good ideas generated by the process are also explicitly consolidated into the system, and are associated with the context in which they were found to be relevant. Thus, over time, previously unfamiliar contexts become familiar to the system and ideas that formerly required cognitive effort come more automatically – a well known phenomenon in cognition [31].

We require that the system should be able to do the following:

- Automatically produce familiar ideas in familiar contexts.
- Search efficiently for good ideas in unfamiliar contexts.

- When required, search for novel ideas even in familiar contexts.
- Learn unfamiliar contexts so that they become familiar over time.

In our model, the generation of ideas is seen fundamentally as a *search process* through semantic space. The central principle is that this search should *emerge* from the natural dynamics of the system under appropriate stimulus situations, i.e., when the system is cued with some task/problem context, its natural response should be to search for and generate relevant ideas. Ideally, all aspects of this search process must emerge from the neural dynamics, and should not depend on off-line databases, buffers, or supervisors.

To be effective, the emergent search process must be both systematic and flexible, striking an appropriate and adaptive balance between the *exploitation* of prior knowledge and the *exploration* of new possibilities. Purely random search is clearly not useful, but nor is a rigid, stereotypical search that only follows previously known trajectories. The search should also not get trapped easily in sub-optimal situations, or ignore useful information that is available.

To meet these imperatives, we model the search for ideas as a two-level process:

- **Dynamic Configuration of an Effective Search Domain (ESD):** At this level, the system identifies a subset within the overall semantic space on which to focus its search in the expectation that this subspace is disproportionately dense in good ideas for the given context. This effective search domain (ESD) is instantiated as a biasing projection onto the space of concepts, and emerges dynamically from the system's inherent functionality over the course of the search process. Details of this are discussed in a later section.
- **Itinerant Search within the ESD:** At any point in time, the search for ideas occurs within the concepts included in the ESD. If the ESD is dense in good ideas, the search is much more effective than if it had occurred over the entire semantic space.

The overall effectiveness of the search process thus depends on configuring a productive ESD and on generating good search trajectories within it. Three sources of information shape the search process at both levels:

- Prior learning about semantic relationships, e.g., associations between concepts
- Similarity between the given context and already known contexts.
- Real-time evaluative feedback from the internal critic.

These interact with inherently stochastic activity among concepts and modulated global inhibition among categories (see below). As a result, the search, while broad and stochastic, has a strongly directed systematic component. This component is relaxed adaptively when it fails to yield good ideas, or when novel ideas are required.

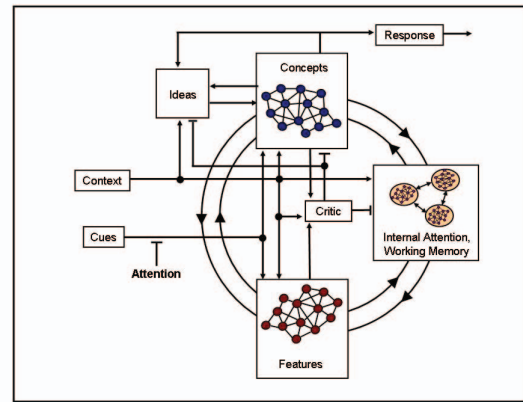


Fig. 1. Architecture of the computational model.

III. MODEL DESCRIPTION

The computational model consists of the following components:

- An *input* or stimulus layer that encodes the contextual information describing the task/problem. This layer provides a stimulus for the context-associated units in the category, idea and feature subsystems.
- A *concept* subsystem that provides an amodal encoding of individual semantic entities termed concepts. The concept units in the concept subsystem are connected to one another recurrently based on their past joint utility. As a result, units that have been usefully coactive in the past tend to become active together.
- A *feature* subsystem, whose units encode sensory and functional attributes of concepts. If activated, a concept unit also activates its associated features. Some features are also associated with the input units, and are activated by specific contexts based on pattern recognition. Features can also be activated by ideas from the ideas subsystem or category units from the category subsystem.
- A *category* subsystem, each of whose units is associated with a salient subset of features. Each category unit is also linked with the concepts that possess the features it encodes, thus creating a cluster of concepts with shared attributes in a subspace of the feature space. Category units are recurrently connected based on learned joint utility, and groups of category units that are strongly connected to one another form modules. In addition to the connections between units, modules have excitatory and inhibitory connections with each other, which are gated by interneurons. Opening or closing these gates allows active modules to merge and switch, dynamically shaping the search space.
- An *ideas* subsystem consisting of units representing ideas known to be useful in specific contexts. Each idea unit is associated with the context patterns for which it is relevant. It is also linked reciprocally with the concept units that are its components, and with features that

characterize the idea, which need not be the union of the features of its constituent concepts. For example, a meal idea comprising a hamburger, non-diet soda ad fries might have the feature “unhealthy” whereas each of its three components only have the feature “high calorie”. Ideas units have recurrent connections with other similar idea units, and activate associated ideas.

- A *critic* that evaluates the ideas produced by the system, and provides a graded reward signal. This produces commensurate reinforcement between units that produced the reward, and association of these units with the goal. The critic is the only part of our system that is not yet implemented as a neural system. Rather, it is a rule-based oracle with complete domain knowledge that participates in the overall system dynamics. In a sense, the task for the ideation system is to discover ideas that are acceptable to the critic using only system dynamics and reinforcement learning. As such, the critic “grounds” the system in the current version.

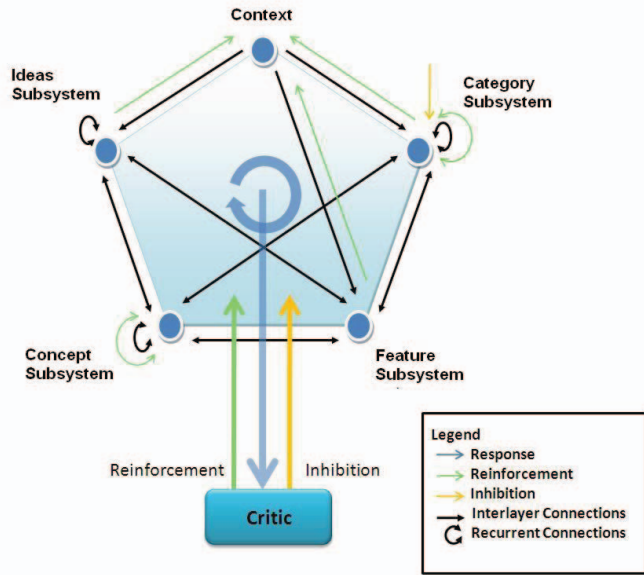


Fig. 2. Connectivity pattern of the computational model.

A. System Dynamics

The search for a productive ESD arises mainly from the dynamics of the category subsystem. It is organized into modules with dense connectivity within the modules and sparse connectivity between them. In addition, there are excitatory and inhibitory gated connections between modules, mediated by interneurons. The input to the category unit i in the category subsystem is:

$$s_i(t) = \zeta_{SG} \sum_j W_{ij}^{SG} y_j^G(t-1) + \zeta_{SS} \sum_j W_{ij}^{SS} y_j^S(t-1) + \zeta_{SC} \sum_j W_{ij}^{SC} y_j^C(t-1) + \zeta_{SF} \sum_j W_{ij}^{SF} y_j^F(t-1) + \lambda_I \sum_B g_{BA}^I Y_A(t) + \lambda_E \sum_B g_{BA}^E Y_A(t) - U^S(t) + N \quad (1)$$

where ζ^{SL} is the gain from layer L where $L \in \{G = \text{context/goal}, S = \text{category layer}\}$ to the category layer, y_j^G , y_j^S , y_j^C and y_j^F are the outputs of unit j of the context, category, concept, and feature layers, respectively, U^S is the inertia in the category subsystem and N is noise.

The third and fourth terms in the equation represent the effect of excitatory and inhibitory gating between modules. λ_I and λ_E are the excitatory and inhibitory gates between modules in the category subsystem and g_{BA}^I and g_{BA}^E are the excitatory and inhibitory gains from module A to module B in the category subsystem respectively, and $Y_A(t) = \sum_w y_w(t-1)$, where $y_w(t)$ is the output from unit $w \in$ module A at time t . When inhibition between two modules is reduced and excitation increases, the modules merge forming a ‘super module’. In contrast, first reducing and then increasing the inhibition causes a switch of activity from one module to the other. These effects are used by the dynamics to continually reconfigure the ESD toward better productivity as described below.

In the concept subsystem, the units are connected recurrently based on previous joint utility in good ideas. The activation equations are similar to those for the category subsystem, except that there is no modular structure or gating between modules. The subsystem does have inertia and noise variables. Activity in the concept subsystem is competitive, with k -of- N units allowed to be active provided their excitation meets a threshold. A unit can only remain active for a limited time of random duration (between 4 and 10 time steps), after which it becomes refractory for a random interval. This ensures that the itinerant dynamics in the concept subsystem maintains variety.

The idea subsystem has equations similar to those of the category and concept subsystems. Idea units are recurrently connected based on similarity, and are associated with contexts in which they were found to be useful. Thus, a given context activates a set of candidate ideas based on previous experience, which then prime more ideas. An itinerant dynamics – again driven by refractoriness in idea units – proceeds and projects to concept units, adding an idea-based dynamic component to the ESD in the concept subsystem.

The feature subsystem does not have recurrent activity, inertia or noise. Its activity pattern provides a passive representation of the current idea at the lowest semantic level, which can be seen as “grounding” the idea. It should be noted that the feature subsystem still participates in the core dynamic loop of the system, and affects the flow of activity in this loop.

Connectivity between the subsystems is summarized in Figure 2.

B. The Search Process

To start off, the external context activates the idea and category subsystems, which bias the concept subsystem to define an initial ESD. The context also activates permanent context-specific features. The input to these features is high enough to keep them on regardless of other activity in the system.

During the initial part of the search, the influence of the context on the category units, of the category and idea units on concepts and features, and of the activated concepts on the features is high, while the other connections are relatively ineffective. As the search proceeds without reward, the global inhibition on the category subsystem is lowered, reducing constraints on the number of category units that are allowed to be active simultaneously. Also, the influence of the concept units on the category and idea units, and the influence of the features on the concept, category and idea units is allowed to go up. Thus, initially, the top down influence, i.e. influence from the context, is the main factor driving the search. As the search gets more relaxed without reward, influence from the internal structure of the system – encoding previously inferred semantic relationships – increases. The net result is an expansion of the ESD by recruiting a larger set of active category units, with category modules merging or switching in and out of the active set.

Once some good ideas begin to be generated, the critic produces reward signals, which have three consequences:

- 1) Every time an idea is rewarded, the active category units associated with its constituent concepts potentiate the connection between themselves. We hypothesize that, in the brain, this potentiation could arise from dopamine-mediated increase in excitability of the category units, which are then fired by the recurrent input from their associated active concept units, leading to Hebbian potentiation.
- 2) When an idea is rewarded, the connections between its constituent concepts are potentiated slightly, making them more likely to be co-active in the future. Over time, this establishes the joint utility-based connectivity of the concept subsystem.
- 3) As reward accumulates, the global inhibition on the category units increases again, reducing the number of active units. However, due to effect (1) above, the category units that survive this are likely to be the ones that jointly produced recently rewarded ideas.

The interplay between effects (1) and (3) above is crucial: The initiation of the search space is based on similarity of the context with previously known ones, its pattern of expansion is based on prior semantic knowledge encoded in the system, and its pattern of narrowing is determined by the real-time rewards produced by the ongoing search. Going back and forth over the expansion-contraction cycle has a ratchet-like (or annealing-like) effect, leading to the configuration of better and better ESDs.

If, however, this process still does not yield success, the system relaxes to an unsystematic stage. The influence of the idea and category subsystems is greatly diminished, which effectively also removes the influence of the context, since that is not directly connected to the concept subsystem. The search proceeds in the concept subsystem alone as a free-association process. Remnant concepts from the previous phase start off the search which is now directed only by the association of concepts, in the presence of noise.

C. Critic

The critic models a thinker's internal judgement of the value of an idea within a context. Ideas are evaluated qualitatively by the critic, and a reward signal is generated.

This evaluation is based on four criteria:

- 1) **Admissibility (A) - does the idea make sense in the given context?** This metric is binary-valued. The system checks if the features of the generated idea satisfy the requirements of the context. If they do, $A = 1$, else 0.
- 2) **Quality (Q) - how good or relevant is the idea in the given context?** This is a number between 0 and 1. For each context a list of desirable features is identified and their contribution to the quality specified. Quality is assigned to the idea based on how many of the concepts in the idea have these features, and to what degree.
- 3) **Novelty (N) - has this idea been seen before?** Since connections between concepts are strengthened when ideas are rewarded, the strengths of connections between the component concepts of an idea provides a natural measure of familiarity. Novelty is calculated from this as a reciprocal quantity, and has a value between 0 and 1.
- 4) **Efficiency (E) - is the idea sufficient without being redundant?** The efficiency measure is the difference of two terms: 1) the percentage of essential elements that are there in the idea, and 2) the percentage of unnecessary components that are included. For example, a **vacation** would consist of a *transportation mode*, a *place to stay*, and some *activities*. The efficiency of an idea would be 1 if all these are present, and no redundant concepts (for example two transportation modes) are included. In the concept subsystem, admissible idea needs an efficiency of 1 to be reinforced. However, if the quality is above 0.8, then the efficiency is less stringent, it has to be above 0.6 for reinforcement.

Reward is generated using the above criteria as follows. First, a *raw reward value* is calculated as:

$$r = A(aQ + bN) \quad (2)$$

where a and b are parameters that can be adjusted based on whether it is the quality or the novelty of the idea that is being sought. This reward is used to adjust inhibition and potentiate the connections in the category subsystem. However, for potentiating the connections in the concept subsystem, an adjusted reward value, r_c , is computed as follows:

- If $Q > 0.8$ and $E > 0.6$, $r_c = r$.
- If $Q \leq 0.8$ and $E = 1$, $r_c = r$.
- Else, $r_c = 0$.

If $r_c = 0$, concept subsystem connections are not modified.

D. Reinforcement

As discussed above, when ideas are rewarded, the units in the different subsystems that were responsible for the reward-winning idea are reinforced. This occurs in two different ways.

- 1) Strengthening of connections between active units in a subsystem - the synaptic strengths between active units that constituted the idea are strengthened, so that the chances of their firing together in future searches is increased. This mode of reinforcement occurs in the category subsystem and the concept subsystem.
 - In the concept subsystem, connections between all active units are reinforced.
 - In the category subsystem, only connections between units that are associated with the active units in the concept subsystem are reinforced.

This reinforcement in the concept subsystem changes weights between concept units as follows:

$$\Delta W_{ij}^C(t) = 0.5r_c G_{CC}(1 - W_{ij}^C(t-1)) \quad (3)$$

$$W_{ij}^C(t) = W_{ij}^C(t-1) + \Delta W_{ij}^C(t) \quad (4)$$

where W_{ij}^C is the strength of the synaptic connection from unit i to unit j in concept subsystem, r_c is the magnitude of reward in the concept subsystem, and G_{CC} is the recurrent gain, which quantifies the influence of the concept subsystem units on themselves.

The weights in the category subsystem are changed as:

$$\Delta W_{ij}^S(t) = 0.5r G_{SC}(1 - W_{ij}^S(t-1)) \quad (5)$$

$$W_{ij}^S(t) = W_{ij}^S(t-1) + \Delta W_{ij}^S(t) \quad (6)$$

where W_{ij}^S is the strength of the synaptic connection from unit i to unit j in category subsystem, r is the magnitude of reward, and G_{SC} is the category to concept gain, which quantifies the influence of the category subsystem on the concept units. The latter is 1 under normal search conditions, and 0 during unsystematic search, in which case the ideas generated are not influenced by the category units and should not affect synaptic modification among them.

- 2) Strengthening of connections between active units and the current context - connections between active units and the current context input are strengthened, so that these units are associated with that context for the future. When the system receives the same context input again, these units can now be directly turned on by it. This mode of reinforcement occurs in the category, idea and feature subsystems.

The formula for the change of weights is:

$$\Delta W_{gx}^{GX}(t) = 0.5r \sum_{c' \in C^*} W_{xc'}^{XC} G_{XC}(1 - W_{ij}^X(t-1)) \quad (7)$$

$$W_{gx}^{GX}(t) = W_{gx}^{GX}(t-1) + \Delta W_{gx}^{GX}(t) \quad (8)$$

where W_{gx}^{GX} is the strength of the synaptic connection from input unit g to unit x in subsystem X where $X \in \{S, I, F\}$, $W_{xc'}^{XC}$ is the strength of the synaptic connection from unit x in subsystem X to unit c' in the concept subsystem, C^* is the set of all active concepts, r is the magnitude of reward, and G_{XC} is the gain from subsystem X to the concept subsystem.

In the feature subsystem, the feature units that are turned on often enough by the active concepts are rewarded.

The formula of reward for the feature unit f is:

$$r_f = r \sum_{c' \in C^*} W_{c'f}^{CF} \quad (9)$$

$$r_f^* = \begin{cases} 1, & \text{if } r_f > 0 \\ 0, & \text{otherwise;} \end{cases} \quad (10)$$

where $W_{c'f}^{CF}$ is the strength of the synaptic connection from unit c' in the concept subsystem to unit f in the feature subsystem, and C^* is the set of all active concepts.

The formula for reinforcement for the feature units is:

$$W_{gf}^{GF}(t) = \alpha r_f(t-1) + W_{gf}^{GF}(t-1) - \beta(1 - r_f^*) \quad (11)$$

where $W_{gf}^{GF}(t)$ is the strength of the synaptic connection from input unit g to the feature unit f at time t . α , the rate at which reinforcement is increased is much higher than β , the rate at which reinforcement is decreased.

In addition, the inhibition increases in the category subsystem based on cumulative reward as:

$$\Delta i(t) = 0.8(r - i(t-1)) \quad (12)$$

$$i(t) = i(t-1) + \Delta i(t) \quad (13)$$

where $r(t)$ is the magnitude of reward at time t , Δi is the increase in inhibition, and i is the inhibition. Thus, the increase in inhibition is less if the inhibition is already high, so as to not quickly zero in on the rewarded attractor unless the reward is high enough to do so

E. Injection of Variety into the Search

A key requirement for an effective search process is inherent variation, i.e., the tendency not to consider the same idea repeatedly. In our model, there are three processes that ensure this inherent variation:

- 1) When a concept unit is active for several consecutive time steps, it switches off, and then remains off for a certain refractory period. Both the duration of activity – which can be seen as a burst – and the duration of the refractory period vary randomly within a range. This generates an itinerant trajectory through the semantic space that explores a large number of potential ideas within the ESD.
- 2) When two concept units are active together over a longer period of time, their connections temporarily diminish in strength, reflecting a habituation phenomenon. This allows weaker associations to exercise their influence, and generate more novel conceptual combinations.
- 3) Over a much longer period of time, each concept unit builds up inertia, reflecting a decrease in excitability due to extended activity. Thus, even if it gets more input, the inertia makes it less likely to fire, allowing fresh units to win the competition for activation. In this way, the system is able to look at concepts that have not been considered before.

IV. SIMULATIONS

The model described above has been instantiated in an artificial environment comprising three semantic domains - vacation, meals and furnishings. Contexts are defined as specific goals or tasks within these domains. There are 118 concepts, 80 features, 30 categories and 24 pre-existing ideas in the system. In addition, the system has 19 familiar contexts that it is assumed to have learned.

A recent set of experiments in our group studies the effects of priming on ideation [32]. In the experiment, individual human participants are given a specific task and asked to generate as many unique relevant ideas as possible in a session of fixed duration. During the course of the session, the brainstormer is provided ideas from an audiotape as cues. Two types of cuing (or priming) ideas are considered: 1) Ideas relevant to the given task; and 2) Ideas of which only a fraction are relevant to the given task. Two experimental protocols are used: 1) A small number of cues is provided; and 2) A large number is provided. Data is also obtained on a control case without cues. These experiments are designed to mimic real brainstorming in groups where participants often spend part of their time discussing – or being distracted by – irrelevant ideas.

Figure 3 shows the results of the experiment [33]. As expected, more priming produces more ideas, and priming with entirely relevant cues is more helpful than priming with irrelevant cues included. This result is not surprising in itself, but demonstrates the distracting effects of irrelevant cues in group interactions. However, this leaves open the question of

what effect irrelevant cues alone might have on the ideation process in novel contexts. We used our model to explore this issue and generate predictions for future experiments.




Fig. 3. Experimental data on the effect of priming with relevant and irrelevant cues on the generation of ideas (redrawn from [34]).

We considered three distinct situations for our study: 1) When the given context is similar to a familiar context and requires similar ideas; 2) When the novel context is only slightly similar to a familiar context and requires substantially different ideas; and 3) When the novel context has no similarity to any familiar context.

For the first situation, we used the context *Romantic Vacation*, which was similar to the familiar context *Physically Active Vacation*. This is intended to simulate the fact that, in most brainstorming situations, the brainstormer is at least somewhat familiar with the problem at hand. Each trial in the simulation consisted of 1000 concept time steps (a total of 1000 conceptual combinations were produced). For each task, 5 trials were run - a control where there were no cues, and the other four where relevant and irrelevant cues were provided to the system at both high (10 cues/trial) and low (3 cues/trial) frequencies. Each cue consisted of a combination of concepts (typically 5), and was presented to the system by directly activating the corresponding concept units. The strength of this forced activation was substantially higher than the average input into the concept units during normal activity to ensure that the cue was “received”. Once activated, a cue stayed on for five time steps, after which, it turned off, reflecting the fact that cues in the actual experiments probably persist in working memory only for a finite time.

The results of the simulation, shown in Figure 4, shows that in this situation: 1) A large number of good ideas are generated; 2) Relevant primes increase productivity over the unprimed case; and 3) Even irrelevant primes can increase the number of good ideas generated, albeit less than relevant primes. Though all these results are consistent with the experimental data the last result is the most interesting and somewhat counter-intuitive. We conjecture that irrelevant primes improve productivity because, while they are not

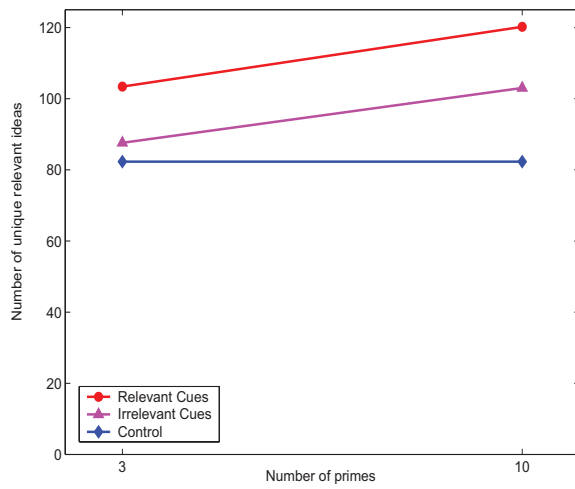


Fig. 4. Simulation results for the effect of relevant and irrelevant cues on idea generation in a novel context (Romantic Vacation) that is quite similar to a familiar one (Physically Active Vacation) Note that, as in the experimental data, both types of cues increase productivity.




Fig. 6. Simulation results for a completely novel context. Only a small number of unique relevant ideas is generated, and a small number of irrelevant cues only provide a modest boost.

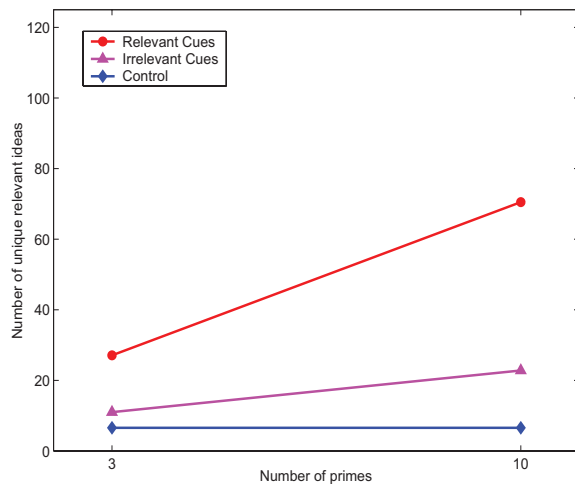


Fig. 5. Simulation results for a novel context with low similarity to any familiar context. A few irrelevant cues are not enough to produce much improvement, but a large number do so.

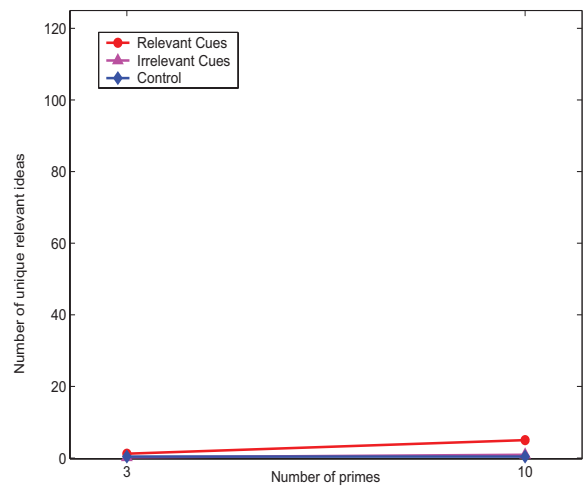


Fig. 7. Simulation results for another completely novel context. Again, only a small number of unique relevant ideas is generated, and a small number of irrelevant cues provides no boost.

good ideas, they often include concepts that are part of other good ideas. This partial, accidental overlap allows the irrelevant cues to initially *exploit* the semantic structure already available for the similar familiar context. This enables the plasticity mechanisms in the search process to find productive categories for the novel context faster than they would have been found without priming, though not as efficiently as with relevant primes, which can leverage the existing structures more systematically.

To build further upon this, we considered what would happen as the novel context became less similar to any that the system knows about. In Figure 5, the novel context *Light Meal on Relaxing Vacation* is quite dissimilar to the closest familiar one *Fast Food*. The results show that irrelevant cues provide only a minimal boost in productivity – presumably because the concepts primed by these cues are not relevant to

the context at hand. Relevant cues still provide a substantial boost, though not as much as in the case of a familiar context.

In Figures 6 and 7, we give the system two completely unfamiliar contexts for which it has no prior experience. In both cases, there is no benefit from irrelevant cues even when the number of such cues is large. This is because, without any prior information about semantic relationships in the given contexts, irrelevant cues are simply noise. Such noise can be useful when it excites semantic structures that are already latent in the system – as with novel contexts similar to familiar ones (Figures 4), but not when such latent structure does not exist.

It is also noticeable that when the given context is not similar to a familiar one, the total number of ideas generated drops significantly even when the cues are relevant. Basically, the relevant cues generate some relevant ideas by directly exciting ideas very similar to themselves, but in the absence

of any semantic structure relevant to the given contexts, this boost does not propagate further. In contrast, for contexts that are similar to familiar ones, each relevant cue *ignites* a whole chain of new relevant ideas by unmasking the relationships already latent in the semantic structure. In a real sense, the cues serve a true priming function only when they *resonate* with existing semantic relationships embedded in the system. Otherwise, they simply act as local disturbances that may shake loose a few ideas, then disappear.

V. CONCLUSION

The detailed mechanisms by which priming affects ideation productivity are still obscure, but our model provides some clues about this, and suggests systematic ways to build on these clues. In particular, we are investigating how the structure and content of the category space – e.g., the presence or absence of category units combining unusual groups of features – affects productivity. It has been hypothesized [35], [23] that especially creative individuals are characterized by intuitive “short-cut” linkages between features, concept or categories that are typically not linked together – perhaps as a result of the involvement of right hemisphere brain regions. Our model provides a concrete way to test this hypothesis, and others in the same vein. Those results will be reported in future papers.

REFERENCES

- [1] P. Paulus and M. Dzindolet, “Social influence processes in group brainstorming,” *Journal of Personality and Social Psychology*, vol. 64, pp. 575–586, 1993.
- [2] H. Coskun, Paulus, V. P.B., Brown, and J. Sherwood, “Cognitive stimulation and problem presentation in idea generation groups,” *Group Dynamics: Theory, Research, and Practice*, vol. 4, pp. 307–329, 2000.
- [3] P. Paulus and V. Brown, “Enhancing ideational creativity in groups: Lessons from research on brainstorming,” in *Group Creativity*, P. Paulus and B. Nijstad, Eds. New York: Oxford University Press, 2003, pp. 110–136.
- [4] K. Dugosh and P. Paulus, “Cognitive and social comparison processes in brainstorming,” *Journal of Experimental Social Psychology*, vol. 41, pp. 313–320, 2005.
- [5] B. Nijstad and W. Stroebe, “How the group affects the mind: A cognitive model of idea generation in groups,” *Personality and Social Psychology Review*, vol. 3, pp. 186–213, 2006.
- [6] V. Brown, M. Tumeo, T. Larey, and P. Paulus, “Modeling cognitive interactions during group brainstorming,” *Small Group Research*, vol. 29, pp. 495–526, 1998.
- [7] S. Diboldi and A. Minai, “A neural network model of creative idea generation in groups,” in *In Proceedings of the Ninth International Conference on Cognitive and Neural Systems, Boston, 2005*.
- [8] V. Brown and S. Diboldi, “A neural network simulation of interactive group brainstorming,” in *Eighteenth Annual Convention of the Association for Psychological Science*.
- [9] S. Diboldi, A. Minai, and V. Brown, “Adaptive dynamic modularity in a connectionist model of context-dependent idea generation,” in *Proceedings of the International Conference on Neural Networks, Orlando, Florida, 2007*, pp. 2183–2188.
- [10] L. Iyer, A. Minai, S. Diboldi, and V. Brown, “Modularity and self-organized functional architectures in the brain,” in *Proceedings of the 7th International Conference on Complex Systems, Boston, MA, 2007*.
- [11] —, “Modeling exploration and exploitation in creative idea generation,” in *In Proceedings of the 19th Midwest Artificial Intelligence and Cognitive Science Conference, 2008*.
- [12] E. Warrington and T. Shallice, “Category specific semantic impairments,” *Brain*, vol. 107, pp. 829–854, 1984.
- [13] A. Martin, “The representation of object concepts in the brain,” *Annual Review of Psychology*, vol. 58, pp. 25–45, 2007.
- [14] M. Kellenbach, M. Brett, and K. Patterson, “Large, colourful or noisy? attribute- and modality-specific activations during retrieval of perceptual attribute knowledge,” *Cognitive, Affective and Behavioral Neuroscience*, vol. 1, pp. 207–221, 2001.
- [15] K. Patterson, P. Nestor, and T. Rogers, “Where do you know what you know? the representation of semantic knowledge in the human brain,” *Nature Rev. Neurosci.*, vol. 8, pp. 976–987, 2007.
- [16] A. Caramazza and J. Shelton, “Domain-specific knowledge systems in the brain: the animate-inanimate distinction,” *Journal of Cognitive Neuroscience*, vol. 10, pp. 1–34, 1998.
- [17] A. Caramazza and B. Mahon, “The organization of conceptual knowledge: The evidence from category-specific semantic deficits,” *Trends in Cognitive Sciences*, vol. 7, pp. 354–361, 2003.
- [18] A. Damasio, “The brain binds entities and events by multiregional activations from convergence zones,” *Neural Computation*, vol. 1, pp. 123–132, 1989.
- [19] H. Damasio, T. Grabowski, D. Tranel, R. Hichwa, and A. Damasio, “A neural basis for lexical retrieval,” *Nature*, vol. 380, pp. 499–505, 1996.
- [20] H. Damasio, D. Tranel, T. Grabowski, R. Adolphs, and A. Damasio, “Neural systems behind word and concept retrieval,” *Cognition*, vol. 92, pp. 179–229, 2004.
- [21] E. Bowden, M. Jung-Beeman, J. Fleck, and J. Kounios, “New approaches to demystifying insight,” *Trends in Cognitive Sciences*, vol. 9, pp. 322–328, 2005.
- [22] K. Heilman, S. Nadeau, and D. Beversdorf, “Creative innovation: possible brain mechanisms,” *Neurocase*, vol. 9, pp. 369–379, 2003.
- [23] W. Duch, “Intuition, insight, imagination and creativity,” *IEEE Comput. Intell.*, pp. 40–52, 2007.
- [24] A. Graybiel, “Building action repertoires: memory and learning functions of the basal ganglia,” *Current Opinion in Neurobiology*, vol. 5, pp. 733–741, 1995.
- [25] J. Houk, “Agents of the mind,” *Biol. Cybern.*, vol. 92, pp. 427–437, 2005.
- [26] W. Schulz, “Multiple reward signals in the brain,” *Nature Reviews Neuroscience*, vol. 1, pp. 199–207, 2000.
- [27] P. Apicella, “Leading tonically active neurons of the striatum from reward detection to context recognition,” *Trends in Neuroscience*, vol. 30, pp. 299–306, 2007.
- [28] C. Aston-Jones and J. Cohen, “An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance,” *Ann. Rev. Neurosci.*, vol. 28, pp. 403–450, 2005.
- [29] A. Minai, L. Iyer, D. Padur, and S. Diboldi, “A dynamic connectionist model of idea generation,” in *submitted to IJCNN 2009*.
- [30] I. Tsuda, “Towards an interpretation of dynamic neural activity in terms of chaotic dynamical systems,” *Behav. Brain Sci.*, vol. 24, pp. 793–847, 2001.
- [31] S. Dehaene, 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*, vol. 95, pp. 14 529–14 534, 1998.
- [32] K. Dugosh, P. Paulus, E. Roland, and H.-C. Yang, “Cognitive stimulation in brainstorming,” *Journal of Personality and Social Psychology*, vol. 79, pp. 722–735, 2000.
- [33] E. Roland, “The effects of stimulation and distraction in brainstorming,” *Unpublished master’s thesis, University of Texas, Arlington*.
- [34] V. Brown, J. Sherwood, P. Paulus, P. Dossett, K. Dugosh, and E. Roland, “Predicting semantic priming effects in group brainstorming: A cognitive modeling approach,” *Preprint*.
- [35] M. Schilling, “A small-world network model of cognitive insight,” *Creativity Res. J.*, vol. 17, pp. 131–154, 2005.