

Knowledge Integration in Creative Problem Solving

Sébastien Hélie (helies@rpi.edu)

Rensselaer Polytechnic Institute, Cognitive Science Department, 110 Eighth Street,
Troy, NY 12180 USA

Ron Sun (rsun@rpi.edu)

Rensselaer Polytechnic Institute, Cognitive Science Department, 110 Eighth Street,
Troy, NY 12180 USA

Abstract

Most psychological theories of problem solving have focused on modeling explicit processes that gradually bring the solver closer to the solution in a mostly explicit and deliberative way. This approach to problem solving is typically inefficient when the problem is too complex, ill-understood, or ambiguous. In such a case, a ‘creative’ approach to problem solving might be more appropriate. In the present paper, we propose a computational psychological model implementing the Explicit-Implicit Interaction theory of creative problem solving that involves integrating the results of implicit and explicit processing. In this paper, the new model is used to simulate insight in creative problem solving and the overshadowing effect.

Keywords: Creative problem solving, insight, problem solving, psychological modeling, creativity.

Introduction

Many psychological theories of problem solving have emphasized a role for implicit cognitive processes (e.g., Evans, 2006; Reber, 1989; Sun, 1995). For instance, implicit processes are often thought to generate hypotheses that are explicitly tested (Evans, 2006). Also, feature similarity among concepts has been shown to affect rule-based reasoning through processes that are mostly implicit (Sun, 1995). Yet, most psychological theories of problem solving have focused on modeling explicit processes that gradually bring the solver closer to the solution in a mostly “rational”, deliberative way (Martindale, 1995). However, when the problem to be solved involves an ambiguous, ill-understood, or overly complex state space that is difficult to search, the solution is usually found by ‘insight’ (Durso, Rea, & Dayton, 1994; Pols, 2002) and regular problem solving models are for the most part unable to account for this apparent absence of “rational” strategy.

A complementary research effort in psychology, creative problem solving, has tried to tackle ‘insight’ problems (Pols, 2002). However, the lack of detailed mechanistic theories and computational models of creative problem solving in psychology has resulted in a limited impact on the field of problem solving in general. In contrast, several computational models of creative problem solving have been proposed in computer science (for a review, see Rowe & Partridge, 1993). However, they are not psychologically motivated and have not been applied to any fine-grained simulation of psychological data. Hence, their psychological

relevance has not been established. Here, we attempt to fill this gap by proposing a detailed computational model implementing the Explicit-Implicit Interaction theory of creative problem solving, a psychological theory developed earlier in Hélie & Sun (submitted). This theory, and the proposed model, focus on the beneficial effect of integrating the results of both implicit and explicit processing (Evan 2006; Sun, Slusarz, & Terry, 2005). The new model is used here to simulate insight in problem solving (Durso et al., 1994) and the overshadowing effect (Schooler, Ohlsson, & Brooks, 1993) found in the psychological literature.

Creative Problem Solving

Human creative problem solving is a four-stage process (Wallas, 1926). The first stage, preparation, is the accumulation of knowledge that can allow directly solving the problem at hand (i.e., regular problem solving). However, this approach to problem solving sometimes fails to provide a satisfactory solution and an impasse is reached. Surprisingly, as shown by numerous psychological experiments, the problem is more likely to be later solved if it is put aside for a while (Smith & Dodds, 1999). This phenomenon (which has also been the subject of countless anecdotes in the history of science) is what Wallas called ‘incubation’ (the second stage). Because the solution often comes as a surprise following incubation, the next stage was called illumination (or ‘insight’). During this third stage, the problem solver may feel that the solution completely elucidates the problem. However, the true quality of the solution remains to be evaluated by the fourth stage: the validation process.

When a problem is clearly defined and can be solved using regular problem solving strategies, only the first and fourth stages are involved (preparation and verification). However, when the problem is unclear (e.g., ambiguous, overly complex, or difficult to search), all the stages are likely to be required. Many psychological theories have been proposed to individually explain each stage of creative problem solving (for reviews, see Pols, 2002; Smith & Dodds, 1999). However, prior to the Explicit-Implicit Interaction theory, none proposed an integrated explanation for all the stages.

The Explicit-Implicit Interaction Theory The Explicit-Implicit Interaction theory is a psychological theory that relies mainly on the following principles: (1) there are two

types of knowledge, namely explicit and implicit (a well established distinction in psychology: see, e.g., Reber, 1989; Seger, 1994), that are simultaneously involved in most tasks (Sun et al., 2005). (2) Explicit and implicit knowledge are often “redundant” (Sun, Merrill, & Peterson, 2001), and the results of explicit and implicit processing are often integrated to provide a response (Sun et al., 2005). In psychology, implicit knowledge refers to knowledge that affects behavior without the awareness of the participant while explicit knowledge may (potentially) come with the accompanying awareness (Reber, 1989). (For details and justifications of the principles, see Hélie & Sun, submitted; Sun et al., 2005).

In the Explicit-Implicit Interaction theory, the preparation and verification stages of problem solving are mostly captured by explicit processing while implicit processing corresponds mostly to incubation. Wallas’ (1926) third stage, insight, involves an “explicitation” process and consequently produces a solution in an explicit form. In addition, the Explicit-Implicit Interaction theory assumes the presence of activation thresholds that control the internal confidence level necessary to come up with a solution to the problem. If a threshold is not crossed after the results of explicit and implicit processing have been integrated, another iteration of processing is performed. This simple yet powerful iterative theory has been formalized into a connectionist model that was used to simulate a variety of tasks. The model, Explicit-Implicit Interaction with Bayes Factor (EII-BF), is described in the following section.

The EII-BF Model

The general structure of the model is shown in Figure 1. The model is composed of two major modules, representing explicit and implicit knowledge respectively. These two modules are connected by using bidirectional associative memories (i.e., the **E** and **F** weight matrices; Kosko, 1988). In each trial, the task is simultaneously processed in both modules, and the outputs (response activations) are integrated in order to determine a response distribution. Once this distribution is specified, a response is stochastically chosen and the Bayes factor (Kass & Raftery, 1995) is computed (using the two most likely responses). If this measure is higher than a predefined threshold, the chosen response is output; otherwise, another iteration of processing is done in both modules, using the chosen response as the input.

In EII-BF, explicit knowledge is captured using a two-layer linear connectionist network while implicit knowledge is captured using a non-linear attractor neural network (*NDRAM*: Chartier & Proulx, 2005). The inaccessible nature of implicit knowledge may be captured by distributed representations in an attractor neural network, because units in a distributed representation are capable of accomplishing tasks but are less individually meaningful (Reber, 1989). This characteristic corresponds well with the relative inaccessibility of implicit knowledge (Sun et al., 2005). In contrast, explicit knowledge may be captured in

computational modeling by localist representations, because each unit in a localist representation is more easily interpretable and has a clearer conceptual meaning. This characteristic captures the property of explicit knowledge being more accessible and manipulable (Sun et al., 2005). This difference in the representations of the two types of knowledge leads to a dual-representation, dual-process model. Theoretical arguments for such models are presented in Evans (2006), Reber (1989), Sun (1995), and Sun et al. (2005).

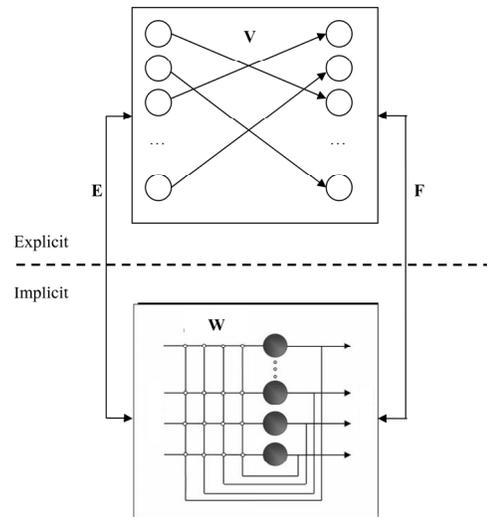


Figure 1: General architecture of the EII-BF model. The top level contains explicit knowledge while the bottom level contains implicit knowledge. Upper-case letters (**V**, **W**, **E**, **F**) represent the weight matrices.

Specifically, explicit knowledge is localistically represented in the top level using binary activation. The left layer in Figure 1 (denoted **x**) is composed of *n* units while the right layer (denoted **y**) is composed of *m* units. These layers are connected using the binary weight matrix **V**, and the information is transmitted using the standard dot product (i.e., $\mathbf{y} = \mathbf{V}\mathbf{x}$).¹

In the bottom level, implicit knowledge is represented using *r* bipolar units (denoted **z**). Precisely, $\mathbf{z} = \mathbf{t}_1 + \mathbf{t}_2$, where **t**₁ represents the first *s* units in **z**, which are connected to the left layer in the top level using the **E** weight matrix. Meanwhile, **t**₂ represents the remaining *r* – *s* units in **z**, which are connected to the right layer in the top level using weight matrix **F**. In words, the **E** and **F** weight matrices are used to ‘translate’ explicit knowledge into implicit knowledge (i.e., ‘implicitation’) and vice-versa (i.e., ‘explicitation’).

¹ In EII-BF, all the weight matrices are learned using Hebbian learning. This type of learning has the advantage of psychological and biological plausibility. The **V**, **E**, and **F** weight matrices were trained using regular Hebbian learning (i.e., the outer matrix product). The bottom-level weight matrix (**W**) was trained using a contrastive Hebbian learning rule (Chartier & Proulx, 2005).

Bottom-level activation (\mathbf{z}) is modified through a settling process using the NDRAM transmission rule (Chartier & Proulx, 2005):

$$\mathbf{z}_{[t+1]} = f(\mathbf{W}\mathbf{z}_{[t]}), \quad f(z_i) = \begin{cases} +1 & , z_i > 1 \\ (\delta+1)z_i - \delta z_i^3 & , -1 \leq z_i \leq 1 \\ -1 & , z_i < -1 \end{cases}$$

where $\mathbf{z}_{[t]} = \{z_1, z_2, \dots, z_r\}$ is the bottom-level activation after t iterations in the network, \mathbf{W} is the bottom-level weight matrix, and $0 < \delta < 0.5$ is the slope of the transmission function. This settling process amounts to a search through a soft constraint satisfaction process, where each connection represents a constraint and the weights represent the importance of the constraints (Hopfield & Tank, 1985). Note that it was estimated psychologically that each iteration in this network takes roughly 350 ms of psychological time (for justifications, see Hélie & Sun, submitted).

Once the response activations have been computed in both levels, they are integrated using the *Max* function:

$$o_i = \text{Max} \left[y_i, \lambda \sum_{j=1}^r f_{ij} z_j \right]$$

where $\mathbf{o} = \{o_1, o_2, \dots, o_m\}$ is the integrated response activation, $\mathbf{y} = \{y_1, y_2, \dots, y_m\}$ is the result of top-level processing, λ is a scaling parameter, and $\mathbf{F} = [f_{ij}]$ is a weight matrix. The integrated response activation is then transformed into the response distribution:

$$P(o_i) = e^{o_i/\alpha} \left(\sum_j e^{o_j/\alpha} \right)^{-1}$$

where α is a noise parameter (i.e., the temperature). Note that low noise levels tend to exaggerate the probability differences, which lead to a narrow search of possible responses and favors stereotypical responses. In contrast, high noise levels tend to minimize the probability differences, which leads to a more complete search of the response space.

Table 1: Algorithm of the EII-BF model

- | |
|---|
| <ol style="list-style-type: none"> 1. Observe the current state of the environment; 2. Compute the response activations; 3. Compute the integrated response activation and the resulting response distribution; 4. Stochastically choose a response and compute the Bayes factor: <ol style="list-style-type: none"> a. If the Bayes factor is higher than ψ, output the response; b. Else, if there is time remaining, go back to step 2. |
|---|

A response is stochastically chosen and the Bayes factor (i.e., the ratio of the probabilities of the two most likely responses) is computed. This measure represents the relative

support for the most likely response. According to Kass & Raftery (1995), Bayes factors smaller than 3.2 are not worth mentioning, while Bayes factors higher than 10 are considered strong evidence (and Bayes factors higher than 100 are *decisive*). In the current model, the chosen response is output if the Bayes factor is higher than a free parameter (ψ); otherwise, the search process continues with a new iteration using the chosen response to activate the left layer ($\mathbf{x} = \mathbf{V}^T \mathbf{o}$; $\mathbf{z} = \mathbf{E}\mathbf{x}$). The algorithm specifying the complete process is summarized in Table 1.

Insight in Creative Problem Solving

Many theories of insight assume that this phenomenon is the consequence of some form of knowledge restructuring (as reviewed in Pols, 2002). Because knowledge is often modeled using knowledge graphs, Durso et al. (1994) hypothesized that insight could be observed by constructing and comparing participants' knowledge graphs before and after insight had occurred. To test this hypothesis, the participants were asked to explain the following story: "A man walks into a bar and asks for a glass of water. The bartender points a shotgun at the man. The man says 'thank you', and walks out."

The participants' task was to explain why the sight of the shotgun seemed to have satisfactorily replaced the man's need for a glass of water (i.e., the fact that he had the hiccup). To explain this substitution, the participants had two hours to ask "yes – no" questions to the experimenter. After this questioning period, the participants were split into two groups (solvers and non-solvers) and asked to rate the relatedness of pairs of concepts. These ratings were used to build the solvers' and non-solvers' knowledge graphs using a scaling algorithm (see Figure 1 in Durso et al., 1994).

The solvers' knowledge graph differed from the non-solvers' by twelve edges. These differences reflected a shift of the graph focal points (i.e., center and median) from 'Bartender' to 'Relieved'. Furthermore, the correlation between the two graphs was essentially zero (Durso et al., 1994).

Simulation Setup

In this task, the Explicit-Implicit Interaction theory posits that participants' hypotheses are generated implicitly and evaluated explicitly (Hélie & Sun, submitted; see also Evans, 2006). Hence, in EII-BF, hypotheses were generated by the bottom level and sent to the top level.

More precisely, each concept in the story was represented by a separate unit in each layer of the top level of EII-BF ($n = m = 14$). The non-solvers' graph was assumed to represent common prior knowledge and was pre-coded in the top level of EII-BF (the \mathbf{V} weight matrix). In the bottom level, each concept was associated to a random set of ten units ($r = 140$, $s = 0$). A distributed representation was randomly generated to represent each top-level unit in the bottom level. Furthermore, the non-solvers' graph was redundantly coded as implicit associations in the bottom level by

training the attractor neural network with the distributed representations.²

Hypothesis generation was initiated by random activation in the bottom level of EII-BF. This activation was processed by the attractor neural network until the convergence of the network. The resulting stable state was sent bottom-up to activate the right layer to compute the integrated response activation ($\lambda = \infty$) and the response distribution. A response was stochastically chosen to activate the left layer ($\mathbf{x} = \mathbf{V}^T \mathbf{o}$). The left layer was then also transformed into a Boltzmann distribution, and a response in that layer was stochastically chosen. If the Bayes factor was high enough ($\psi = 10$), a question was put to the simulated experimenter concerning the existence of a connection between the two activated responses. If a connection was present in the solvers' graph (see Figure 1 in Durso et al., 1994), a 'yes' answer was given by the simulated experimenter and a connection was added into the weight matrix \mathbf{V} (representing explicit knowledge); otherwise, the corresponding connection was removed from the \mathbf{V} weight matrix. If the Bayes factor was too small to output a question to the experimenter, the explicit knowledge remained unchanged.

In all cases, the activation present in the left layer was sent top-down for another iteration of processing in the bottom level ($\mathbf{z} = \mathbf{E}\mathbf{x}$). This ongoing process ended if the \mathbf{V} weight matrix was identical to the adjacency matrix of the solvers' graph or 20 571 iterations had occurred in the bottom level (because $20\,571 \times 350\text{ ms} \approx 2\text{h}$).

This simulation can also be used to test another prediction made by many insight theories: A more diffused search in memory is more likely to yield an insight solution (e.g., Martindale, 1995; Pols, 2002). This phenomenon is modeled in EII-BF by the noise parameter in the construction of the response distribution (α). The present simulation was also used to test the adequacy of this modeling approach. Hence, noise was varied between 10^{-2} and 10^5 (with an increment of 1 for the exponent). 100 simulations were run with each of these noise levels.

Simulation Results

The mean performance by noise level is shown in Figure 2. As can be seen, EII-BF was generally able to modify its explicit representations by simulating "yes / no" questions to the experimenter. A between-subject analysis of variance on the noise levels was highly significant $F(7, 792) = 1336.77$, $p < .01$, showing that the number of edges differing between the responses proposed by EII-BF and the solvers' graph was negatively related to noise level. More precisely, Tukey post hoc analyses showed that low noise levels resulted in poor performance, and that each increment between 10^{-2} and 10^0 significantly improved the performance of the model ($p < .01$). From that point on,

increasing the noise level did not significantly improve the performance of the model ($p > .01$). Overall, the noise parameter in EII-BF changes the probability of correct responses in insight problems, which is in line with associationist theories of insight (e.g., reviewed in Pols, 2002).

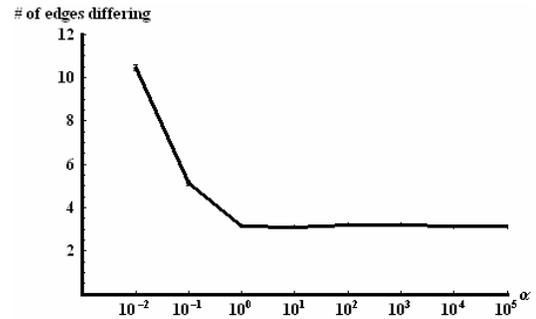


Figure 2: Mean number of edges differing between EII-BF's explicit knowledge graph and the solvers' graph in Durso et al. (1994).

Discussion

The preceding analysis showed that the performance of EII-BF improved when the noise level was increased. This result is consistent with evolutionist theories of creativity (Martindale, 1995), which equate creativity to a more thorough exploration of the response space because infrequent responses are responsible for producing creative insight. This can be captured in artificial neural networks using a noise parameter to sample infrequent responses.

Overshadowing in Insight Problem Solving

The implicit learning literature (e.g., Reber, 1989; Seger, 1994; Sun et al., 2001) has shown that explicitly looking for rules and regularities can impair performance when the participants are unable to find them. This *overshadowing* effect of explicit thinking on implicit learning is robust and also affects performance in insight problem solving (Schooler et al., 1993). According to the Explicit-Implicit Interaction theory (Hélie & Sun, submitted), this experimental phenomenon follows from the inhibition of the incubation phase (mostly implicit processing) in favor of overly extended preparation and verification phases (mostly explicit processing).

The presence of overshadowing has been shown using many "insight" problems (Schooler et al., 1993). For instance: "A dealer of antique coins got an offer to buy a beautiful bronze coin. The coin had an emperor's head on one side and the date 544 B.C. stamped on the other. The dealer examined the coin, but instead of buying it, he called the police. Why?"

The participants had two minutes to explain this problem initially. Following this initial period, the participants were interrupted and half were assigned to an unrelated task while the remaining were asked to verbalize their problem solving

² 100 epochs, $\eta = 0.001$, $\zeta = 0.9999$, $\delta = 0.4$, and $p = 3$. Note that these are the default values in the NDRAM model. For a detailed explanation of the parameters, see Chartier and Proulx (2005).

strategies. In both cases, this interruption lasted 90 seconds and was followed by another four-minute attempt to explain the problem. The dependent variable was the proportion of explained problems.

The results were as shown in Figure 3 (grey bars). After the participants spent the interruption period working on an unrelated task, 45.8% of the insight problems were correctly explained. In contrast, only 35.6% of the problems were correctly explained after the participants had to verbalize their problem solving strategies. According to Schooler et al. (1993), this statistically significant difference, $t(74) = 2.13, p < .05$, was the consequence of an overly verbal mode of problem solving, which prevented the participants from reaching an insight.

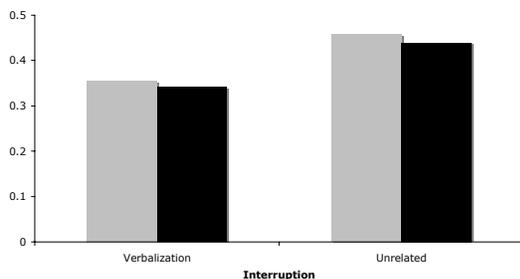


Figure 3: The grey bars represent the proportion of problems solved by humans (Schooler et al., 1993) and the black bars represent the proportion of correct outputs by EII-BF.

Simulation Setup

Each of the features of the problem (i.e., coin material, carved pattern, coin date, dealer's decision) was represented by two units in the left layer of the top level of EII-BF ($n = 8$): the first two units represented the date (good, bad), the two following units represented the material (good, bad), units five and six represented the carved pattern (good, bad), and the last two units represented the dealer's decision (buy, don't buy). In the right layer, eight responses (abstract explanations) were locally represented ($m = 8$).³ The dependency between the dealer's decision and the features of the coin was pre-coded into the top level of EII-BF (i.e., the V weight matrix): when all the features of the coin were *good*, the antique dealer chooses to buy the coin. Otherwise, the dealer does not buy the coin.

In the bottom level, each top-level unit was represented by a distributed representation (randomly generated). The

³ This number of explanations was found by using the Cartesian product of the first three features (which were assumed to be independent). Adding more explanations is also possible but would require the consideration of other features of the problem by the model. As argued elsewhere (Evans, 2006; Hélie & Sun, submitted), the preparation phase, while unable to explain the problem, serves to contextualize subsequent processing and eliminate irrelevant features.

sizes of the distributed representations were made to reflect the constraints of the task. For instance, the dealer's decision was to be explained and should not be changed by the settling process. Hence, the dealer's decision was represented by more units in the bottom level than the other features. In contrast, the coin date might be problematic, so it was represented using fewer units in the bottom level.⁴ Eight training stimuli were generated (i.e., one for each response) and were used to train the attractor neural network to redundantly encode the top-level rules in the bottom level.⁵

To simulate this task, the initial stimulus represented a good date, a good material, a good carved pattern, but a refusal of the dealer to buy the coin (i.e., the problem). This stimulus did not correspond to any exemplar used to pre-train the bottom level, nor to the pre-coded rules in the top level. The same parameter values were used as in the previous simulation, except for $\lambda = 325$ and $\alpha = 1$. Because each iteration in the bottom level takes approximately 350 ms of psychological time, the first period of problem solving lasted a maximum of 343 iterations (2 minutes).

The interruption following the initial problem solving period lasted 257 iterations (90 seconds). During this time, the participants who were assigned to an unrelated task were assumed to continue generating implicit hypotheses to explain the problem. To represent this continuity in hypothesis generation, the simulations representing these participants were not modified in any way. In contrast, the verbalization group was not allowed any implicit processing during this period; hence, the scaling parameter was set to zero in the simulations representing these participants (i.e., $\lambda = 0$).

Finally, for both conditions, there was another 4-minute period of implicit and explicit processing for explaining the problem (a maximum of 686 iterations). During this final problem solving period, the scaling parameter in the verbalization condition was reset to its initial value ($\lambda = 325$; which was the same as the value in the other condition). 41 simulations were run in each condition. The dependent variable was the proportion of correct responses.

Simulation Results

The simulation results are shown in Figure 3 (black bars). Only 34.2% of the simulations in the verbalization condition output the correct response, whereas 43.9% of the simulations in the unrelated problem condition output the correct response. This difference is statistically reliable $B(41, 0.342) = 18, p < .05$.⁶ Unfortunately, we were unable

⁴ More precisely, the attractor neural network representing the bottom level contained 280 units ($r = 280$), 230 of which were used to represent the coin features [i.e., material (50) + pattern (50) + date (30) + decision (100)]; the remaining 50 units were used to represent the explanations ($s = 230$).

⁵ The parameters were the same as in the previous simulation except for 25 epochs and $p = 1$.

⁶ A binomial statistic was used here because each simulation produced a single output classified as *correct* or *incorrect*.

to put the human and simulated data in an ANOVA (with data source as a factor), because we did not have access to the human data. However, the fit is quite good and the reliable statistical difference is reproduced. EII-BF thus successfully captured the overshadowing data.

Discussion

In this simulation, the bottom level was more likely to generate the correct explanation because the top level could only produce stereotypical responses that are direct consequences of its pre-coded rules. However, the bottom level is a soft constraint satisfaction process (Hopfield & Tank, 1985) that can weigh some constraints (e.g., the dealer's decision) more heavily than others (e.g., the coin date). The bottom-level activation converged toward existing attractors that satisfied the stronger constraints, which constituted an intuitively appealing interpretation of the phenomenon.

This result can also be interpreted as a test of the effect of the scaling parameter (λ): when the result of implicit processing was weighed more heavily in knowledge integration, insight was more likely to happen (up to a point). However, when bottom-level processing was disregarded in knowledge integration, EII-BF could no longer perform soft-constraint satisfaction and was thus less likely to find the correct response. This interpretation of the simulation results is in line with the constraint theory of insight (reviewed in Pols, 2002).

Conclusion

In this paper, a new connectionist model of creative problem solving has been proposed (inspired by CLARION, but different; see Sun et al., 2005). EII-BF is simple and yet powerful enough to explain difficult problems and capture psychological data related to insight problem solving (Durso et al., 1994) and the overshadowing effect (Schooler et al., 1993). These simulations suggest that, in line with existing psychological theories and human data, the probability of reaching an insight solution is positively related with the noise level (e.g., Martindale, 1995) and the amount of implicit processing (Schooler et al., 1993).

Other simulations supporting the presence of other stages of Wallas' (1926) analysis of creative problem solving have been run using the EII-BF model (e.g., incubation). The results are promising. Future work should be devoted to the simulation of many more such problems, as well as the simulation of regular problem solving to further substantiate the model.

Acknowledgments

This research was supported (in part) by a postdoctoral research fellowship from *Le Fonds Québécois de la Recherche sur la Nature et les Technologies*, and by the *Army Research Institute* contract W74V8H-05-K-0002 and the *Office of Naval Research* grant N00014-08-1-0068.

References

- Chartier, S. & Proulx, R. (2005). NDRAM: Nonlinear Dynamic Recurrent Associative Memory for learning bipolar and non-bipolar correlated patterns. *IEEE Transaction on Neural Networks*, 16, 1393-1400.
- Durso, F.T., Rea, C.B., & Dayton, T. (1994). Graph-theoretic confirmation of restructuring during insight. *Psychological Science*, 5, 94-98.
- Evans, J.B.T. (2006). The heuristic-analytic theory of reasoning: Extension and evaluation. *Psychonomic Bulletin & Review*, 13, 378-395.
- Helie, S. & Sun, R. (submitted). Incubation, insight, and creative problem solving: A unified theory and a connectionist model. *Submitted to a Journal*.
- Hopfield, J.J. & Tank, D.W. (1985). "Neural" computation of decisions in optimization problems. *Biological Cybernetics*, 52, 141-152.
- Kass, R.E. & Raftery, A.E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773-795.
- Kosko, B. (1988). Bidirectional associative memories. *IEEE Transactions on Systems, Man, and Cybernetics*, 18, 49-60.
- Martindale, C. (1995). Creativity and connectionism. In S.M. Smith, T.B. Ward, & R.A. Finke (Eds.) *The Creative Cognition Approach*. Cambridge, MA: MIT Press.
- Pols, A.J.K. (2002). *Insight Problem Solving*. Doctoral Dissertation, Department of Psychology, University of Utrecht, Netherlands.
- Reber, A.S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, 118, 219-235.
- Rowe, J. & Partridge, D. (1993). Creativity: A survey of AI approaches. *Artificial Intelligence Review*, 7, 43-70.
- Schooler, J.W., Ohlsson, S., & Brooks, K. (1993). Thoughts beyond words: When language overshadows insight. *Journal of Experimental Psychology: General*, 122, 166-183.
- Seger, C. (1994). Implicit learning. *Psychological Bulletin*, 115, 163-196.
- Smith, S.M. & Dodds, R.A. (1999). Incubation. In M.A. Runco & S.R. Pritzker (Eds.) *Encyclopedia of Creativity*. San Diego, CA: Academic.
- Sun, R. (1995). Robust reasoning: Integrating rule-based and similarity-based reasoning. *Artificial Intelligence*, 75, 241-295.
- Sun, R., Merrill, E., & Peterson, T. (2001). From implicit skills to explicit knowledge: A bottom-up model of skill learning. *Cognitive Science*, 25, 203-244.
- Sun, R., Slusarz, P., & Terry, C. (2005). The interaction of the explicit and the implicit in skill learning: A dual-process approach. *Psychological Review*, 112, 159-192.
- Wallas, G. (1926). *The Art of Thought*. New York: Franklin Watts.