Expertise, Text Coherence, and Constraint Satisfaction: Effects on Harmony and Settling Rate[1]

Bruce K. Britton
Institute for Behavioral Research
University of Georgia
Athens, GA, 30602-3013
IBR@UGA.cc.Uga.edu

F. James Eisenhart
A.I. Programs
BGSRC
Room 111
University of Georgia
Athens, GA 30602
Jeisen@AISUNI.AI.uga.edu

Abstract

This paper reports three experiments showing that 17 experts' mental representations had significantly higher harmony and faster settling rates than 638 novices' when activation was spread through the representations in a simulation of thinking; that when coherent texts were read by novices, they produced mental representations with significantly higher harmony and faster settling rates than less coherent texts; and that novices whose representations matched the experts' mental representations had significantly higher harmony and faster settling rates. The results were found for declarative experts in history and procedural experts in literary interpretation, for novice groups including U.S. Air Force recruits and undergraduates, and for both history texts and literary texts. These results were consistent with our hypothesis that the quality of a person's prior knowledge determines the harmony and settling rates of their representations and that these can be measured by simulating the spread of activation through the person's mental representation of a subject matter domain. Harmony may also be used as a metacognitive signal.

In these studies, we investigated the quality of mental representations. To do this, first we measured each subject's mental representation for a domain, and then we simulated the subject's thinking process about it. We then compared the end product of the thinking process with the original mental representation; a good match indicated a high quality representation, in the sense that the constraints in the original representation had been satisfied. Another index that we used for the quality of the representation was the number of processing cycles required for the simulated thinking process to settle.

We measured the mental representations of experts and novices by asking them to rate the relatedness of all possible pairs of 12 terms from a text. The result was a fully interconnected network of associations. To simulate the thinking process, we spread activation through this fully interconnected network of associations until it settled into a stable state. Besides observing the time it took to settle (here called settling rate) we also calculated a harmony score, which is a measure of the "goodness" of the mental representation, based upon the degree to which the settled representation satisfies the constraints inherent in the original ratings of relatedness (McClelland & Rumelhart, 1988). Our hypothesis was that subjects with better prior knowledge would produce mental representations which better satisfied the constraints, leading to higher harmony scores and faster settling rates.

The experiential counterpart to our simulation would be a person thinking about a subject matter area that they knew something about. After completing their thinking process, the person would compare the end

---

[1] This research was supported by AFOSR (LAMP Grant AFOSR-89-0515), Office of Naval Research Grant No. N00014-92-4-1977, and the National Reading Research Center under Dept. of Education Grant 117A20007.

product of it to the constraints inherent in the network of ideas s/he started with. If they matched perfectly, then the person would conclude that s/he had a "good" representation, corresponding to a harmony score of 1.0. But if they did not match perfectly, then the representation would be judged less good, with a harmony score below 1.0; harmony can go as low as -1.0. The person could use the feeling of "goodness" reflected in the harmony score as an index of how well s/he understood the subject-matter area, or for other purposes.

In our simulation, the person's knowledge about the subject matter area was represented by an associative network in the computer program, with the network's links, and the strengths of those links, constructed from the person's ratings of relatedness of the 12 terms. The person's thinking about this network of association was simulated by spreading activation along the links in accord with their strengths until the network settled.

What are Harmony and Settling Rate?

Some simple illustrations of the spreading activation process will clarify the notions of harmony and settling rate. Consider first the two simplest possible networks, those of two nodes, A and B, connected by a single link, which can be either excitatory, represented by a +1, or inhibitory, represented by a -1, shown in Figure 1. An excitatory link is equivalent to a positive association, corresponding to the notion that thinking of one idea leads to thinking of another idea: If one thinks of A, one will think of B, and vice versa. An inhibitory link is a negative association, corresponding to the notion that thinking an idea leads to not thinking of another idea: If one thinks of A, one will be less likely to think of B, and vice versa. If we start to spread activation from A to B and B to A, and allow the network to settle into a stable state, both of these networks will settle very quickly, and their final state will satisfy the constraints perfectly, leading to a harmony score of +1.0. That is, for the excitatory network both nodes will be on after the network settles, entirely consistent with the constraint that both nodes should be in the same state, expressed by the +1. So the harmony score will be at its maximum of +1 because the final state of the network is completely consistent with the constraints.

Figure 1
Two-Node Networks

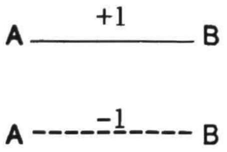

For the inhibitory network, one node will be on and the other will be off, consistent with the constraint that the nodes should be in different states, expressed by the -1. Again the harmony score will be at its maximum of +1 because the final state matches the constraints perfectly.

Harmony scores of less than one can occur in three-node networks. For example, consider the network at the top of Figure 2. Here A is excitatorially linked to B and C, but B and C are inhibitorially linked to each other. The thought that is equivalent to this network can be expressed in words as "Two things that are different from each other (i.e., B and C) can both be the same as something else (i.e., A)." In trying to simulate the thinking of this thought, the spreading activation process may start out with A in some state, say S. Then B will tend to be in the same state S, because of the excitatory link between A and B represented by the +1. Then C will tend to be in a different state, say Not-S, because of the inhibitory link between B and C. But if C is in the state Not-S, then A will tend to be in the same state, namely Not-S, because of the excitatory link between C and A represented by the +1. But if A is Not-S, it will tend to make B be Not-S, which will tend to make C take on the state S, which will make A take on the state S, and so on. It is obvious that it is impossible for the final state of the network to match the constraints perfectly, and this leads to a harmony score below 1. Also, when activations are started at 1.0 for all units, the network takes a relatively long time (ten cycles) to settle to the relatively unharmonious solution it arrives at.

The experiential counterpart of this is that it is difficult for the human thinking process to come up with an entirely satisfactory interpretation of the thought "Two things that are different from each other can both be the same as something else". That the spreading activation process has this same difficulty bodes well for its adequacy as a simulation. In the simulation, this difficulty in thinking satisfactorily about that thought is reflected quantitatively in low harmony scores and slow settling rates. In the human mind, there may be some cognitive counterpart to these quantities.

In contrast, consider the network at the bottom in Figure 2, corresponding to the thought: "Things that are the same as each other (B and C) can both be different from something else (A)." Experientially, this is relatively easy to think about, and in the simulation the spreading activation process is quickly able to settle to a state that satisfies the constraints perfectly. That is, if A is in state S, then B is in Not-S, as is C, and A is in state S, and all these are entirely consistent with the constraints, leading to a harmony score of +1.0, and taking only 2 cycles to settle.

Figure 2
Three Node Networks

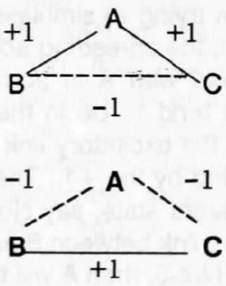

The Present Studies

Our hypothesis is that experts in a subject matter area, like history, will have arrived at an internally consistent mental representation of historical events in their area. Therefore the constraints in their network of associations will be easily satisfied, leading to high harmony scores and fast settling. Similarly, experts in the skill of literary interpretation will be well-practiced in arriving at internally consistent mental representations of literary texts. Therefore the constraints in their network of associations will be easily satisfied, leading to high harmony scores and fast settling.

In contrast, novices in history and literary criticism will be less likely to come up with internally consistent representations of texts in those subject matter areas, and so the constraints in their network of associations will be less likely to be easily satisfied, leading to lower harmony scores, and slower settling for novices than for experts.

But even if novices don't have prior knowledge representations in an area, there are several ways in which they can construct internally consistent mental representations in a subject matter area. One way is by reading a coherent text that provides an internally consistent textual representation of the subject matter area. If that textual representation is learned, the result can be that the novice has acquired an internally consistent representation of the subject matter area, leading to relatively higher harmony scores and faster settling rates for those novices who have read coherent texts. But if the text is less coherent, the chances of the novice learning an internally consistent representation are lower, leading to relatively lower harmony scores and slower settling rates.

Another way in which novices can arrive at an internally consistent mental representation of a subject matter area, even if they are reading an incoherent text, is by using their reading abilities and domain prior knowledge to supplement what is provided in the text, so as to produce an internally consistent representation. So novices who are good readers and/or have relevant prior knowledge should have relatively higher harmony scores and faster settling rates.

In the experiments reported below, we tested these hypothesis by gathering mental representations from experts and novices who had read coherent or incoherent history or literary texts, and then simulated their thinking process by spreading activation through our copy of their mental representations, measuring harmony and settling rate. We expected higher harmony and faster settling rates for experts, for novices who read coherent texts, and for novices who were good readers.

Experiment 1

Method

Subjects. Eight subject matter experts were tested, including an ambassador, four military historians, and 3 military employees. For the comparison of the Original to the

Revised texts, the novices were 357 Air Force recruits, 179 of whom read the Original text, and 178 the Principled Revision, which had been revised to be more coherent. For the study of novices who were good or poor readers, the subjects were 211 Air Force recruits who had read the Original text.

Materials. The Original text was a 1000 word self-contained section from an Air Force ROTC training document. The Revised version had been modified to improve its coherence (Britton & Gulgoz, 1991).

Tests. The 12 most important terms from the text were selected by consensus between the author and a graduate student. All possible pairs of the 12 terms were constructed and presented to all subjects for rating on a seven point scale labeled from "very closely related" to "very distantly related." For purposes of the associative network used in our analysis, responses of "very closely related" were interpreted as excitatory links of +1 and "very distantly related" as inhibitory links of -1; the intermediate values were interpreted as +.67, +.33, 0, -.33, and -.67, respectively.

Analyses. The ratings of each subject were entered into a network in a program based on the Construction-Integration model of Kintsch (1988). Initial activation values were set to 1.0 for the data reported here. (Additional simulations with initial activations differing non-systematically gave the same pattern of results.) Harmony was calculated by taking the sum of the product of the correlations between the final activation of the units and the initial weight matrix, and scaling it by the sum of the absolute values of the weights. That is, first we rescaled the activation vector: rescaled activation = 2(old activation -.5). Then we multiplied the activation vector by its transpose. Finally, we multiplied the resulting correlation matrix by the weight matrix, added the resulting values, and divided by the sum of the absolute values of the weights. The number of cycles to settle was recorded by the program.

Results

The top row of Table 1 shows the harmony and settling rates for the average expert and the next row shows them for the average novice in the groups who read the original text and the text revised to be more coherent. The harmony was significantly higher and the settling rates significantly lower for the experts than for the novices in either group ($t(184) = 6.72$ and 7.48, for Original text harmony and settling time, and $t(185) = 5.66$ and 8.07 for the Revised text, all $p < .05$.

Table 1

Harmony and Settling Rates in Experiment 1

	Experts (n=8)
Harmony	.789*
Settling Rate	4.00*

	Novices who read Original Text (n=179)	Novices who read Revised Text (n=178)
Harmony	.286	.306*
Settling Rate	9.39	9.23

	Novices with Expert Representation (n=68)	Remaining Novices (n=143)
Harmony	.469*	.262
Settling Rate	8.10*	10.25

*$p < .05$

Also, the average harmony was significantly higher for the novices who read the revised text than those who read the original text, $t(353) = 3.01$, $p < .01$ and the settling rate was faster, although the latter was not significant.

A separate group of 211 novices who read only the original text were factor analyzed, and three factors were identified (as in Britton & Tidwell, 1991). The mental representation of novices in one of these factors was significantly correlated with the representation of the average expert (as in Britton & Tidwell, 1991). These novices were presumably the more successful readers. The bottom row of Table 1 shows the average harmony and settling rate for subjects in that factor, and for the remaining subjects. The novices with the expert representation had significantly higher harmony and faster settling rate, consistent with the hypothesis that they had more internally consistent mental structures than the remaining subjects, $t(209) = 6.37, p < .001$, $t(209) = 2.13, p < .05$.

Discussion

These results are consistent with our expectations that experts, readers of coherent text, and good readers would have mental representations that were more internally consistent than novices, readers of less coherent texts and poor readers. We conclude that the quality of the testee's prior knowledge may be measurable by calculating harmony and settling rates. The remaining studies replicated this finding with a different type of expertise, and also with different materials, experts, and novices.

Table 2

Results of Experiments 2 and 3

Experiment 2

	Literary Versions	Literal Versions
Harmony	.170*	.114
Settling Rate	14.99	15.75

Experiment 3

	Literary Versions	Literal Versions
Harmony	.230*	-.050
Settling Rate	12.00	20.00

Experiment 2

A well-written literary text can be given a coherent literary interpretation. A well written informational text can be given a coherent informational interpretation. But suppose a reader attempts to give an informational interpretation to a literary text: Attempting to interpret a text differently from the way its author intended will presumably be more difficult, and the resulting mental representation will likely be less internally consistent. If so, harmony will be lower when readers are asked to interpret a literary text as a literal one, and settling rate will be slower. In Experiment 2 we asked novices to provide literal interpretations for four literary texts; novices in the control condition were asked to produce literary interpretations for the same four texts.

Subjects. Seventy undergraduates were tested, half in the Literary Condition and half in the Literal Condition.

Text. Four literary texts were selected from an international anthology of prose poems; all were by well-known authors, three by winners of the Nobel Prize for literature. One text described an elm log which had been sawn a year before, yet had produced a fresh green shoot; a second text described the struggles of a fly who had become stuck in flypaper; a third was about a dream of a chess game, and the other was about ants returning to a burning log. For each text, a literary title was provided for the literary group and a literal title for the literal group. The literary title was constructed to cue a literary interpretation: for example, for the Elm Log passage it was "Rebirth of the Spirit"; for the Flypaper passage it was "How Man Ages and Dies". The corresponding literal titles were "Description of an Elm Log" and "A Description of how a Fly becomes Tangled in Flypaper". Titles were constructed similarly for the other two passages.

Tests. Twelve terms were selected, including six terms representing perennial literary themes -- love, death, life, hate, happiness and sadness (these terms were the same for all passages) -- and six terms from each passage; for example, for the elm log passage these were log, green, branch, saw, tree, chop; for the fly passage they were flypaper, old age, attempt, give up, helpless, man. For each text, about half of the terms were more susceptible to a literal interpretation and half to a more literary interpretation. (For some passages, we judged that one or more perennial theme terms could more easily be interpreted literally in the context of that passage, so the same number of other terms were selected to be more susceptible to a literary interpretation.) All possible pairs were constructed and presented to each subject after s/he read the text, and each subject rated them on the same seven-point scale as used in Experiment 1.

Results and Discussion

Table 2 shows the results. The novices who provided the literary interpretation had higher harmony and faster settling rate than those who provided the literal interpretation, consistent with our hypotheses, but only the harmony difference was statistically reliable, $t(138) = 2.42, p < .02$.

Experiment 3

This experiment examined the ability of experts to provide a literal interpretation of a literary text.

Subjects. The experts were 9 faculty members professionally concerned with literary interpretation, from the Departments of English, Comparative Literature, Romance Languages, and Philosophy at the University of Georgia.

Materials. The texts and tests were the same as in Experiment 2.

Procedure. Each expert provided both a literary and a literal interpretation for each text, on two occasions separated by at least one month.

Results

Table 2 shows the results. The experts had higher harmony and faster settling rates when they provided the literary interpretations than the literal ones, but the effect was significant only for harmony scores $t(8) = 2.67, p < .05$. That the average harmony for the experts' interpretations of the literal versions is negative is particularly notable. It appears that literary experts have exceptional difficulty providing a literal interpretation for a literary text.

General Discussion

These results show that the quality of a person's mental representation can be measured by spreading activation through it, and observing its settling rate and harmony. Validation of this was obtained by comparing (a) mental representations whose quality was known to be high *a priori*, i.e., those of experts, those of novices who arrived at representation similar to experts, and both experts and novices who read texts that permitted a coherent interpretation; to (b) mental representations which were known to be of lower quality *a priori*: novices, especially novices who did not arrive at representations like those of experts, and both experts and novices who read texts less susceptible to a coherent interpretation. In all cases the results were consistent with expectations.

The validity of this measure suggests it as a solution to a long-standing problem in the theory of metacognition: How does the person come to realize that there is something wrong with his mental representation, wrong enough to trip off fix-up mental activities, such as rereading the text to try to understand it better, asking questions, etc.? One way to know that you don't understand is to compare your understanding with the correct understanding, but this is not possible unless you already have the correct understanding. In contrast, it is always possible to calculate harmony, and the self-calculation of harmony can provide a feedback signal, which may correspond experientially to the feeling that your understanding is incomplete. This signal could then be used to trip off mental activities that will induce fix-up actions.

Harmony also provides a measure of quality that is independent of comparisons to experts, which may prove useful when experts have differing views, as often happens in complex subject matter areas, or when qualified experts are not available. For example, in evaluating the representation of a person who is the only expert in a field, or in which there are no experts, harmony may be the only measure available.

References

Britton, B. K. & Gulgoz, S. (1991). Using Kintsch's computational model to improve instructional text: Effects of repairing inference calls on recall and cognitive structures. *Journal of Educational Psychology*, *83*, 329-345.

Britton, B.K. & Tidwell, P. (1991). Shifting novices' mental representation toward experts': Diagnosis and repair of mis- and missing conceptions. In *Proceedings of Thirteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review*, *95*, 163-182.

McClelland, J. L. & Rumelhart, D. E. (1988). *Explorations in parallel distributed processing*. Cambridge, MA: MIT Press.