Nokes, T. J., & Ohlsson, S. (2001). How is abstract generative knowledge acquired? A comparison of three learning scenarios. In J. D. Moore and K. Stenning (Eds.), *Proceedings of the Twenty Third Annual Conference of the Cognitive Science Society* (pp. 710-715). Mahaw, NJ: Erlbaum.

How is Abstract, Generative Knowledge Acquired? A Comparison of Three Learning Scenarios

Timothy J. Nokes (**tnokes@uic.edu**)
Stellan Ohlsson (**stellan@uic.edu**)

Department of Psychology
The University of Illinois at Chicago
1007 West Harrison Street (M/C 285)
Chicago, IL 60607, U.S.A.

Abstract

Several theories of learning have been proposed to account for the acquisition of abstract, generative knowledge including schema theory, analogical learning and implicit learning. However, past research has not compared these three theories directly. In the present studies we instantiated each theory as a learning scenario (i.e., direct instruction, analogy training and implicit training) and then tested all three training groups on a common problem. Results show that the analogy training groups and one of the direct instruction groups performed significantly better than the other groups on problem solving performance. The findings are interpreted in terms of opportunity to practice generating a response of the relevant type.

Theories of Deep Learning

In order to solve complex, novel problems one must be able to retrieve previously learned information from memory and apply it to the current situation. For instance, students learning geometry need to be able to apply mathematical formulas acquired during study to novel problems encountered at test. Although surface features of the problems change (e.g., specific values: a=5 to a=15) the abstract operators used to solve the problems stay the same (e.g., the Pythagorean theorem: $a^2 + b^2 = c^2$). Thus, in order for the knowledge gained from study to be helpful on the test it must be both abstract and generative. How such deep knowledge is acquired remains a central question for researchers in psychology, philosophy and education.

Several theoretical explanations have been proposed as to the origin of such abstract, generative knowledge including: schema theory (Marshall, 1995; Thorndike, 1984), analogical learning (Gentner, 1983; Holyoak & Thagard, 1988) and implicit learning (Reber, 1989).

Research on schema theory has shown that abstract knowledge is constructed during various types of higher-order cognitive activities including text comprehension (e.g., Kintsch & van Dijk, 1978; Thorndike, 1977), problem solving (e.g., Marshall, 1995) and direct instruction (e.g., Ohlsson & Hemmerich, 1999; Ohlsson & Regan, in press). For the purposes of this paper we are not concerned with the induction hypothesis of schema acquisition but instead with whether a schema can be taught directly (Ohlsson & Hemmerich, 1999). Although schema theory has provided much insight into the nature and form of abstract knowledge representations (e.g., Bobrow & Collins, 1975) it has done little to articulate a specific theory for how abstract schemas are acquired.

A second major theoretical proposal is the analogical learning hypothesis. Research on analogical learning suggests that one acquires deep knowledge through a systematic process in which a person retrieves an analog from memory and maps the underlying conceptual structure to a novel problem (Gentner, 1983; Holyoak & Thagard, 1988). In a typical analogical learning experiment participants first solve a source problem (e.g., story problems: Gick & Holyoak, 1983; or algebra problems: Reed, 1987) and then solve a test problem that has different surface features (i.e., a different context) but retains the deep relational features of the source problem. When participants are given the hint to use the source problem to solve the test problem they perform better than a control group who did not receive prior training, indicating that explicit knowledge of the prior solution procedure facilitates subsequent problem solving.

In contrast to the prior two theories research on implicit learning suggests that knowledge acquisition is a passive, inductive process that is independent of any intention to learn (Reber, 1989; Seger, 1994). In the training phase of artificial grammar learning – a typical implicit learning paradigm – the participants memorize letter strings that are generated from an artificial grammar. Participants are not informed of the rule-based nature of the memorization strings until after the training phase. In the test phase, the participants are given a classification task in which they are asked to judge whether or not new letter strings, half generated by the relevant grammar and half violating one or more of the rules, are like those memorized during the training phase. A large amount of evidence (Reber 1989; Seger; 1994; Stadler & Frensch, 1998) shows that participants perform better than chance in the test phase, indicating that they have acquired some knowledge of the underlying grammar.

These three theories present a complicated if not contradictory picture of knowledge acquisition. Each theory has a history of empirical support, experimental paradigms and explanatory problems associated with it. In order to

further explore and understand the relationship between these theories comparative empirical studies are essential. In the present studies we examine all three theories via a single experimental paradigm. This allows us to compare and examine the type of the knowledge generated by the various learning scenarios associated with each theory. How does that knowledge function on subsequent tasks? Is the knowledge representation abstract? Is it generative?

We instantiate each theory as a different type of learning scenario (i.e., direct instruction, analogical training and implicit training) and then test all three training groups on a common problem, Thurstone's sequence extrapolation task (Thurstone & Thurstone, 1941).

Sequence extrapolation problems have been studied from a cognitive perspective by Kotovsky and Simon (1973) and Simon (1972), among others. In this type of problem, the problem solver is given a sequence of symbols (usually letters) generated in accordance with some complex pattern and asked to extrapolate it. In order to solve the problem, he or she must first *discover* the pattern in the given segment of the symbol sequence and then articulate that pattern to *generate* the next N positions of the sequence.

In the current studies, experiments were divided into a training phase and a test phase. Participants in the implicit training condition memorized letter strings that had the same abstract pattern as that used in the test phase. Participants in the analogy training condition solved source problems that had the same abstract pattern as that used in the test problems. Participants in the instruction condition read a general tutorial on how to solve sequence extrapolation problems and studied the abstract rules of the pattern for each target problem.

In the test phase of the experiments all participants solved two types of extrapolation problems, a target problem and a transfer problem. The *target problem* had the same deep structure (i.e., the relations between the elements of the pattern) as that used in the sequences of the training procedures. The *transfer problem* was generated from the target problems by 'stretching' relations between letters, e.g., "forward 1 step in the alphabet" becomes "forward 2 steps" and "backwards 1" becomes "backwards 2" (see Table 1).

The goal of the present studies is to examine the nature and function of the knowledge created by each of the three training scenarios. If the knowledge gained is generative (i.e., can articulate a sequence of temporally related actions) and abstract (i.e., not bound to surface features) then performance should be facilitated on the target task. If the knowledge is of a higher level of abstraction it should facilitate problem solving performance on the transfer problem.

Experiment 1: Analogy vs. Implicit Training

Experiment 1 compares the effect of different types of analogical training (i.e., solving 1 problem three times vs. solving 3 structurally similar problems once) to different levels of implicit training (i.e., memorizing 6 vs. 18 instances of the pattern) on problem solving task performance.

Method

Participants One hundred and twenty seven undergraduate students from the University of Illinois at Chicago participated in return for course credit.

Materials The target tasks were two sequence extrapolation problems with a periodicity of six items for problem 1 and seven items for problem 2. Each task was instantiated as both a target and transfer problem; see Table 1. Target and transfer problems were related in that they contain similar over-arching pattern types but differed in the particular instantiation of the relations (i.e., manipulating some of the relations of the target pattern by a magnitude of 2 for transfer problems). To enable participants to detect the pattern, the given segments were 12 items long for task 1 and 14 items long for task 2. That is, they covered two complete periods of the underlying pattern. The extrapolation problems were created specifically for this experiment with a design similar to the problems used by Kotovsky and Simon (1973).

Table 1. Two sequence extrapolation problems and their associated transfer problems.

Problem Type	Given letter or number sequence & the correct 8-step extrapolation
Problem 1	
Target	E F D G C O F G E H D P G H F I E Q H I
Transfer	E G D I C O G I F K E P I K H M G Q K M
Problem 2	
Target	A C Z D B Y Y D F X G E W W G I V J H U U J
Transfer	A E Z G C X X G K V M I T T M Q R S O P P S

There were also three extrapolation *training* problems for each target task. The three training problems followed the exact same pattern as the associated target problem; see Table 2a for an example. Training problems were constructed so they do not overlap (i.e., do not share any of the surface features) with each other or the target problems. The single analog group was trained on the first of the three training problems and the multiple analog group was trained on all three.

In addition, there were 36 letter training strings consisting of 12 letters for task 1 and 14 letters for task 2, eighteen strings for each problem. The eighteen strings associated with each problem followed the exact same

pattern as the given sequence for that problem; see Table 2b for an example. The low implicit participants were trained on six strings per task and the high implicit participants were trained on eighteen strings per task.

Table 2. Two training sequences for Problem 1.

Example	
A. Source Problem:	I J H K G S J K I L H T
B. Training String:	M N L O K W N O M P L X

Each participant was tested on a Macintosh computer with a 14' color monitor, standard keyboard and mouse. All stimuli were presented in black 30pt font in the center of the screen. The experiment was designed and generated using PsyScope software. Target and transfer problems and associated training stimuli were counter-balanced across all conditions.

Design and procedure The participants were randomly assigned to one of four groups: *single analog* (n = 25), *multiple analog* (n = 23), *low implicit* (n = 26), *high implicit* (n = 23). In addition, a separate *control group* (n = 30) was tested on the target and transfer problems to provide a measure of baseline performance.

In the analogy training groups, participants solved letter sequence extrapolation training problems that conformed to the same patterns as those used in the target problems. The single analog group solved one and the same training problem three times and the multiple analog group solved three different training problems once each. Each of the multiple analog problems had different surface features but they all shared the same underlying pattern. In both implicit learning groups, participants memorized letter strings that conformed to the same patterns as those in the target extrapolation problems. The low implicit group memorized six training strings and the high implicit group memorized eighteen training strings. In the control group participants received no training.

General procedure. Participants were tested in groups of 1-4 people. The procedure consisted of *two cycles*. Each cycle was composed of problem training followed by solving target and transfer problems.

Procedure for analogy groups. Participants were first given general instructions on how to solve sequence extrapolation problems. Next, they were presented with the first extrapolation training problem. They were given 6 minutes to solve each problem. After participants had finished solving a problem or max time had elapsed, they were presented with the next problem. After participants solved all three training problems they were presented with the target problem instruction. Target problem instructions were the same as the training instructions except that they added the hint that if participants noticed a pattern on any of the prior problems it would help them solve the target problem. They were then presented with the target problem and were given 6 minutes to solve it. Finally, they were given the transfer problem and instructed to solve it in the same manner as the target problem. The second cycle proceeded in the same manner. The entire procedure took between 60-80 minutes.

Procedure for implicit learning groups. The participants were instructed to memorize and recall each letter string one by one; six strings for the low implicit group and eighteen strings for the high implicit group. They were then given 45 seconds to memorize each string. After 45 seconds the string disappeared and they were given 30 seconds to recall and type in the string. After they finished recalling the string or 30 seconds elapsed, participants were presented with the next string. This procedure was continued through all of the training strings. Next, participants were instructed to write down whether or not they noticed a pattern in the memorization strings. If they noticed a pattern they were instructed to describe it as best they could. Participants were then given general instructions on how to solve the sequence extrapolation problems. They were presented with the target sequence extrapolation problem and given the hint that if they noticed a pattern from the memorization strings it might help them on problem solving. They were given 6 minutes to solve the problem. Finally, they were given the transfer problem and were instructed to solve it in the same manner as the target problem. The second cycle proceeded in the same way. The entire procedure took approximately 70-90 minutes.

Results

The central question of interest is whether or not the various training scenarios facilitated performance on the target and transfer problem tasks. The problem solving score was the number of letters correctly extrapolated in each problem solving task. Because the participants were asked to continue the sequence to eight places their problem solving scores varied between 0 and 8.

Initial analyses revealed non-significant differences within both the implicit and analogy groups and all subsequent analyses collapsed across them, F (1, 46) = .922, ns and F (1, 47) = .02, ns respectively. Figure 1 shows the mean problem solving scores for the analogy, implicit and control groups on target and transfer problems.

A 3 (treatment: analogy vs. implicit vs. control) by 2 (problem-type: target vs. transfer) mixed analysis of variance (ANOVA) revealed a main effect for both treatment and problem-type, F (1, 124) = 7.88, MSE = 12.96, $p < .05$ and F (1, 124) = 14.26, MSE = 1.86, $p < .05$ respectively. The interaction was not significant, F (1, 124) = .32, ns. The main effect of problem-type shows that the participants performed better on the target problems than on the transfer problems. Follow up comparisons on treatment showed that the analogy group performed significantly better than both implicit and

control groups, F (1, 95) = 13.77, MSE = 12.35, p < .05 and F (1, 77) = 8.39, MSE = 14.23, p < .05. The implicit group did not significantly differ from the control F (1, 76) = 0.02, ns.

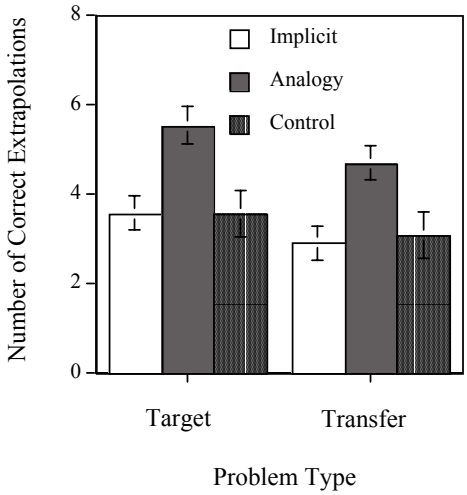

Figure 1: Problem solving as a function of training condition.

In addition, a simple comparison of target vs. transfer was conducted for the control group to provide a measure of baseline performance for problem-type. The analysis revealed a non-significant difference indicating that control participants performed equally well on both target and transfer problems, F (1, 29) = 2.77, ns.

We next compare analogy training to direct instruction.

Experiment 2: Analogy vs. Direct Instruction

Both analogy and implicit learning are indirect training methods. Is it possible to teach a sequential schema directly, by simply telling the participants what the pattern is? Experiment 2 compares different types of analogy training (single vs. multiple analog training) to different levels of direct instruction (high vs. low) on problem solving task performance.

Method

Participants One hundred and nineteen undergraduate students from the University of Illinois at Chicago participated in return for course credit.

Materials The test problems and the analogy training problems were exactly the same as those used in experiment 1.

In addition, there were two extrapolation problem tutorial booklets, one for low instruction training (12 pages) and one for high instruction training (14 pages). Both tutorials consisted of general instructions for how to *find* pattern sequences as well as detailed descriptions of the component relations of patterns (e.g., forward, backward, repeat and identity relations). The high instruction tutorial had two additional pages of general instruction describing how to *extrapolate* or *continue* sequence patterns. There was also a general tutorial test that consisted of four recall questions and two comprehension questions.

In addition, there were two diagrammatic illustrations of the underlying pattern relations for each of the target problems as well as two blank diagrammatic recall sheets (see Table 3 for an example of a diagrammatic pattern illustration). There were also two distractor tasks that consisted of three multiplication problems each.

Table 3. A sample diagrammatic pattern illustration.

Example: *square boxes represent letter positions*

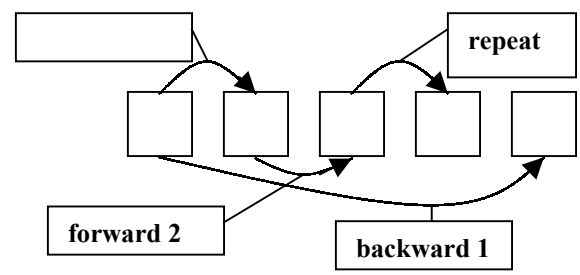

The test problems and analogy training problems were presented via computer with the same specifications as experiment 1. Direct instruction training material was presented on sheets or in booklet form. Target and transfer problems and associated training stimuli were counter-balanced across all conditions.

Design and procedure The participants were randomly assigned to one of four groups: *single analog* (n = 30), *multiple analog* (n = 31), *low instruction* (n = 28), *high instruction* (n = 30). The same *control condition* (n = 30) was used from experiment 1 as a measure of baseline performance.

In the instruction training conditions participants first read general tutorials, then memorized and recalled the abstract patterns for each target task. The only difference between high and low instruction groups was that the high instruction participants were given two additional pages in the tutorial which provided specific step by step instructions for how to extrapolate a problem.

Procedure for analogy groups. Procedure was exactly the same as in experiment 1.

Procedure for direct instruction groups. Participants were tested in groups of 1-4 people. The procedure consisted of *two cycles*, a training phase and a test phase. Before the training phase cycle all participants were given the general tutorial text to read (max time allowed = 18 minutes) after which they were given the tutorial test (max time = 6 minutes). At the beginning of the training cycle participants were given 3 minutes to memorize the first diagrammatic pattern illustration. Next, participants were presented with

the diagrammatic blank recall sheet and instructed to recall and write down the relations of the pattern (max time = 3 minutes). Participants were then given the distractor task (max time = 3 minutes). Next, participants were presented with the general instructions for the test problems. They were then given 6 minutes to solve the target problem. Finally, they were given the transfer problem and were instructed to solve it in the same manner as the target problem. The second cycle proceeded in the same way. The entire procedure took approximately 70-90 minutes.

Results

Again, the question of interest pursued here is whether or not training facilitated problem solving performance on the test tasks. The problem solving score was calculated in the same manner as experiment 1.

Initial analysis revealed a non-significant difference between analogy groups and all subsequent analyses collapsed across them, F (1, 59) = .51, ns. Figure 2 shows the mean problem solving scores for the analogy, high instruction, low instruction and control groups on target and transfer problems.

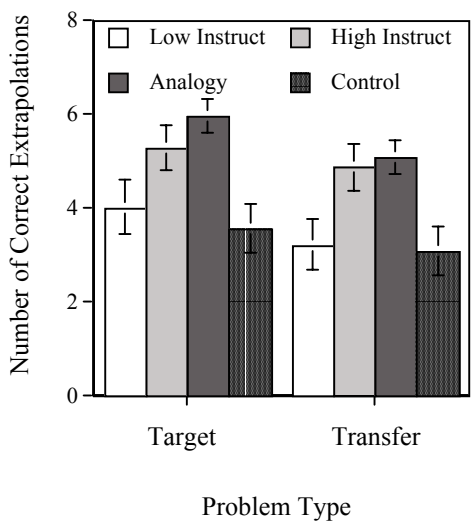

Figure 2: Problem solving as a function of training condition.

A 4 (treatment: analogy vs. high instruction vs. low instruction vs. control) by 2 (problem-type: target vs. transfer) mixed ANOVA revealed a main effect for both treatment and problem-type, F (1, 145) = 6.18, MSE = 14.45, $p < .05$ and F (1, 145) = 18.56, MSE = 1.49, $p < .05$ respectively. The interaction was not significant, F (1, 145) = .70, ns. The main effect of problem-type shows that the participants performed better on the target problems than on the transfer problems. Follow up comparisons on treatment showed that both the analogy and high instruction group performed significantly better than the low instruction and control groups, F (1, 87) = 9.19, MSE = 15.13, $p < .05$, F (1, 89) = 13.11, MSE = 14.77, $p < .05$, and F (1, 56) = 4.37, MSE = 13.94, $p < .05$, F (1, 58) = 6.78, MSE = 13.43, $p < .05$ respectively. The analogy group did not significantly differ from the high instruction group and the low instruction group did not significantly differ from control, F (1, 89) = .60, ns and F (1, 56) = .15, p = ns respectively.

Discussion

So how is abstract, generative knowledge acquired? The present study suggests that there are at least two ways to acquire such knowledge, one through analogical problem solving and the other through direct instruction.

Experiments 1 and 2 showed that participants in the analogy and high instruction training conditions performed better than participants in the implicit, low instruction and control conditions on both target and transfer problems. Target problem performance indicates that the knowledge acquired from analogy and high instruction training was both *generative*, in that the representation could be used to continue a sequence of temporally related actions, and *abstract*, in that the knowledge was flexible and could be applied to novel stimuli. This result supports typical findings on analogical transfer in problem solving (e.g., Gentner & Markman, 1998; Gick & Holyoak, 1983; Reed, 1987).

The analogy and high instruction groups also performed better than implicit, low instruction and control groups on the transfer problems indicating that the knowledge representation was generalizable to similar types of problem structures. However, there was a main effect for problem-type showing that analogy and high instruction participants performed better on target than on transfer problems. In contrast, the control group performed no differently on target than on transfer. These results show that the difference in performance on the target and transfer problem was a function of knowledge gained from training and not of differences in problem stimulus.

These results suggest at least two plausible explanations. One possibility is that some of the participants in the analogy and high instruction groups acquired a knowledge representation of a higher-level abstraction that facilitated their performance on the transfer problem whereas the others did not. Individual differences within the groups would account for the acquisition of a more abstract representation for only a portion of the participants. A second possibility is that participants in addition to learning the abstract pattern from the training stimuli also learned general problem solving heuristics for solving sequence extrapolation problems (i.e., employing specific pattern finding strategies such as searching for repetitions or backward relations). In this case, although knowledge of the specific abstract pattern facilitated performance on target problems it failed to transfer to the transfer problems and participants resorted to more general (and less accurate) problem solving heuristics.

Why did analogy and high instruction training facilitate problem solving and the other training conditions fail? The

prior analysis of the properties of a successful knowledge representation – abstraction and generativity – also reveals the potential components for failure in problem solving including failures of generativity and abstraction.

The failure of implicit training can be explained by either of the above components. Previous research does not provide definitive support for either component. For example, in a prior study we investigated implicit learning in sequence extrapolation problems and found that the knowledge created from the training procedures was of a moderate level of abstraction but was also potentially generative (Nokes & Ohlsson, 2000). Further research is needed to differentiate between each of these components.

The reason for the failure of low instruction can be investigated by examining the differences between the high and low instruction training materials. The only difference between the two training scenarios was that the high instruction tutorial had two additional pages of instruction describing in detail how to *extrapolate* sequence patterns. This description included one example problem that was worked through step by step in detail. This difference in training materials suggests that the low instruction group failed to construct a knowledge representation that was generative.

This hypothesis is also supported by other results in the literature. For example, Ohlsson and Regan (in press) used an intervention paradigm to teach participants several abstract concepts relating to the structure of DNA and facilitated subsequent use of those concepts on a discovery problem. They had participants practice *generating* their own concrete examples of the concepts in the training phase in addition to being given an example from the experimenter. Since low instruction participants in the current study never practiced articulating the abstract schema this component of the knowledge representation was not strengthened.

Thus a commonality that ties both implicit and low instruction together is the lack of practice in *articulating* the abstract schema. Although both learning scenarios had participants memorizing and recalling exemplars, whether they were letter strings or abstract rules, the participants never studied or practiced using these knowledge representations. In contrast, the analogy groups practiced generating pattern sequences on three separate occasions and the high instruction group studied pattern articulation in depth. It is proposed here that it is the practice of pattern articulation of an abstract schema that gives the analogy and high instruction groups their advantage over the other two learning scenarios.

References

Bobrow, D., & Collins, A. (Eds.), (1975). *Representation and understanding*. New York: Academic Press.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy, *Cognitive Science, 7*, 155-170.

Gentner, D., & Markman, A. B. (1998). Structure mapping in analogy and similarity. In P. Thagard (Ed.), *Mind readings* (pp. 127-156). Cambridge, MA: MIT Press

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15*, 1-38.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13*, 295-355.

Kintsch, W., & van Dijk, T. (1978). Toward a model of text comprehension and production. *Psychological Review, 85*, 363-394.

Kotovsky, K., & Simon, H. (1973). Empirical tests of a theory of human acquisition of concepts for sequential patterns. *Cognitive Psychology*, *4*, 399-424.

Marshall, S. (1995). *Schemas in problem solving*. Cambridge, UK: Cambridge University Press.

Nokes, T. J., & Ohlsson, S. (2000). An inquiry into the function of implicit knowledge and its role in problem solving. In L. R. Gleitman and A. K. Joshi, (Eds.), *Proceedings of the Twenty Second Annual meeting of the Cognitive Science Society* (pp. 829-834), Mahaw, NJ: Erlbaum.

Ohlsson, S., & Hemmerich, J. (1999). Articulating an explanatory schema: A preliminary model and supporting results. In M. Hahn and S. Stoness, (Eds.), *Proceedings of the Twenty First Annual meeting of the Cognitive Science Society* (pp. 490-495), Mahaw, NJ: Erlbaum.

Ohlsson, S., & Regan, S. (in press). A function for abstract ideas in conceptual learning and discovery. *Cognitive Science Quarterly*.

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, *118*, 219-235.

Reed, S. K. (1987). A structure mapping model for word problems. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 13*, 127-139.

Seger, C. A. (1994). Implicit Learning. *Psychological Bulletin, 115*, 163-196.

Simon, H. (1972). Complexity and the representation of patterned sequences of symbols. *Psychological Review*, *79*, 369-382.

Stadler, M., & Frensch, P, (Eds.), (1998). *Handbook of implicit learning*. Thousand Oaks, CA: SAGE.

Thorndike, P. W. (1977). Cognitive structures in comprehension and memory of narrative discourse. *Cognitive Psychology, 9*, 77-110.

Thorndike, P. W. (1984). Applications of schema theory in cognitive research. In J. R. Anderson & S. M. Kosslyn (Eds.), *Tutorials in learning and memory: Essays in honor of Gordon Bower* (pp. 167-192). New York: Freeman & Company.

Thurstone, L. & Thurstone, T. (1941). *Factorial studies in intelligence*. Chicago: University of Chicago Press.