

**DO NOT QUOTE WITHOUT PERMISSION
(contact: psmath@lsu.edu)**

**Implicit and Explicit Processes in the Development of
Cognitive Skills: A Theoretical Interpretation with Some
Practical Implications for Science Education**

Ron Sun

Rensselaer Polytechnic Institute

Robert C. Mathews

Sean M. Lane

Louisiana State University

“The most exciting phrase to hear in science, the one that heralds new discoveries, is not 'Eureka!' (I found it!) but 'That's funny ...' “

Isaac Asimov (Science fiction novelist & scholar: 1920 - 1992)

As the quote above illustrates, there are facets of the scientific process that are prompted by feelings of intuition. Although this notion may seem closer to stereotypes of art than science, the notion that scientists can rely upon knowledge that is generated without their awareness has an empirical foundation in the literature on implicit learning (e.g. Reber, 1989; Lewicki, Czyzewska, & Hoffman, 1987; Mathews, et al., 1989). So, although the subjective experience of intuition may seem “magical,” researchers have begun to characterize the characteristics of such knowledge. One major theme of this chapter will be that *implicit knowledge* (knowledge gained directly from experience) and *explicit knowledge* (e.g. the explicit facts that are typically acquired in science instruction) may be fruitfully combined (and clearly are in mature scientists). To use the example above, the intuition that alerts one to the detection of an anomaly (e.g., “That’s funny”) is likely to spur a more explicit reasoning process that involves trying to locate the source of the anomaly and understand its implications.

Most educational settings focus on teaching conceptual (explicit) knowledge rather than setting up an opportunity for gaining substantial experiential (mostly implicit) knowledge. While this may be appropriate for some subject areas, other subjects areas may require learning information (e.g., features of complex systems or categories) that are better learned (at least initially) through extensive hands-on experience than with lectures or textbooks alone (that is, with explicit learning alone). This issue has a significant bearing on the question of the role of laboratory research experiences in science education (see, for example, the chapter by Taraban et al. in this volume). In general, lab experiences help to promote implicit and/or procedural learning, while classroom lectures and textbooks often promote explicit learning of explicit, conceptual knowledge. While emphasizing the importance of explicit knowledge, we should not downplay the significance of implicit learning and implicit knowledge in the overall learning experience. For example, to be a scientist is not simply a matter of accumulating a large amount of specific and explicit knowledge; it is, more importantly, a matter of developing intuitive (implicit) understanding and intuitive (implicit, procedural) skills for approaching and exploring scientific issues and for conducting scientific experiments (Sahdra & Thagard 2003). We believe that a major goal of science education (and education more broadly) is to develop skills that involve the ability to access well-integrated implicit and explicit knowledge in the course of solving problems. The epitome of such a goal is a person who can perform a complex task (e.g., run an experiment, perform surgery, utilize complex equipment) and hopefully also be able to describe to others the conceptual facts and “tricks of the trade” that guide their performance. Such people are commonly recognized as outstanding teachers and collaborators. Viewed in this light, understanding the interaction of implicit and explicit learning becomes a critical issue.

Given the focus of this volume, the goal of this chapter is to provide a theoretical and empirical overview of implicit and explicit learning processes (focusing on how they interact) and to draw out implications of research in this area for science teaching. To do so, we begin by articulating the characteristics of implicit and explicit learning. Next, we focus on how implicit and explicit learning can interact by first describing a computational theory/model (namely CLARION; see Sun, 2002; 2003) of human cognitive processes. Further, we describe the results of several research studies that illustrate the potential interactions between implicit and explicit learning along with computer simulations of the data demonstrating the power and utility of the CLARION theory/model. Finally, we conclude the chapter by discussing what the results of this research have to say about how implicit and explicit learning can successfully be combined in the science classroom.

I. Implicit and Explicit Learning

Substantial research (e.g., Berry & Dienes, 1993; Kolers & Roediger, 1984; Lewicki, 1986; Mathews, et al., 1989; Reber, 1993; Sun 2002; Sun, Slusarz, and Terry, 2005) has documented two different and complementary types of cognitive processes (and subsequent representations) that are involved in the acquisition of cognitive skills. These two processes, implicit and explicit learning, differ on a number of characteristics that are highlighted in Table 1. Implicit knowledge is acquired relatively directly from the environment and requires substantially less mental effort during the learning process than explicit learning. For example, you might learn geographic information about an area simply by driving to a different city even without having an intention to learn. Learning that same information by consulting a map would be more effortful and you would likely be consciously aware of trying to commit the information to memory. Despite its relatively effortless manner, one common misconception about implicit learning is that it does not require attention to the subject matter. Participants in studies of implicit learning *are* attending to a task; they are just not deliberately trying to acquire certain aspects of the task stimuli (e.g., its underlying structure) that are the focus of later testing. Using our example of the driver learning geography, the driver (hopefully) is attending to the task of driving safely even though he or she is not intentionally trying to remember the landscape. Another important aspect of implicit learning is that it is relatively more error tolerant than explicit learning (e.g., Sallas, et al., 2006). By this we mean that implicit learning is more sensitive to the variability of stimuli (e.g., among members of a category) in the real world and is more useful for detecting less salient features of the environment. Explicit learning is less tolerant of error or noise because people typically try to form simple hypotheses that explain a particular phenomenon. When the phenomenon is fairly stable and the relations between variables fairly obvious, simple hypothesis testing can be fairly accurate but hypothesis testing is less effective when relations are made less salient by noise. By noise, we are talking about the natural variation of phenomenon in the world such as the subtle differences between members of a species or how the motion of a tennis ball off a racket might be slightly different due to subtle differences from one moment to the next (e.g., wind speed, fatigue).

Table 1. Characteristics of implicit and explicit learning.

Characteristics	Implicit Learning	Explicit Learning
<i>Effort</i>	Easy	Hard
<i>Learning</i>	Unaware	Aware
<i>Robustness</i>	Error tolerant	Error intolerant
<i>Knowledge</i>	Difficult-to-Verbalize	Easy-to-Verbalize
<i>Type of Cognition</i>	Hot (emotional)	Cool
<i>Speed</i>	Fast	Slow
<i>Control</i>	Cue-driven (unconscious)	Conscious
<i>Solutions</i>	Heuristic	Algorithmic
<i>Representation</i>	Holistic	Analytic

The knowledge that results from implicit and explicit learning also differs in a number of important

ways. One well-known difference is that explicit knowledge is much easier to verbalize than implicit knowledge (while verbalizing implicit knowledge may be impossible without explicitation first). For example, it is quite easy to talk about the names for the bones in the human hand that you just memorized for your anatomy class, but it is likely quite difficult for you to explain the specific movements you just used to pick up your coffee cup. Consequently, explicit knowledge plays a major role in the ability to communicate to others what is known. In a related way, implicit knowledge is much more likely to reveal itself with emotional responses ("hot" cognition) than is explicit knowledge. For example, the feeling you get when you are walking in a dangerous area ("the hair on the back of your neck") can be an implicit response to cues in the environment. In the context of scientific practice, such affective responses can alert a scientist to anomalies before he or she is aware of the source of the response. Implicit knowledge can be deployed much faster than explicit knowledge, and is often triggered automatically by cues in the environment. Because of this, implicit knowledge may be particularly good at providing quick answers that have worked well in past situations (a heuristic), but are not guaranteed to do so in the present moment. Conversely, explicit knowledge is retrieved much more slowly, typically after deliberate effort to do so. For example, think of your own efforts to remember lecture material when trying to answer test questions. Explicit knowledge can be used to fashion highly accurate, dependable and effortful ("algorithmic") solutions to problems. Examples of algorithms would include a procedure that describes how to compute a function, or a recipe for making gumbo. Finally, implicit representations are generally holistic (e.g., through a general impression or reaction) while explicit representations allow one to focus more analytically on individual features. To use an example from Melcher & Schooler (1996), any experienced wine drinker will be able to tell you whether they like a particular wine or not, but only a trained wine taster typically has the analytic understanding and vocabulary to tell you which features of the wine contribute to whether they like it or not (or which aspects they like and which they do not).

II. Interaction of Explicit and Implicit Learning

In the first section we outlined the theoretical qualities of the two systems that guide human thinking. In this section, we present results from three studies demonstrating these qualities and different ways in which the two systems can interact. Sometimes such interactions increase performance as compared to either system operating alone. In other situations activating the explicit system may impair implicit learning or have no effect. It is important to understand these interactions to inform educators on best learning practices. Various experimental manipulations can shift participants' activity more toward one system or the other. For example, telling participants to reflect or figure out the rules while performing a task increases explicit learning processes (e.g., hypothesis testing).

We use Ron Sun's model (cognitive architecture) CLARION (Sun, 2002; 2003) to simulate results from each experiment. This model is a cognitive architecture designed to simulate a wide range of human cognitive processes. One of its most unique features and the one that is most important for present purposes, is that it is one of very few computer models that posits the two systems and allows them to interact in both directions, what we will refer to as bottom-up learning (implicit knowledge becoming explicit) and top-down learning (explicit knowledge affecting implicit learning). The purpose of the simulations is to test the adequacy of our theoretical ideas concerning the interactions of implicit and explicit learning. In most other work on cognitive architectures, the focus has been almost exclusively on "top-down" learning (that is, learning first explicit knowledge and then implicit knowledge on the basis of the former). The bottom-up direction (that is, learning first implicit knowledge and then explicit knowledge) or learning both in parallel has been largely ignored. However, there are a few studies that do demonstrate the parallel development of the two types of knowledge or the extraction of explicit knowledge from implicit knowledge (e.g., Rabinowitz & Goldberg, 1995; Karmiloff-Smith, 1986; Stanley et al., 1989; Sun, Slusarz & Terry, 2001; Sun, 2002).

1. A General Framework

The role of implicit learning in skill acquisition and the distinction between implicit and explicit learning have been widely recognized in recent years (see, e.g., Reber, 1989; Stanley et al., 1989; Willingham et al., 1989; Proctor & Dutta, 1995; Sun, 2002; Sun et al., 2005). Although implicit learning has been actively investigated, the complex and multifaceted interactions between the implicit and the

explicit and the importance of this interaction have not been universally recognized; to a large extent, such interactions have been downplayed or ignored, with only a few noteworthy exceptions. Research has been focused on showing the lack of explicit learning in various learning settings (see especially Lewicki et al., 1987) and on the controversies stemming from such claims. Similar oversight is also evident in computational simulation models of implicit learning (with a few exceptions).

There is increasing recognition that it is difficult, if not impossible, to find a situation in which only one type of learning is engaged (Reber, 1989; Seger, 1994; but see Lewicki et al., 1987). Our review of existing data (see Sun et al. 2001; 2005) has indicated that, while one can manipulate conditions to emphasize one or the other type, in most situations, both types of learning are involved, with varying amounts of contributions from each (see, e.g., Sun et al., 2001, 2005; Stanley et al., 1989; Willingham et al., 1989).

In the work described below, we will focus on interactions between the two systems and resulting effects on performance. The guiding hypothesis of this work is that the interaction between implicit and explicit knowledge is the key to understanding acquisition of skills. Further, we believe this work has implications for understanding how best to teach the implicit and explicit skills associated with science.

2. CLARION: A Computational Theory and an Integrative Model

We next discuss a computational theory (namely CLARION) of human cognitive processes that incorporates both implicit and explicit processes. We do so because we believe that the best applications are those which are informed by theory and theory-driven empirical research.

Representation. Implicit knowledge is like pattern matching; we know who it is we are looking at, or how to react to a situation, but we are not aware of how we know. For example, when you recognize a friend's face you probably are not aware of what features are used to identify him—you just know it is "George." In computational modeling, this relatively inaccessible nature of implicit knowledge may be captured by subsymbolic distributed representations provided by a connectionist (i.e., "backpropagation") network (Rumelhart et al., 1986). This is because representational units in a distributed representation are capable of accomplishing tasks, but are subsymbolic and generally not individually meaningful (see Rumelhart et al., 1986; Sun, 1995). Although it is not true that distributed representations are never accessible to conscious awareness, they are definitely less accessible (e.g., more difficult to verbalize or tell someone else how to perform the task), than symbolic/localist representations. Distributed representations may only be accessed through indirect, transformational processes. Such processes may be compared to watching your own behavior and discovering how you perform a task. For example, when asked by his young daughter how to balance a bicycle while starting off, one of the authors had to get on a bike and perform the task to recognize that turning the front wheel was the key.

In contrast, explicit knowledge may be captured in computational modeling by symbolic or localist representations (Clark & Karmiloff-Smith, 1993) in which each unit is easily interpretable and has a clear conceptual meaning. This is similar to memorizing steps in a procedure or recipe (i.e., an algorithm). Each step is meaningful and is described to tell another how to do it. This characteristic captures the property of explicit knowledge being accessible and manipulable (Sun, 1995).

This difference in the representations of the two types of knowledge led to a two-level cognitive architecture CLARION (which stands for Connectionist Learning with Adaptive Rule Induction ON-line; proposed in Sun 1997; 2002), whereby each level using one kind of representation captures one corresponding type of process (either implicit or explicit). Greater detail regarding the theoretical arguments for such two-level models can be found in work by Sun (1995; 1997; 2002).

Learning. Implicit knowledge is tuned to our environment over time through feedback from its effects (positive or negative) each time it is employed. We experience much of this feedback process through emotional reactions to our success or failure when trying to accomplish a task. In computational modeling, this continuous learning process (at the bottom level of CLARION) can be captured in a variety of ways consistent with the nature of distributed representations. In learning settings where the correct responses are not explicitly provided, reinforcement learning (Watkins, 1989, especially Q-learning) can be implemented using backpropagation networks. Such learning methods are theoretically and empirically justified: For example, Shanks (1993) showed that human instrumental conditioning (a simple type of skill learning) was best captured by associative models (i.e., connectionist networks) when

compared with a variety of rule-based models. Further, Cleeremans (1997) has argued that implicit learning could not be captured by models using only symbolic representations.

Whereas implicit learning is a gradual, continuous process, explicit learning can occur on one trial (at the top level of CLARION). For example, if I tell you that you must stop at a red light, a new driver can immediately adopt this behavior without a need for continuous feedback. Explicit learning (at the top level) can also be modeled in a variety of ways using symbolic/localist representations. While complex mental models can and are adapted over time, CLARION uses one-trial learning to create top-level explicit rules. Such rules may be forgotten or discarded over time, but rule creation is thought to be a one-time process (the rule is either present or not, see Nosofsky, Palmeri, & McKinley, 1994; Sun, 1997). For example, if you tell a class that “all living things interact with their environment”, we can assume that for many students (those paying attention) this new rule is now explicitly encoded in their memories and it is available to be used perhaps to distinguish between living and non-living things. Of course, this is not to say it was represented in every student’s head (we wish!) or that it will be remembered when needed. Nonetheless, with such learning, individuals explore the world, and dynamically acquire explicit representations and modify them as needed, reflecting the dynamic (on-going) nature of skill learning (Sun, 1997; Sun et al., 2001).

The implicit knowledge already acquired in the bottom level can be utilized in learning explicit knowledge (through so called “bottom-up learning”; Sun et al., 2001). This is accomplished in CLARION by having the top level examine a positive outcome from using implicit knowledge and generating a rule that encompasses the current context and the implicitly controlled response. Such bottom-up learning is also likely to happen following failure of explicit (top level) rules. When this occurs the explicit rule is rejected (deleted) and a new rule is created through CLARION “observing” its own implicitly controlled behavior (e.g., Dienes & Fahey, 1995; Nosofsky et al., 1994). This rejection of failed rules occurs in CLARION when a measure of a rule’s performance falls below the deletion threshold and the rule is deleted.

Whenever all the rules of a certain form are deleted, a new set of rules of a different form are hypothesized, and the cycle repeats itself. In hypothesizing rules, CLARION progresses from the simplest rule form to the most complex (just as humans have been shown to do; Berry & Broadbent, 1988; Stanley et al., 1989).

As in real life, CLARION often faces situations where the two systems suggest different actions. Should I go with my gut (implicit) feelings or my head (my explicit plan)? CLARION makes this decision probabilistically based on a **weighting parameter** which biases the system toward using the bottom (implicit) or top level (explicit) in a given situation. This weighting parameter is one of the primary variables we changed to simulate different task instruction effects in the following experimental simulations. A second parameter, the **rule threshold** parameter is also changed to simulate increased activity in generating rules at the top level. All other parameters are fixed based on a long standing attempt to build a computer architecture complex enough to resemble the human mind in some way (for further details see Sun, 2002; 2003).

III Human Experiments and Model Simulations

One paradigm we have used to examine interactions between implicit and explicit learning involves learning an artificial language (Reber, 1989). These experiments typically involve seeing many exemplars of letter strings that are created by a set of rules (an artificial grammar). Just like learning a real language, this paradigm offers the opportunity to study the effects of experiencing exemplars (mostly implicit knowledge) and teaching (or self-induction) of the rules of the grammar (i.e., its underlying structure). The basic elements of this task resemble, for example, situations where students learn categories from exposure to exemplars of the category (e.g., when they learn about the category “crawfish” by seeing different specimens).

1. Task Description of Artificial Grammar Experiment


The data reported here come from a published study by Domangue, Mathews, Sun, Roussel and Guidry (2004). The finite-state grammar used by Mathews et al. (1989) was used in this experiment (see Panel B of Figure 1). This grammar generates 177 exemplars (“words”) ranging in length from 5 to 11

letters. There were two 1-hour sessions conducted one week apart. Each session began with a 20-minute training phase requiring participants to perform a training task, followed by a 20-minute cued-generate test. One hundred twenty undergraduate students participated in the experiment.


Participants in the EP (implicit) training groups were instructed to copy as many of the 88 instances (exemplars) as possible into the response sheets (top of Figure 1) in 20 minutes. Each letter of each exemplar was to be copied into the appropriate circle on the sheet. Thus these participants experienced many exemplars without a need or time to reflect on the grammar rules.

The GR (explicit) training condition required participants to observe a copy of the artificial grammar (Panel B of Figure 1) for 2 ½ minutes then turn the diagram over. For another 2 ½ minutes participants reproduced the artificial grammar diagram from memory by drawing it on a blank sheet of paper. This was repeated four times for a total of 20 minutes training time. The goal of the GR task was to teach an explicit representation of the grammar without showing many valid letter strings that could stimulate implicit learning. However, it was essential that participants understood how to use the diagram to generate strings. Therefore, prior to the first session, three test cues were used to demonstrate how to generate strings using the diagram.

Participants in the ED (exemplar diagramming) groups were instructed to trace as many of the 88 exemplars through the diagrams (bottom of Figure 1) on their response sheets as possible in 20 minutes. They were instructed to copy each letter of each exemplar into the corresponding transition box until the exemplar was completed. The rationale for this task was to have participants process exemplars within the context of the grammatical structure. Thus this task provided integrated training both in the structure of the grammar and in seeing exemplars generated by the grammar.



Panel A



Panel B

Figure 1

During the test phase, participants were provided with a set of blanks with two letters (cues) filled in. Their job was to fill in missing letters in a string (a cued generation task). If the letter string generated



by the participant did not match at least 70% of the letters of the closest not-yet-generated instance, all non-matching letters were erased from the screen and the participant had to try again. This process was continued until at least 70% of the letters typed by the participant matched an instance. When the 70% criterion was achieved, the computer retrieved and displayed the closest not-yet-generated instance.

It was predicted that participants having only the purely implicit (EP) training would generate strings the fastest, using only fast implicit processes. It was expected that the purely explicitly (GR) trained group would be the most accurate, but the slowest, using only explicit knowledge. The integrated (ED) training was expected to fall in between the two pure groups, employing some fast implicit processes combined with slower explicit knowledge.

2. Results of Human Experiment

Performance measures for both sessions are provided in Figure 2. Figure 2 illustrates the performance of various training tasks on three dependent measures.

The following measures were used: (1) Achievement is measured in terms of the number of acceptable strings (matching at least 70% of the letters of an instance) generated per minute during the test phase. (2) Accuracy is a measure of the proportion of attempts that generate perfect letter strings. (3) Speed of responding was measured in terms of number of attempts per minute during the test phase.






Figure 2. Results of Artificial Grammar Experiment

Recall that the goal of the test was to complete as many strings as possible by getting close (70%) to a valid string. Success in this goal is measured by achievement. Note that purely implicit training resulted in the best achievement of this goal by the end of session two. This training also led to the fastest performance (see speed in Figure 3). This shows that purely implicit training can be very effective when speed is needed, but a high level of accuracy is not.

Strictly explicit training (GR) led to the highest accuracy and the slowest responding on the cued-generate test. The integrated training was in between, having higher accuracy and lower speed than EP, and lower accuracy but higher speed compared to GR. Unfortunately, the integrated training did not lead to the best of both worlds—both high speed and accuracy. However, as will be mentioned in the discussion, we have since found a way to do just that!

3. Simulation Using CLARION

We simulated the above human data with CLARION (Sun et al., 2001, 2005; Sun, 2002). As mentioned before, the model (the cognitive architecture) consists of two levels: the top level encodes explicit knowledge and the bottom level encodes implicit knowledge. Experienced strings (as presented to participants or sampled from presented grammar diagrams) are used to train an associative memory made up of a backpropagation network. The network maps input to output; in this particular case, it maps some partial strings (each of which is a part of an experienced string) to the full experienced string. This associative mapping allows implicit grammatical knowledge to develop.ⁱⁱ

At the top level, experienced strings are encoded as rules. For example, if a string "SCP V" is experienced, the following three rules may be encoded there: S->C, C->P, P->V.

The outcome from the model can be either from the bottom level or from the top level. However, the bottom-level implicit processes are significantly faster than the top-level explicit processes (see Schneider & Oliver, 1991; Sun & Zhang, 2004). In CLARION, response time is determined by parameters that specify the time lag of each step of associative memory retrieval at the bottom level, and the time lag of each step of rule application at the top level.

For the explicit (GR, GR) group, the top level is mainly responsible for generating the outcome during test. This is because the system was configured in such a way that mainly the top level is used, due to the fact that this experimental setting encourages an explicit mode because of the presentation of grammar diagrams to participants during training. (The top vs bottom weighting parameter was set during training in a way that supports this configuration.) During test, the top level uses learned rules to attempt

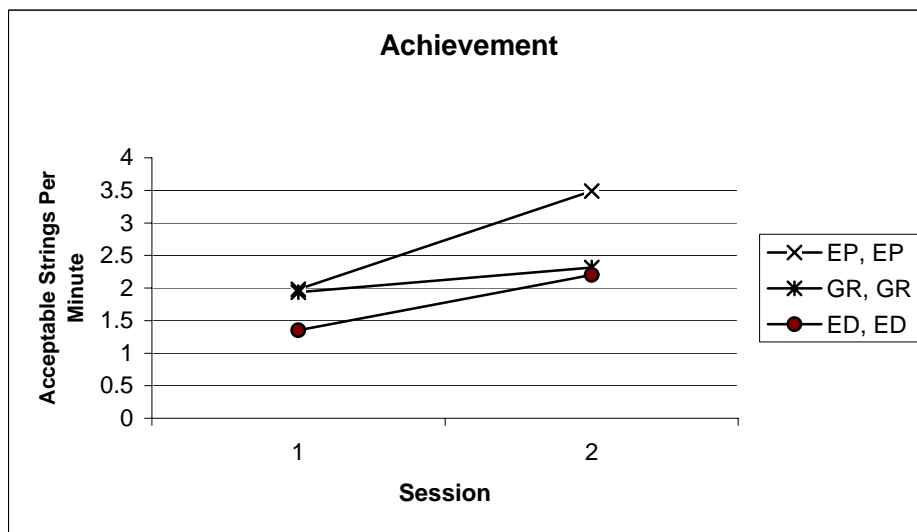
to complete each given partial string. That is, given the test cue, it searches for a possible completion guided by the rules at the top level, using depth-first search (with backtracking). For example, given a partial string ``S __ V'', the search has to go through all the rules in the form of S->x, or in the form of x->V, where x can be any letter, and many other similar rules (e.g., concerning the relation between the second and third letters). This search process is slow, but the outcome from the top level is rather accurate. ⁱⁱⁱ

For the implicit (EP, EP) group, during test, the bottom level is responsible for generating the outcome. This is because the system is configured in such a way that mainly the bottom level is used, due to the fact that this experimental setting during training encourages an implicit mode by repeatedly presenting training instances (Sun, 2002). (The top vs bottom weighting parameter was set during training in a way that supports this configuration.) During training, the bottom level uses an associative memory (in the form of a backpropagation network) to map a partial string to a full string that is a likely completion of the partial string. This way of capturing implicit learning during training is appropriate, considering the fact that participants in this task marked experienced strings on a bubble sheet, which naturally led to multiple partial strings. The bottom level is, generally speaking, less accurate but much faster.

For the integrated training (ED,ED) group, a combination of the two levels was used, because the experimental settings involve both implicit training and explicit training, due to the use of both repeated presentation of strings and the presentation (and tracing) of the grammar diagram. During test, the combination process of the two levels proceeds this way: The bottom level generates candidate completions of partial test strings; then the top level checks each of these strings using the rules already learned at the top level. The check by top-level rules is carried out through straightforward applications of relevant rules, without any backtracking. For example, if ``S C P V'' was suggested by the bottom level, at most three rules may be applied: S->C, C->P, P->V, if these rules do exist at the top level. ^{iv}In this case, although the bottom level works at a fast pace, the top level is slower. But because there is no full-blown depth-first search with backtracking, the top level is not as slow as in the case of the explicit (GR, GR) group. However, due to multiple applications of rules it is slower than the bottom-level implicit processes alone. So, the final outcome is, on average, at a speed somewhere between the implicit group and the explicit group.

The purpose of the simulation was to see if a model using dual representational structures could capture and explain the key features of the data. No attempt was made to fine tune the fit of the model by varying parameters, because at this stage we are only interested in the overall features of the data.

The key feature we were trying to capture in the simulation was that exposure to a diagram of the grammar either through grammar replication (GR) or exemplar diagramming (ED) would enhance accuracy and efficiency but such exposure would reduce speed. In addition, a high level of achievement could be accomplished through implicit processing (EP,EP) alone, without exposure to a diagram of the grammar. The simulation results are shown in Figure 4. Let us examine the data from the test phase of session 2 below.



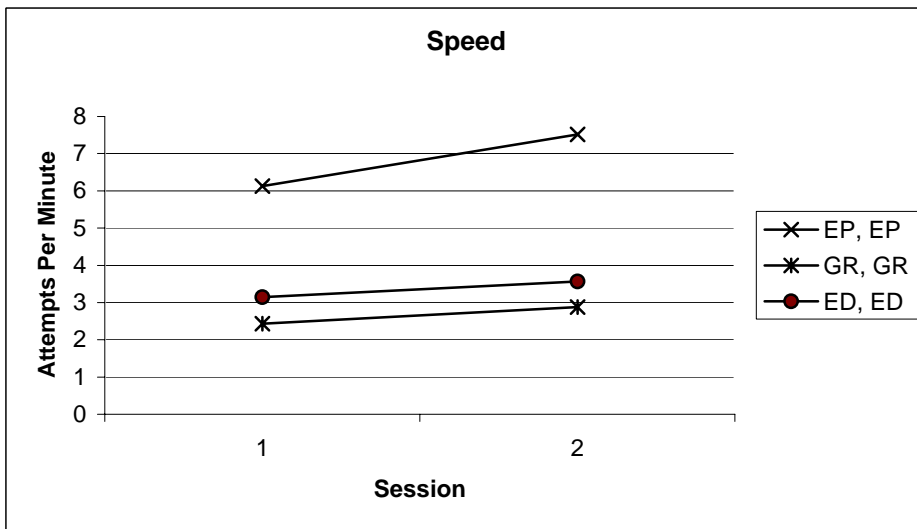
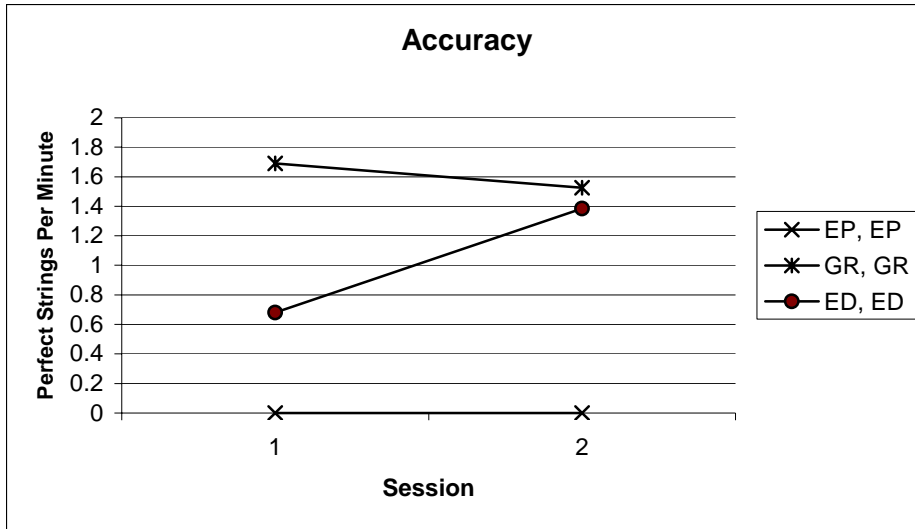


Figure 4. Artificial grammar learning simulation illustrates the performance of the CLARION model capturing the human data.

The simulation results for the accuracy data (from the test phase of session 2) are quite similar to the human data. In session 2, the two highest groups in both the simulation and the human data were those exposed to a diagram of the grammar. (However, in the human accuracy data, the grammar replication (GR,GR) group was far superior to all other groups, while in the simulation it was slightly better.)

However, exposure to the diagram also reduced the speed of string generation. In both the human and the simulation data, the group only exposed to implicit training (EP,EP) were the fastest during session 2. Finally, the implicit only group (EP,EP) had the highest level of achievement in both the simulation and the human data during session 2.

Looking at the accuracy and the speed data (the two contrasting measures), during session 2, it is clear that the integrated training group (ED,ED) was somewhere between the implicit group (EP,EP) and the explicit group (GR,GR), in both the human data and the simulation data. This is an important match between the human and simulation data, validating our earlier prediction concerning the effect of integrated training.

Thus, the simulation results support the notion that a dual representational model can account for these data. Our simulation using CLARION has produced some interesting interpretations (explanation) of the human data. These interpretations are embodied in our simulation setups as described earlier.

They described a plausible mechanistic underpinning of human performance in this task based on the implicit-explicit interaction. In particular, they provide an explanation of why the integrated training group performed the way it did.

4. The Effect of Hints and Reflection on Learning a Process Control Task

Another paradigm often used to study implicit learning is the process control task (e.g., Berry and Broadbent, 1988). In this type of task, an input variable is changed over trials by learners trying to keep an output variable at a specific level. In this section we present and then simulate the results of one such well known study by Berry and Broadbent (1988). As will be seen, the results of this study suggest that simultaneous use of implicit and explicit processes can be useful in learning an "implicit" task. However, when the task becomes more difficult, reflection is actually harmful to learning.

The Task. Subjects interacted with a computer simulated "person" whose behavior ranged in twelve levels of affection going from "very rude" (level 1) to "loving" (level 12) and the task was to maintain the behavior at "very friendly" (level 6) by controlling his/her own behavior (which could also range in the same 12 levels, from "very rude" to "loving"). Letting "P" stand for the level of the computers response on the next trial and "W" the learner's response on a given trial, In the salient version of the task, the behavior of the computer "person" was determined by the immediately preceding input of the subject (W). In the nonsalient version output was determined by the input made two trials ago (denoted as W1, instead of the current input W). In addition, random variation (i.e., noise, denoted as N) was added to both equations determining the response (output) of the computer "person" so that there was a chance of being up or down one level (a 33% chance respectively). So in the salient version the equation relating inputs and outputs was $P = W - 2 + N$. In the non-salient version, it was $P = W1 - 2 + N$. Four groups of subjects were used: salient experimental, salient control, non-salient experimental, and non-salient control. The experimental groups were given explicit rule search instructions after the first set of 20 trials, and after the second set of 20 trials were given explicit how-to instructions in the form of indicating the relevant input that determined the computer responses (W or W1). The control subjects were just told to try to keep the computer's response at a friendly level (level 6).

The Data. The exact target value plus/minus one level (that is, "friendly", "very friendly", or "affectionate") was considered "on target." The average number of trials on target was recorded for each subject for each set of 20 trials.

Figure 5 shows the data for the four groups of subjects for the three sets of trials. Analysis showed that on the first set, neither of the two experimental groups differed significantly from their respective control groups. However, on the second set, the salient experimental group scored significantly higher than the salient control group ($p < 0.01$), but the non-salient experimental group scored significantly lower than the non-salient control group ($p < 0.05$). On the third set, both experimental groups scored significantly higher than their respective control groups ($p < 0.01$). The results clearly showed (1) a benefit of explicit rule search in the salient condition and a negative effect of rule search in the non-salient condition; (2) Providing a hint about how inputs were related to outputs in the final set of trials for the experimental groups facilitated performance in both the salient and non-salient conditions. Finally there was a main effect of salience in the second set of trial, as both salient groups performed better than both non-salient groups.

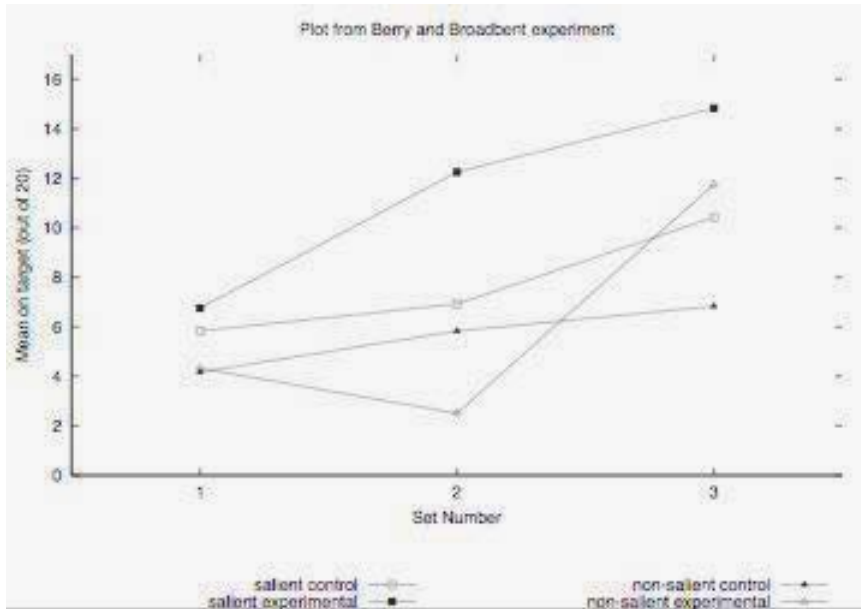


Figure 5: The data of Berry and Broadbent (1988).

The Model Setup. To capture in CLARION the explicit rule search condition (during the second training set), the rule generation threshold in CLARION (as mentioned before) was raised to cause more activity (rule generation) at the top level. To capture the explicit hint condition in the third session, only rules that related the given critical variable were generated in the third session for the experimental groups.

The Match. The results of the CLARION simulation are shown in Figure 2. The data demonstrated clearly the explicit how-to instruction effect (improving performance in all conditions), and showed to some extent the explicit search effect (improving performance in the salient condition and worsening performance in the non-salient condition), as well as the salience difference effect along with the explicit search effect. The data showed the extent and the limit of synergy (in that the non-salient condition discouraged synergy).

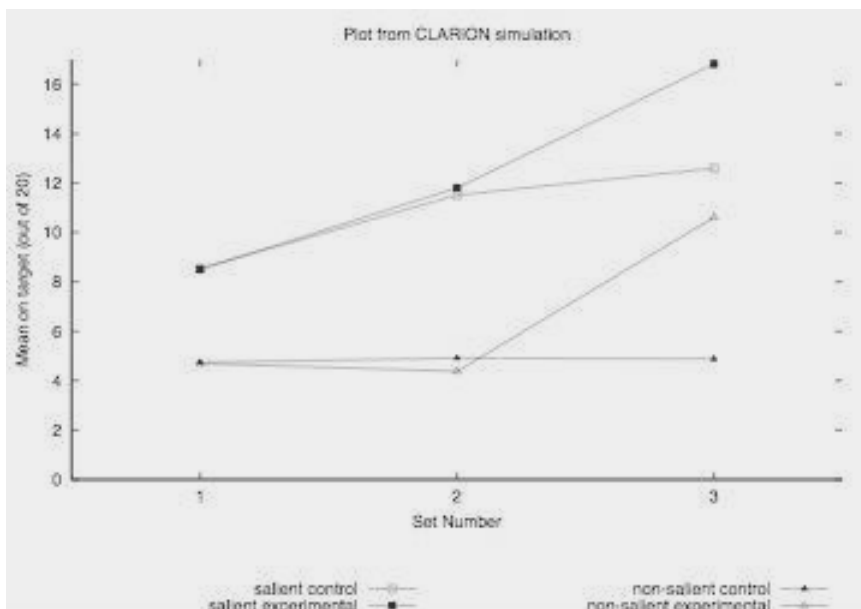


Figure 2: The simulation of Berry and Broadbent (1988).

5. Synergy in a Process Control Task

In the Berry and Broadbent (1988) study, there was evidence that providing hints was effective in both the salient and non-salient tasks. Before considering implications for teaching science, we present one more study (Stanley, et al., 1989). However, as will be seen, the critical result from this study is that explicit learning processes can enhance learning of a difficult implicit task if explicit information is provided (in the form of a hint) or people reflect on their performance *after having experience with the task* (delayed reflection).

The Task. Two versions of the process control task were used in Stanley et al. (1989). In the "person" version, subjects were to interact with a computer simulated "person" whose behavior ranged from "very rude" to "loving" (over a total of 12 levels) and the task was to maintain the behavior at "very friendly" by controlling his/her own behavior (which could also range over the 12 levels, from "very rude" to "loving"). In the sugar production factory version, subjects were to interact with a simulated factory to maintain a particular production level (out of a total of 12 possible production levels) by adjusting the size of the workforce (which has 12 levels). In either case, the behavior of the simulated system was determined by this equation $P = 2 * W - P1 + N$, where P was the current system output, P1 was the previous system output, W was the subjects' input to the system, and N was noise. Noise (N) was added to the output of the system, so that there was a chance of being up or down one level (a 33% chance respectively). Beside the noise, this task is difficult because increasing workers does not always lead to increased output (it depends on the current level of output (P1)). Thus, the system often behaves in counter-intuitive ways.

There were four groups of subjects. The control group was not given any explicit how-to instruction and not asked to verbalize. The "original" group was required to verbalize: Participants were asked to verbalize *after* each block of 10 trials. Other groups of participants were given explicit instructions in various forms, for example, "memory training", in which a series of 12 correct input/output pairs was presented to participants, or "simple rules", in which a simple heuristic rule ("always select the response level half-way between the current production level and the target level") was given to participants. All the participants were trained for 200 trials (20 blocks of 10 trials).

The Data. The exact target value plus/minus one level (that is, "friendly", "very friendly", or "affectionate") was considered on target. The mean scores (numbers of on-target responses) per trial block for all groups were calculated. Analysis showed the delayed verbalization effect: The score for the "original" group was significantly higher than the control group ($F(1, 73) = 5.20; p < 0.05$). Analysis also showed the explicit hint effect: The scores for the memory training group and for the simple rule group were also significantly higher than the control group. See Table 2.

Table 2: The human data for the process control task from Stanley et al (1989).

	Human Data	
	Sugar Task	Person Task
Control	1.97	2.85
Original (Delayed Reflection)	2.57	3.75
Memory Training	4.63	5.33
Simple Rule	5.91	4.00

The Model Setup. To capture in CLARION the delayed verbalization condition, the rule creation threshold in CLARION was increased in the verbalization groups (the "original" group), implying that verbalization leads to increased rule generation at the top level. To capture the explicit instruction conditions, in the "memory training" condition, each of the 12 examples given to participants was represented as rules in the top level of CLARION; in the "simple rule" condition, the simple rule (as described earlier) was represented at the top level of CLARION.

The Match. Our simulation captured the delayed verbalization effect in the human data well (for

further details, see Sun et al., 2005). Compare Tables 2 and 3. We used a t-test to compare the "original" group with the control group in the model data, which showed a significant improvement of the original group over the control group ($p < .01$), the same as in the human data.

Table 3: The model data for the process control task from Stanley et al (1989).

	Model Data	
	Sugar Task	Person Task
Control	2.276	2.610
Original (Delayed Reflection)	2.952	4.187
Memory Training	4.089	5.425
Simple Rule	4.073	5.073

Our simulation also captured the explicit instruction effect, as shown in Table 2. We used pairwise t-tests to compare the "memory training" and "simple rule" groups with the control group in the model data and found significant improvements of these two groups over the control group ($p < .01$, respectively), the same as in the human data.

Both effects point to the positive role of the top level of the model. When the top level is enhanced, either through delayed verbalization or through externally provided hints, performance is improved, although such improvement is not universal (e.g., see Sun et al., 2001). In a sense, both findings demonstrate synergy between top-level explicit processes and bottom-level implicit processes.

6. Summary and Discussion

We would summarize the results of the presented empirical results in the following way. First, the artificial grammar study (Domague, et al., 2004) demonstrated that implicit learning strategies often lead to fast, but relatively inaccurate performance, while explicit learning strategies often lead to highly accurate, but slower performance. This study also illustrated that attempting to combine these two types of strategies sometimes leads to compromised performance. Second, the Berry and Broadbent (1988) study illustrates that reflection (i.e., thinking about the possible rules underlying the task) *while learning the task* is only helpful if the task is very easy and straightforward (otherwise, it can impair performance). The study also illustrated the potential utility of hints to implicit learning. Finally, the Stanley et al. (1989) study illustrates that reflecting on task rules can sometimes be helpful when the reflection is done after substantial experiential learning (i.e., when one does not attempt to simultaneously learn the task and reflect on the rules). The study also demonstrated that task hints (explicit information) can improve learning if the information directs attention to important task features or strategies.

Note that although implicit learning is a controversial topic, the existence of implicit processes in skill learning is not in question. What is in question is their extent and importance. CLARION allows for the possibility that both types of processes and both types of knowledge coexist and interact with each other to shape learning and performance, so it goes beyond the controversies and the studies that focused mostly on the minute details of implicit learning.

The work reported thus far highlights the importance of the interaction of implicit and explicit processes in skill learning. It captures the interaction through a model (a cognitive architecture) that includes both types of processes. This model reveals something new in the existing data (cf. Gibson et al., 1997; Lebiere, Wallach, & Taatgen, 1998): The contribution of this model lies in interpreting human data in skill learning through the interaction of the two types of processes, and also in demonstrating the computational feasibility and psychological plausibility of bottom-up learning (Sun et al., 2001). Note that many other simulations using CLARION have similarly shown the interaction between implicit and explicit knowledge during skill learning (see Sun, 2002 for details).

Although theoretical and empirical rigor as demonstrated above is critical to our understanding of the processes underlying skill learning, ultimately the goal of many of our readers concerns the utility of this information. In the next section, we attempt to distill the results of research in this area for science education.

IV Applications to science education

In this section, we will note the implications of CLARION as well as the related vast literature on explicit and implicit learning for the teaching of science. To do so, we first contrast the characteristics of implicit and explicit learning, noting how these characteristics might bear on science education. However, implicit and explicit learning are hardly ever found in “pure” form in real-world tasks (Mathews, et al., 1989, Sun 2002) and, more importantly, the ultimate goal of science education is to produce students who can use both implicit and explicit knowledge to solve problems they encounter. Thus, we also discuss how explicit and implicit learning processes interact, and highlight the fact that this interaction can sometimes interfere and sometimes enhance knowledge and skill acquisition. Finally, we will focus on some suggestions regarding how science instruction might be structured to increase the likelihood that students will emerge with rich implicit and explicit knowledge (Sahdra and Thagard 2003).

As discussed earlier in the chapter, one strength of implicit learning is that it is more error tolerant than explicit learning (e.g., Lane, Mathews, Sallas & Sun, 2006). More generally, we mean that implicit learning is more likely to represent the natural variability of a category than explicit learning (which involves constructing potential rules). For example, one common task in a biology course is to identify the genus and species of a given organism. When learning this task, students typically see a variety of related organisms that differ, subtly and not-so-subtly, from each other. In most natural categories, the presence or absence of a single feature does not discriminate between related category members but rather they share a “family resemblance” (i.e., they share a number of characteristic features). However, explicit learning often focuses on trying to identify simple rules (e.g., If specimen contains X, then it is a Y). Implicit learning processes are more likely to lead to a representation that contains the variability within the category. Thus, it is clear that experience with the members of the category under study can provide important information that is of considerable use during the task of identification. After all, it is entirely possible for students to confidently state rules they learned in class without being able to apply them (e.g., using that knowledge to identify a target organism). However, one common misconception about the use of implicit learning processes in the classroom is that it involves no direction on the part of the teacher (i.e., it simply comes from experience). An important thing to keep in mind is that implicit learning is exquisitely tuned to the examples to which it is exposed. The more varied and representative the examples are for the category being learned, the more robust the ultimate representation of that knowledge. Thus, one implication for science teaching is that the type of examples to which students are exposed will impact the knowledge gained from experience. For example, it is important to not only expose students to typical members of a category, but also less typical members as well as members of other categories that are commonly mistaken as being from the target category.

Choosing good, representative examples in science teaching is also important for another reason. It has often been claimed that the representations created by explicit learning are more flexible than those which result from implicit learning (e.g., Dienes & Berry, 1997). More generally, the claim is that explicit learning leads to knowledge that transfers more easily to new situations than implicit learning. Given that transfer is one of the major goals of education, this might be worrisome to those wishing to integrate more experiential aspects of learning into their science instruction. However, there are several reasons to call this claim into question. First, people are not particularly good at transferring *explicitly* learned information to new situations (e.g., Gick & Holyoak, 1980). Second, research on implicit learning has found evidence that the knowledge acquired with one set of task constraints can be applied to a task with new constraints (e.g., Mathews et. al., 1989). In other words, implicit learning can lead to transfer. Finally, note again that implicit knowledge is tuned to the range of experience to which it is exposed. To the extent a student might see only a small, limited range of examples, one might expect that his or her knowledge of the category would be limited as well.

Although implicit knowledge is an important and often neglected aspect of the learning process, ultimately science practice involves the integration of both implicit and explicit knowledge (Sahdra and Thagard 2003). For example, we have argued that pattern matching (a central feature of implicit knowledge) is important in nearly every area of science. Biologists are able to recognize organisms at scales ranging from viruses to large mammals. Astronomers distinguish between various types of celestial phenomena. Radiologists are required to make subtle distinctions when viewing visual images. Although each of these examples clearly demonstrate the essential nature of perceptual skill in scientific pursuits, this type of implicit knowledge alone is not enough. Instead, explicit knowledge and reasoning processes must also be brought to bear. For instance, one important aspect of science practice is the ability to detect anomalies (e.g., Popper, 1972); the feeling that some piece of data “does not fit.”

Although the initial feeling may be the output of an implicit process, it typically triggers an explicit process of attempting to understand the anomaly (e.g., “Perhaps this is a new strain of bacteria.”). In addition, the ability to communicate what one knows (explicit knowledge) is central to the scientific enterprise for a multitude of reasons including the issues of replicability (e.g., conveying procedures), the practice of peer review and mutual critiquing, and the increasingly social nature of scientific work (e.g., groups of scientists planning, running and interpreting the results of experiments). Thus, one cannot become a scientist or even an educated student of science without being able to deploy both implicit and explicit scientific knowledge in an integrated way.

Research has revealed that integration between implicit and explicit knowledge does not always happen automatically. For example, a study of “naïve physics” revealed that even students who successfully completed a college-level physics course nevertheless still wrongly predicted the trajectory of a ball in a simple task (McCloskey & Kohl, 1983). Readers who are science teachers will similarly recognize the lack of integration either in those students who are able to perform procedures without being able to articulate the rationale or in those who are able to articulate concepts on an exam but are unable to use those concepts in a laboratory task. Further, attempts to integrate implicit and explicit learning at the same time have sometimes been harmful or ineffective (e.g., Lane, et al., 2006; Roussel, 1999) and sometimes helpful (e.g., Sun, Merrill, and Peterson, 2001; Sallas, Mathews, Lane, & Sun, 2006). Although there is still much to be learned about how to maximize the integration, there are nevertheless a number of lessons to be learned from past research that are applicable to science instruction.

One common way of attempting to integrate explicit and implicit knowledge is to have people reflect on their understanding during the process of learning a task. In order to understand the impact of reflection, it is important to keep in mind that its effects can differ based on the type of knowledge to be acquired, the nature of the task, and the timing within the knowledge acquisition process (i.e., how far along in the learning process) when it is engaged. For example, reflecting (e.g., self-explanations) while reading a text describing the circulatory system is effective in increasing comprehension and repairing misconceptions (e.g., Chi, 2000; Chi, de Leeuw, Chiu, & LaVancher, 1994). In this case, the material is largely conceptual and declarative in nature, and reflection is clearly beneficial. However, there are a number of circumstances, where it can be harmful. For example, verbalizing one’s thoughts about difficult-to-verbalize aspects of one’s knowledge can impair performance (i.e., *the verbal overshadowing effect*, Schooler & Engstler-Schooler, 1990; Schooler, 2002). Indeed, verbalization has been demonstrated to impair insight problem solving (Schooler, Ohlsson, & Brooks, 1993), analogy retrieval (Lane & Schooler, 2004), and memory for faces (see Meissner & Brigham, 2001 for a meta-analysis). Similarly, our own research has found that attempting to reflect on one’s understanding during the process of learning tasks that contain non-salient or complex rules can sometimes harm knowledge acquisition (e.g., Lane, et al., 2006; Roussel, 1999; see also, Berry & Broadbent, 1988). The timing of reflection also appears to be critical. For example, in contrast to concurrent reflection, reflection that is delayed to the end of a block of trials does not appear to impair learning (e.g., Roussel, 1999). Further, the level and type of knowledge attained at a particular point in time can also have an impact. The verbal overshadowing literature (e.g., Melcher & Schooler, 1996) suggests that reflection may be less harmful when people already have both verbal and non-verbal (mostly implicit) knowledge of a domain than it is for people who have only non-verbal knowledge. CLARION has been used to explain these phenomena (Sun, 2002; Sun et al., 2005)

Although work in this area is continuing within the framework of CLARION, research suggests a number of important considerations for teachers. First, not surprisingly, when attempting to teach conceptual (explicit) aspects of an area, reflecting while learning may aid the acquisition of this knowledge. Second, having students reflect when first learning tasks that are largely perceptual (e.g., learning to categorize samples) or involving difficult-to-verbalize stimuli is likely to impair knowledge acquisition. However, if reflection is utilized in these situations, it should be used sparingly and non-concurrently (e.g., during breaks between seeing extensive examples). Our findings suggest the possibility that concurrent reflection may tie up working memory resources that are typically utilized during the learning process. Third, the negative effects of reflection appear to be reduced after initial learning and may be beneficial after students have begun to develop both implicit and explicit knowledge of a domain. In conclusion, we believe that the existing research suggests, for initially learning relatively experiential aspects of a domain, having students concurrently reflect about what they are learning is not likely to lead to integration and may actually impair acquisition of that knowledge. Synergy is more likely

to occur when implicit and explicit knowledge are initially acquired separately and then integrated over time. For example, some studies suggest that learning in these types of tasks might be best facilitated by an initial period where students are exposed to good examples of category (mostly implicit learning) before they receive more explicit instruction (Domangue, et al. 2004; Mathews, et al., 1989). In other words, rather than attempting to integrate explicit and implicit knowledge by co-activating both types of learning processes (using reflection + experiential learning), this suggestion involves a period of implicit learning prior to explicit learning.

Although the subject of much research, reflection is only one way that explicit knowledge and learning processes can be brought to bear on learning difficult experiential tasks. We have found a number of other ways that explicit and implicit learning can lead to synergy. For example, a number of studies have indicated that giving participants explicit hints about successful task completion improves performance on an implicit task (Roussel, 1999; Sun & Mathews, 2005) relative to participants who only learn the task experientially. In general, we argue that such hints direct attention to important aspects of the task without burdening working memory resources. Our simulation with CLARION has confirmed this (Sun et al., 2005). In addition, recent work in our laboratory has discovered a way that explicit knowledge can be introduced in ways that aid the implicit learning of an artificial grammar task. Specifically, this research (Sallas, et al., 2006) found that participants who saw example "words" in the context of an animated diagram of the grammar performed more accurately on a test and were just as fast as participants who saw the "words" alone (in contrast to the Domangue, et al., 2004 study discussed in an earlier section). It is important to note that it is not animation, per se, that improves learning in this task as the benefit is only conferred when combined with diagram information. Instead improved performance appears to result from the provision of explicit information in a form that reduces the working memory load of participants while directing attention to important task information. From the viewpoint of teaching, both studies suggest that experiential learning can be aided by the judicious use of information that points students in the right direction.

Our work also has relevance for more general strategies that are employed in the teaching of science. In many of our research studies, "experiential" participants learn the task on their own after being given a general goal (e.g., "learn how to control the nuclear reactor so it stays at 6000"). Given that these participants often learn the task quite well, one might surmise that trial-and-error is an efficient strategy for learning implicit tasks. However, our work (e.g., Prattini, 2006) has documented that participants who are explicitly told to use a trial-and-error strategy do quite poorly. Instead, it appears the nature of the task and the general goal lead participants to engage in a relatively systematic (somewhat explicit) exploration of the problem domain. Trial-and-error participants, on the other hand, do not apparently learn enough from their exposure to examples to constitute sufficient coverage of the domain. The data from this and other studies, along with the research on hints noted above, suggests that a guided instruction approach to the teaching of experiential aspects of science knowledge may be superior to "purely" experiential approaches, which is directly related to our earlier point about synergy; for more general arguments on this issue pertaining to science instruction, see Klahr & Nigam (2004). Although we believe a more guided instruction approach to be superior, we would emphasize that such guidance should be subtle rather than extensive. For example, guidance could come in the form of brief instructions (e.g., pay attention to feature X and how it relates to feature Y), by choosing examples carefully, and by structuring the task to provide suitable feedback to students.

We believe that teaching skills requires teachers to be actively involved in teaching scientific concepts as well as actually using the knowledge doing science. It also requires a willingness to not only tell students the rules (explicit conceptual knowledge) but how you actually do it. Just telling the explicit rules hides the implicit side of science. For example, the gold standard in psychological research is inferential statistics (ANOVAs, multiple regression, etc), that basically tell you which dots on your graph you can "trust" as being truly different. This is a very explicit, analytic approach to interpreting data. However, one of us (Robert Mathews) likes to take a more holistic (and more implicit) look at the big picture of the data (e.g., look at a table of means before worrying about significant differences). Often this approach leads to a different method of analyzing the data, or a change in experimental procedures for the next experiment. This skill can be recommended to students but you cannot expect that students will immediately apply this strategy in ways that are equivalent to many years of experience. Expertise takes time; knowledge integration takes time. In addition to the suggestions we have offered above, we believe teachers can speed up the process by being willing to expose their deviations from standard (explicitly recommended) strategies. Most of us are not even aware of our hidden implicit knowledge (the art of our

science). The path to discovering this knowledge lies in being actively involved in doing science and carefully observing your own implicit reactions as you do it.

V Conclusion

Developing the ability to bring both conceptual and experiential aspects of science knowledge to bear on a problem is an important goal of science education. We have argued throughout this chapter that these types of knowledge have different characteristics and can be usefully combined. We also described the results of research that address the problem of knowledge integration, and the utility of having a powerful theoretical model (CLARION) in understanding the complex relationship between implicit and explicit processes. We believe that this model and further research on this issue are central to understanding how humans learn complex skills and will continue to provide valuable lessons for education and training. Finally, we suggest that science teachers attempting to integrate experiential and conceptual knowledge avoid ineffective strategies (such as concurrent reflection) and instead focus on strategies that allow students to systematically explore the examples of the domain under study while providing appropriate guidance and feedback.

Acknowledgements

We would like to thank Xi Zhang and Paul Slusarz for helping with simulating these tasks using CLARION. We would like to thank Paul Gade, Jonathan Kaplan, and Michael Drillings at ARI for their support (through contracts DASW01-00-K-0012 and W74V8H-04-K-0002). Thanks are also due to Roman Taraban for his invitation to contribute this chapter and his comments on the draft of this chapter.

References

- Berry, D. C., & Broadbent, D. E. (1988). Interactive tasks and implicit-explicit distinction. *British Journal of Psychology*, *79*, 251-272.
- Berry, C. C., & Dienes, Z. (1993). *Implicit learning: Theoretical and empirical issues*. East Sussex, UK: Erlbaum.
- Chi, M.T.H. (2000). Self-explaining expository texts: The dual processes of generating inferences and repairing mental models. In R. Glaser (Ed.), *Advances in Instructional Psychology*, pp. 161-237. Mahwah, NJ: Lawrence Erlbaum Associates.
- Chi, M.T.H., de Leeuw, N., Chiu, M.H., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science*, *18*, 439-477.
- Clark, A & Karmiloff-Smith, A. (1993). The cognizer's innards. *Mind and Language*, *8*, 487-519.
- Cleeremans, (1997). Principles for implicit learning. In: D. Berry (Ed.), *How implicit is implicit learning?*, 195-234. Oxford University Press, Oxford, England.
- Dienes, Z., & Berry D. (1997). Implicit learning: Below the subjective threshold. *Psychonomic Bulletin & Review*, *4*, 3-23.
- Dienes, Z., & Fahey, R. (1995). Role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 848-862.
- Domangue, T. J., Mathews, R. C., Sun, R., Roussel, L. G., and Guidry, C. E. (2004). Effects of model-based and memory-based processing on speed and accuracy of grammar string generation. *Journal of Experimental Psychology: Learning Memory, and Cognition*: *30*, 1002-1011.
- Gibson, F., Fichman, M., & Plaut, D. (1997). Learning in dynamic decision tasks. *Organizational Behavior and Human Decision Processes*, *71*, 1-35.
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, *15*, 1-38.
- Karmiloff-Smith, A. (1986). From meta-processes to conscious access: Evidence from children's metalinguistic and repair data. *Cognition*. *23*. 95-147.
- Klahr, D., & Nigam, M. (2004). The equivalence of learning paths in early science instruction: Effects of direct instruction and discovery learning. *Psychological Science*, *15*, 661-667.
- Kolers, P. A., & Roediger, H. L. (1984). Procedures of mind. *Journal of Verbal Learning and Verbal*

- Behavior*, 23, 425-449.
- Lebiere, C., Wallach, D., Taatgen, N. (1998). Implicit and explicit learning in ACT-R. *Proceedings of ECCM'98*, 183-189. Nottingham University Press.
- Lane, S. M., Mathews, R. C., Sallas, B., & Sun, R. (2005). Reflecting vs experiencing: Factors that facilitate or impair learning. Manuscript in preparation.
- Lane, S. M. & Schooler, J. S. (2004). Skimming the surface: The verbal overshadowing of analogy retrieval. *Psychological Science*, 15, 715-719.
- Lewicki, P. (1986). Processing information about covariations that cannot be articulated. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 135-146.
- Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 13, 523-530.
- Mathews, R. C., Buss, R. R., Stanley, W. B., Blanchard-Fields, F., Cho, J. R., & Druhan, B. (1989). The role of implicit and explicit processes in learning from examples: A synergistic effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 1083-1100.
- Mathews, R. C., Lane, S. M., Sallas, B., Prattini, B. & Sun, R. (2005). Hints are effective for increasing the acquisition of memory-based knowledge. In preparation.
- Meissner, C.A., & Brigham, J. C. (2001). A meta-analysis of the verbal overshadowing effect in face identification. *Applied Cognitive Psychology*, 15, 603-616,
- Melcher, J. M., & Schooler, J. W. (1996). The misremembrance of wines past: Verbal and perceptual expertise differentially mediate verbal overshadowing of taste memory. *Journal of Memory & Language*, 35, 231-245.
- McCloskey, M., & Kohl, D. (1983). Naive physics: The curvilinear impetus principle and its role in interactions with moving objects. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 9, 146-156.
- Nosofsky, R., Palmeri, T., & McKinley, S. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53-79.
- Popper K. (1972). *Conjectures and Refutations: The Growth of Scientific Knowledge*. London: Routledge.
- Prattini, R. (2006). *Knowledge representation acquired in a dynamic process control task*. A masters thesis presented to Louisiana State University
- Proctor, Robert W.; Dutta, Addie; (1995). *Skill acquisition and human performance*. Thousand Oaks, CA: Sage Publications.
- Rabinowitz, M., & Goldberg, N. (1995). Evaluating the structure-process hypothesis. In Weinert, F. & Schneider, W. (Eds.) *Memory Performance and Competencies*. Lawrence Erlbaum Associates. Hillsdale, NJ.
- Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, 118, 219-235.
- Reber, A. S. (1993). *Implicit learning and tacit knowledge: An essay on the cognitive unconscious*. New York: Oxford University Press.
- Roussel, L. G. (1999). *Facilitating knowledge Integration and Flexibility: The effects of reflection and exposure to alternative models*. Unpublished doctoral dissertation. Louisiana State University.
- Rumelhart, D., McClelland, J., & PDP research group (1986). *Parallel distributed processing*. MIT Press. Cambridge, MA.
- Sahdra B. & Thagard P. (2003). Procedural knowledge in molecular biology. *Philosophical Psychology*, 16(4), 477-407.
- Sallas, B., Mathews, R. C., Lane, S. M., & Sun, R. (2006). *Integrating memory-based and model-based knowledge of an artificial grammar using animation*. Manuscript in preparation..
- Schooler, J.W. (2002). Verbal overshadowing produces a transfer inappropriate processing shift. *Applied Cognitive Psychology*, 16, 989-997.
- Schooler, J.W., & Engstler-Schooler, T.Y. (1990). Verbal overshadowing of visual memories: Some things are better left unsaid. *Cognitive Psychology*, 17, 36-71.
- Schooler, J. W., Ohlsson, S., & Brooks, K. (1993). Thoughts beyond words: When language overshadows insight. *Journal of Experimental Psychology: General*, 122, 166-183.
- Schneider, W. and Oliver (1991), An intractable connectionist/control architecture. In: K. Van Lehn (Ed.), *Architectures for Intelligence*. Lawrence Erlbaum Associates: Hillsdale, NJ.
- Seeger, C. (1994). Implicit learning. *Psychological Bulletin*, 115, 163-196.
- Shanks, D. R. (1993). Associative versus contingency accounts of category learning. *Journal of*

- Experimental Psychology: Learning, Memory, & Cognition, Vol 19, 1411-1423.*
- Stanley, W. B., Mathews, R. C., Buss, R. R., & Kotler-Cope, S. (1989). Insight without awareness: On the interaction of verbalization, instruction, and practice in a simulated process control task. *Quarterly Journal of Experimental Psychology, 41*, 553-577.
- Sun, R. (1995). Robust reasoning: Integrating ruled based and similarity based reasoning. *Artificial Intelligence, 75*, 241-296.
- Sun, R. (1997). Learning, action, and consciousness. *Neural Networks, 10*, 1317-1331.
- Sun, R. (2002). *Duality of the mind*. Lawrence Erlbaum Associates, Mahwah, NJ.
- Sun, R. (2003). *A Tutorial on CLARION 5.0*. <http://www.cogsci.rpi.edu/~rsun/sun.tutorial.pdf>
- Sun, R. and Mathews, R. (2005). *Exploring the Interaction of Implicit and Explicit Processes to Facilitate Individual Skill Learning*. Technical Report TR-1162, Army Research Institute for the Social and Behavioral Sciences, Arlington, VA.
- Sun, R., Merrill, E., & Peterson, T. (2001). From implicit skill 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*(1), 159-192.
- Sun, R. & Zhang, X. (2004). Top-down versus bottom-up learning in cognitive skill acquisition. *Cognitive Systems Research, 5*, 63-89.
- Watkins, C. (1989). *Learning with Delayed Rewards*. PhD Thesis, Cambridge University, Cambridge, UK.
- Willingham, D., Nissen, M., & Bullemer, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 1047-1060.

Footnotes:

ⁱ Backpropagation is the term that describes a common method of “teaching” a computer network using feedback. Starting from the first trial, the output of the network is compared to the target output. The difference between the current and the target output is used to adjust the network settings to minimize this difference. Over time, this leads the network to learn how to generate the target output.

ⁱⁱ This method of training can be justified based on the fact that such associative learning can be easily performed from observing a given string and it can provide the needed implicit grammatical knowledge (as embedded in the network weights).

ⁱⁱⁱ When a completion of a partial string is found, and it is completely consistent with the rules available, the completed string is used as output. However, if a completion is impossible using given rules at the top level (due to the lack of applicable rules), the model attempts to complete as many positions as possible (it compares different partial completions and chooses the most complete one). Then, the bottom level is used. The partially completed string generated thus far by the top level is used as input to the bottom level to come up with a full string. Then, this (guessed) completion is used as output.

^{iv} If all the relevant rules are available and consistent with the candidate completion of the given partial string as generated by the bottom level, then that completion is used as output. If any of these rules are absent, an alternative rule will be used, which corrects the position that failed validation.