Undoing One's Learning

Magda Osman (m.osman@ucl.ac.uk)
Department of Psychology, University College London. Gower Street
London, WC1E 6BT England

Abstract

Buchner, Funke, & Berry, (1995) claim that in order to ensure good control performance in complex dynamic control tasks it is necessary to maintain similar learning instances during training and limit the exploration of the system. From this, they make the prediction that the diversity of learning experiences is inversely related to control performance. Contrasting this position, studies of goal specificity effect (i.e. Non-goal orientated instead goal-orientated learning leads to global processing of task information and successful transfer of knowledge) shows that focused Specific Goal (SG) learning in which the learning experiences are restricted leads to poor control performance and poor transfer of control performance compared to diverse learning experiences gained under Non Specific Goal (NSG) learning (e.g., Burns & Vollmeyer, 2002; Vollmeyer, et al., 1997). The aim of the present study was to investigate these competing claims. In order to achieve this, a novel procedure was employed in which participants solved two control tasks under NSG based learning conditions. For half, the learning phase from the first task was recorded and replayed in the second task (i.e. restricted learning experiences). For the remaining half a different learning phase from their first was presented in the second task (i.e. diverse learning experiences). The findings showed decrements in control performance under conditions in which the diversity of the learning experiences was restricted.

Complex Dynamic Control Tasks

Complex dynamic control tasks were designed to study the acquisition and transfer of skill based learning in complex interactive environments. Many (e.g., Berry, 1991; Dienes & Berry, 1997) have argued that the ability to control and maintain dynamic systems is implicit: that is, the knowledge that is acquired and utilized is beyond the individual's conscious awareness. The reason for this view is that the knowledge necessary to perform complex control task is procedural and the representations encoded are not arranged in a way that allows them to be manipulated, decomposed, and analyzed intentionally (bottom-up learning). Others claim that skilled problem solving is an example of intentional top-down processing that includes hypothesis testing which is dependent on representations being available for conscious manipulation (top-down learning). Studies investigating the influence on top-down processes in control tasks have reported an effect referred to as the goal specificity effect. When people solve a problem following a specific goal they examine information that is only relevant to the identified goal, at the cost of having a superficial understanding of the problem (Sweller, 1988). In contrast, when solving a task under a non-specific goal problem solving search is unrestricted, and so people have been shown to gain a broader and deeper understanding of the problem (Miller, Lehman, & Koedinger, 1999).

An Illustration of a Complex Dynamic Control Task

Typically complex dynamic control tasks involve several input variables which are continuous (e.g., concentration levels of salt, carbon, and lime) and that are connected via a complex causal structure or rule to several output variables that are also continuous (e.g., Chlorine concentration, Oxygenation levels, temperature) (See Figure 1). The example used here is taken from Burns and Vollmeyer's (2002) task which was originally based on a water tank purification system. In their system the starting values of the inputs were set to 0, and those of the outputs are: Oxygen = 100, Chlorine concentration = 500, Temperature = 1000, and the objective was to learn the causal structure and numerical relationship between the inputs and outputs; which is described as linear, but with constant value added to each connection.

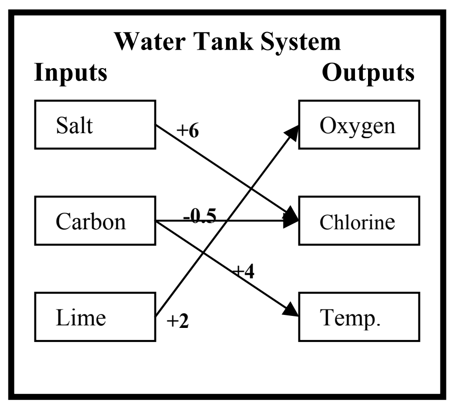

Figure 1: Water tank system

Thus, learning about the system and then attempting to control it requires that participants accurately incorporate the continuous feedback they receive on the output variables while changing the input variables. For instance, if on the first trial a participant decides to change the level of the input variable Carbon to 100 units, the value for the output Temperature will be 1104 (i.e., 1000 (which is the starting value) + 100 (value of Carbon) + 4 (the added value)). Because the input Carbon belongs to a common effect causal structure, the output Chlorine Concentration is also affected, and its value will be 599.5 (i.e., 500 (which is the starting value) + 100 (value of Carbon) + -0.5 (the added value)). If on trial 2 the value of the input Salt is changed, the output values of Temperature and Oxygenation will remain the same as the previous trial and only Chlorine Concentration will change because it is the only output connected to the input Salt.

Conflicting theoretical issues

Where both research interests (implicit learning & goal specificity) in the control task paradigm meet is in the discovery of the transfer limitations of knowledge gained during induction. Studies of control tasks devised to examine implicit learning outline the task environments that enable transfer of knowledge (i.e. perceptual similarity of transfer tasks). Studies of goal specificity detail the properties of the instruction that can lead to successful transfer (i.e., NSGs leads to global processing of task information and successful transfer of knowledge), and in both these research domains, instance based knowledge or its equivalent SG based knowledge, does not transfer beyond situations similar to those in which it was acquired.

Where both research interests (implicit learning & goal specificity) differ, is the level of self insight that they posit individuals have of the processes they use in control tasks. Berry (Berry, 1991; Dienes & Berry, 1997) claims that people only have access to the end products of the processing but not the actual process involved in generating them (e.g., Berry & Broadbent, 1984, 1986; 1988; Dienes & Fahey, 1995, 1998). The parallels between action (procedural learning) and implicit learning, and observation and explicit learning have also been explored (e.g., Berry, 1991). Consistent with this, there is evidence to suggest that there is good control performance in complex dynamic control tasks under conditions of active learning and poorer performance when learning takes place through observation (e.g. Berry, 1991; Lee, 1995; Zang et al, 2004).

This is at odds with evidence from studies of the goal specificity effect which show a correspondence between verbal self-reports of problems solving behavior and behavioral evidence of problem solving ability (e.g., Burns & Vollmeyer, 2002; Miller, Lehman, & Koedinger, 1999). In addition, performance on direct and indirect measures is associated (Burns & Vollmeyer, 2002); in particular, this has been reported in complex control systems that adopt a non-salient input-output structure with multiple causal links from inputs to outputs (e.g., Burns & Vollmeyer, 2002; Vollmeyer, et al, 1996).

Buchner et al's (1995) attempt to reconcile the dissociations between implicit and explicit learning reported in studies of implicit learning is also inconsistent with evidence from studies of the goal specificity in control tasks. Rather than evidence of dissociations between implicit and explicit learning processes, Buchner et al (1995) claim that conflicts between good performance on indirect (implicit) measures of knowledge and poor performance on direct (explicit) measures reflects differences in the range of learning experiences gained in a control task. Focused learning experiences lead to the accumulation of homogenous instance based knowledge which leads to good performance on indirect measures but poor explicit knowledge tapped by direct measures. Unrestricted learning experiences lead to good performance on direct test of knowledge because a wider range of instances of input-output states are accumulated, but at the cost of poor performance on indirect tests, because accurately controlling the system requires repeated exposure to similar learning experiences.

Buchner et al. 1995 claim that to ensure good control performance it is necessary to maintain similar learning instances during training and limit the exploration of the system and predict that the diversity of learning experiences is inversely related to control performance. The corollary to this is that the diversity of learning experiences should also be inversely related to the transfer of control performance to control tasks that are perceptually and structurally similar to those in which knowledge was acquired. This is at odds with evidence that shows focused SG learning leads to poor control performance and poor transfer of control performance compared to unrestricted NSG learning (e.g., Burns & Vollmeyer, 2002; Vollmeyer, et al., 1997).

Present Experiment

The purposes of the present study are twofold. First, the contradiction between the two approaches studying inductive processes in control tasks leaves open the question of whether unconstrained learning that generated diverse instances of a task lead to poor control performance and poor transfer compared to constricted learning. The present study addresses this issue by examining the effect on control performance in situations in which participants' problem solving is based on a NSG. The innovation of the present study is to contrast the effects of NSG learning on the transfer of control performance across perceptually dissimilar but structurally similar control tasks, and when learning experiences are repeated compared with learning experiences that are diverse. By doing this, the present study examines the prediction that the diversity of learning experiences should be inversely related to control performance and transferability of this skill to similar task domains.

The second purpose is to examine the effects on performance when learning was action-based or observation-based. To investigate this, participants either observed the learning phase in second task (Observe-self, Observe-other), or they actively interacted with the system (Act-on-self, Act-on-other). If procedural learning is necessary for good performance in control tasks, then the performance of observation-based learning conditions would be compromised compared to the performance of action-based learning conditions. It was also hypothesized that if participants in action-based learning conditions were learning implicitly, then they would be unable to accurately detect their own from another's learning compared with the observation-based learning conditions; this is because the representations acquired by participants in observation-based learning conditions would be explicit.

Method

Complex Dynamic Control Tasks Used

The two control tasks used in the present study were based on Burns and Vollmeyer's (2002) Water tank problem solving task. The water tank task version involved learning to control a linear system (Water tank system) which consisted of three inputs (substances: salt, carbon, lime) that were connected to three outputs (measures: oxygenation, chlorine concentration, temperature) (see Figure 1). Furthermore, there was a constant value added to each input-output link. The second control task was structurally identical to the first with the exception that the context and visual layout of the task was changed. Participants were told that they were newly recruited ghost hunters and that they had returned from a field experiment. It was their job to learn the relationship between three pieces of equipment that had been used in the field and the phenomena that each machine detected. The three machines (i.e. GGH meter, Anemometer, TrifieldMeter) represented the three inputs, and the three phenomena that were detected (Electro Magnetic Waves, Radio Waves, Air Pressure) represented the three outputs. The control element of the task was to modify the levels of the readouts of phenomena by manipulating the dials on each machine.

Participants Forty-eight students from University College London volunteered to take part in experiment and were paid £6 for their involvement. They were randomly allocated to one of the four conditions (Observe-self, Observe-other, Act-on-self, Act-on-other) with twelve in each. Participants were tested individually. The order of presentation of the two dynamic control tasks (i.e. the Water tank system, Ghost hunting task) was randomized for each participant.

Design The experiment was a 2x2 between subjects design with two types of exploration phase (self, other) and two types of learning condition (action, observation). All participants were required to solve two complex dynamic control tasks. The order in which participants were presented the two control tasks was randomized, (hence forth the control tasks will be referred to as Problem 1 [i.e. the first control task participants solved] and Problem 2 [i.e. the second control task participants solved].

Each problem comprised a learning phase referred to as the *Exploration phase* which consisted of 12 trials, and two control tests referred to as *Control Test 1* and *Control Test 2*; in these test participants were required to demonstrate their ability to control the system to pre-specified criteria.

The exploration phase in Problem 1 was experienced in the same way by all participants taking part in the study. The critical manipulation in the present study was in the format in which the exploration phase was presented (observational, procedural) and whether the learning experiences were identical to Problem 1 (self) or different to Problem 1 (Other).

In Problem 1 in the *Exploration phase* for each trial participants were free to change as many inputs as they liked. Once they were satisfied with their changes to the inputs, participants would click on a button labeled 'output readings' and this would reveal the values of all three outputs. This procedure would complete a trial. For action based conditions (Act-on-self, Act-on-other) the inputs changed and the values they were changed by were recorded for each trial of the exploration phase for each participant. In Problem 2, prior to starting the exploration phase, participants were presented with a record sheet with the trial history from the exploration phase of Problem 1. Participants were instructed to change inputs on each trial according to the value on the record sheet they were given.

The Act-on-self, were presented with the trial history of their own exploration phase from Problem 1, the Act-on-other condition were yoked to a participant in the Act-on-self condition, and therefore were presented with the trial history of an exploration phase different to that experienced by them in Problem 1.

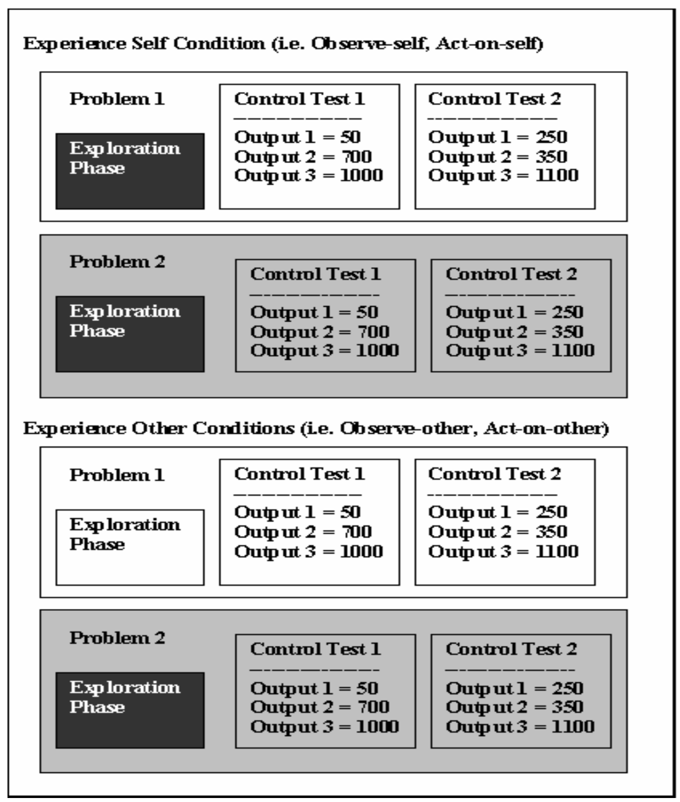

Figure 2: Experimental Design

For the observation conditions (Observe-self, Observe-other) in the exploration phase of Problem 2 participants were unable to directly intervene by selecting the inputs they wanted to change and the values they wished to change them by. Instead observation-based learning conditions tracked changes to input values and the effect on the output values for each trial. The Observe-self condition tracked the changes to the inputs and the input value changes that they had chosen themselves in the exploration phase of Problem 1. Observe-other condition was yoked to a participants from the Observe-self condition and so tracked the changes to the inputs and the input value changes that were different those they had made themselves in the exploration phase of Problem 1.

Figure 2 depicts the design of the experiment. The large white and grey boxes depict each problem that was solved. In Figure 2 the exploration phase of self experience conditions (i.e. Observe-self, Act-on-self) is shaded in the same color in both Problem 1 and Problem 2. For these conditions the problem content of Problem 1 and 2 differed, but exploration phases of each problem were matched. In Figure 2 the exploration phase of experience other conditions (i.e. Observe-other, Act-on-other) is shaded differently in both Problem 1 and Problem 2. For these conditions the problem content of Problem 1 and 2 differed as did the exploration phases of each problem. It should also be noted that the shading of the exploration phase in Problem 2 is identical to that of the exploration phases of the experience self conditions.

Thus, in the exploration phase the input values that participants had selected on each trial in the exploration phase of Problem 1 were recorded along with the corresponding output values, and for Experience Self conditions (Observe-self, Act-on-self) these were presented to them in the exploration phase of Problem 2. The experience other conditions (Observe-other, Act-on-other) were yoked to participants in the experience self conditions, that is, in the exploration phase of Problem 2 the experience other conditions were presented with the exploration phase of a participant from the experience self condition. Thus, all participants made their own decisions in the first exploration phase, but in the second exploration phase, half of the participants were presented with their own exploration phase again, while the remainder experienced a different exploration phase to their self generated one.

Procedure Participants were told that they would be taking part in a problem solving task. They were also assured that throughout the task, were they to encounter any difficulties, they should ask the experimenter for assistance. On completion of the Problem 1 participants were then told that they would be presented with another problem solving task. At no stage throughout the experiment were participants informed that the two problems were similar, or were they given advanced warning that they would be performing two problems.

Exploration phase: All participants were presented with the same computer display which was similar to Figure 1 with the exception that the connections and added contrast value of each connection were not presented.

Control test 1: In this phase all participants were required to change the input values to achieve the following output values: Oxygen = 50, Chlorine concentration = 700, Temperature = 900. Participants were allocated 6 trials in which they were to reach and then maintain the output values given.

Control test 2: As with Control Test 1, all participants were now required to change the input values to achieve given output values. However, the required output values were now: Oxygen = 250, Chlorine concentration = 350, Temperature = 1100.

Posttest question: Before being debriefed, at the end of the experiment participants were informed that there were two conditions in the experiment and that they were randomly allocated to one. It was explained that in one condition participants were presented with their exploration phase from Problem 1 again in Problem 2, or they were presented with the exploration phase of another participants in the Problem 2. Participants were asked which of the two conditions they took part in, and what they based their decision on.

Scoring

Control Test Error Scores: The procedure Burns and Vollmeyer used in scoring performance in the Control tests was adopted here. Success in achieving the goal states in each phase was computed as the sum of the absolute differences between each goal value and the value produced for each of the three outputs. All analyses of these scores are based on the mean error, for Control test 1 averaged over all six trials and averaged across all three output variables. A log transformation (base 10) was applied to the error scores of each individual participant for each round to minimize the skewedness of the distribution of scores. Control test 2 scores were calculated in same way as the Control test 1 scores. For Control test error scores in each phase a lower score indicates better performance.

Results

Figure 3 shows that overall error scores in Control Test 1 appear to be lower than error scores in Control Test 2. Figure 3 also indicates that for Control test 1 and Control test 2 the error scores of the Observe-self and Act-on-self appear to have decreased in Problem 2 compared to Problem 1. In contrast, the error scores of the Observe-other and Act-on-other conditions appear to be stable across Problem 1 and 2. To examine the possible interaction between the diversity of learning experiences and control performance across Problem 1 and 2 the following analyses were conducted.

Control Test Error Scores: The following analyses were conducted on the mean error scores calculated for each participant. A 2x2x2x2 ANOVA was carried out with Control Test (Control Test 1, Control Test 2) and Problem (Problem 1, Problem 2) as the within subject factors, and exploration phase (Self, Other) and learning format (action, observation) as between subject factors.

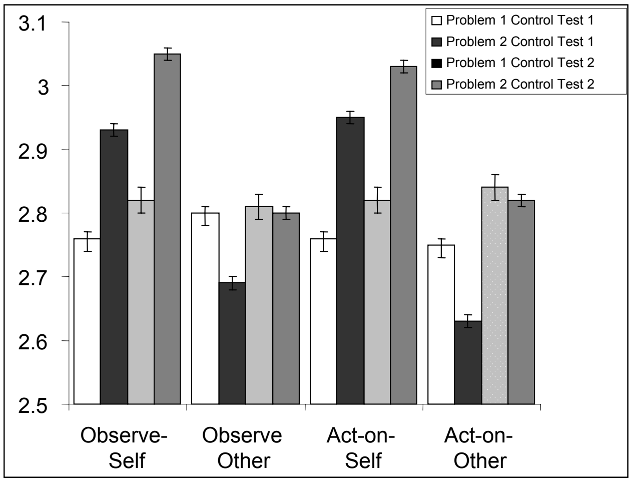

Figure 3: Mean Error scores for Control Test 1 and Control Test 2 for Problem 1 and 2 by Condition

The analysis revealed a significant main effect of Control Test, $F(1, 44) = 11.561$, MSE = 0.39, $p < 0.002$. There was no significant main effect of Problem $F(1, 44) = 1.88$, MSE = 0.1, and no main effect of learning format $F(1, 44) = 0.04$, MSE = 0.004. There was a significant main effect of exploration phase on error scores $F(1, 44) = 7.06$, MSE = 0.72, $p < 0.05$, and a significant Exploration phase x Problem interaction $F(1, 44) = 21.08$, MSE = 1.14, $p < 0.0005$. Thus the results confirmed the pattern of findings indicated in Figure 3 that overall, there was a difference in performance between Control test 1 and 2. There was no overall difference between error scores in Problem 1 and 2, and no evidence that the learning format influenced the pattern of error scores. To locate the source of the interaction, tests of simple main effects were carried out. Because learning format was not found to have a significant effect on error scores, error scores were collapsed across Observe-self and Act-on-self conditions, and across Observe-other and Act-on-other conditions. The significant decrease in performance across problems for the experience self conditions was confirmed by planned comparisons of error scores in Control Test 1 $t(23) = -23.23$, $p< 0.0005$, and Control Test 2 $t(23) = -17.77$, $p< 0.0005$. The increase in performance across problems for the experience other conditions was not statistically confirmed by planned comparisons of error scores in Control Test 1 $t(23) =1.91$, $p= 0.067$, and Control Test 2 $t(23) =1.98$, $p< 0.061$, although both tests approached significance.

Table 1: Proportion of accurate responses to posttest recognition question

Exploration phase	Observe	Act
Self	91.7	66.7
Other	75	83.3

Table 1 presents the proportion of correct responses to the posttest question. Pearson's chi-squared analysis revealed that there was no significant difference in the accuracy of participants response to the recognition question, $\chi^2(1) = 3.43$, $p> 0.05$. Responses in each group indicated that participants were able to accurately identify which experimental condition they took part in.

Discussion

The evidence from this study did not confirm Buchner et al. (1995) prediction that the diversity of the learning experiences is inversely related to control performance in dynamic control tasks. Instead the findings from the present study revealed that by limiting the learning experiences to those previously experienced control performance is adversely affected to the extent that it produces a negative learning effect. However, this pattern of findings is also not predicted by the position that restricted goal learning reduces control performance relative to unrestricted goal learning which produces a varied set of learning experiences. However, consistent with this position, there was evidence to suggest that increasing the range of learning experiences does lead to improved control performance.

There was no evidence to suggest that control performance was adversely affected by the observation–based learning format of the exploration phase compared to the procedural based learning version. In addition, contrary to previous studies using control tasks to examine implicit learning processes, participants demonstrated accurate self insight and were able to respond accurately in the posttest awareness test.

The findings raise three questions: Why was there no difference between action-based and observation-based learning conditions? What caused a negative learning effect in the Observe-self and Act-on-self conditions? Why were problem solvers able to accurately identify their own learning experience in a task that is commonly thought to elicit implicit processes?

Why was there no difference in performance between action-based and observation-based learning conditions? Consistent with previous evidence (Osman & Heyes, 2005), the findings from the present study suggested that was an advantage of action-based learning over observation-based learning One reason that observers and active learners performed similarly is that in previous studies using complex dynamic control task paradigms (e.g., Berry, 1991; Less, 1995) observation based learning conditions are passive, and there is a lack of engagement with the task which may in turn lead to a lack of motivation for observers compared to active learners. In contrast to this, participants in the observation conditions in the present study were motivated to examine the information on screen, particularly because during the first phase of the experiment changes to input values on a given trial were not accompanied by the subsequent output values until participants pressed a button to reveal them. By separating the stages in which people observed input values from the stage when the effects on the output values was revealed allowed participants to actively process input information and predict the effects that occurred. Thus, for each trial in the exploration phase participants would have to mentally calculate the relationship between the input and output values in order to track the differences between them. Thus, the high attentional and motivational demands of the observation-based learning produced learning that was equivalent to procedural-based learning.

Why was performance of the observe-self and act-self groups adversely effected? One reason that there was a decline in performance after re-exposure to their own prior learning was because the range of learning experiences were limited to those encountered earlier. If participants had difficulty determining what the underlying causal structure of the control task was after solving the task for the first time, then being presented the same values selected in an early learning phase again in a different task environment (even if the underlying causal structure is the same) may have interfered with participants' ability to examine the learning environment objectively.

Why were problem solvers able to accurately identify their own learning experience in a task that is commonly thought to elicit implicit processes that are inaccessible to conscious awareness? The accuracy with which participants responded to the posttest recognition test is inconsistent with the commonly held claim that control tasks elicit implicit learning processes. Moreover, previous evidence (e.g., Berry & Broadbent, 1984) suggests that participants are unable to verbalize their knowledge despite demonstrating the ability to control complex dynamic problem solving tasks. One reason for the discrepancy between previously reported findings and those of the present study is simply that under forced choice conditions participants are able to accurately recall information that is relevant to make a discrimination as to whether the source of the learning experience is their own or that of another's. The accessibility of knowledge, and whether a task is considered to elicit implicit processes, is based on the particular test used to examine it. The results in the present study suggest that with sufficient scrutiny, participants are self aware under conditions that are commonly believed to promote implicit learning. Moreover, the present study found that there was no difference in accuracy between action and observation based learning conditions. This suggests that even in action-based learning, which is regarded as procedural and implicit, participants demonstrate awareness of their learning behavior.

Conclusions

The finding from this study demonstrated that participants were able to accurately identify their own learning experiences in a complex dynamic control task, suggesting that contrary to numerous studies, people are aware of the learning process they are engaging whilst problem solving. In addition, the findings also revealed that there was no difference in control performance when learning to solve a control task under an observation-based condition compared to an equivalent action-based version.

Acknowledgments

Preparation for this article was supported by Economic and Social Research Council ESRC grant RES-000-27-0119. The support of the Economic and Social Research Council (ESRC) is gratefully acknowledged. The work was also part of the programme of the ESRC Research Centre For Economic Learning and Human Evolution.

References

Berry, D. (1991). The role of action in implicit learning. *Quarterly Journal of Experimental Psychology, 43*, 881-906.

Berry, D., & Broadbent, D. E. (1984). On the relationship between task performance and Associated verbalizable knowledge. *Quarterly Journal of Experimental Psychology, 36*, 209-231.

Burns, B. D., & Vollmeyer, R. (2002). Goal specificity effects on hypothesis testing in problem solving. *Quarterly Journal of Experimental Psychology, 55*, 241-261.

Dienes, Z., & Berry, D. (1997). Implicit learning: Below the subjective threshold. *Psychonomic Bulletin & Review*, *4*, 3-23.

Lee, Y. (1995). Effects of learning contexts on implicit and explicit learning. *Memory and Cognition*, *23*, 723-734.

Miller, C.S., Lehman, J. F., & Koedinger, K. R. (1999). Goals and learning in microworlds. *Cognitive Science, 23*, 305-336.

Newell, A., & Simon, H. A. (1972). *Human problem solving.* Englewood Cliffs, NJ: Prentice Hall.

Osman, M., & Heyes, C. (2005). Practice doesn't always make perfect: Goal induced decrements in the accuracy of action-and –observation based problem solving. *In Proceedings of the 27th Annual Meeting of the Cognitive Science Society, Stresa, Italy, August.*

Sweller, J. (1988). Cognitive load during problem solving: Effects of learning. *Cognitive Science, 12*, 257-285.

Vollmeyer, R., Burns, B. D., & Holyoak, K. J. (1996). The impact of goal specificity and systematicity of strategies on the acquisition of problem structure. *Cognitive Science, 20*, 75-100.

Zang, A., Richard, J., & Tijus, C. (2004). Implicit learning in rule induction and problem solving. *Thinking and Reasoning, 10*, 55-83.