The Pennsylvania State University

The Graduate School

College of the Liberal Arts

MAINTENANCE OF ACTIVATION

WITHIN GOAL HIERARCHIES

A Thesis in

Psychology

by

Christopher A. Stevens

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science

August 2011

The thesis of Christopher A. Stevens was reviewed and approved* by the following:

Richard A. Carlson
Professor of Psychology
Associate Head of the Department of Psychology
Thesis Adviser

Nancy A. Dennis
Assistant Professor of Psychology

Cynthia Huang-Pollock
Assistant Professor of Psychology

Melvin M. Mark
Professor of Psychology
Head of the Department of Psychology

*Signatures are on file in the Graduate School.

Abstract

Maintenance of goal activation is important to the successful completion of cognitive tasks. However, it is still unclear how this activation is maintained. The hypothesis that superordinate goals provide a source of activation for their subgoals was tested in two experiments. Across experiments, the presence of task cues was manipulated as a way of measuring activation loss in subgoals. A benefit of task cues would suggest activation loss. In Experiment 1, participants completed a task with a superordinate goal - multi-step arithmetic problems. Task cues did not impact performance, suggesting that superordinate goals may have been restoring subgoal activation. In Experiment 2, the presence of a superordinate goal was manipulated. Task cues improved accuracy regardless of goal condition. However, task cues did not affect the rate of slowing within problems. Therefore, the evidence for activation loss is mixed. Overall, the results fail to provide clear evidence for or against the hypothesis.

Table of Contents

List of Figures……v

INTRODUCTION…… 1
The Interactive Activation Model…… 1
The Memory for Goals Model…… 3
Superordinate Goals and Goal Maintenance……4
Experiment Overview and Hypotheses……5

EXPERIMENT 1…… 8
Method…… 8
Results……10
Discussion……11

EXPERIMENT 2…… 13
Method…… 15
Results……18
Discussion……23

GENERAL DISCUSSION……25
Future Directions……29
Conclusion…… 31

References……33

Appendix: ANOVA Summary Tables……36

List of Figures

Figure 1. Paradigm for Experiment 1……………………………………………………………10
Figure 2. Response times by problem step and task cue in Experiment 1…………………… 11
Figure 3. Paradigm for superordinate goal condition of Experiment 2………………………16
Figure 4. Paradigm for no superordinate goal condition of Experiment 2.…………………...17
Figure 5. Accuracy data for the no superordinate goal condition in Experiment 2…………...20
Figure 6. Response time by problem type and superordinate goal in Experiment 2................21
Figure 7. Response time by problem step, superordinate goal, and task cue
in Experiment 2…………………………………………………………………22

Introduction

Recent theories of cognitive control postulate that goals – representations of actions or action sets - have associated activation levels. Activation refers to the strength of a given representation in memory and thus the likelihood that it will be retrieved during a search of memory. The higher the activation level of a representation, the faster it will be retrieved (Anderson & Lebiere, 1998). Activation also determines the ability of a goal to control behavior (Altmann & Trafton, 2002; Byrne & Bovair, 1997; Norman, 1981). However, the way in which this activation is maintained over time is still very much an open question.

Goals are often assumed to exist as part of a hierarchy, so a candidate mechanism for maintaining goal activation is the goal hierarchy itself. Control flows from abstract, higher-level goals to basic actions at the bottom level (e.g. Norman & Shallice, 1986). Because of this flow of control, it has been postulated that higher-level goals maintain activation in lower goals in the hierarchy. Current theories differ regarding how higher-level goals affect activation levels in lower-level goals. Some theories suggest that upper-level goals continuously strengthen lower-level goals, shielding them from interference from other goals (Byrne & Bovair, 1997; Cooper & Shallice, 2000). Others suggest that upper level goals merely initialize lower level goals, providing an activation boost during planning only (Altmann & Trafton, 2002; Altmann & Gray, 2008). The aim of this thesis is to elucidate this question and explore its implications for performance.

Superordinate Goals and Subgoals

According to several current theories, goals are organized hierarchically into layers of superordinate goals and subgoals (Altmann & Trafton, 2002; Byrne & Bovair, 1997; Norman, 1981). A superordinate goal represents a complex action or action sequence. Each superordinate goal has one or more subgoals that must be completed before it can be achieved. For example, making coffee is a superordinate goal. Its subgoals include adding coffee to the coffeemaker, adding water, and turning on the coffeemaker. If these subgoals are not completed, the superordinate goal, making coffee, will not be achieved.

The Interactive Activation Model

Norman (1981) proposed the Activation-Trigger-Schema (ATS) theory. According to this theory, goals are represented as a hierarchy of schemas. Each goal in the hierarchy has an associated activation level. This quantity determines whether the goal can be accessed and executed by the system. If a goal is not active enough, the system will fail to execute it. This activation has two main sources: the environment and other goals. All superordinate goals (parent schemas) have associative links to their respective subgoals (child schemas). When selected, superordinate goals spread activation to their subgoals through these links (Norman, 1981).

The ATS framework was later implemented in a computational model known as the Interactive Activation (IAN) model (Cooper & Shallice, 2000; Cooper & Shallice, 2006). According to this model, when a superordinate goal reaches an activation threshold, it is selected by the system. As long as the superordinate goal remains selected, it continuously spreads activation to its subgoals.

A similar theory that also assumes continuous maintenance of subgoals is Byrne and Bovair's (1997) computational model of post-completion errors. A post-completion error is the failure to complete a procedure after the superordinate goal of that procedure has been satisfied. An example of this type of error is forgetting one's keys in a door lock. Here the superordinate goal of the procedure, opening the door, is complete. Removing the keys is merely a "clean up" step. Like the IAN model, the Byrne and Bovair model assumes that, as long as they are still active, superordinate goals provide continuing activation for their subgoals. However, once the superordinate goal is satisfied, it ceases to sustain the subgoals. These subgoals then rapidly decline in activation. As a result, there is an increased chance that the remaining subgoals will fall below threshold before they can be executed.

The Memory for Goals Model

A somewhat different approach to goal memory was taken by the ACT-R family of theories (Anderson & Lebiere, 1998). Goals were assumed to be arranged in a stack. In this framework, the system begins with one or more goals placed in the stack. The system always executes the top goal on the stack. Each new goal is placed on the top of the stack. Once the top goal is completed, it is removed from the stack and the next goal is executed. Thus goals are executed in a first-in last-out manner. One weakness of this approach is that goal memory is assumed to be perfect. Once a goal is placed on the stack, it remains there until it is executed and removed. Although it eliminates the need to postulate a mechanism of goal maintenance, this assumption leads to difficulty in predicting certain serial order patterns in behavior, particularly sequence errors (Altmann & Trafton, 2002; Byrne & Bovair, 1997). Anderson and Douglass (2001) later rejected the perfect goal memory assumption, demonstrating that retrieving suspended goals incurs a time cost. They concluded that goals have associated activation levels that impact the likelihood of goal retrieval and the time cost associated with doing so.

Altmann and Trafton (2002) expanded on the findings of Anderson and Douglass (2001) by developing a theory of goal memory called the Memory for Goals (MFG) theory. The theory is based primarily on two counterintuitive results from the task switching literature: within-run slowing and within-run error increase (Altmann, 2002; Altmann & Gray, 2008). When participants complete the same task repeatedly (i.e. a run of trials) their reaction times and error rates increase linearly with each repetition. However when a task cue is presented, performance is improved. The MFG theory argues that these results are caused by a loss of goal activation. The theory assumes that goals are representations in episodic memory. During every processing cycle, the system selects and executes the most active goal in memory. Over time, the activation level of each goal decays. This makes goal retrieval more time-consuming and error-prone due to increased interference from old goals. Information from the environment (e.g. a task instruction) raises the activation level of associated goals.

Superordinate goals are also stored in episodic memory. Unlike in the IAN model, superordinate goals raise activation levels of subgoals only during planning rather than continuously. Here planning refers to a state in which the system is selecting and preparing to execute a superordinate goal. Once a superordinate goal is selected, the system focuses attention on this goal to raise the activation level of its subgoals. Once this process is complete, the focus of attention is moved away from the superordinate goal and the system relies on retrieval of the subgoals to direct behavior.

The MFG theory has been successful in predicting performance in several domains. Models based on the MFG theory have successfully simulated performance in the Tower of Hanoi task (Altmann & Trafton, 2002), sequence errors (Trafton, Altmann, & Ratwani, 2011), and recovery from interruption (Altmann & Trafton, 2007; Hodgetts & Jones, 2006). Though these tasks all have hierarchical goal structures, they do not lend themselves to studying the temporal dynamics of individual subgoals because there are often many distinct subgoals for each superordinate goal and switching between these subgoals is frequent. In the two-choice task switching paradigms studied by Altmann (2002) and Altmann & Gray (2008), within-run slowing and within-run error increase can be used as indexes of the activation loss of individual subgoals, but these tasks lack a hierarchical structure. The task reported in this thesis was developed to possess both a hierarchical structure while still allowing for measurements of activation loss.

Superordinate Goals and Goal Maintenance

An important theoretical distinction between the MFG model and the IAN model is thus the role of superordinate goals in memory for subgoals. In the MFG model, the system can focus attention on superordinate goals before task execution begins. Focusing on a superordinate goal causes it to spread activation to its respective subgoals so that the subgoals can be retrieved at the correct time. Once the person begins executing the task, behavior is guided solely by the subgoals. Superordinate goals are then used only when the system fails to retrieve an appropriate subgoal and must infer the next step. Conversely, the IAN model suggests that superordinate goals remain in working memory and keep their respective subgoals active. Thus the two theories make diverging predictions regarding the effect of task structure on performance.

The IAN model predicts that superordinate goals should provide protection from activation loss in subgoals by supplying a continuous source of activation. Although the model has not been applied to speed of performance, it can be extended to do so by assuming higher activation leads to faster execution of a goal. Therefore, the model predicts that the presence of a superordinate goal should offset the effects of activation loss observed by Altmann (2002) and Altmann and Gray (2008).

By contrast, the MFG model treats superordinate goals similarly to external task cues. That is, when attention is focused on a superordinate goal, it spreads activation to its subgoals.

However, as soon as attention is moved away from the superordinate goal, the subgoals begin to lose activation (Altmann & Trafton, 2002). Consequently, according to the MFG model, activation loss should occur regardless of the presence of a superordinate goal.

Experiment Overview and Hypotheses

The purpose of the following experiments was to clarify the way in which goal hierarchies direct the moment-to-moment control of cognition. Specifically, do they serve as an ongoing source of activation (as in the IAN model) or do they simply activate subgoals during planning (as in the MFG model)? To answer this question, the presence of a superordinate goal was manipulated in the context of a multiple-step arithmetic task. The IAN model predicts that the superordinate goal should offset the subgoal's activation decay, and therefore slowing should be reduced. Whereas the MFG model predicts that goal-activation decay should be evident to the same extent regardless of the presence of a superordinate goal.

Previous studies in task-switching indicate that task cues boost performance even if the participant already knows the goal (Altmann, 2002; Koch, 2003). The putative reason for this is that the goal is losing activation, but remains above threshold. If the subjects' goals are losing activation, then task cues should restore that activation and reduce slowing. If the goals are already fully active, then task cues should have no effect. Therefore, I manipulated the presence of these cues. In each experiment, half of the participants saw task cues at every step and half did not. Thus the IAN model would predict that task cues should reduce slowing only when there is no superordinate goal present. Conversely, the MFG model predicts that task cues should reduce slowing in both conditions.

Testing the predictions of the two models requires a manipulation of the goal structure of the experimental task. A task is required in which a superordinate goal can be added without changing the essential features of the task. I chose a serial arithmetic task because of its flexible goal structure. Individual operations can be done to solve a larger problem (i.e. to complete a superordinate goal) or they can stand alone (Carlson & Lundy, 1992; Lundy, Wenger, Schmidt, & Carlson, 1994).

In this task, each trial consisted of combining two numbers through addition or subtraction to compute a result. In the no superordinate goal condition, these trials were all independent. Solving them did not contribute to the completion of a superordinate goal. In the superordinate goal condition, the trials were combined into larger problems, so that solving each trial was necessary to solve the larger problem. In this case, the subgoals were to solve the trials (e.g. compute-sum-trial1, compute-sum-trial2). The superordinate goal was to compute the answer for the larger problem.

In Experiment 1, I tested the hypothesis that goals lose activation over time in the presence of superordinate goals. The presence of this decay would suggest that superordinate goals do not sustain activation in their subgoals, consistent with the MFG model. If this hypothesis is correct, then task cues should reduce the amount of slowing of performance. Alternatively, a lack of this effect would suggest that goals are not losing activation, consistent with the predictions of the IAN model. Then, in Experiment 2, I performed a stronger test of the superordinate goal hypothesis by directly comparing tasks with a superordinate goal to those without one. Further, I explored the possibility that the impact of goal structure may be affected by an individual's working memory capacity.

Experiment 1

The aim of Experiment 1 was to determine whether goal activation decay occurs in a task with a superordinate goal. In the task, participants solved a series of multi-step arithmetic problems. Here, the superordinate goal was to provide the final answer to each problem. Completing an individual step constituted a subgoal. Half of the participants saw task cues (+, -) and half did not. Based on the MFG model, I expected that, without task cues, goal activation would decay and performance would slow down in later steps of each problem. When task cues were present to boost goal activation, I expected the slowing to be reduced or even eliminated. Therefore I hypothesized an interaction effect between step and task cue such that response times should increase with step and that this increase should be sharper in the no task cue condition.

Method

Design. The experiment was a 2 x 5 mixed factorial with task cue (present, not present) as a between-subjects variable and step (2 – 6) as a within-subjects variable. The dependent measures were proportion of problems correct and median response time for each step.

Participants. Seventy Penn State University undergraduate students took part in the experiment in return for partial course credit. All participants gave informed consent in accordance with Penn State procedures

Procedure. Participants were asked to solve a series of multiple-step arithmetic problems. Each problem was either an addition problem or a subtraction problem. After completing 2 blocks of practice problems, participants completed 9 blocks of 6 problems. In each block, three problems were addition and three were subtraction. Order of problem type was randomized within-blocks.

Task. Each problem consisted of 6-8 steps. The purpose of the variable problem length was to reduce the participants' ability to anticipate the end of the problem as this might cause them to slow down in anticipation of a possible task switch. The addends or subtrahends were always 2, 3, or 4, selected randomly with the following constraints: an addend or subtrahend could not be selected twice in a row, and the result of the step could never end in zero. Participants were instructed that the tens digit of the number was not important and were required to enter only the ones digit of the answer. Participants were also instructed that the answers would always be positive.

Figure 1 illustrates the task. All stimuli were displayed on a computer screen in white font inside of a small black box in the center of an off-white screen. In each problem, participants saw a two-digit starting number and were instructed to add or subtract a series of one-digit numbers starting with that number to compute a final result. For half of the participants, task cues indicating the current problem type ('+' for add and '-' for subtract) were presented between steps. After viewing the starting number for as long as they wished, participants pressed the spacebar to see the first addend. Participants then saw either a blank screen or a task cue for 200 milliseconds, depending on their experimental condition. This display was followed by a presentation of the addend or subtrahend (also 200 milliseconds). After the addend or subtrahend disappeared, the cue or blank screen returned and remained until the participant pressed the spacebar. Once the spacebar was pressed the next addend was displayed for 200 milliseconds and then replaced once again with the cue or blank screen. This cycle repeated until the end of the problem, at which time the participant was asked to enter the total.

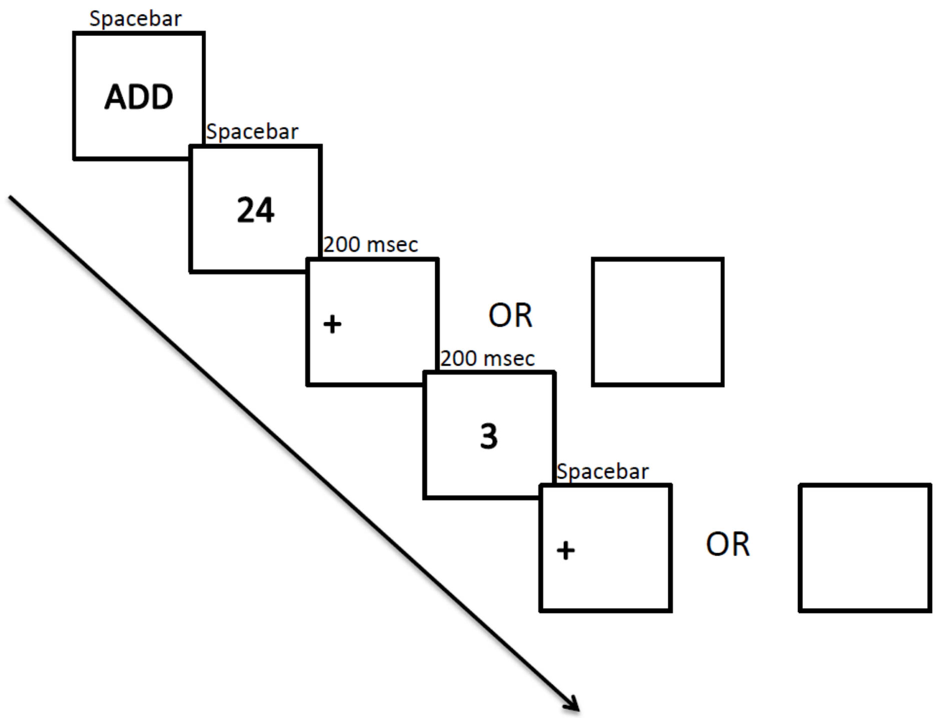

Figure 1. Paradigm for Experiment 1.

Results

Accuracy. The proportion correct scores (M = .893, SD = .007) were submitted to a repeated measures ANOVA with task cue (present, absent) as a between-subjects variable and problem type (addition, subtraction) as a within subjects variable. This analysis yielded no significant effects (all ps > .1).

Response Time. Response times for each step were measured from the onset of the task cue to the spacebar press. One subject was excluded from the response time analysis for failing to maintain an accuracy of at least 75 percent. Once again, task cue and problem type were entered as variables in the ANOVA. Step (2-6) was included as an additional within subjects variable. Only steps 2-6 were included because no computations occurred on step 1 and steps 7 and 8 were not present in every trial. This analysis revealed a significant main effect of step ($F(4, 268) = 6.2$, $p < .0005$, $\dot{\eta}^2 = .475$) and the linear contrast of step was significant ($F(1, 67) = 22.8$, $p < .0005$, $\dot{\eta}^2 = .254$). Response times were significantly longer on subtraction problems

(M = 1403 ms, SE = 35.7) than on addition problems (M = 1243 ms, SE = 33) ($F(1, 67) = 60.5$, $p < .0005$, = .085). This effect did not interact with any other variables (ps > .1). The task cue by step interaction predicted by the MFG model was not significant ($F(4, 268) = .324$, $p = .862$, $\dot{\eta}^2 = .005$). Figure 2 illustrates this interaction.

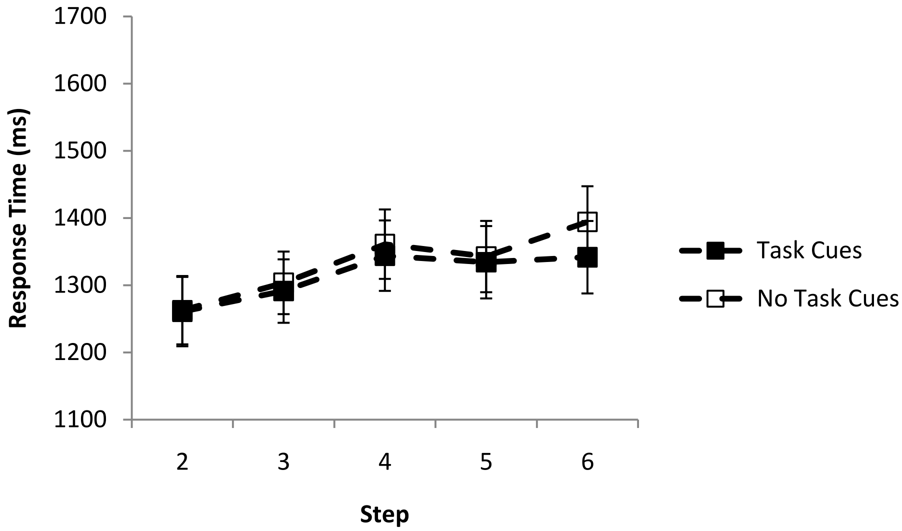

Figure 2. Response times by problem step and task cue in Experiment 1.

Discussion

The MFG model's hypothesis of a steeper slowing curve in the task cue absent condition was not supported. Although I did replicate the linear slowing effect observed previously in task-switching paradigms (Altmann & Gray, 2008), the lack of an effect of task cue suggests that this slowing was not due to decay of goal activation. Consistent with the IAN model, these results suggest that the superordinate goals were restoring the activation lost by the subgoals due to decay.

The main effect of step may be due to proactive interference of subtotals in working memory. For example, research by Lustig, May, and Hasher (2001) indicates that old memories remain active in working memory, making the retrieval process slower and more error-prone.

According to the IAN model, schemas remain active after they are executed. These lingering schemas can decrease the activation level of the current schemas by lateral inhibition (Cooper & Shallice, 2000). Perhaps in this case older subtotals inhibited newer subtotals. As a result, the new subtotals may have required more time to reach a selection threshold, causing slower performance.

Additionally, the finding that participants were slower on subtraction problems seems likely to be due to lower resting activation levels of subtraction facts due to lower frequency of use. A similar explanation has been offered for the problem size effect – that solving problems with larger numbers is slower than solving problems with smaller numbers (Ashcraft, 1992).

Experiment 2

Consistent with the IAN model, Experiment 1 showed no evidence that task cues affected the rate of slowing, suggesting that a superordinate goal may have been restoring the activation lost by decay. Experiment 2 provides a stronger test of the model's predictions by including a comparison with an arithmetic task with no superordinate goal. If the predictions based on the IAN model are correct, there should be evidence of goal activation loss when superordinate goals are absent, but not when they are present.

In addition, Experiment 2 addresses several alternative explanations of the results of Experiment 1. First, it could be that the task cues provided too little activation to the goals to affect performance. Miyake, Emerson, Padilla, and Ahn (2004) found that cue transparency is important to the effectiveness of a task cue. Transparency refers to the strength of the relationship between the cue and the goal. Because word reading is relatively automatic, word cues tend to be more transparent than symbolic cues (Miyake et al., 2004). Therefore, word cues may provide sufficient activation to overcome decay in this paradigm.

It is also possible that the way in which the task cues were displayed in Experiment 1 was the reason for the lack of an effect. The task cue and digit were never on the screen at the same time. Rather the display alternated between the two, which may have discouraged the participants from associating them together. This alternating display may have also distracted participants. In Experiment 2, these issues are addressed by displaying the task cue first, then adding the digit to the display so that both are present simultaneously.

In order to make the display of the task cue condition more similar to the no task cue condition, I replaced the blank screen condition with an irrelevant cue condition. The irrelevant cue - a set of pound signs - occurred at the same place and time on the display as the task cues in the relevant cue condition. This made it less likely that any observed effect of task cue could be due to the task cue functioning as a signal to orient attention to the digit.

Another possible explanation of the lack of an effect of the task cues in Experiment 1 is that the effect was obscured by individual differences in working memory. A variety of work indicates that goal maintenance is a crucial component of working memory capacity (Duncan, Emslie, Williams, Johnson, & Freer, 1996; Duncan et al., 2008; Kane & Engle, 2003; McVay & Kane, 2009). This research suggests that individuals with low working memory capacity (i.e. low spans) should have a more difficult time keeping goals active across trials. Further, these individuals should benefit more from the presence of task cues than those with high capacity (i.e. high spans). To address this possibility, I included an operation span task (Turner & Engle, 1989) as a measure of working memory capacity in Experiment 2.

Previous research suggests that high spans are better able to manage goal sets than low spans (Duncan et al., 1996). Thus it is possible that high spans are better able to utilize internal goal hierarchies to maintain subgoal activation than low spans. If this is true, then working memory span scores should positively correlate with the benefit of superordinate goals on slowing.

The IAN and MFG models make different predictions regarding the results of this experiment. In terms of accuracy, the IAN model would predict an interaction between task cue and superordinate goal such that the absence of task cues results in lower accuracy when no superordinate goal is present. By contrast, the MFG model predicts a benefit of task cues regardless of the presence of a superordinate goal.

As for response time, both models predict the same results for the no superordinate goal condition. In this case, task cues should speed up response time and this benefit should increase with step. In the superordinate goal condition, the IAN model predicts no benefit of task cues. However, the MFG model predicts the same effect of task cues as in the no superordinate goal condition. Overall, the IAN model predicts a three-way interaction between superordinate goal, task cue, and step where the benefit of task cue on slowing is present in the no superordinate goal condition but not in the superordinate goal condition. By contrast, the MFG model predicts only a two-way interaction between task cue and step.

Method

Design. The design was a 2 x 2 x 5 mixed factorial with superordinate goal (present, absent) and task cues (present, absent) as between-subjects variables and step (2-6) as a within-subjects variable. Originally, Experiment 2 was planned as two separate experiments so the superordinate goal condition and the no superordinate goal condition were run sequentially rather than simultaneously. The dependent measures were proportion of problems correct and median response time for each step.

Participants. One hundred sixty-seven undergraduate students from Penn State University participated in the experiment in exchange for partial course credit. All participants gave informed consent in accordance with Penn State procedures.

Procedure. After giving informed consent, participants completed the experiment on a desktop PC running E-Prime software. Participants first completed the serial arithmetic task and then the operation span task. After completing both tasks, the participants were debriefed and dismissed. The maximum length of the experimental session was 60 minutes.

Serial Arithmetic Task. The superordinate goal version of the task was very similar to that used in Experiment 1. However, in this task, the symbolic cues (+, -) were replaced with verbal cues (plus, minus). Moreover, in the cues-absent condition, pound signs (#####) were displayed instead of a blank screen. Finally, the cue/blank screen was not continuously present between steps. Rather, it was displayed briefly (200 ms) before the digit for each step. Once again, the superordinate goal was to provide the answer to the multiple step problem and the subgoals were to solve the steps of that problem. Figures 3 and 4 illustrate the task.

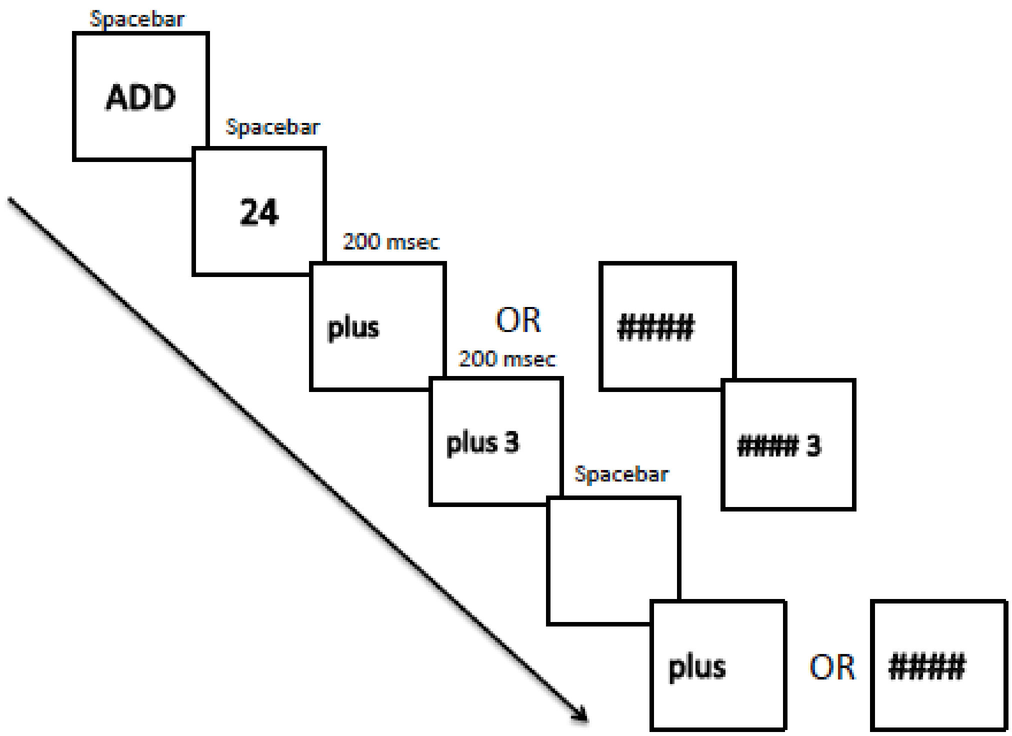

Figure 3. Paradigm for superordinate goal condition of Experiment 2.

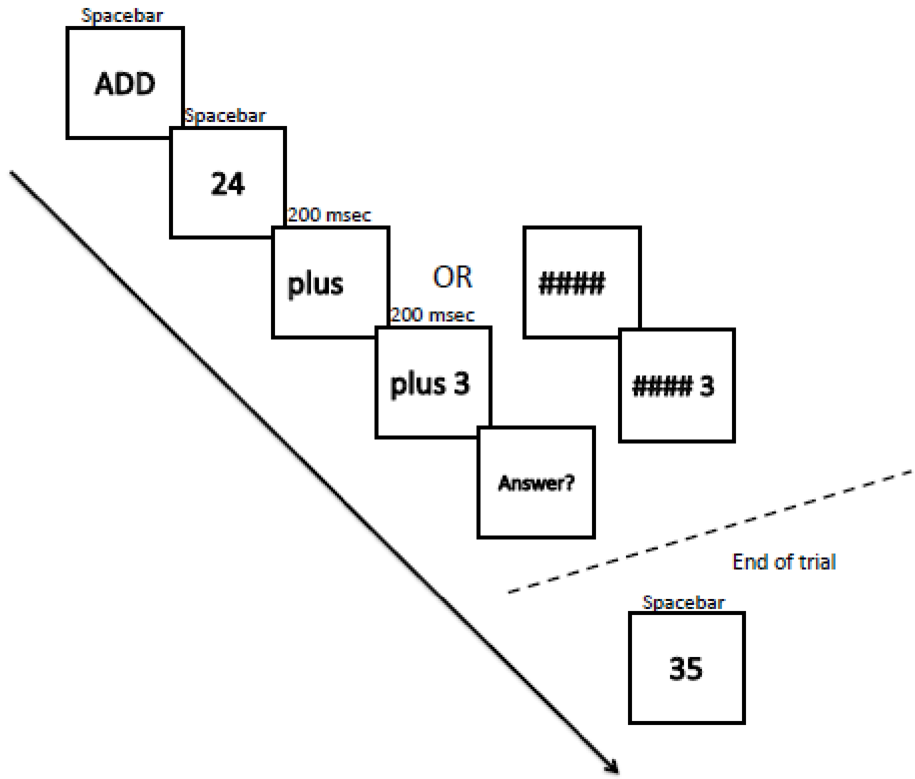

Figure 4. Paradigm for no superordinate goal condition of Experiment 2.

In the no superordinate goal condition, participants completed sets of 6-8 addition or subtraction problems instead of one addition or subtraction problem with 6-8 steps. Thus the superordinate goal of completing the multi-step addition or subtraction problem was removed. Each set of problems began with an instruction indicating the type of problem (ADD, SUB). After the participant pressed the space bar, the program displayed a two-digit starting number for the first problem. The participant then pressed the space bar to show the cue or blank screen followed by the digit screen. Following the offset of the digit, the program prompted the participant for the ones-digit of the answer. After entering the answer, the participant could change the answer if desired by pressing the backspace key. To confirm the answer, the participant pressed the spacebar. Doing so displayed the starting number for the next problem. The process repeated until all of the problems in the set were answered. A feedback screen indicating proportion correct for the set and for the experiment followed every set of problems.

Operation Span Task. To measure working memory capacity, I employed a computerized version of the operation span task (Turner & R. W. Engle, 1989). In this task, participants must verify whether given mathematical equations are correct (e.g. $4/2 - 1 = 1$). After verifying each equation participants are given a word to remember until the end of the block. Participants reported words by typing them on the keyboard after being prompted by the program. In order to offset recency effects, the program instructed the participants not to enter the final word first. Blocks consist of 2 to 6 trials and grow increasingly longer until the end of the procedure. Each participant's working memory score is defined as the total number of words recalled correctly throughout the experiment.

Results

One subject was excluded from the analysis because he had participated in a pilot version of the experiment. Subjects who failed to maintain an accuracy level of at least 75 percent, or who had average response times at least 3 standard deviations from the mean were excluded from all analyses. This resulted in the exclusion of 6 participants.

Accuracy. Accuracy data (M = .936, SE = .004) were submitted to a 2 x 2 x 2 repeated measures ANOVA with problem type (addition, subtraction) as a within-subjects variable and superordinate goal (present, absent) and task cues (present, absent) as between-subjects factors. The analysis indicated that participants in the no superordinate goal condition (M = .964, SE = .005) were more accurate than those in the superordinate goal condition (M = .907, SE = .006) ($F(1,156) = 53.5$, $p < .0005$ $\dot{\eta}^2 = .255$). It should be noted, however, that the no superordinate goal condition was somewhat easier in terms of accuracy than the superordinate goal condition. In the superordinate goal condition, answering a problem correctly required computing the correct answer for each step in the problem. In the no superordinate goal condition participants could get the wrong answer for one step but still get the others correct. Additionally, participants with task cues ($M = .946$, $SE = .006$) were more accurate than those without ($M = .925$, $SE = .006$) ($F(1, 156) = 6.7$, $p = .01$, $\eta^2 = .042$). This effect is consistent with decay of goal activation when no task cues are present. No other contrasts were significant (all $ps > .1$).

If the task cues were preventing goal activation decay, then the participants who did not have task cues should have become gradually less accurate in later steps than in earlier steps (Altmann, 2002; Altmann & Gray, 2008). To determine if this was the case, I performed a follow-up 2 x 5 repeated-measures ANOVA on the accuracy data in the no superordinate goal condition (These data were not available for the superordinate goal condition because answers were not collected at every step). The analysis indicated that, although there is an overall main effect of task cue ($F(1,86) = 9.5$, $p < .003$, $\eta^2 = .10$), there was no main effect of step ($p > .1$) or interaction between step and task cue ($p > .1$). Thus, Experiment 2 failed to replicate the within-run error increase effect reported in previous work (Altmann, 2002; Altmann & Gray, 2008). These data are plotted in Figure 5.

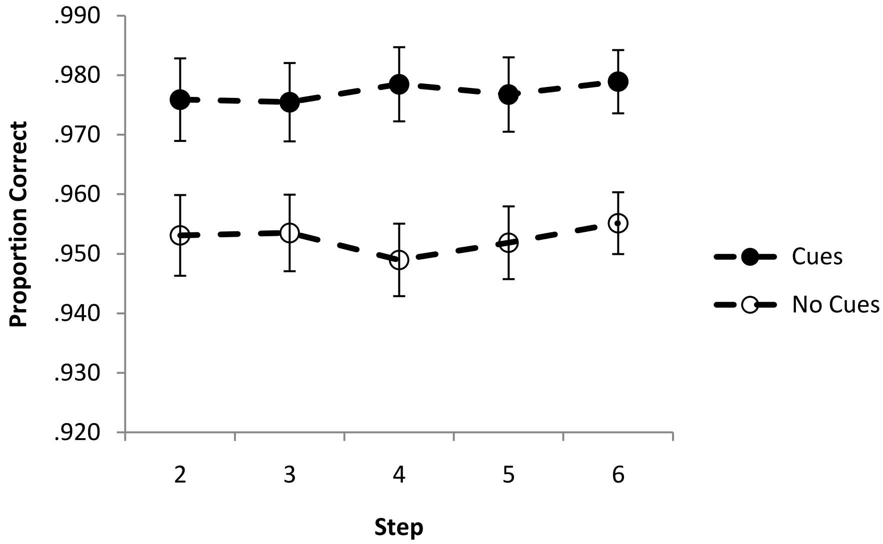

Figure 5. Accuracy data for the no superordinate goal condition in Experiment 2.

Response Time. In the superordinate goal condition, response times for each problem step were measured from the offset of the previous digit until the spacebar press. In the no superordinate goal condition, response times were measured from the onset of the starting number until the participant confirmed the trial answer with the spacebar press. Response time data were analyzed using a 2 x 2 x 2 x 5 repeated measures ANOVA with problem type (addition, subtraction) and step (2-6) as within-subjects factors and superordinate goal (present, absent) and task cues (present, absent) as between-subjects factors. The main effect of step was significant ($F(4, 624) = 6.6$, $p < .0005$, $\dot{\eta}^2 = .041$) and once again, the linear contrast of step was significant ($F(1,156) = 20.8$, $p < .0005$, $\dot{\eta}^2 = .118$). The main effect of problem type was also significant ($F(1,156) = 93.2$, $p < .0005$, $\dot{\eta}^2 = .374$), with subtraction (M = 1480 ms, SE = 29 ms) taking longer than addition (M = 1344, SE = 26 ms). Response times were also significantly longer in the no superordinate goal condition (M = 1537 ms, SE = 36) than in the superordinate goal condition (M = 1287, SE = 39) ($F(1, 156) = 22.4$, $p < .0005$, $\dot{\eta}^2 = .125$). This may reflect the fact that participants had to enter an answer at every step only in the no superordinate goal condition.

The ANOVA also revealed several significant interaction effects. The step by superordinate goal interaction was marginally significant ($F(4, 624) = 2.1$, $p = .086$, $\dot{\eta}^2 = .013$). This appears to reflect a larger overall rate of slowing in the superordinate goal condition, perhaps due to working memory load. The superordinate goal by problem type interaction was also significant ($F(1,156) = 14.7$, $p < .0005$, $\dot{\eta}^2 = .086$). This interaction appears to be driven by a larger benefit of superordinate goal for addition than subtraction, as illustrated by Figure 6. Neither of the hypothesized interactions was significant. Task cues did not interact with step ($F(4, 624) = 1$, $p = .406$, $\dot{\eta}^2 = .006$) and the three-way interaction between task cues, step, and superordinate goal was not significant ($F(4, 624) = .8$, $p = .553$, $\dot{\eta}^2 = .005$). Figure 7 depicts response time as a function of these variables.

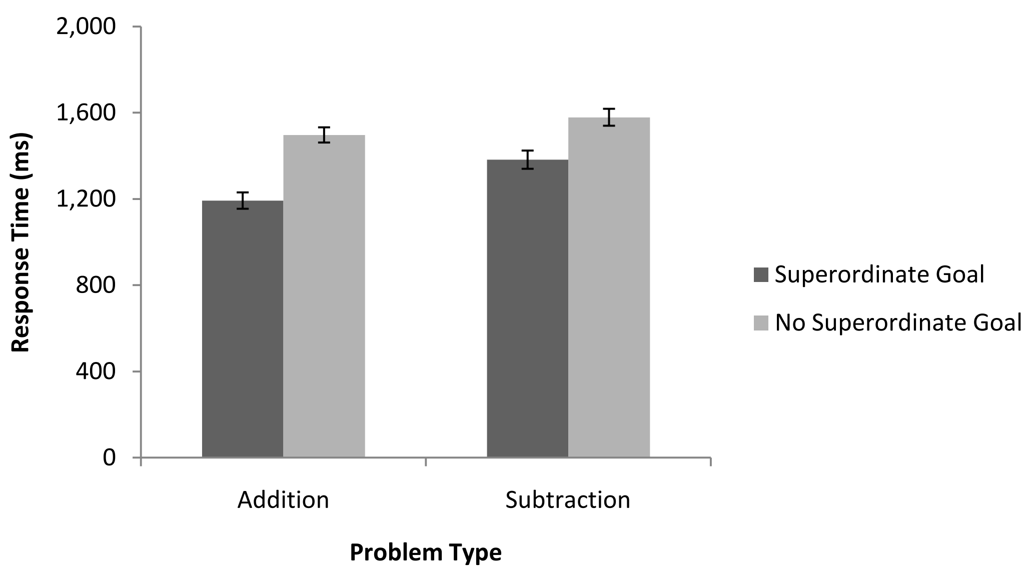

Figure 6. Response time by problem type and superordinate goal in Experiment 2.

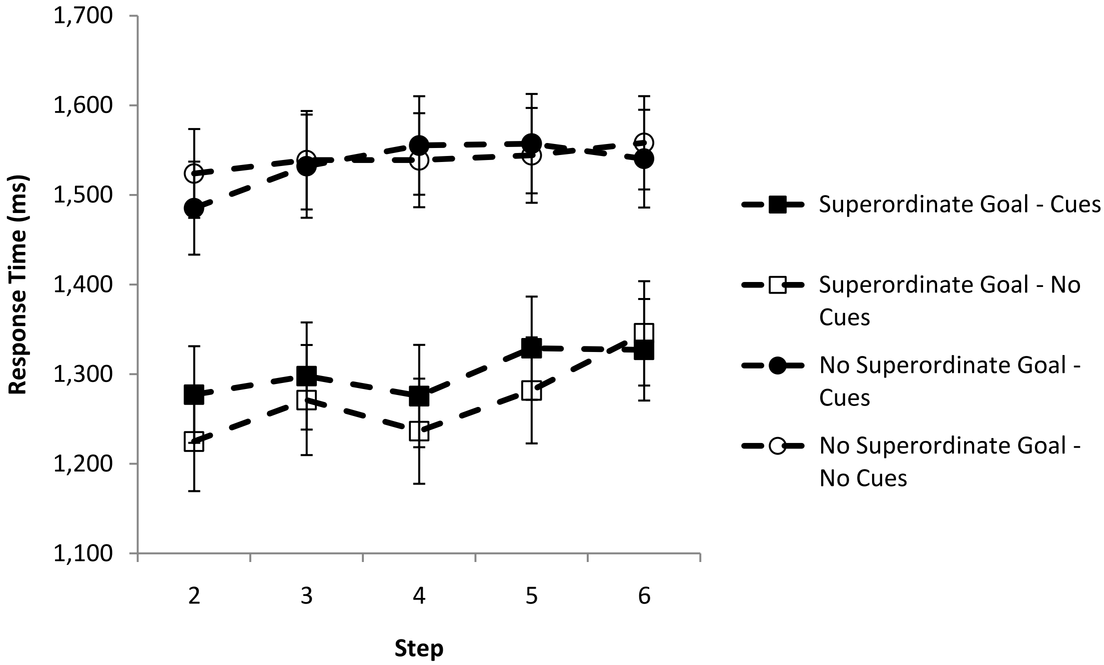

Figure 7. Response time by problem step, superordinate goal, and task cue in Experiment 2.

Operation Span. For each participant, an accuracy score on the mathematics task was calculated and all participants achieving less than 70 percent accuracy were excluded. I did not use the 80 percent cutoff used by Turner and Engle (1989) because doing so would have resulted in discarding a very large proportion of the data. Even with this more lenient criterion, 22 subjects were excluded. Also, as a result of a program error, data from 13 subjects were lost. The final sample size for this analysis was 125.

Each participant's working memory score was computed using the "Total Memory Span" method reported in Turner and Engle (1989). Participants received one point for every correctly recalled word. No points were awarded for a word if it was the final word in the set and it was reported first. Further, misspelled words counted as correct unless the misspelling resulted in a different word. For example, if the word was "rain," participants received a point if they typed "rane" but received no point for "train." Finally, no points were given or taken away for intrusions.

Each participant's response time data were plotted as a function of step and fitted with a linear regression function. The slopes of these lines were used as a measure of slowing (M = 5.8, SE = 4.6). The response time slopes were then entered as a dependent variable in a linear regression analysis with working memory score (M = 53.6, SE = .4), task cue (-1 for absent, 1 for present) and superordinate goal (-1 for absent, 1 for present), and all two and three-way interactions of these variables as predictors. This regression was not significant ($F(7,117) = .82$, $p = .574$). This suggests that, at least in this limited sample, differences in working memory capacity did not moderate the effects of superordinate goals.

Discussion

Overall, the data do not provide clear support for either Altmann and Trafton's Memory for Goals (MFG) model or Norman and Shallice's Interactive Activation (IAN) model. The finding that task cues boosted accuracy regardless of superordinate goal partially supports the MFG model because it suggests that the subgoals required additional activation. However, the fact that this benefit was constant across steps is puzzling from the MFG perspective because goal activation should have been decaying with increasing step. One possible explanation is that reading the task cues caused a speed-accuracy tradeoff. Although there was a nonsignificant trend such that those with task cues (M = 1418 ms, SE = 38 ms) were slightly slower than those without task cues (M = 1406 ms, SE = 37 ms), the small size of the trend makes this explanation questionable.

The response time data replicate the data of Experiment 1. When a superordinate goal is present, no benefit of task cue was observed for response time. However, the null effect of task cue in the no superordinate goal condition is puzzling and not predicted by either model. In the

MFG model, activation levels are closely tied to response time, so the continuous activation provided by the task cues should have speeded performance. A similar result can be expected from the IAN model because no superordinate goal is increasing goal activation levels. One possible explanation is that the goal encoding process is time-consuming (which partially accounts for task-switching effects, see Altmann & Gray, 2008). Therefore, on every step, the system must spend time encoding a new task cue when it already has an appropriate cue in memory. However, this explanation does not account for the improvement in accuracy when task cues are present. If task cues are slowing down processing then they should also be interfering with encoding the digit (which is presented for only 200 ms).

Participants were faster and less accurate in the superordinate goal condition, but they also had to enter fewer answers in this condition. So in this case, the effect of goal structure is confounded with task demands. It is also possible that the difference is the result of a speed/accuracy tradeoff. Perhaps participants were less inclined to closely monitor answer accuracy at each step when they did not have to report the answer.

The superordinate goal variable interacted with problem type in an interesting and unexpected way. Participants were faster in the superordinate goal condition for both problem types, and this benefit was larger for addition problems than subtraction problems. People tend to be more skilled in addition than subtraction (Campbell & Xue, 2001), possibly due to higher levels of practice. It is possible that this higher amount of practice resulted in streamlined sets of addition goals that can be more easily activated by superordinate goals. It is difficult to say for certain because the IAN model does not specify how organization in such schematic knowledge develops (Botvinick & Plaut, 2004).

General Discussion

The central question pursued in this thesis concerns whether loss of activation in subgoals is mitigated by superordinate goals. The Interactive Activation (IAN) model, a hierarchical schema model, suggests that superordinate goals should restore lost activation in subgoals and thus offset the effects of decay (Cooper & Shallice, 2000). By contrast, the Memory for Goals (MFG) model assumes that superordinate goals boost subgoal activation only during planning and therefore should not offset decay. In both experiments, a task cue manipulation was employed as a measure of goal activation decay. If subgoals were losing activation, then task cues should improve performance. If no activation was being lost, then no task cue effect should be observed.

The results do not provide a definitive answer to this question. In Experiment 1, performance slowed down linearly with problem step. However, there was no effect of task cue, suggesting slowing was not due to decay of goal activation. Similarly, the response time data showed no evidence of goal activation loss, as task cue had no effect on slowing regardless of the goal structure. In Experiment 2, task cues improved accuracy. But the benefit to accuracy was constant across steps, which does not support the hypothesis that the task cues were restoring decayed activation. Experiment 2 replicated the stepwise slowing observed in Experiment 1. Once again, no effect of task cue on response time was observed. Problem type interacted with superordinate goal, suggesting that the superordinate goal provided a larger advantage to response time in addition problems than that of subtraction problems.

The presence of a main effect of task cue on accuracy in Experiment 2 but not Experiment 1 conceptually replicates Miyake et al. (2004), which found that verbal cues are more effective at reducing task switching costs than symbolic cues. It is somewhat puzzling that the benefit is not greater in later steps than in earlier steps. If the subgoals are losing activation with every step, then accuracy should become worse on later steps when task cues are not present. However, this interaction was not present. Therefore, this result is not consistent with Altmann and Trafton's Memory for Goals (MFG) model. It is possible that the effect is due to a speed/accuracy tradeoff where those with task cues are performing slightly more slowly to maintain higher accuracy. This could be caused by the additional cognitive resources required to process the verbal cues at every step.

In terms of response time, the no superordinate goal condition in Experiment 2 failed to replicate previous studies of goal activation decay (Altmann, 2002; Altmann & Gray, 2008). That is, task cues did not mitigate within-run slowing. There are several possible reasons for why this occurred. One possibility is that participants were more focused on performing accurately than quickly. It is possible that a strategy that emphasizes accuracy may diminish the effect of goal activation on response time. However, there is evidence to suggest that participants will strive to reduce response time (even by milliseconds) in laboratory tasks (Gray & Boehm-Davis, 2000). Additional work might rule out this possibility by placing a greater emphasis on response times in the instructions.

An important difference between the reported task and those modeled by Altmann and Gray (2008) is the number of available response alternatives. Increasing the number of response alternatives decreases the amount of activation received by each individual response from a given source – a phenomenon known as the fan effect (Anderson & Reder, 1999; Anderson, 1974). The spreading activation mechanisms that underlie this effect are implemented in both the IAN and MFG models (Altmann & Trafton, 2002; Cooper & Shallice, 2000). The task-switching paradigms employed by Altmann and Gray (2008) required simple, two-choice decisions. In the serial arithmetic task, by contrast, participants had to select one of ten choices

(1-9). In this case, each goal ("Add" for instance) has a certain quantity of activation, N, that it can spread to its responses. Since there are 10 possible responses, each response receives N/10 units of activation in a given cycle. Conversely, a parity judgment (Altmann & Gray, 2008) has only two responses, so each one receives N/2 units of activation per cycle. This additional activation would decrease the probability that all of the responses are below threshold on a given cycle, thus reducing response time.

Moreover, numbers, like words, may tend to automatically cue certain goals when perceived (Stroop, 1935). Specifically, while words activate procedures for reading, numbers may activate mathematical operations or arithmetic facts. In the experiments discussed in Altmann and Gray (2008), participants made semantic judgments about numbers. However, there is evidence suggesting that numerical stimuli prime their associated arithmetic facts (Rusconi, Priftis, Rusconi, & Umilta, 2006). Such automatic activation could interfere with processing the current task by activating facts associated with the competing task (e.g. addition facts primed during the execution of the subtraction task). If this explanation is true, it would suggest possible boundary conditions on the within-run slowing and within-run error increase phenomena modeled in Altmann and Gray (2008).

It may at first seem strange that task cue effects were observed for accuracy and not for response time. This is perhaps due to the fact that accuracy and speed are determined by different mechanisms in both models. In the MFG framework, slowing results when no goals are above threshold. Accuracy errors, by contrast, result when the incorrect goal is selected.   In these experiments, the incorrect goal may be receiving activation from the stimulus, so it should be incorrectly retrieved more often. These two mechanisms – low activation and faulty selection – occur in the IAN model as well.

The interaction between problem type and superordinate goal on response time is interesting and raises questions for further study. This interaction appears to be driven by a larger benefit of superordinate goals for addition problems than for subtraction problems, suggesting that the benefits of superordinate goals may be limited to highly practiced skills. Perhaps the IAN model is more applicable to tasks in which a person is highly skilled. This makes sense because it was originally applied to routine activities like making coffee (Cooper & Shallice, 2000). Conversely, the MFG model may be more appropriate for tasks in which skill is low and no firm goal hierarchy has been established. This result also suggests that the amount of activation that superordinate goals provide to subgoals is lower than I hypothesized. Perhaps superordinate goals are providing only a small boost of activation rather than the large amount required to completely offset decay. Alternatively, perhaps the fact that superordinate goals are themselves subject to decay is a limiting factor on the amount of activation they can provide. As superordinate goals decay, they may spread less activation to their subgoals.

The lack of any effect of task cue on response time suggests that the linear slowing observed in both experiments is not the result of goal activation decay. Rather, it may have been due to proactive interference from previous subtotals. This account explains why a reaction time interaction was observed between superordinate goal and step in Experiment 2. In the superordinate goal condition, participants had to maintain subtotals from step to step. Every time a new total was introduced, the older subtotals competed with it for selection. In the no superordinate goal condition, subtotals would still interfere with later subtotals. However, this interference was weaker because people did not have to maintain the subtotals after they give each response.

Although it is possible that fatigue may have contributed to the slowing within problems, this seems unlikely because the problems were relatively short. Further, there was not a substantial break between problems, so fatigue that accumulated on one problem could have easily carried over to the next. Another possible source of slowing is anticipation of an upcoming task switch. Although steps were taken to reduce this, it was possible to predict the end of a problem because the probability that the problem would end on a given step increased in later steps. Participants may have been slowing down in order to prepare for the beginning of the next problem. They may have even been slowing down in preparation to report the final answer. One might argue that the failure to find support for the IAN model is because the task does not require a distinct superordinate goal for addition and subtraction. Rather, both tasks may share a single superordinate goal (e.g. "compute total" rather than "compute sum"). If this was true, then the activation from the superordinate goal would be divided among the two tasks, and would therefore provide no net benefit. This shared-superordinate goal hypothesis is also inconsistent with the data. If the currently irrelevant goal is receiving activation from the superordinate goal, then more interference between the subgoals should occur in the superordinate goal condition. Therefore, this hypothesis would predict a larger task cue effect in the superordinate goal condition than in the no superordinate goal condition.

It may also appear that the superordinate goal and no superordinate goal tasks have too many differences in terms of timing and response requirements to be comparable. This is an issue only insofar as it could "mask" the effect of task cues on slowing in one of the conditions and not the other. If task cue effects had appeared in only one of the conditions, follow up experiments might have been necessary to ensure that one of these differences did not confound the results. However, because task cues did not reduce slowing in either condition, the overall interpretation remains the same.

Future Directions

A future task that could address the issues raised above is a binary logic task. Previous work in our lab suggests that binary logic tasks behave similarly to arithmetic tasks in terms of goal activation (Sohn & Carlson, 1998). The task would involve the application of simple logical operators with two possible outcomes, “0” and “1.” For instance, given a starting digit and a new digit, the participant would have to apply a logical rule to obtain a new result. The new result would then be immediately reported or carried over to the next step. This task has fewer response options than the serial arithmetic task, which should reduce fan effects. The binary logic task also solves the problem of automatic cuing of mathematical operators by numerical stimuli because the logical operators will be relatively unfamiliar to the participants. Similar to serial arithmetic, the binary logic task has a flexible goal structure, because each result can easily serve as input for an additional step.

People rely on multiple types of cues to determine how goals should be organized within a hierarchy. Both theories discussed here assume that goals which share a common purpose (e.g. making coffee) are grouped together under a superordinate goal. However there are other ways in which this grouping can be achieved. Lien and Ruthruff (2004) found that tasks that have close temporal or spatial proximity tend to be grouped together. This is evidenced by group-level switch costs which overshadow task-level switch costs. Schneider and Logan (2006) demonstrated that performing tasks in a consistent sequence induces similar group-level switch costs, even without close temporal or spatial proximity. Thus it is possible that the questions raised in this thesis could be answered using one of the methods reported in Lien and Ruthruff (2004) and Schneider and Logan (2006) to impose a hierarchy onto the tasks employed by Altmann (2002) and Altmann and Gray (2008). This similarity would help to rule out many of the extraneous factors (such as proactive interference of subtotals) present in the above data.

Imposing hierarchical structure on the tasks modeled in Altmann and Gray (2008) might be achieved by simply arranging the tasks in two different sequences and training participants on those sequences. In order to observe within-run effects, it would be necessary that these sequences contain runs of the same task at least four trials long (e.g. AAAABB) because this is the shortest run-length that has produced the effects (Altmann, 2002; Altmann & Gray 2008). Random task sequences would be used as a control. In this paradigm, the IAN model would predict an absence of within-run effects in the sequence condition but not in the random condition, whereas the MFG model would predict within-run effects in both.

If the results support the MFG model and not the IAN model, the possibility remains that a stronger hierarchy is necessary to offset subgoal activation decay. This hypothesis could be addressed using a modified serial arithmetic task. Each problem in this task would consist of two multiple step sub-problems. At the end of each problem, participants would be required to combine the two answers to produce a final result. For instance, sub-problems 1 and 2 would each require adding 4 separate numbers together. At the end of the problem, participants would subtract the two subtotals and report the result. The goal structure of this task better fits the definition of hierarchical tasks outlined in Carlson and Lundy (1992). Perhaps the more explicit hierarchical structure of this task would be enough to cause a hierarchical goal structure to form. If these experiments still fail to produce a benefit of superordinate goal, then the activation dynamics of the IAN model may have to be reconsidered.

Conclusion

The reported studies are inconclusive concerning the exact nature of the impact of superordinate goals on the maintenance of subgoals in cognitive tasks. Although participants slowed down in later steps of the arithmetic problems, this slowing was not mitigated by task cues, suggesting that goal activation decay was not the cause. Task cues resulted in better accuracy in Experiment 2, but this effect was constant across steps. This lack of interaction makes a goal activation interpretation questionable. The results do not clearly support either the IAN model or the MFG model, as both models predicted a loss of goal activation in the no superordinate goal condition. Follow-up studies are needed to clarify the results with respect to these two models. One follow-up would be to use a simpler paradigm in order to isolate the effects of interest and remove the extraneous task features that may be obscuring the results. In addition, a task with a more explicit hierarchical structure may be needed to cause the formation and use of superordinate goals. Further experimentation is necessary to understand the function of goal hierarchies in the maintenance of goal activation.

References

Altmann, E. M. (2002). Functional decay of memory for tasks. *Psychological Research*, *66*, 287-297.

Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science*, *26*, 39-83.

Altmann, E. M., & Trafton, J. G. (2007). Timecourse of recovery from task interruption: Data and a model. *Psychonomic Bulletin & Review, 14*(6), 1079-84.

Anderson, J. R, Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological review*, *111*(4), 1036-60.

Anderson, J. R., & Reder, L. M. (1999). The fan effect: New results and new theories. *Journal of Experimental Psychology: General*, *128*(2), 186-197.

Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. *Cognitive Psychology*, *6*(4), 451-474.

Ashcraft, M. H. (1992). Cognitive arithmetic: A review of data and theory. *Cognition*, *44*, 75-106.

Botvinick, M., & Plaut, D. C. (2004). Doing without schema hierarchies: a recurrent connectionist approach to normal and impaired routine sequential action. *Psychological Review*, *111*(2), 395-429.

Byrne, M., & Bovair, S. (1997). A working memory model of a common procedural error. *Cognitive Science*, *21*(1), 31-61.

Campbell, J. I. D., & Xue, Q. (2001). Cognitive arithmetic across cultures. *Journal of Experimental Psychology: General*, *130*(2), 299-315.

Carlson, R. A., & Lundy, D. H. (1992). Consistency and restructuring in learning cognitive procedural sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(1), 127-141.

Cooper, R., & Shallice, Tim. (2000). Contention Scheduling and the Control of Routine Activities. *Cognitive Neuropsychology*, *17*(4), 297-338.

Duncan, J., Emslie, H., Williams, P., Johnson, R., & Freer, C. (1996). Intelligence and the frontal lobe: the organization of goal-directed behavior. *Cognitive Psychology*, *30*(3), 257-303.

Duncan, J., Parr, A., Woolgar, A., Thompson, R., Bright, P., Cox, S., et al. (2008). Goal neglect and Spearman's g: competing parts of a complex task. *Journal of Experimental Psychology. General*, *137*(1), 131-48.

Gray, W. D., & Boehm-davis, D. A. (2000). Milliseconds matter : An introduction to microstrategies and to their use in describing and predicting interactive behavior. *Journal of Experimental Psychology: Applied*, *6*(4), 322-335.

Hodgetts, H. M., & Jones, D. M. (2006). Interruption of the Tower of London task: Support for a goal-activation approach. *Journal of Experimental Psychology: General, 135*(1), 103-15.

Kane, M. J., & Engle, Randall W. (2003). Working-memory capacity and the control of attention: The contributions of goal neglect, response competition, and task set to Stroop interference. *Journal of Experimental Psychology: General*, *132*(1), 47-70.

Koch, I. (2003). The role of external cues for endogenous advance reconfiguration in task switching. *Psychonomic Bulletin and Review*, *10*(2), 488–492.

Lien, M. C., & Ruthruff, E. (2004). Task switching in a hierarchical task structure: Evidence for the fragility of the task repetition benefit. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *30*(3), 697-713.

Lundy, D. H., Wenger, J. L., Schmidt, R. J., & Carlson, R. A. (1994). Serial step learning of cognitive sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(5), 1183-1195.

Lustig, C., May, C. P., & Hasher, L. (2001). Working memory span and the role of proactive interference. *Journal of Experimental Psychology: General*, *130*(2), 199-207.

McVay, J. C., & Kane, M. J. (2009). Conducting the train of thought: Working memory capacity, goal neglect, and mind wandering in an executive-control task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(1), 196-204.

Norman, D A. (1981). Categorization of action slips. *Psychological Review*, *88*, 1-15.

Rusconi, E., Priftis, K., Rusconi, M. L., & Umilta, C. (2006). Arithmetic priming from neglected numbers. *Cognitive Neuropsychology*, *23*(2), 227-39.

Schneider, D. W., & Logan, G. D. (2006). Hierarchical control of cognitive processes: Switching tasks in sequences. *Journal of Experimental Psychology: General*, *135*(4), 623-40.

Sohn, M., & Carlson, R. A. (1998). Procedural Frameworks for Simple Arithmetic Skills. *Journal of Experimental Psychology: Learning Memory and Cognition*, *24*(4), 1052-1067.

Turner, M. L., & Engle, R. W. (1989). Is working memory capacity task dependent? *Journal of Memory and Language*, *28*(2), 127-154.

APPENDIX

Experiment 1 Accuracy

Within-Subjects Effects

Source	Sum of Squares	df	Mean Square	F	p	Partial Eta Squared
ptype	.006	1	.006	1.033	.313	.015
ptype * Group	.000	1	.000	.001	.977	.000
Error	.389	68	.006			

Between-Subjects Effects

Source	Sum of Squares	df	Mean Square	F	p	Partial Eta Squared
Intercept	111.493	1	111.493	12625.547	.000	.995
Group	.003	1	.003	.292	.591	.004
Error	.600	68	.009			

Experiment 1 Response Time

Within-Subjects Effects

Source	Sum of Squares	df	Mean Square	F	p	Partial Eta Squared
ptype	4410153	1	4410153	60.523	.000	.475
ptype * Group	8326	1	8326	.114	.736	.002
Error(ptype)	4882133	67	72868			
step	1038433	4	259608	6.236	.000	.085
step * Group	53997	4	13499	.324	.862	.005
Error(step)	11157584	268	41633			
ptype * step	223907	4	55977	1.801	.129	.026
ptype * step * Group	143248	4	35812	1.152	.332	.017
Error(ptype*step)	8329528	268	31080			

Within-Subjects Contrasts

Source	Contrast Type	Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared
step	Linear	884787	1	884787	22.806	.000	.254
	Quadratic	65389	1	65389	1.343	.251	.020
step * Group	Linear	32406	1	32406	.835	.364	.012
	Quadratic	7290	1	7290	.150	.700	.002
Error(step)	Linear	2599367	67	38797			
	Quadratic	3261246	67	48675			
ptype * step	Linear	124614	1	124614	3.265	.075	.046
	Quadratic	23622	1	23622	.834	.365	.012
ptype * step * Group	Linear	54353	1	54353	1.424	.237	.021
	Quadratic	3903	1	3903	.138	.712	.002
Error(ptype*step)	Linear	2556798	67	38161			
	Quadratic	1898638	67	28338			

Between-Subjects Effects

Source	Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared
Intercept	1208267346	1	1208267346	1616.120	.000	.960
Group	58908	1	58908	.079	.780	.001
Error	50091517	67	747635			

Experiment 2 Accuracy

Within-Subjects Effects

Source	Sum of Squares	df	Mean Square	F	p	Partial Eta Squared
ptype	.001	1	.001	.417	.519	.003
ptype * Group	.000	1	.000	.068	.795	.000
Error(ptype)	.361	158	.002			

Between-Subjects Effects

Source	Sum of Squares	df	Mean Square	F	p	Partial Eta Squared
Intercept	281	1	281	43415.486	.000	.996
Group	< 1	1	< 1	4.305	.040	.027
Error	1	158	< 1			

Experiment 2 Response Time

Within-Subjects Effects

Source	Sum of Squares	df	Mean Square	F	p	Partial Eta Squared
step	790789	4	197697	6.620	.000	.041
step * Group	119590	4	29898	1.001	.406	.006
step * Experiment	245099	4	61275	2.052	.086	.013
step * Group * Experiment	90542	4	22636	.758	.553	.005
Error(step)	18635302	624	29864			
ptype	7323893	1	7323893	93.223	.000	.374
ptype * Group	41812	1	41812	.532	.467	.003
ptype * Experiment	1153842	1	1153842	14.687	.000	.086
ptype * Group * Experiment	99145	1	99145	1.262	.263	.008
Error(ptype)	12255796	156	78563			
step * ptype	44703	4	11176	.439	.781	.003
step * ptype * Group	70078	4	17520	.688	.601	.004
step * ptype * Experiment	163359	4	40840	1.603	.172	.010
step * ptype * Group * Experiment	173900	4	43475	1.706	.147	.011
Error(step*ptype)	15899503	624	25480			

Within-Subjects Contrasts

Source	Contrast Type	Sum of Squares	df	Mean Square	F	p	Partial Eta Squared
step	Linear	696194	1	696194	20.803	.000	.118
	Quadratic	4	1	4	.000	.991	.000
step * Group	Linear	6994	1	6994	.209	.648	.001
	Quadratic	79180	1	79180	2.982	.086	.019
step * Experiment	Linear	59762	1	59762	1.786	.183	.011
	Quadratic	119439	1	119439	4.498	.036	.028
step * Group * Experiment	Linear	66425	1	66425	1.985	.161	.013
	Quadratic	6537	1	6537	.246	.620	.002
Error(step)	Linear	5220675	156	33466			
	Quadratic	4142810	156	26556			
step * ptype	Linear	5030	1	5030	.209	.648	.001
	Quadratic	16	1	16	.001	.977	.000
step * ptype * Group	Linear	5188	1	5188	.216	.643	.001
	Quadratic	4737	1	4737	.238	.626	.002
step * ptype * Experiment	Linear	19089	1	19089	.794	.374	.005
	Quadratic	131334	1	131334	6.595	.011	.041
step * ptype * Group * Experiment	Linear	1723	1	1723	.072	.789	.000
	Quadratic	22484	1	22484	1.129	.290	.007
Error(step*ptype)	Linear	3749861	156	24038			
	Quadratic	3106549	156	19914			

Between-Subjects Effects

Source	Sum of Squares	df	Mean Square	F	p	Partial Eta Squared
Intercept	3167782369	1	3167782369	2838.145	.000	.948
Group	51233	1	51233	.046	.831	.000
Experiment	24949461	1	24949461	22.353	.000	.125
Group * Experiment	129713	1	129713	.116	.734	.001
Error	174118653	156	1116145			