6-13-2012

Individual Differences in Function Learning as They Relate to the Learning of Conceptual Information

Cynthia Fadler
Washington University in St. Louis

WASHINGTON UNIVERSITY IN ST. LOUIS

Department of Psychology

Dissertation Examination Committee:
Mark A. McDaniel (Chair)
David A. Balota
Heather J. Rice
Henry L. Roediger, III
R. Keith Sawyer
James V. Wertsch

Individual Differences in Function Learning as They Relate to the Learning of Conceptual Information

by

Cynthia Lynn Fadler

A dissertation presented to the
Graduate School of Arts and Sciences
of Washington University
in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy

August 2012

Saint Louis, Missouri

Abstract

Individual differences have not often been considered within the problem-solving or concept-learning literatures despite the indication that some individuals are better able to transfer to novel problems and that manipulations in strategy can effectively increase the ability to transfer (Gick & Holyoak, 1983). Research in the function-learning domain indicates that there may be two qualitatively different types of learners: those who remember distinct example associations (exemplar learners) and others who abstract rules that govern each association (rule learners; DeLosh, Busemeyer, & McDaniel, 1997). Data from two unpublished studies (McDaniel, Cahill, Robbins, & Trumpower, 2012; Fadler, Lee, Scullin, Shelton, & McDaniel, 2012) have demonstrated the stability of these two types of learning across a variety of different higher order problem-solving, concept-learning, and cognitive tasks. However, it remains to be seen whether these differences between learners have implications for the type of conceptual material often used in classrooms.

In the current project, this issue was addressed through two experiments. During Experiment 1, participants were first identified as exemplar or rule-based learners on the basis of function learning transfer performance. Each group then read several passages and answered questions about the passages that ranged in their degree of transfer. Rule learners performed better than exemplar learners on each question type and the two types of learners also demonstrated qualitatively different processing during function learning training

and on a test of analogical transfer. The data from Experiment 2 showed that rule learners behaved qualitatively differently from exemplar learners during function learning training but failed to replicate the passage data from Experiment 1. However, a benefit was found on recognition memory for exemplar learners on a concept-learning task.

The current study is the first to show differential benefits for exemplar and rule-based processing. It also provides evidence that function-learning tendency can be used to predict differences on concept-learning tasks and that only rule learning is associated with abstraction ability. The findings suggest that individual differences should be considered both in current hybrid models of categorization, but also potentially in classrooms that might rely heavily on problem solving, where the differences in types of learners may have an impact on student performance and understanding.

Acknowledgements

First and foremost, I would like to thank my advisor and committee chair, Mark McDaniel, for taking a chance on me and for his enthusiastic advice and support over the past three years. I also must express my appreciation to Heather Rice, who has served as a constant source of encouragement and guidance throughout my graduate career. I would also like to acknowledge Dave Balota, Roddy Roediger, Keith Sawyer, and Jim Wertsch for serving on my dissertation committee and providing helpful comments on the project.

I am extremely grateful to the past and present members of the Memory and Complex Learning Laboratory for their warmth, advice, thoughts on data, and unfailing support. The environment that they have established is unique and cherished.

I also thank my husband, Trevor, who has kept all the other plates spinning throughout this process, and the rest of my immediate family: my parents, Ed and Jayne Wooldridge, and my three best friends and dear sisters, Julie, Jessie, and Mickey. Without the continual love and encouragement of my family and friends, this document could not have come to fruition.

This research was supported by the James S. McDonnell Foundation 21st Century Science Initiative grant, Applying Cognitive Psychology to Enhance Educational Practice: Bridging Brain, Mind, and Behavior Collaborative Activity Award.

Table of Contents

List of Tables

List of Figures

List of Appendices

Individual Differences in Function Learning as They Relate to the Learning of Conceptual Information

Within the problem-solving domain, research on individual differences is relatively rare. A few studies have shown that some individuals are able to transfer solutions to novel problems, while others do not transfer (Gick & Holyoak, 1983; Novick, 1988). These findings suggest that the tendency to transfer may vary across individuals, but these potential differences are rarely discussed in the literature. Instead of investigating potential qualitative differences in learners, researchers generally attempt to explain variations in transfer ability through a continuous dimension of intelligence (see Wenke, Frensch, & Funke, 2005, for review). Although it is possible that intelligence might account for these differences, empirical work has not been conducted to demonstrate that intelligence is a reliable predictor of transfer ability. McDaniel, Cahill, Robbins, and Trumpower (2012) instead argued that there are two qualitatively different types of learners: those that retain specific example-response associations (exemplar learners) and those that abstract an underlying rule that governs each association (rule learners).

The goal of the present research was to demonstrate that these qualitative differences in learning tendency extend to other domains. Specifically, learners who remember example-response associations and those who abstract rules may perform differently in a classroom setting with educational materials. Experiment 1 was designed to examine a potential interaction in conceptual material such that exemplar learners perform better on tests of retention, while rule learners are better able to transfer to novel situations. Experiment 2 was designed to replicate and extend the findings of Experiment 1 as well as to explore the effects of function learning tendency in other concept-learning tasks.

In order to provide a foundation for the project, I will first briefly describe models of rule and exemplar processing and then outline the research demonstrating individual differences in problem solving and specifically in function learning. I will then describe several unpublished studies that examine the stability of individual differences in function-learning tendency and how these might have educational implications.

Rule and Exemplar Processing

Historically, there has been a debate in the categorization literature between those who argued that categorization was supported fully by either rule-based processing (Reed, 1972; Einhorn, Kleinmuntz, & Kleinmuntz, 1979; Nosofsky, Palmeri, & McKinley, 1994) or exemplar-based processing (Estes, 1994; Kruschke, 1992; Medin & Shaffer, 1978; Nosofsky, 1984; Nosofsky & Palmeri, 1997; Nosofsky & Johansen, 2000; DeLosh et al., 1997). Generally, proponents of rule-based models explained that individuals use a controlled cognitive process in order to abstract underlying information that governs classification of all (or most, see Nosofsky et al., 1994) items. When encountering a novel item, it could then be compared to the abstraction to determine its category membership. This theory is contrasted with exemplar models, which described learning as occurring through a memorization process such that each exemplar is represented in memory. Upon encountering a novel item, the item is compared with the stored representations of exemplars and the novel item would be categorized according to the most similar memorized exemplar (or an accumulation of information from weighted averages; see Juslin et al., 2003, for more details).

More recently, several hybrid models have been proposed to account for the discrepancies between rule and exemplar models (Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Erickson & Kruschke, 1998; Anderson & Betz, 2001; Bott & Heit, 2004). These hybrid models describe rule and exemplar processing as two qualitatively different types of cognitive processing, and these models draw support from data showing that rule and exemplar processing rely on different neural substrates (Smith, Patalano, & Jonides, 1998). Within hybrid models such as the ACT-R model (Anderson & Betz, 2001), the particular type of processing is chosen based on which is most appropriate in the given situation. Anderson and Betz explained, "participants track how well each basis is working on the stimulus set and select each method in rough proportion to its past history of success" (p. 630). Specifically, materials that contain few exemplars, pictorial stimuli, and nominal feedback are more likely to encourage exemplar processing (Juslin, Olssen, & Olssen, 2003). Thus, for hybrid models, the type of processing used depends on the task and materials, but these models do not account for individual differences in the type of processing that might be chosen, particularly in the absence of strong task-based determinants of processing.

Individual Differences in Concept Learning and Problem Solving

As described above, there is little empirical work examining potential qualitative differences between individuals who are able to transfer problem solutions and those who are not. In a study of analogical transfer, Gick and Holyoak (1983) gave half of their subjects two problems with different cover stories, but that required the same solution, while the other half of subjects received only one of these problems. Later, subjects were given a transfer problem that required the same solution. If subjects had previously been given two problems, they were significantly more likely to recognize the utility of the previous solution and therefore answer the transfer problem. Giving subjects the additional problem led learners to recognize that a single schema could govern multiple problems with different surface features and therefore changed the way they approached the novel criterial problem. However, only 52% of subjects were able to transfer even given the two problems with differing surface structure. These data may indicate that learners who were able to transfer may be naturally adopting a different strategy than those who were unable to transfer. Gick and Holyoak explained that the potential mechanism by which individuals were able to transfer was their ability to develop a schema associated with training examples. That is, subjects needed to be aware of the underlying structure of a problem in order to transfer to a new problem with a similar structure. It is possible that those subjects who were able to complete the task did so by using this more successful strategy, while others may have been using a strategy that was qualitatively different.

While few researchers have considered these differences in transfer ability, one early study did examine qualitative differences in learning. Medin, Altom, and Murphy (1984) trained subjects to learn various examples from a category, such that the rule for category membership was not very clear. Medin et al. asked learners to specifically describe the strategies that they used to accomplish the task. While there was not complete data for all subjects, the responses indicated that some learners were categorizing based on an abstracted prototype, while others had memorized the exemplar-category pairs to determine which items were part of the category. Thus it appears that within the problem-solving and concept-learning domains, there exist qualitatively different strategies for learning and transfer.

Individual Differences in Function Learning

In line with Medin et al. (1984), evidence for qualitative differences in transfer also has been found in the function-learning domain. DeLosh, Busemeyer, and McDaniel (1997) trained participants on a set of input-output values that followed a function that was unknown to the participants. After training, subjects were tested on items that were similar to training values, but also on items that fell outside the training range (extrapolation). When using a quadratic function, a single hybrid model (which incorporated both rule-based and exemplar-based strategies) predicted the performance of most learners. However, there were some individuals who showed noticeably different extrapolation patterns. Some learners appeared to continue the function and closely followed the correct quadratic curve, indicating that these individuals had abstracted the underlying function rule and were applying it to the novel input values. Other subjects instead used output values that were very similar to those at the extreme ends of the training range, indicating that these individuals had retained specific input-output associations and were using the closest retained example to produce their responses, thus relying on memorized points instead of a general rule (see Figure 1).

Following the results of DeLosh et al. (1997), McDaniel et al. (2012) examined individual differences in function learning more closely. In Experiment 1a, subjects were given input-output pairs that followed a V-shaped function (described in detail in Experiment 1 below). Participants were then classified as either rule or exemplar learners based on their extrapolation errors. Individuals showing extrapolation that was close to the function (as opposed to flat extrapolation) were classified as *rule learners* for having abstracted the underlying function. Individuals who did not show this functional understanding were classified as *exemplar learners*. These classifications correlated with working memory capacity such that individuals with high working memory were more likely to abstract the underlying function. Within the rule learners, Ravens Advanced Progressive Matrices (RAPM, a non-verbal measure of fluid intelligence or the ability to abstract relations between items; Raven, Raven, & Court, 1998) correlated negatively with the rate at which subjects learned during training. That is, the higher a rule-learner's ability to abstract the less time that it took them to abstract the rule. However, the correlation between RAPM and rate of learning did not extend to the exemplar learners, indicating that their strategy was qualitatively different from rule learners and not based on abstraction ability. In addition, overall RAPM scores did not differ between the two types of learners. As described above, most researchers explain individual differences in transfer ability through differences in intelligence (see Wenke, Frensch, & Funke, 2005 for review). These data collectively indicate that any differences found between learners could not be accounted for by intelligence alone.

In Experiment 1b of McDaniel et al., subjects categorized novel animals into two distinct categories. This method was based on a study by Regehr and Brooks (1993) who initially used the animals to show that natural stimuli can both have underlying rules for categorization, and can also be classified with the use of memory for idiosyncratic features. After training, individuals were given a test in which they categorized repeated training items, as well as a lure that was the "twin" of each training item. That is, for each lure, the idiosyncratic features were almost identical with the training "twin", with only one changed feature. For half of the lures, the changed feature caused the lure to be placed in a different category than its training "twin". For the other half, the changed feature did not change the category membership. McDaniel et al. identified individuals who were classified as exemplar learners (on the function-learning task) who also learned well during training. For these exemplar learners, it appeared that they based their novel categorization on their nearest training example, which caused them to place half of the lures into the wrong category. Rule learners, on the other hand, performed similarly on the two types of lures, indicating that their categorization was based on a qualitatively different strategy. McDaniel et al. again showed the stability of function learning tendencies across materials in Experiment 2, in which subjects were trained to place exemplars into an appropriate abstract coherent category. Subjects were shown four features of a fictitious machine and were asked to classify the machine as a “morkel” or “krenshaw”. While never made explicit to the subject, “morkels” were made up of features which were coherent (made sense together) and “krenshaws” were made up of features that were incoherent (did not make sense; e.g. rolled on wheels in water; see Table 2). Again, subjects who abstracted a rule during the function-learning task were significantly more likely to classify novel exemplars into the appropriate category during the transfer portion of the study.

These data led McDaniel et al. to propose that there are two qualitatively different approaches to concept learning displayed by learners. This theory contradicts models that have argued that concept learning is supported singularly either by exemplar retention (e.g. Kruschke, 1992) or through rule abstraction (e.g. Koh & Meyer, 1991). Instead, McDaniel et al. argued that concept learning might actually be governed by both processes and that differences may exist among individuals in the extent to which they rely on each process. While some learners appear to learn by memorizing specific exemplars associated with a category or response, others tend to abstract an underlying rule that governs the exemplar associations. Importantly, the differences in function-learning tendency were used to predict performance on other higher-order cognitive tasks, indicating that these learning strategies are relatively stable.

While the data supporting individual differences contradict solely rule or exemplar models of concept learning, it is not entirely inconsistent with hybrid models, which incorporate both types of processing. McDaniel et al. demonstrated that individual differences on a task that is most likely to encourage rule-based processing (i.e. function-learning task, because it has a large number of exemplars and uses continuous as opposed to binary feedback; see Juslin et al. for discussion) could be used to predict performance on a task most likely to encourage exemplar-based processing (i.e. from Experiment 1b and 2; again see Juslin et al., 2003). McDaniel et al. explained that, while it is true that certain materials might lend themselves to a particular strategy, within a given set of materials, individuals differ in the way they approach the task, with some relying more on exemplar-based learning and others on abstraction. Indeed, even proponents of hybrid models have described that individual differences in the use of rule versus exemplar strategies have hindered their ability to find task-specific effects (Juslin et al., 2003), yet these individual differences have remained unaccounted for in the literature.

Applying Function Learning Tendency To Education

Individual differences in function-learning tendency appear to be stable across different concept-learning and problem-solving tasks. However, it is unclear whether or not these differences in learning will have implications for other materials. Specifically, the skills that are required in a problem-solving task may differ considerably from the skills needed for learning conceptual material within the classroom. In the function-learning tasks used by DeLosh et al. (1997) and McDaniel et al. (2012), subjects were required to learn input-output pairs along a continuous mathematical function. It is possible that a particular skill underlying this type of learning may be general mathematic ability. That is, rule learners might be those individuals who are better able to abstract a particular mathematical function and exemplar learners, lacking the same ability, must rely on the memorization of input-output pairs to succeed in the task. On the concept-learning tasks used by McDaniel et al., there was no mathematical function required. However, some amount of problem-solving ability might explain the differences between rule and exemplar learners on these tasks as well as the function-learning task. It is possible that the type of skills needed in a mathematical or problem-solving task might not readily map onto the discourse and text processing that is necessary for learning from lectures and textbooks in a typical classroom environment.

However, other tasks have shown differences in learners that may be more comparable to the type of reasoning required in the classroom. In the analogical reasoning studies discussed above (e.g. Gick & Holyoak, 1983), learners must develop a schema for the problem in order to map on the underlying structure to other analogous problems. In order to understand prose material, students may need to develop a schema for the information in order to gain a coherent understanding of the material and to map the information learned onto other novel materials. It is possible that rule learners abstract a schema associated with material presented in the classroom, which allows them to make connections between various materials and gain a better understanding of concepts within the classroom. Exemplar learners, on the other hand, may focus on the surface details of the given information (as in an analogical reasoning task) and therefore be able to retain information presented, but lack a more conceptual understanding of that information, leaving them less able to transfer to novel stimuli. Interestingly, both types of learners might be successful within the classroom. Rule learners may be able to draw upon information learned throughout their school careers to understand the current material, while exemplar learners may have excellent retention of information conveyed, often producing excellent test scores.

While exemplar learners may be successful in achieving high marks on tests, a common goal of education is transfer of knowledge and not solely memorization of facts. Because rule learning produces better extrapolation, it may also produce better transfer of concepts for prose materials. Therefore if the differences in learning tendency were associated with performance on the classroom materials, there would be strong implications for educators who may want to encourage rule-based processing in order to promote better transfer to novel situations.

Following these theoretical ideas, my colleagues and I recently conducted a study in which subjects were given a function-learning task and then were trained on supply and demand economics problems (Fadler, Lee, Scullin, Shelton, & McDaniel, 2012). In this way we attempted to determine whether function learning tendencies would predict differences on similar, but ecologically valid materials. Individuals were told to learn information about the supply and demand of a company in order to give advice to the company following training. After being trained on problems that follow four supply and four demand rules (never made explicit), subjects were given a test in which they answered similar supply and demand problems as well as problems that required them to combine supply and demand information (equilibrium problems). In order to answer these equilibrium problems, subjects had to have a clear conceptual understanding of how and why supply and demand curves shift. Preliminary results indicate that there is an interaction between function learning tendency and scores on different types of problems. Specifically, exemplar learners perform better on items that are very similar to training problems, but rule learners are superior on the novel equilibrium problems. While these data indicate that function learning may have implications for other domains and specifically classrooms which rely heavily on problem solving, the materials are composed of problems and may therefore use skills that are closer to the function-learning task (e.g. mathematical ability) than skills that are needed for learning of conceptual material. It therefore remains to be seen whether differences in function learning can be used to predict differences in more conceptual information.

The Present Research

While there are many potential avenues for examining differences in function learners, the current research focused on retention and transfer of conceptual material (prose passages) similar to what might be learned within a classroom. In addition to narrowing the materials in the current study, it is important to define the constraints of transfer that were used. Barnett and Ceci (2002) explained that there are many dimensions of transfer that make it difficult to classify and test (see Figure 2). Transfer can vary across different levels of learned skills, memory demands, and distance of transfer. In the current study, the materials required different levels of learned skills, such that some items were designed to require a representation of information while others required an underlying principle. The memory demands of the questions differed as well. Some questions required subjects to remember the information that was explicitly requested while other questions required subjects to additionally determine which information they needed to retrieve and to apply it appropriately. Finally, the materials varied from very near transfer to far transfer (to a different knowledge domain; see Appendixes B and C for full materials). An example of a factual question (explicitly stated in the passage) was, "What happens to bats physiologically when in a torpor state?" (Answer: *When in a torpor state, a bat's metabolism slows down, reducing biological activity and conserving energy.*) An example of an inferential question (required transfer) was, "The U.S. Military is looking for inspiration in developing a new type of aircraft that promotes increased maneuverability. How would this new type of aircraft differ from traditional aircrafts like fighter jets?" (Answer: *Traditional aircrafts are modeled after bird wings, which are rigid and good for providing lift. Bat wings are more flexible, and thus an aircraft modeled on bat wings would have greater maneuverability.*) The materials were specifically designed such that the lowest level of question was most likely to address the type of processing engaged by exemplar learners. That is, if exemplar learners tend to use a learning strategy that involves simply committing information to memory, they are likely to perform better on questions that require surface-level information (Van Dijk & Kintsch, 1983). Rule learners, on the other hand, may abstract schemas for the text that can enhance their understanding of the material, leading to better performance on items that require understanding or application to novel domains.

Overview of Experiments

There were three primary purposes of the current study. The unpublished studies described above (McDaniel et al., 2012; Fadler et al., 2012) provide evidence that there are indeed two qualitatively different and identifiable learning tendencies. However, identification of these distinct differences requires replication to demonstrate the strength of the effect. Therefore, the current study was designed to identify rule and exemplar learners using the classification methodology developed by McDaniel et al. and to determine the extent to which these differences have implications on tasks that are unrelated to the problem-solving or concept-learning domains, specifically to analogical reasoning and learning of conceptual material. In addition, the study was designed to determine whether a benefit could be found for exemplar-based processing on recognition memory in a concept-learning task where only a rule-based advantage for categorization had been previously observed. A final exploratory purpose was to determine whether function-learning tendency might be associated with other demographic characteristics of individuals. For example, learners could choose different courses depending on the way in which they approach problems and so we could find that rule and exemplar learners differ in academic major.

Experiment 1

The first goal of Experiment 1 was to identify two qualitatively different strategies for learning through the use of a function-learning paradigm developed by DeLosh et al. (1997) and to replicate the differences in extrapolation profiles reported by McDaniel et al. (2012). Rule learners were expected to abstract the underlying V-shaped function and continue the function in extrapolation, while exemplar learners were expected to base their extrapolation on the most similar exemplar and display fairly flat extrapolation.

The second goal of Experiment 1 was to determine whether these learning strategies could be used to predict differences in performance on other materials; specifically on a test of analogical reasoning and on an assessment of conceptual learning. Differences in analogical transfer might be governed by the ability to abstract an underlying schema across items (Gick & Holyoak, 1983). On the function-learning task, rule-based learners are characterized specifically by abstracting the underlying function when given multiple examples. Therefore, they might also be better able to abstract underlying similarities between analogical transfer items, leading to better performance on novel problems requiring the same solution. Exemplar learners, conversely, may focus on the surface features of each problem, committing it to memory. Therefore, on the novel problem, they may not recognize the shared underlying schema and fail to transfer.

Similarly, we might also see differences in learners on an assessment of conceptual learning. When subjects are asked to read multiple passages, rule learners might make connections among the passages (or abstract underlying schemas) while exemplar learners work to memorize each passage as a separate exemplar. Therefore, when answering questions that are factual in nature, exemplar learners may outperform rule learners, but when answering questions that require subjects to make connections between the passages, rule learners may outperform exemplar learners. In addition to these types of questions, subjects were also asked inferential questions, which required them to apply their knowledge of passage material to a new domain. If rule-learning tendency represents a general tendency to make connections between material, rule learners may outperform exemplar learners on these items, as they are more aware of the need to connect what they learned from the passages to novel information. Exemplar learners may also perform well on these items because they are aware that this is a transfer context and are able to recall the critical information needed to transfer.

The third and final goal of Experiment 1 was to explore any potential demographic differences between the two types of learners. While no differences were expected on race or sex, other differences were possible. Because the current sample was drawn from a population of college students, it is possible that age could differ between learning types. Specifically, early in their college careers, students might find that an exemplar strategy is sufficient for classroom achievement, as most introductory classes require rote memorization of material. However, more advanced courses might require abstraction of underlying concepts. It is therefore possible that general learning strategy could shift over time. In addition to age, college major might differ between learning tendencies. If rule-based learning is associated with mathematical ability, rule learners might perform better in natural science courses and therefore choose majors in those disciplines. Alternatively, individuals in natural science courses may learn that rule-based learning is beneficial and adopt this strategy more generally.

There were other measures used in Experiment 1 to assess whether they might be related to the tendency to rely on exemplars versus abstract underlying rules. The Kolb Learning Styles Inventory (Kolb LSI; Kolb, 1993) was used to classify learners along two dimensions: *taking in experience* and *dealing with experience*. Taking in experience is described by the extent to which someone relies on concrete experience or abstract conceptualization, which could be related to the rule-learning tendency. That is, if rule learners are those individuals who tend to abstract underlying information, they may score high on abstract conceptualization. Exemplar learners may shy away from abstraction and therefore rate high on concrete experience and low on abstract conceptualization. In addition to the Kolb LSI, fluid intelligence was measured using Raven's Advanced Progressive Matrices (RAPM: Raven, Raven, & Court, 1998). This task requires subjects to view a visual display and determine the rule that governs the relationship between items in order to select a stimulus that correctly completes the display. Therefore, RAPM might predict those individuals who are able to abstract a rule and those that are not. However, the difference between rule and exemplar learners is not necessarily described by the *ability* to abstract a rule, but rather the tendency to rely on exemplar or rule-based processing in a learning situation that allows for either type of processing to be successful. That is, even if one is *able* to abstract a rule, they may instead choose to adopt an exemplar strategy. It is therefore more likely that rule and exemplar learners will not show differences in RAPM, but that, within the rule learners, RAPM will be correlated with their ability to learn the rule as assessed by their rate of learning. Indeed, this pattern was found by McDaniel et al. (2012) and would represent a replication of those data.

Experiment 1 consisted of two sessions to accommodate all of the materials. During session 1, participants completed a demographics questionnaire, the function-learning task, the analogical reasoning convergence problems, and an abbreviated version of RAPM. During session 2, participants read twelve passages, completed the Kolb LSI, and took a 30-item quiz over the passages. Rule learners were expected to perform better than exemplar learners on the analogical reasoning problems and on the inferential and connecting questions on the quiz and to show a relationship with RAPM. Exemplar learners were expected to perform better on the factual questions on the quiz and show no relationship with RAPM.

Method

Participants. Eighty-six participants were recruited from the Department of Psychology human subject pool and received either credit towards completion of a research participation requirement or cash payment ($5 for each half hour of participation). Four participants did not return for the second session of the experiment or did not complete participation due to time constraints. Three participants failed to comply with instructions during the function-learning task and were therefore excluded from analyses. In addition, seven participants did not demonstrate adequate learning of the input-output pairs during training (mean absolute error on the final training block >10) and were excluded from further analysis in accordance with the methods used by McDaniel et al. (2012). Therefore, the final sample consisted of 72 participants.

Procedure. Participants were tested in small groups in two sessions, two days apart. The procedures for the two sessions are depicted in Figure 3.

Session 1. Participants first completed a demographics questionnaire that assessed age, sex, race, grade point average, SAT/ACT scores, and college major. Participants then completed the function-learning task used by McDaniel et al. (2012). During this task, subjects were trained on a set of input-output pairs that made up a continuous function. The function was V-shaped with the vertex at 100, but this function was never made explicit to the subjects. For input values less than 100, the function followed the equation $f(x) = 230 – 2.2x$; for input values greater than 100 the function followed the equation $f(x) = 2.2x – 210$. Subjects were given a total of 200 training trials composed of 10 blocks of 20 randomly ordered numbers. Within each block, the input values were all odd numbered integers between 80 and 120.

Participants were told that they would be working for NASA examining data printouts about a newly discovered organism on Mars. They were given the task of determining how much of a particular element ("Beros") the organism released based on the amount of another element ("Zebon") that was absorbed. On each trial, participants were shown three bars (see Figure 4): a bar that displayed the input value ("Zebon Absorbed"), a bar that displayed each participant's predicted value ("Your Prediction"), and a bar that displayed the correct answer ("Beros Released"), which served as feedback for the participants. Subjects were given unlimited time to respond by using the arrow keys and were given immediate feedback. Feedback consisted of the output value displayed on the "Beros Released" bar and a sentence stating, "Your prediction was __ units off." Feedback appeared on the screen for 4 seconds before the computer automatically moved on to the next trial. In addition, subjects received feedback at the end of each block giving their mean accuracy for the given block.

At the end of training, subjects immediately began the test. The test was composed of 60 trials: 20 repeated training trials, 20 trials that were within the range of training trials but were not previously seen (even integers, termed *interpolation* trials), and 20 trials that extend beyond the range of the training trials (10 odd integers above and 10 below the training range, termed *extrapolation* trials), presented in a single random order to all subjects. The visual display was identical to that presented during training (see Figure 4), but no feedback was provided.

After the function-learning test, participants read the first two convergence problems (taken from Gick & Holyoak, 1983). Subjects were told that they were being tested on their reading comprehension and to simply read each story carefully because they would be asked questions about it later. The order of the two stories was randomly assigned across participants. Participants were given three minutes to study each story (see Appendix A) and were then asked to summarize the story. The second story was presented immediately after the first. Each story was removed during summarization and participants were not told that there was any connection between the two stories.

After reading the two stories, participants completed an abbreviated version of RAPM. On each trial, participants were shown eight boxes arranged into a 3 x 3 grid with the bottom right block missing. Subjects were instructed that they were to choose, from eight different options, the block that would complete the pattern both vertically and horizontally. Subjects were given a total of 12 trials, consistent with the short form version of the RAPM (Bors & Stokes, 1998; Set II), and were given unlimited time to complete each item.

After completing RAPM, participants were given the criterial convergence problem (see Appendix A) and told that there were many possible correct answers, so they should write down as many answers as possible in the allotted time. Participants were then given three minutes to solve the problem, consistent with the methods used by Gick and Holyoak (1983). However, participants were not given any hints to use the previous stories to help solve the current problem. After the three minutes, participants were dismissed and reminded to return for session 2 two days later.

Session 2*.* During session 2, participants first read 12 passages on a variety of topics. Six of the passages were taken directly from Butler (2010), while the other six were created to be similar in content and length. The six new passages were each paired with one of the passages from Butler (2010), such that there was an overlapping piece of information that connected the two passages. For example, the “bread” passage stated, “Whereas yeast takes two to three hours to produce its leavening action, <u>a dry chemical leavening agent like baking powder is instantaneous</u>.” The connected “volcanoes” passage stated, “A common mistake in making a model volcano is using baking powder instead of baking soda. Baking powder does not react with vinegar as quickly as pure baking soda, and <u>baking powder can also start reacting on its own because it contains the acid and base needed for the production of the carbon dioxide</u>.” All passages were developed from three online sources (www.en.wikipedia.org, www.encyclopedia.com, and www.howstuffworks.com). Each passage was approximately 500 words in length, separated into four paragraphs (see Appendix B).

Before each passage, participants were presented with the title of the passage and asked to press the space bar when they were ready to begin. The passages were each displayed on the screen for three minutes and were presented in a single random order to all participants. After each passage participants were presented with the title of the next passage, allowing them time for a short break if it was needed. After the final passage, subjects were asked how many of the passages they finished in the allotted time on a scale from 1 (all of the passages) to 4 (only a few of the passages).

After reading all the passages, subjects completed the Kolb LSI, which has been shown to have moderate to high test-retest reliability (r = .90, Veres, Sims, & Locklear, 1991; r = .54, Ruble & Stout, 1991). The inventory consists of 12 sentence stems (e.g. “I learn best when:”) followed by four response options which participants are told to rank from 1 (*least* like you) to 4 (*most* like you). Each response option is associated with one of four “learning modes” (concrete experience, reflective observation, abstract conceptualization, and active experimentation). The rankings associated with each learning mode are summed and the scores are used to compute two dimensions of learning. The dimension of *taking in experience* is calculated by subtracting the concrete experience score from the abstract conceptualization score and the dimension of *dealing with experience* is calculated by subtracting the reflective observation score from the active experimentation score.

The final task was to complete the test over the passages. For the six passages used by Butler (2010), two *fact* questions per passage were taken directly from the Butler (2010) materials (called “conceptual” questions by Butler). The answers to the fact questions could be answered with information that was explicitly stated in the passage and were therefore assumed to rely on surface details only (see Appendix C). Two *inferential* questions per passage were adapted from the “inferential conceptual (different domain)” questions developed by Butler (2010). These questions were altered such that no hint was given as to which passage should be used to answer the question. For example, a question from Butler read, “The U.S. Military is looking at bat wings for inspiration in developing a new type of aircraft. How would this new type of aircraft differ from traditional aircrafts like fighter jets?” The question was revised to say, “The U.S. Military is looking for inspiration in developing a new type of aircraft that promotes increased maneuverability. How would this new type of aircraft differ from traditional aircrafts like fighter jets?” Finally, for the six remaining passages, one question was created per passage that required information from one of the six original passages created by Butler (2010), but also information from one of the newly created passages in order to correctly answer (*connecting* questions; see example about “baking soda” provided above). The final test therefore consisted of 12 fact items, 12 inferential items, and 6 connecting items, for a total of 30 items, which were randomly ordered. The test was cued recall and participants were asked to answer every question even if they had to guess in order to maintain a constant response criterion across participants and avoid any floor effects. No feedback was provided. After completing the test, participants were debriefed about both sessions and dismissed.

Results

Function learning classifications. Mean absolute errors (MAE) were calculated for each participant for first and last training blocks, interpolation trials, and extrapolation trials. As indicated above, participants whose MAE > 10 on the last training block (N = 7) were considered non-learners and were excluded from further analyses. Extrapolation MAE was then used to classify the remaining participants as either *rule learners* or *exemplar learners.* In this particular task, flat extrapolation would produce an MAE of 34.72 (indicative of an exemplar model; see DeLosh et al., 1997). If participants are using rule-based information, their MAE should be significantly less than 34.72 because they should deviate from flat extrapolation in favor of the function. Therefore, 95% confidence intervals were computed for each participant's extrapolation MAE and those participants with confidence intervals that fell entirely below 34.72 were classified as *rule learners* with the remainder of individuals classified as *exemplar learners* (with five exceptions, described below). As seen in Figure 5a, rule learners showed extrapolation patterns that closely follow the underlying function, while exemplar learners did not appear to extrapolate their learning with any clear pattern that would be predicted given the training values. These patterns are consistent with models that incorporate exemplar and rule learning as separate mechanisms, (e.g. DeLosh et al., 1997), such that exemplar learners performed in a manner consistent with exemplar (associative learning) models and rule learners performed in a manner consistent with rule-based models.

Five individuals demonstrated extrapolation patterns that followed an oscillating pattern instead of the V-shaped function (see Figure 5a). The MAE for these individuals was above the 34.72 criterion, which would classify them as exemplar learners, but because an oscillating (sine-like) function is a reasonable abstraction from the training points (i.e. a possible abstracted function), these individuals were considered rule learners (see Bott & Heit, 2004). Each of their MAEs was calculated for a sine function and their 95% confidence intervals were compared to a criterion MAE of 24.09 (flat extrapolation with respect to the sine-like function). The confidence intervals for all five subjects fell entirely below 24.09 and all five were therefore classified as rule learners. After classifying all learners, the final sample included 34 exemplar learners and 37 rule learners.

Mean absolute errors for each training block are displayed in Figure 6. By block 3, rule learners on average had reached criterion and then steeply dropped off, maintaining very low error. Exemplar learners, instead, did not reach criterion until block 5 and then gradually reduced error through block 10, F (9, 612) = 2.373, p < .05 for the interaction. These data indicate that rule learners learned the rule and then displayed low error, while exemplar learners had a slower rate of learning as they learned each of the points. In addition, rule learners (M = 6.21) showed lower error overall than exemplar learners (M = 9.09), F (1, 68) = 18.67, MSE = 77.42, p < .001, but this difference was driven by the lower error on blocks 4 through 10, all F's (1, 69) > 13.55, all p's < .001. When reducing the analysis to only the first and last blocks, rule learners (M = 10.24) again had lower MAE overall than exemplar learners (M = 12.37), F (1, 69) = 11.58, MSE = 13.82, p = .001. However, the interaction term was non-significant, F (1, 69) = 2.24, MSE = 12.60, indicating that, although rule learners performed better than exemplar learners overall and the learning rates were different, by the end of training both groups had learned the items equivalently.

As described by McDaniel et al. (2012), it could be argued that exemplar learners simply become confused when seeing new items and this would be reflected in poor performance on all novel items, both interpolation and extrapolation. While rule learners (M = 2.55) did have lower MAE than exemplar learners (M = 5.01) on interpolation trials, F (1, 70) = 24.98, MSE = 4.29, p < .001, both groups made predictions that closely followed the function for these points (see Figure 5a, lower panel). A 2 (learner type) x 2 (trial type: interpolation vs. extrapolation) mixed ANOVA showed that the difference between the two learner types on extrapolation (M_{diff} = 23.18) was significantly larger than the difference on interpolation (M_{diff} = 2.46), F (1, 69) = 29.37, MSE = 129.49, p < .001, for the interaction. Therefore the two groups differed primarily on their extrapolation MAE, showing very similar performance on training and on interpolation trials, which is consistent with models demonstrating that both exemplar and rule models perform well on interpolation but differ on extrapolation (DeLosh et al., 1997).

Conceptual passages and test. After reading all of the passages, participants indicated how many passages they were able to read. Sixty-two participants (88.6%) reported that they were able to read all or most of the passages in the allotted time. The number of passages read (all, most, some, only a few) was equally distributed across learner types, χ^2 (3, N = 71) = 2.84. The subsequent analyses were analyzed after excluding the participants who indicated that they were unable to read all or most of the passages (N = 8) and the pattern of data remained the same. Therefore, results for the full sample are reported below.

The test data were graded according to the grading criteria in Appendix C. Each item was scored in three ways—strict scoring, lenient scoring, and a score for whether the correct passage was used to answer the question. A 2 (learner type) x 3 (question type: factual, inferential, connecting) mixed ANOVA was conducted on each type of scoring. For the strict scoring, the interaction was not significant ($F < 1$), but there was a significant effect of question type ($F(2, 138) = 41.34$, $MSE = .022$, $p < .001$), such that individuals scored highest on factual questions ($M = .58$) and lowest on connecting questions ($M = .35$), with inferential questions in between ($M = .47$; see Figure 7a). In addition, collapsing across question type, rule learners ($M = .51$) performed better on the test overall than exemplar learners ($M = .42$), $F(1, 69) = 4.58$, $MSE = .092$, $p < .05$. When reducing the analysis to only the factual and connecting questions (where the interaction was most expected), the interaction was still not significant, $F < 1$. The lenient scoring produced the same pattern of results, with the main effect of learner type dropping to marginal significance, $F = 3.35$, $MSE = .054$, $p = .07$.

When analyzing the correct passage scoring, factual items were not considered because all answers were associated with the correct passage. Therefore, for correct passage scoring, a 2 (learner type) x 2 (question type: inferential, connecting) mixed ANOVA produced a significant effect of question type, $F(1, 69) = 140.35$, $MSE = .02$, $p < .001$, such that participants used the correct passage more often on inferential items ($M = .56$) than on connecting items (M = .29), where they were required to use information from the two connected passages to get credit (see Figure 7b). However, there was no effect of learner type, F (1, 69) = 1.09, MSE = .04, and no interaction, F < 1.

Analogical reasoning. Each participant was given a binary score of "1" if they mentioned the convergence solution in their response to the criterial problem and a score of "0" if they did not. Rule learners (M = .57) did not significantly differ from exemplar learners (M = .56) on their use of the convergence solution, F < 1. However, within the rule learners, there was a significant correlation between RAPM and use of the convergence solution (r = .35, p < .05), but not within the exemplar learners (r = .06).

Other measures. Rule learners (M = .60) did not significantly differ from exemplar learners (M = .57) on RAPM, F < 1. In addition, within the exemplar learners, RAPM was not significantly correlated with rate of learning (r = .10) as defined by the training block in which the participant reached a learning criterion of MAE < 10 (lower block number associated with faster rate of learning). There was, however, a correlation that trended toward significance between RAPM and rate of learning for the rule learners (r = -.27, p = .11), such that rule learners who took fewer trials to reach learning criterion scored higher on RAPM. There was also no correlation between RAPM and any of the test scores (all r*'s* < .19). Learner type did not predict any of the Kolb LSI learning modes, or the two dimensions of learning (all F < 1).

Demographics. As seen in Table 1, the two groups did not significantly differ in age, (F (1, 70) = 1.011, MSE = 4.54), grade point average (F < 1), or

ACT scores ($F < 1$; because some students reported only SAT scores, SAT scores were converted to ACT scores using the ACT/College Board concordance tables, 2008, in order to standardize the data). Rule and exemplar learners also did not differ on sex (χ^2 (1, N = 71) = .006), or race (χ^2 (3, N = 69) = 2.77). In order to examine academic major, participants were divided into those who indicated *STEM* (science, technology, engineering, or mathematics) majors, *non-STEM* majors, or both (i.e. double-majors, one STEM and one non-STEM). Individuals who were undecided (N = 5) were removed from this analysis. The two groups did not significantly differ in academic major when STEM, non-STEM, and double majors were included, χ^2 (3, N = 71) = 2.55, nor when double majors were removed, χ^2 (1, N = 50) = .1.52.

Function learning classifications revisited. One potential issue with the above analyses lies in the manner in which the function learning classification was conducted. Specifically, the confidence interval approach to classification selects participants on the basis of how closely they mirror the underlying V-shaped function and compares these individuals to everyone else. However, it is possible that other individuals may abstract part of a rule and need more training to abstract the entire rule, or that individuals are using a combination of exemplar and rule-based processing. It is also possible that participants could use a rule plus exception model (Nosofsky et al., 1994) such that they might abstract a generally positive linear trend and only items in the 80-100 training range would be exceptions and therefore memorized points. These types of learning are more difficult to distinguish with an MAE or confidence interval approach but ignoring these possibilities could dampen effects that we see on other tasks. In order to address this issue, a classification guide was used (see Appendix D) to distinguish rule patterns from exemplar patterns of extrapolation, removing ambiguous patterns of extrapolation from analyses (used by Fadler et al., 2012). This type of classification can therefore be considered an extreme groups classification, such that only the clearest rule learners and the clearest exemplar learners are included in the analysis.

Using the classification guide, two independent raters divided the learners into three groups: rule, exemplar, or ambiguous. The two raters agreed on 87.8% of the subjects' classifications and all discrepancies were resolved before proceeding. Using this approach participants were divided into 38 rule learners and 13 exemplar learners. The patterns of results were similar to those described above, but there were no significant effects, presumably due to the few exemplar learners in the sample.

Discussion

Rule learners and exemplar learners were successfully identified in a manner that replicated McDaniel et al. (2012). While rule learners had lower MAE overall, both groups learned equivalently by the end of training. However, rule learners learned faster, indicating a qualitatively different strategy for learning. Importantly, the two types of learners also showed similar interpolation profiles and diverged specifically on extrapolation trials, consistent with formal models of each type of processing (DeLosh et al., 1997). In addition, within the rule learners, RAPM trended toward association with rate of learning, indicating that abstraction ability could be related to speed of abstracting the rule. However, no such correlation existed within the exemplar learners, indicating that exemplar learners might be using a qualitatively different strategy, unrelated to abstraction. These data were again similar to McDaniel et al. where a significant correlation between RAPM and rate of learning was found within the rule learners but not the exemplar learners.

There was also evidence in Experiment 1 of stability of function learning tendency across materials. While there were no differences found on the analogical reasoning problems, this could be because both rule and exemplar processing could support success on these items. That is, rule learners could have drawn on abstraction of the underlying schema to promote success. Gick and Holyoak (1983) demonstrated that using two stories with differing surface features caused subjects to be significantly more likely to abstract the underlying schema and therefore abstract at similar levels as in the current study (M = .52). Exemplar learners, on the other hand, might have tried to recall the most similar problem they encountered in the current experiment. Because there was a short (approximately 10 minute) delay, exemplar learners may have been able to easily recall one of the stories they previously read and be more successful with mapping the appropriate features. There is some evidence for this conclusion as, within the rule learners, RAPM was significantly correlated with performance on the criterial convergence problem, but no correlation was found within the exemplar learners. As above, this indicates that the *ability* to abstract was associated with successful transfer in the rule learners, but that the exemplar learners were using a qualitatively different strategy to promote transfer.

While the analogical reasoning task did not produce direct differences between groups, there were distinct differences on the assessment of conceptual learning. Rule learners outperformed exemplar learners on all question types, contrary to the hypothesis that exemplar learners should perform better on factual items. However, these items were still conceptual in nature and utilized information that was needed for comprehension of the passage. It is therefore possible that rule learners were making connections with prior knowledge in order to understand factual information, which in turn strengthened their memory for factual material and allowed them to perform well on these items.

Finally, it appeared that learning tendency represented a unique type of assessment that could not be explained by a general ability. That is, there were no differences between rule and exemplar learners on RAPM or ACT scores. There was also no association between learning tendency and scores on the Kolb LSI or on more exploratory demographic items (sex, race) and these differences do not seem to change over time (as indicated by age) or mandate academic major selection.

Experiment 2

While rule learners outperformed exemplar learners on all question types on the assessment of conceptual learning in Experiment 1, the expected interaction (learner type by question type) did not emerge. This effect may have occurred because rule learners were making connections between factual material and prior knowledge. Converging evidence for the tendency for rule learners to make connections comes from the inferential items in which rule learners outperform exemplar learners. Therefore, the first goal of Experiment 2 was to replace the factual items with questions that asked about unnecessary details in each passage. An example of one of these questions was, “What is the name of the thin membrane of skin found on a bat’s wing?” (Answer: *Patagium.*) These items involved examples of conceptual information that did not affect comprehension of the passage as a whole. Briefly, Kintsch (1988) explained that text is processed at three distinct levels*: surface, propositional,* and *situational levels.* The surface level involves the verbatim words and linguistic structure of a given sentence. Sentences are then converted into propositions, which contain the meaning of the text. The situational level then contains the overarching context of the text—the circumstances directly related to the information described in the text.

On the function-learning task, exemplar learners use a strategy in which they memorize the input-output pairs. If that type of strategy extends to text processing, exemplar learners should focus on memorization of the surface features of the text and might not create many links between propositions and to situation models. More specifically, exemplar learners do not abstract the relational information in function learning. In text processing, they also may not relate the propositions with each other or with, for example, prior knowledge. Rule learners, on the other hand, approach the function-learning task by trying to abstract an underlying rule. If they approach text processing similarly, they might develop considerable links between propositions and particularly with situation models to guide understanding. Rule learners might therefore abandon the surface level features in favor of the gist of the text. The two types of learners might therefore process the same text in different ways and these processes might lend themselves differentially to the three question types.

The fact items from Experiment 1 were taken from sentences that were strongly linked with the rest of the passage in meaning and comprehension and would have been easily integrated and connected to existing situation models. Rule learners may have more strongly encoded this information than did exemplar learners who may have focused on the verbatim text instead of the meaning and associations. The new *example* questions used in Experiment 2 contained information that was somewhat irrelevant for meaning and comprehension. These items would still have been converted into propositions but would have been less easily integrated with the propositions from the rest of the passage and certainly less integrated with existing situation models. If rule learners are focused on information that *can* be readily integrated into situation models, they may have paid considerably less attention to example information. Exemplar learners may not necessarily discriminate between areas of more of less import for text comprehension, instead attempting to commit the surface level features to memory. If this is the case, exemplar learners should be better able to recall these verbatim examples than are rule learners. It was therefore predicted that on the new surface-level example questions, exemplar learners should outperform rule learners, having retained more of the surface-level features, while rule learners should remain superior on inferential and connecting questions.

The second goal of Experiment 2 was to replicate and extend the findings in McDaniel et al. (2012). In Study 1b, McDaniel et al. trained subjects to categorize novel animals and then tested them on repeated training items and novel animals. Performance at the end of training was not perfect (*M* = 74%), but when isolating effects to perfect learners, exemplar learners performed in a manner consistent with an exemplar strategy. In the present study, a simpler rule was adopted. McDaniel et al. used an additive rule, in which any combination of two out of three critical features had to be present to be considered a *builder*, while in the current study a conjunctive rule was used, such that big animals with spots were considered *builders* and all others were considered *diggers*. In addition, the number of training blocks was increased in order to encourage greater learning by the end of training.

In addition, the testing procedure was changed for the current study to reflect the differences in learning strategy. Specifically, during training, once rule learners have adopted a rule, they may be less likely to pay attention to the other features of each animal, as those features are irrelevant to the current task. However, if exemplar learners are memorizing the animals, they may be more likely to notice changes in idiosyncratic features that are irrelevant to the rule because they may be using all of the features in order to memorize the exemplar as a unique item. These differences would be clear on categorization of novel animals (as seen in McDaniel et al), but also in recognition, such that exemplar learners would be more likely to discriminate between training and transfer items. If such an interaction emerged, it would indicate that the neither type of learning is inherently “better” than the other, but rather that the success of a given processing strategy depends on the goals of the current task.

In the current study, two different types of lures were created for the testing phase, which reflect differences in processing (see Appendix E for examples of each type of lure). Each lure was paired with a training item, such that the two were the same on most idiosyncratic features. Recognition lures (called “Good Transfer”’ by McDaniel et al.) were then created by making small changes to features that were not critical for learning the rule. These items may therefore be most difficult during recognition and disproportionately so for rule learners who may have disregarded the features that were not critical for the rule. Categorization lures (calls “Bad Transfer” by McDaniel et al.) were instead created by making changes to one of the critical features and therefore changing the category of the item. If exemplar learners are using their closest memorized exemplar in order to categorize, they should perform poorly on these items. However, if rule learners have abstracted the rule, they should be able to categorize these items comparably to all other items. The current study therefore extended the concept-learning paradigm used by McDaniel et al. by increasing the likelihood of learning the rule and by adding a recognition component in order to explore potential benefits of exemplar learning.

McDaniel et al. also used an abstract coherent categories task, which was implemented in the current study. Unfortunately, due to a programming error, there were only four unique items used during the training portion of this task and it does not stand as a direct replication of McDaniel et al. (who used eight training items). This method may have encouraged an exemplar strategy by both types of learners because of the relatively small amount of effort required to memorize so few items and the difficulty in abstracting a general rule with so little variability.

As in Experiment 1, demographic information was collected to explore the extent to which individual differences might correlate. Ravens Advanced Progressive Matrices was collected to replicate the findings of Experiment 1. However, as there were few differences between learners on analogical reasoning or the Kolb LSI, these measures were dropped from Experiment 2. Finally, there were enough subjects to allow comparisons using the extreme groups approach described above, such that individuals could be classified based on their extrapolation patterns and individuals who did not show a clear trend toward either rule- or exemplar-based processing could be removed. It was expected that rule learners would outperform exemplar learners on categorization in the concept-learning and abstract coherent categories tasks and on the inferential and connecting questions over the passages. It was expected that exemplar learners would outperform rule learners on recognition in the concept-learning task and on the new example questions over the passage.

Method

Participants. Seventy-six participants were recruited from the Department of Psychology human subject pool and received either credit toward completion of a research participation requirement or cash payment ($5 for each half hour of participation). Four participants did not return for the second session of the experiment or did not complete participation due to time constraints. In addition, nine participants did not demonstrate adequate learning of the input-output pairs during training (mean absolute error on the final training block >10) and were excluded from further analysis. Therefore, the final complete sample consisted of 62 participants. However, due to time constraints and technical issues, the Regehr and Brooks task data were missing from 3 participants and the abstract concept categorization data were missing from 1 participant. In addition, participants (N = 14) who did not show adequate learning of the training stimuli in the abstract coherent categories task (final training block accuracy ≤ 75%) were excluded from those analyses.

Procedure. Participants were tested in small groups in two sessions, two days apart. During session 1, participants completed a demographics questionnaire, the function-learning task, an abbreviated version of RAPM, and the concept-learning task (Regehr & Brooks, 1993). During session 2, participants read the twelve passages, completed the abstract coherent categories task, and took a 30-item test over the passages. This procedure is depicted in Figure 3.

Session 1. Participants first completed the demographics questionnaire and the function-learning task, which was identical in type and procedure as in

Experiment 1. Participants then completed the abbreviated version of RAPM, which was again identical to that of Experiment 1. After RAPM, participants completed a modified version of the Regehr and Brooks' (1993) concept learning task. In this task, participants were shown images of fictitious animals that varied on six binary dimensions: body shape (angular or round), leg length (short or long), number of legs (two or six), neck (short or long), spots (spots or no spots), and animal size (big or small). Each animal was classified as either a *digger* or a *builder*, and group membership was determined using a conjunctive rule, such that animals that were big with spots were classified as *builders* and all other animals were classified as *diggers*. Each image was shown on a background such that image size could be judged with reference to the background image and therefore made salient (see Appendix E).

Each training animal had a unique form across the six primary features described above. For example, although some animals had six legs and others had two, the shape of the legs varied across each individual animal. In this way, perceptual distinctiveness was maximized, while still allowing a rule to govern classification. During training, each image was presented until the participant classified the animal as either a builder or digger by pressing designated keys. Participants were not presented with the rule, but were simply instructed to classify the animals into the appropriate category and to do so as quickly and accurately as possible. After the participant made a response, they were told whether they were correct or incorrect. There were a total of 4 animals in each category for a total of 8 stimuli, which were presented in random order within each of 10 blocks, for a total of 80 trials. There were two versions created such that animals that represented builders in one version were changed to diggers in the second version. The two versions were then counterbalanced across participants.

After training, participants completed a test, which consisted of three different trial types. Eight of the items were repeated from the training portion, termed Repeated Training Items. Four items were created by changing idiosyncratic features of the original image that were irrelevant for the rule (eye shape, toe shape, and leg length). These items would be incorrectly identified as old items if individuals were only paying attention to the critical features for the rule (spots and size). These items were therefore termed Recognition Lures, as they should be most difficult during recognition. An additional four items were created by changing critical features associated with category membership (e.g. spots changed to no spots). These items were similar to trained items and would be incorrectly categorized if subjects did not correctly apply the rule (i.e. if they categorized the item in the same way as its most similar trained item). Therefore these items were termed Categorization Lures. The final test therefore consisted of 8 Repeated Training Items, 4 Recognition Lures, and 4 Categorization Lures presented in random order.

During the test, each image was presented on the screen and subjects were required to first make a button press to indicate if the item was old (had appeared in training) or new. While the image was still on the screen, subjects made a second button press, classifying the animal as either a *builder* or *digger*.

Each image remained on the screen until the subject had made the recognition judgment and categorization judgment. After completing the test, participants were reminded about the second session and dismissed.

Session 2. During session 2, participants first read the same 12 passages in the same order as in Experiment 1, but were asked after each passage if they were able to complete their reading of the passage (Yes or No). After reading all the passages, participants completed the abstract coherent categories task to replicate McDaniel et al. (2012). The task was similar in nature to the classification condition used by Erickson et al. (2005, Experiment 3). Participants were presented with a list of four attributes that described a machine and told that they would be classifying the machine as either a *morkel* or *krenshaw*. The four features represented 1) where the machine operated, 2) the action it was used for, 3) what instrument it used, and 4) its means of locomotion, and the four features were presented in this order for all trials. While participants were never explicitly told the rule, *morkels* were comprised of two sets of coherent features that, when combined, formed four features that were also coherent (i.e. made sense together). *Krenshaws* were also composed of two sets of coherent features, but, when combined, no longer yielded a plausible machine (e.g. features 1 and 3 were coherent and 2 and 4 were coherent, but 2 and 3 could not be plausibly combined; see Table 2). Participants were told that two machines of the same type could have different features and that machines of differing types could share some features. Participants had unlimited time to classify each list of features by pressing an associated key and were then given feedback that stated, “Yes, that was a ____.” or “No, that was a ____.” The list of features stayed on the screen during feedback and subjects were given an unlimited amount of time to process feedback before moving on. Four machines (two morkels and two krenshaws) were presented two times per block and randomly ordered within each of 8 blocks, for a total of 64 training trials.

After training, subjects were told that they would see two features of the machines they had just classified and would need to categorize each set of features as belonging to either a morkel or krenshaw. Participants were shown every combination of features except for those that were always presented together (i.e. features 1 and 3 and features 2 and 4) as consistent with the two-feature test used by Erickson et al. (2005) for a total of 16 randomly ordered trials. After each classification trial, participants were asked to rate their confidence in their classification on a scale of 1 (least confident, just guessing) to 7 (certain) and did not receive any feedback.

Participants were then given a novel classification test. Participants were told that they would see new features of machines and would need to classify each set of features as belonging to morkels or krenshaws. The features represented the same type of features (e.g. where it operates) in the same order as training, but with novel features. For this test, morkels were again completely coherent and plausible machines. Krenshaws, instead, were machines in which the location of operation was coherent with instrument used, but not with location and locomotion (see Table 2). There were a total of 12 randomly ordered trials

(6 morkels, 6 krenshaws). As in the two-feature test, there was no feedback and subjects rated their confidence on a scale from 1 to 7.

After completing the abstract coherent categories task, subjects took a test over the passages, which was identical to that of Experiment 1 with the following exceptions. The Fact questions were removed and replaced by a single detail-oriented question (termed *Example* questions) for each of the 12 passages. Therefore the test consisted of 12 Example questions, 12 Inferential questions, and 6 Connecting questions, for a total of 30 randomly ordered items. The procedure for presenting these items was identical to that of Experiment 1. When subjects completed the test, they were debriefed about both sessions and dismissed.

Results

Function learning classification. Subjects were classified as rule or exemplar learners according to the same confidence interval approach described in Experiment 1. There were 9 participants whose MAE > 10 on the last training block and were excluded from further analyses. After computing the MAE on extrapolation trials and classifying participants into groups, there were 33 rule learners (including 3 participants who displayed sine-like function learning) and 34 exemplar learners.

Mean absolute errors for each training block are displayed in Figure 6. A 2 (learner type) x 10 (training block) ANOVA showed that, while rule learners (M = 8.08) had lower MAE overall than exemplar learners (M = 6.38), F (1, 62) =

8.09, p = .006, the two groups learned at approximately the same rate, F (9, 585) = 1.10, MSE = 8.09, for the interaction. When reducing the analysis to only the first and last blocks, rule learners (M = 10.54) no longer displayed lower MAE overall than exemplar learners (M = 11.53), F (1, 65) = 2.09, MSE = 15.80. As in the above analysis, the interaction term was non-significant, F <1, indicating that, although rule learners performed nominally better than exemplar learners overall, the two groups learned at approximately the same rate and displayed equivalent amounts of learning by the end of training.

As in Experiment 1, rule learners (M = 3.02) had significantly lower MAE on interpolation trials than exemplar learners (M = 4.77), F (1, 66) = 6.21, MSE = 8.25, p < .05 but the difference was significantly larger on extrapolation trials (M =25.69), F (1, 65) = 40.36, MSE = 118.89, p < .001 for the interaction (see Figure 5b). Therefore, as in Experiment 1, the patterns are consistent with exemplar and rule-based models in that the groups show similar learning rates and interpolations patterns, but differ considerably on extrapolation (DeLosh et al., 1997).

Conceptual Passages and Test. More than half (56.5%) of the subjects indicated that they finished reading all of the passages and 95% of the subjects indicated that they read more than half of the passages. In addition, there was no difference between rule (M = 10.47) and exemplar learners (M = 11.07) on the number of passages they were able to read, F (1, 61) = 1.11, MSE = 4.96. Table 3 shows the number of participants who were unable to read each passage. In order to avoid any decrement in performance due to these effects, if a participant indicated that they were unable to finish reading a passage, it was removed from further analyses.

The test data were graded according to the same grading criteria as in Experiment 1 (see Appendix C). For the strict scoring, there was a significant effect of question type, F (2, 120) = 4.20, MSE = .04, p < .05, such that participants scored highest on inferential questions (M = .39), followed by connecting questions (M = .36), and lowest on example questions (M = .30). However there was no interaction, F (2, 120) = 1.21, MSE = .04, and no effect of learner type, F < 1 (see Figure 8a). The lenient scoring showed the same pattern of results and will not be further considered. Participants were more likely to mention information from the correct passage when answering the inferential questions (M = .55) than the connecting questions (M = .25), F (1, 60) = 81.35, MSE = .03, p < .001. However, there was again no interaction, F < 1, and no effect of learner type, F < 1 (see Figure 8b).

In order to determine if excluding the unread passages was responsible for changing the pattern of results from Experiment 1, the above analyses were conducted again, including all passages in the analysis. However, the pattern of results was the same as those described above (see Table 4 for means with and without unread passages).

Concept learning task. Individuals who correctly categorized less than 75% of the items during the final training block (3 exemplar and 3 rule learners) were removed from subsequent analyses. Rule (M = .94) and exemplar learners

(*M* = .95) did not differ on their accuracy on the last block of training, *F* < 1, and both groups showed very high accuracy by the end of training.

Recognition. As seen in Figure 9a, exemplar learners (*M* = .92) correctly recognized more of the repeated training items than did rule learners (*M* = .80), $F(1, 60) = 5.11$, *MSE* = .033, $p < .05$. However, when examining recognition of lures, both types of learners appeared to better recognize categorization lures than recognition lures (see Figure 9b). These impressions were confirmed by a 2 (learning type) x 2 (trial type: recognition lures, categorization lures) mixed ANOVA on the number of correct responses (correct rejections), which showed a main effect of trial type, $F(1, 59) = 33.84$, *MSE* = .02, $p < .001$, such that recognition lures were more prone to false alarms (64% reported as old) than categorization lures (40% reported as old). There was no interaction ($F < 1$) and no main effect of learner type ($F < 1$). To look at recognition performance more holistically, *d'* scores were calculated for each individual by taking the standardized proportion correct on repeated training items (hits) and subtracting the standardized proportion incorrect collapsed across the two types of lures (false alarms). As predicted, exemplar learners (*d'* = 2.03) were better able to discriminate between old and new items than rule learners (*d'* = 1.21), $F(1,60) = 5.43$, *MSE* = 1.93, $p < .05$.

Categorization. As seen in Figure 9a, there was no difference between rule (*M* = .88) and exemplar learners (*M* = .86) on categorization of repeated training items, $F < 1$. However, when examining categorization of lures, it appeared that there was no difference between learners on recognition lures, but that exemplar learners better categorized categorization lures (see Figure 9c). It also appeared that recognition lures were much better categorized than categorization lures. A 2 (learner type) x 2 (trial type) mixed ANOVA showed that the interaction was nonsignificant, F (1, 59) = 1.39, MSE = .056, but that recognition lures (M = .83) were indeed categorized significantly better than categorization lures (M = .47), F (1, 59) = 73.40, MSE = .056, p < .001. However there was no effect of learner type, F < 1.

Abstract Coherent Categories. As indicated above (and consistent with McDaniel et al., 2012) individuals who correctly classified fewer than 75% of items were excluded from subsequent analyses. Rule and exemplar learners showed similar rates of learning as demonstrated by a 2 (learner type) x 8 (training block) mixed ANOVA, which showed a non-significant interaction, F < 1. However, both types of learners performed significantly better by the end of training, F (7, 322) = 34.59, MSE = .02, p < .001 and there was no effect of learner type, F (1, 46) = 1.69, MSE = .12. Furthermore, by the end of training, both groups showed very high performance (M > .94 for both groups).

There was no significant difference between learners on the two feature test, F (1, 46) = 1.70, MSE =.03, but the difference on confidence-adjusted scores showed a nonsignificant trend, F (1, 46) = 2.38, MSE = 4.29, p = .13, indicating that rule learners were somewhat more confident in their accurate responses (M = 3.57) than exemplar learners (M = 2.65). The two-feature test was also broken down into two item types (see Figure 10). On half of the trials, there was a functional relationship between the items, such that knowing the rule would allow for correct categorization (e.g. "rolls on wheels" or "slides on skis" paired with "operates on water" or "operates on land"). However, for the other half of the trials, they could not be categorized based on the rule (e.g. "has a shovel" or "has a spongy material" paired with "rolls on wheels" or "slides on skis"). A 2 (learner type) x 2 (item type: rule-based or not) mixed ANOVA showed that there was a significant effect of item type, F (1, 46) = 6.98, MSE = .02, p < .05, such that participants were better able to categorize the rule-based items (M = .81) than the others (M = .73). However, there was no effect of learner type, F (1, 46) = 1.70, MSE = .07, and no interaction, F < 1. On the novel test, there was no difference between learners in overall categorization, F < 1, or on the confidence-adjusted scores, F (1, 46) = 1.34, MSE = 9.46. However, collapsing across learner type, scores on the novel test (M = .61, SD = .20) were significantly above chance, t (48) = 3.62, p < .01.

Demographics. The two groups did not significantly differ in age, grade point average, or ACT scores (all F <1). They also did not differ on sex (χ^2 (1, N = 67) = .022), or race (χ^2 (3, N = 66) = 2.50). Major was analyzed in the same manner as in Experiment 1. There was no significant effect of major when double majors were included (χ^2 (2, N = 56) = 1.05) or removed (χ^2 (1, N = 50) = 1.02) from the analysis (see Table 1).

Ravens Advanced Progressive Matrices. Rule learners (M = .65) did not differ from exemplar learners (M = .66) on RAPM, F < 1. Within the exemplar learners, there was no correlation (r = .15) between RAPM and rate of learning (as indicated by the number of blocks it took to reach criterion; lower number indicates faster learning) during the function learning task, but the correlation approached significance for rule learners (*r* = -.24, *p* = .17). As depicted in Figure 11, when the data were combined across Experiments 1 and 2, the correlation within the rule learners was significant (*r* = -.25, *p* < .05), but not within the exemplar learners (*r* = .06), and the difference between these correlations was marginally significant, *z* = -1.79, *p* = .07. RAPM was significantly correlated with much of the passage data, including correct scores on the inferential questions (*r* = .30, *p* < .05) and connecting questions (*r* = .32, *p* < .01) as well as use of the correct passage on the inferential questions (*r* - .33, *p* < .01) and connecting questions (*r* = .34, *p* < .01).

Function learning classifications revisited. One of the goals of Experiment 2 was to analyze the data using an extreme groups approach according to the guidelines laid out in Appendix D. Two independent raters classified subjects based on their pattern of extrapolation with 85% agreement. All discrepancies were resolved, resulting in 20 exemplar learners and 27 rule learners. All of the above analyses were conducted again and displayed comparable patterns of results with the following exception: rule learners (*M* = .87) categorized items significantly better than exemplar learners (*M* = .75), *F* (1, 33) = 5.36, *MSE* = .02, *p* < .05, on the two-feature test during the concept-learning task.

Discussion

The data from the function-learning task replicate that of Experiment 1. Subjects could be classified into two qualitatively different learning tendencies based on their extrapolation error rates. They showed comparable patterns of learning and interpolation, replicating formal models distinguishing these two types of learning (DeLosh et al., 1997). There was also no difference in RAPM performance between learning types, but within the rule learners, the correlation between RAPM and rate of learning trended toward significance (and reached significance when combined across experiments) but no such relationship existed within the exemplar learners. These data replicate Experiment 1 and indicate that abstraction ability may be related to learning for rule learners, but not exemplar learners, again showing partial evidence that these are two qualitatively different strategies.

While the example questions did reduce the benefit for rule learners as predicted, the strong benefit that rule learners showed on inferential and connecting questions in Experiment 1 failed to replicate in Experiment 2. However, other measures showed a relationship with function-learning tendency. Within the concept-learning task, exemplar learners showed a benefit on recognition of repeated training items as well as greater recognition sensitivity (as measured by *d'*). However, the benefits for exemplar learners on recognition did not extend to categorization. These data collectively indicate that exemplar learners were better able to recognize items they had previously seen, but for the categorization lures, they may have categorized based on the closest exemplar they saw during training. Unfortunately, rule learners also showed a decline on categorization and there are several potential explanations for this finding. Rule learners may have been using a unidimensional rule plus exception strategy (e.g. Learning the rule, “If it has spots it is a builder, and everything else is a digger” and memorizing the single exception to that rule), resulting in excellent training performance, but lower performance on categorization of categorization lures. Indeed, there are several unidimensional rules that would have predicted chance performance on categorization of categorization lures. It is also possible that some rule learners adopted an exemplar strategy during training because they were unable to learn the rule and therefore behaved as exemplar learners during transfer. The observed behavior for these two explanations would be similar and therefore they cannot be disentangled with the current data.

As described at the beginning of Experiment 2, the abstract coherent categories task contained a programming error in which only 4 items were repeated within each block of training. With so few exemplars, it was predicted that the task would be most easily accomplished with an exemplar strategy. Indeed, while training performance was high, there was no difference between learning types on the novel categorization test, indicating that rule learners had not abstracted the rule (or at least no better than exemplar learners). On the two-feature test, there was a benefit for both types of learners on items in which rule-based information could aid in categorization. While on other tasks, the appearance of partial rule-based extrapolation could be explained by prior knowledge, subjects in the current study had no prior knowledge about what constituted a “morkel” or “krenshaw” or that the critical information was the coherence of the machine. There are two other possible explanations for this pattern of data. It is possible that both rule and exemplar learners could have abstracted a small amount of rule-based information. In addition, while there was no difference between learners on the novel test, both types of learners performed above chance, which again indicates some understanding of rule-based information. This could be explained if individuals reached criterion through memorization (using an exemplar strategy) and then proceeded to try to find relationships between the items during the remainder of training. An alternative possibility is that there is a confounding variable within the stimuli such that incoherent information is simply easier to remember. Without subjects' conscious awareness that a particular machine was being classified on the basis of coherence, they might better recall the category in which an incoherent item belongs. However, that explanation could not account for above chance performance on the novel test items, which again indicates that both rule and exemplar learners seem to have abstracted some amount of rule-based information.

As in Experiment 1, there were no differences on any of the demographic characteristics. In addition, while RAPM was not correlated with function learning tendency, it was significantly correlated with several aspects of conceptual learning. While this is different than the pattern of results observed in Experiment 1, it seems that for Experiment 2, the processes underlying problem solving, concept learning, and abstraction are fundamentally different from those that govern the learning of prose passages.

General Discussion

Preliminary evidence exists that function learning tendency is a reliable and stable individual difference that predicts performance on a number of conceptual learning tasks as well as higher order cognitive tasks (McDaniel et al., 2012; Fadler et al., 2012). If these learning styles were shown to be stable across different materials, it could have important implications for education. Specifically, if rule learners are those that are better able to transfer to novel situations, educators may want to encourage this type of learning process. However, no such evidence that individual differences in function learning predict performance on classroom materials currently exists. The present research was therefore designed to determine the extent to which the differences in function learning tendency are associated with the type of conceptual learning often represented in classrooms.

Overview of Findings

There were several important results that emerged in the current study. Rule and exemplar learners were classified based on their extrapolation patterns, which differed considerably, while training and interpolation patterns were very similar across learners. These differences predicted patterns of behavior on analogical reasoning (as indicated by the significant correlation with RAPM in rule learners but not exemplar learners) in Experiment 1, while analogical reasoning performance was comparable. In Experiment 2, function-learning tendency predicted differences in recognition on the concept-learning task, but there were no differences between learners on performance on the abstract coherent categories task. On the conceptual material, rule learners outperformed exemplar learners on each question type in Experiment 1, but this effect did not replicate in Experiment 2. I will first discuss the results from the assessment of conceptual learning and potential explanations for the discrepant results across experiments and then discuss how the current research replicates and extends previous research on function learning tendencies. Finally, I will describe potential implications for education and limitations that can be addressed in future studies.

Application to Conceptual Learning

The primary goal of the current study was to determine the extent to which the above differences in learning tendencies have implications for classroom materials. For the purposes of this study, classroom materials were defined as prose passages and questions aimed at assessing different levels of knowledge. In Experiment 1, a main effect emerged such that rule learners performed better on every type of question on the conceptual test, despite one question type being designed to support exemplar processing. After replacing these questions with items that relied even more on surface-level details, the effect from Experiment 2 failed to replicate, such that there were no differences between learners on any question type. There are a few potential explanations for this discrepancy. One uninteresting explanation is simply that the effects from one of the experiments represents a Type I or Type II error, such that the results were not representative of true effects. Other explanations for the discrepancy may come from the differences between Experiments 1 and 2. These differences were the other tasks completed during the experimental session and the factual questions exchanged for example questions.

If the tasks completed after function learning in a given experiment had pushed participants to use a given strategy, they may have had an effect on the way the conceptual material was processed. The tasks that were unique to Experiment 1 included the analogical reasoning task and the Kolb LSI, with the analogical reasoning task occurring during session 1, but the Kolb LSI immediately preceding the conceptual test. Neither of these tasks seemed to encourage either rule or exemplar processing strategy to complete, although the Kolb may have heightened an individuals' awareness of their own processing preferences (although there was no relationship between function learning tendency and Kolb scores). In Experiment 2, the other tasks included the abstract coherent categories and concept-learning task, both of which can be completed using either a rule or exemplar strategy. The abstract coherent categories task was completed immediately before the conceptual test, but it appeared that all individuals used some amount of both rule and exemplar-based processing on this task. It is possible that the task immediately preceding the test could have led to some amount of priming which affected the way in which participants completed the conceptual test. In this case, because rule and exemplar learners performed similarly on abstract coherent categories, they may have therefore used similar processing on the subsequent test. If these secondary tasks indeed affected the processing used on the conceptual test, it would indicate that the function-learning strategies are extremely flexible, such that individuals might have a general preference, but will adopt a different strategy quite easily. It would also indicate that the function learning tendencies are not only an encoding process, but have distinct retrieval processes, as the secondary tasks would not have affected encoding (the first task encountered during the second session was the conceptual passages).

The final difference between the two experiments is the change from factual questions to example questions. As described above, the information in the factual questions was necessary for comprehension of the passage as a whole. In addition, some of the information in the factual questions was also used in the inferential or connecting questions (Experiment 1). However, this was never the case in the example questions (Experiment 2). It is possible that answering the factual questions may have led to greater activation of information needed to answer the inferential and connecting questions, leading to a potential testing effect (see Roediger & Karpicke, 2006a for review). Most studies demonstrating the testing effect compare a condition in which individuals restudy material to a condition in which individuals take a quiz. On a later test in which the same items are (re)presented, individuals perform better if they had been previously quizzed (e.g. McDaniel & Fisher, 1991; Roediger & Karpicke, 2006b; Carpenter, Pashler, & Cepeda, 2009). However, a few recent studies have demonstrated that retrieval practice can also have a facilitative effect on related information (e.g. Chan, McDermott, & Roediger, 2006; Butler, 2010).

In the current study, because the questions were randomized on the test, it is possible that answering factual questions may have had a facilitative effect on later related inferential or connecting questions. Similarly, if inferential and connecting questions produced increased activation of the underlying factual information, they too could have had a facilitative effect on later factual items. If rule learners are indeed more likely to make connections between information on just one of these item types, it could have produced facilitation on the others. If this explanation were true, it would indeed indicate that rule learners were processing information in a manner different from exemplar learners and in a way that produces better transfer to related information. In Experiment 2, the example questions may not have a facilitative effect on the inferential or conceptual items because these were not as closely related and therefore answering these items may not have produced the same boost in related information. Similarly, answering inferential and connecting questions may not have provided heightened activation of example items because they were unrelated, resulting in no benefit for these items as well. There was some evidence for the dissociation between these items as RAPM was significantly correlated with inferential and connecting question responses in Experiment 2, but not example questions. However, in Experiment 1, there was no relationship between RAPM and any of the passage data, indicating that a different type of processing might account for those data. One might expect that inferential and connecting questions might still have had a facilitative effect on each other, but if these items were somehow too difficult for rule learners, they may not have received the same facilitative effect. It is therefore unclear which of the above is the accurate explanation for the current data.

Replication and Extension of Function Learning Differences

The present research replicates previous findings (McDaniel et al., 2012) demonstrating multiple points of evidence for differences in function learning tendency. In both Experiments 1 and 2, there were distinct differences in patterns of extrapolation, while interpolation and training patterns were very similar across groups. These patterns are consistent with models of exemplar and rule learning (DeLosh et al., 1997). Further data supporting an individual differences approach comes from the relationship between RAPM and rate of learning within the rule learners, but not the exemplar learners in both experiments (approaching significance in both). Because RAPM assesses one's ability to abstract, the relationship indicates that, among individuals who utilize a rule-based learning strategy, those who are better at abstraction to some degree learn the training points (and presumably the function rule governing the points) faster. However, for exemplar learners, ability to abstract (RAPM) is unrelated to learning speed, indicating that these learners are adopting a qualitatively different type of strategy. These differences occur despite the fact that there are no differences in the average RAPM scores for the two groups. Thus, while rule and exemplar learners have the same average *ability* to abstract, only for the rule learners is this ability somewhat associated with learning the function.

These individual differences may clarify why it is difficult to explain all data using a single model of rule or exemplar learning. That is, the current debate in the concept-learning literature between pure exemplar models (e.g. Nosofsky, 1984; Kruschke, 1992) and pure rule-based models of learning (e.g. Bourne, 1984; Koh & Meyer, 1991; Nosofsky et al., 1994) may be resolved by including an individual difference factor. Indeed, the proposed individual differences can be used to supplement current hybrid models of learning (e.g. Anderson & Betz, 2001) where individual differences have been found to reduce task-specific effects and have therefore been considered a hindrance (Jusin et al., 2003). Proponents of hybrid models argue that the type of materials dictate the type of learning that will be used in a given situation. If, on the other hand, a given task can be achieved with an exemplar or rule-based approach, individuals will diverge and adopt one of two qualitatively different strategies. The present data also replicate the findings of McDaniel et al. (2012), which showed evidence that these differences are not isolated to function learning but can be seen across a range of tasks. Specifically, although rule learners were not superior to exemplar learners on transfer performance in the concept-learning task in Experiment 2, the present data provided the first evidence that the more general benefit of rule learning on transfer (as seen in the function learning task) is coupled with a benefit of exemplar learning on recognition of previously seen items (see Figure 9a) and overall sensitivity to whether an exemplar has been seen previously or not (as indicated by *d'* scores). Thus, while it could appear that rule learning is superior to exemplar learning for transfer (particularly in the function learning task), exemplar learning provides certain benefits as well, depending on the goals of the task.

These data are also the first to indicate that there may be a relationship between function learning tendency and analogical reasoning, as seen in Experiment 1. As in function learning training, there was no difference in success on the analogical reasoning task (as measured by the proportion of each type of learner using the convergence solution on the criterial problem). However, there was a relationship between likelihood of using the convergence solution and RAPM, but only within the rule learners. These data again indicate that rule learners who have a higher ability to abstract (i.e. score higher on RAPM) are more likely to transfer in analogical reasoning. However, there is no relationship between abstraction ability and analogical reasoning in the exemplar learners, indicating that they are using a qualitatively different approach, albeit one that is equally successful. That is, while reading the two initial stories rule learners likely abstract the underlying convergence solution schema, which they then apply to the novel problem. However, because of the relatively short delay between the stories and the criterial problem, exemplar learners may read the problem and think back to the closest task they have received to the current problem, which would be one of the stories. They can then map the story onto the current problem and arrive at a solution. If there was a longer delay, or a story in which the surface features more closely matched the criterial problem but with a different solution, exemplar learners may not perform as well on these analogical reasoning problems.

These ideas could also partially explain the data of Gick and Holyoak (1983). When participants were given two stories with the convergence solution and similar features, they performed worse than when given two stories with different surface features. This finding could be primarily because rule learners are unable to abstract the underlying schema. That is, when rule learners try to relate two dissimilar stories, the only similarities are related to the problem solution, which may then be abstracted, but when there are similar surface features they may relate the stories based on these features and not as readily on the solution. When given one convergence problem and one problem unrelated to the convergence solution, performance was at its worst. This finding could be explained if rule learners were again unable to abstract and, in addition, some exemplar learners were choosing the wrong story as the closest exemplar (all stories differed in surface features in that condition). Incorporating stories with differing surface features would help to determine if rule and exemplar learners are indeed using different strategies on these analogical reasoning problems. These data would then indicate that rule learning is not inherently superior to exemplar learning, as both types of learning may be used to accomplish similar goals.

Implications for Education

While the primary goal of the current study was to determine whether individual differences in function learning predict differences on conceptual learning, the data for the conceptual tasks were inconclusive. However, there remain potential implications for education. First and foremost, there are qualitative individual differences in function learning that can be reliably identified. Second, these differences predict behavior in concept learning as well as correlations between RAPM and analogical reasoning. Even if the differences in function learning were demonstrated not to predict differences on conceptual materials, the fact that rule and exemplar learners perform differently on these other tasks indicates that function-learning tendency might have important implications in the classroom. Specifically, learning in the STEM disciplines might be more affected by function learning differences than in non-STEM disciplines because STEM disciplines require more learning from examples, as well as problem solving, both of which might be affected by function learning tendency. Indeed, in a chemistry course at the author's university, function-learning tendencies were assessed and these differences predicted final course grades. Additional empirical work is needed in order to determine whether there are similar implications for conceptual material in non-STEM classrooms and the extent to which STEM disciplines truly are affected by differences in learning tendency.

If these differences are shown to be strong indicators of performance in the classroom, there are many more potential implications for both educators and students. It should be noted that exemplar learning does not necessarily represent an ineffective strategy. While many educators might assert that they want their students to use a rule learning strategy, there are certain courses in which an exemplar strategy may be most effective (e.g. learning foreign language vocabulary, memorizing anatomy). Therefore, first and foremost, educators may need to be aware of the learning strategy that they want to encourage. If we want to encourage rule learning in exemplar learners, it is possible that training in rule learning behavior may be beneficial for typical exemplar learners in classrooms that require this type of learning. In the chemistry course listed above, a type of inquiry learning is being utilized to attempt to do just that, such that exemplar learners might benefit particularly from this type of instruction.

It is also possible that these learning strategies are more flexible than they has been discussed throughout. Educators may be able to simply tell students their course objectives and the way in which students will be assessed in order to prep them to use the most effective strategy. If true, it would be important that educators inform students of which strategy they want them to adopt and then assess accordingly. That is, exemplar learners may have developed their strategy because it has been extremely effective in courses where they are required to recognize a memorized answer (as on most multiple-choice exams) in order to achieve the highest grade possible. The extra effort required to learn rules and underlying concepts may be a wasted one when the assessment does not require this type of behavior.

Potential Limitations and Future Work

There are several possible limitations of the current study. While the differences in function learning may have very real and important implications for education, the current work is limited in ecological validity. Most of the tasks that are used in the current study would not be used in a classroom. Therefore, additional work is needed either within the classroom or at least with classroom materials that might rely specifically on rule-based or exemplar-based processing. In addition, in the current study conceptual material was operationally defined with learning from a set of unrelated prose passages. If rule-based processing is an encoding process, rule learners may not have had any reason to make connections between passages while learning. However, within the classroom, individuals are more likely to learn from a continuous set of materials such as textbook chapters or a continuous lecture. In that case, rule learners might be more likely to attempt to make connections between materials that seem relevant. Therefore, additional work is needed within the laboratory to determine if differences in encoding produce changes in processing of conceptual material.

Conclusions

The current study is among the first to demonstrate the stability of function-learning tendency across a range of tasks, and it is the first to demonstrate differential benefits of exemplar and rule-based processing. The data indicate that individual differences in function-learning tendency have both theoretical import, as a supplement to current hybrid models of categorization, as well as applied implications for classroom materials that might heavily rely on concept-learning or problem-solving. While more evidence is necessary to determine the implications of individual differences in learning tendency, it is clear that this is an area of the concept-learning literature that is deserving of considerably more attention.

References

Anderson, J.R., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin & Review, 8*, 629-647.

Ashby, F.G., Alfonso-Reese, L.A., Turken, A.U., & Waldron, E.M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review, 105*, 442-481.

Barnett, S. M., & Ceci, S. J. (2002). When and where do we apply what we learn?: A taxonomy for far transfer. *Psychological Bulletin, 128*, 612-637.

Bors, D.A., & Stokes, T.L. (1998). Ravens Advanced Progressive Matrices: Norms for first-year university students and the development of a short form. *Educational and Psychological Measurement, 58*, 382-398.

Bott, L., & Heit, E. (2004). Nonmonotonic extrapolation in function learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 30*, 38-50.

Bourne, L.E., Jr. (1974). An inference model of conceptual rule learning. In R.L. Solso (Ed.), *Theories in cognitive psychology: The Loyola symposium* (pp. 231-256). Potomac, MD: Erlbaum.

Butler, A.C. (2010). Repeated testing produces superior transfer of learning relative to repeated studying. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 36*, 1118-1133.

Carpenter, S.K., Pashler, H., & Cepeda, N.J. (2009). Using tests to enhance students' retention of U.S. history facts. *Applied Cognitive Psychology, 23*, 760-771.

Chan, C. K., McDermott, K. B., & Roediger, H. L., III (2006). Retrieval induced facilitation: Initially nontested material can benefit from prior testing. *Journal of Experimental Psychology: General, 135*, 533-571.

DeLosh, E.L., Busemeyer, J.R., & McDaniel, M.A. (1997). Extrapolation: The sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 968-986.

Einhorn, J.H., Kleinmuntz, D.N., & Kleinmuntz, B. (1979). Regression models and process tracing analysis. *Psychological Review, 106,* 180-209.

Erickson, J.E., Chin-Parker, S., & Ross, B.H. (2005). Inference and classification learning of abstract coherent categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 86-99.

Erickson, M.A., & Kruschke, J.K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General, 127*, 107-140.

Estes, W.K. (1994). *Classification and cognition.* New York: Oxford University Press.

Fadler, C.L., Lee, J., Scullin, M., Shelton, J.T., & McDaniel, M.A. (2012). Individual differences in learning tendency predict learning of economics. (Manuscript in preparation).

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

Juslin, P., Olsson, H., & Olsson, A-C. (2003). Exemplar effects in categorization and multiple-cue judgment. *Journal of Experimental Psychology: General*, *132*, 133-156.

Kalish, M.L., Lewandowsky, S., & Kruschke, J.K. (2004). Population of linear experts: Knowledge partitioning and function learning. *Psychological Review, 111*, 1072-1099.

Katona, G. (1940). *Organizing and memorizing*. New York: Columbia University Press.

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

Koh, K., & Meyer, D.E. (1991). Function learning: Induction of continuous stimulus-response relations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17*, 811-836.

Kolb, D.A. (2007). Kolb Learning Style Inventory (Version 3.1). Experience Based Learning Systems, Inc.

Kruschke, J.K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99*, 22-44.

McDaniel, M.A., & Fisher, R.P. (1991). Tests and test feedback as learning sources. *Contemporary Educational Psychology, 16,* 192-201.

McDaniel, M.A., Cahill, M.J., Robbins, M., & Trumpower, D.L. (2012). Individual differences in learning and transfer: Stable tendencies for learning exemplars versus abstracting rules. (Manuscript under review).

Medin, D.L., & Schaffer, M.M. (1978). Context theory of classification learning. *Psychological Review, 85*, 207-238.

Medin, D.L., Altom, M.W., & Murphy, T.D. (1984). Given versus induced category representations: Use of prototype and exemplar information in classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 333-352.

Nosofsky, R.M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 104-114.

Nosofsky, R.M., & Johansen, M.K. (2000). Exemplar-based accounts of “multiple-system” phenomena in perceptual categorization. *Psychonomic Bulletin & Review, 7*, 275-402.

Nosofsky, R.M., & Palmeri, T.J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review, 101*, 53-79.

Nosofsky, R.M., Palmeri, T.J., & McKinley, S.C. (1994). Rule-plus-exception model of classification learning. *Psychological Review, 101*, 53-79.

Novick, L.R. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*, 510-520.

Raven, J.C., Raven, J.E., & Court, J.H. (1998). *Progressive matrices.* Oxford, England: Oxford Psychologists Press.

Reed, S.K. (1972). Pattern recognition and categorization. *Cognitive Psychology, 3,* 382-407.

Regehr, G., & Brooks, L.R. (1993). Perceptual manifestations of an analytic structure: The priority of holistic individuation. *Journal of Experimental Psychology: General*, *122*, 92-114.

Roediger, H. L., III, & Karpicke, J. D. (2006a). The power of testing memory:

Basic research and implications for educational practice. *Perspectives on Psychological Science, 1*, 181-210.

Roediger, H. L., III, & Karpicke, J. D. (2006b). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science, 17*, 249-255.

Ruble, T. L., & Stout, D. E. (1991). Reliability and classification stability, and response set bias of alternative forms of the learning style inventory (LSI-1985). *Educational and Psychological Measurement, 51*, 481-489.

Smith, E.E., Patalano, A.L., & Jonides, J. (1998). Alternative strategies of categorization. *Cognition, 65*, 167-196.

Veres, J. G., Sims, R. R., & Locklear, T. S. (1991). Improving the reliability of Kolb's revised learning style inventory. *Educational and Psychological Measurement, 51,* 143-150.

Wenke, D., Frensch, P.A., & Funke, J. (2005). Complex problem solving and intelligence: Empirical relation and causal direction. In R.J. Sternberg & J.E. Pretz (Eds.), *Cognition & Intelligence: Identifying the mechanisms of the mind.* (pp. 160-187). New York: Cambridge University Press.

Table 1.
Demographic characteristics as a function of learner type in Experiments 1 and 2.

	Experiment 1		Experiment 2	
	Rule	Exemplar	Rule	Exemplar
Age (Mean)	20.4 (1.44)	20.9 (2.68)	20.4 (2.89)	20.0 (1.45)
Sex (Number)				
Female	21	19	20	20
Male	16	15	13	14
Race (Number)				
White	26	21	24	21
Asian	7	9	6	7
African American	2	2	2	5
Hispanic	0	2	0	1
GPA (Mean)	3.52 (.37)	3.53 (.37)	3.55 (.37)	3.53 (.32)
ACT (Mean)	32.38 (1.98)	32.61 (1.58)	32.26 (3.18)	32.33 (2.27)
Major (Number)				
STEM	20	12	19	19
Non-STEM	8	10	4	8
Both	7	9	3	3

Note. Standard deviations are displayed in parentheses where applicable.

Table 2.

Abstract coherent categories stimuli taken from Erickson et al. (2005) and used in Experiment 2.

	Coherent morkel	Incoherent krenshaw
Training Stimuli	Operates on land Works to gather harmful solids Has a shovel Rolls on wheels	Operates on land Works to clean spilled oil Has a shovel Slides on skis
	Operates on the surface of the water Works to clean spilled oil Has a spongy material Slides on skis	Operates on the surface of water Works to gather harmful solids Has a spongy material Rolls on wheels
Novel Test Stimuli	Operates in highway tunnels Works to remove carbon dioxide Has a large intake fan Flies with a propeller	Operates on the seafloor Works to remove broken glass Has a large intake fan Flies with a propeller
	Operates on the seafloor Works to remove lost fishing nets Has a hook Swims with fins	Operates on the beach Works to remove carbon dioxide Has a hook Rolls on a tread

Note. The four features for each morkel item were coherent. The four features for each krenshaw consisted of two pairs of coherent features (location-instrument and pollutant-locomotion), but the two pairs did not fit together to provide a coherent whole.

Table 3.

Proportion of subjects who indicated that they were unable to read each passage.

Passage title	Proportion of subjects
Bats	.05
Tropical Cyclones	.10
Vaccines	.15
Bread	.08
Respiratory System	.19
Internet	.13
Reptiles	.10
Liver	.13
McCarthyism	.18
Volcanoes	.06
Flowers	.06
Balloons	.03

Table 4.

Test performance as a function of item type and learning tendency in Experiment 2 with and without unread passages included.

	All Passages Included		Unread Passages Removed	
	Rule	Exemplar	Rule	Exemplar
Strict scoring				
Example	.30 (.03)	.26 (.03)	.32 (.04)	.28 (.04)
Inferential	.37 (.04)	.40 (.04)	.38 (.04)	.41 (.04)
Connecting	.33 (.04)	.38 (.04)	.33 (.05)	.39 (.05)
Correct passage scoring				
Inferential	.55 (.03)	.56 (.03)	.56 (.03)	.54 (.03)
Connecting	.26 (.04)	.24 (.04)	.26 (.04)	.25 (.04)

Note: Standard deviations are displayed in parentheses.

Figure 1a.

Data from DeLosh et al. (1997) showing individual subjects who extrapolated according to the function.

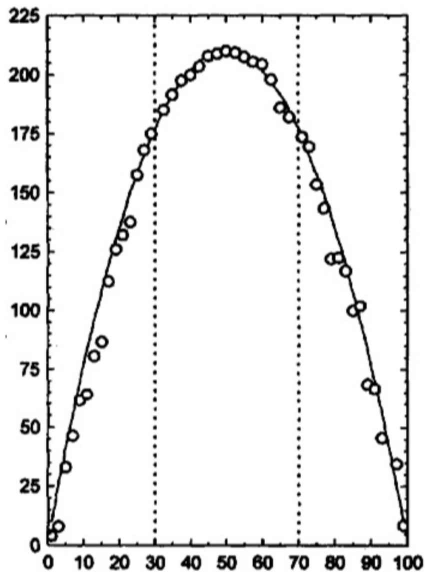

Figure 1b.

Data from DeLosh et al. (1997) showing individual subjects who used output values from extreme points for extrapolation trials.

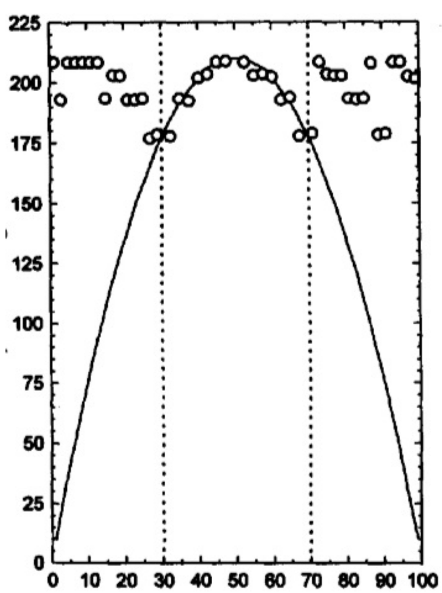

Figure 2.

A taxonomy of transfer proposed by Barnett and Ceci (2002). The upper box represents the content factor, while the lower box represents the context transfer.

A Content: What transferred			
Learned skill	Procedure	Representation	Principle or heuristic
Performance change	Speed	Accuracy	Approach
Memory demands	Execute only	Recognize and execute	Recall, recognize, and execute

B Context: When and where transferred from and to					
	Near ←			→	**Far**
Knowledge domain	Mouse vs. rat	Biology vs. botany	Biology vs. economics	Science vs. history	Science vs. art
Physical context	Same room at school	Different room at school	School vs. research lab	School vs. home	School vs. the beach
Temporal context	Same session	Next day	Weeks later	Months later	Years later
Functional context	Both clearly academic	Both academic but one nonevaluative	Academic vs. filling in tax forms	Academic vs. informal questionnaire	Academic vs. at play
Social context	Both individual	Individual vs. pair	Individual vs. small group	Individual vs. large group	Individual vs. society
Modality	Both written, same format	Both written, multiple choice vs. essay	Book learning vs. oral exam	Lecture vs. wine tasting	Lecture vs. wood carving

Figure 3.

Procedure used in Experiments 1 and 2.

Experiment 1

Session 1:

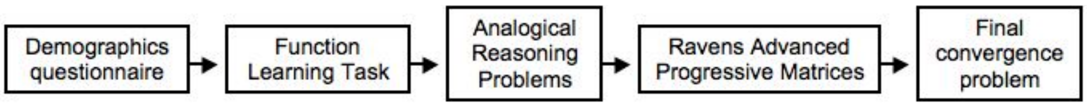

Session 2:

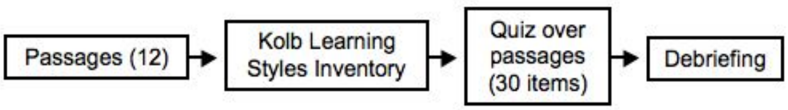

Experiment 2

Session 1:

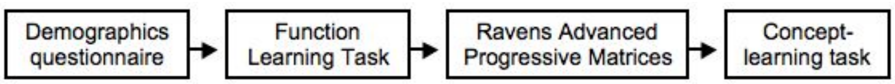

Session 2:

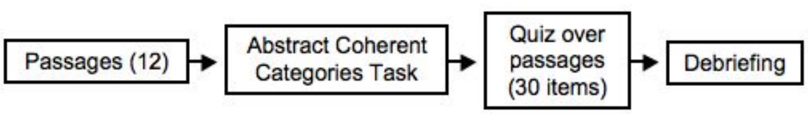

Figure 4.

Sample training screen used during the function-learning task in Experiments 1 and 2.

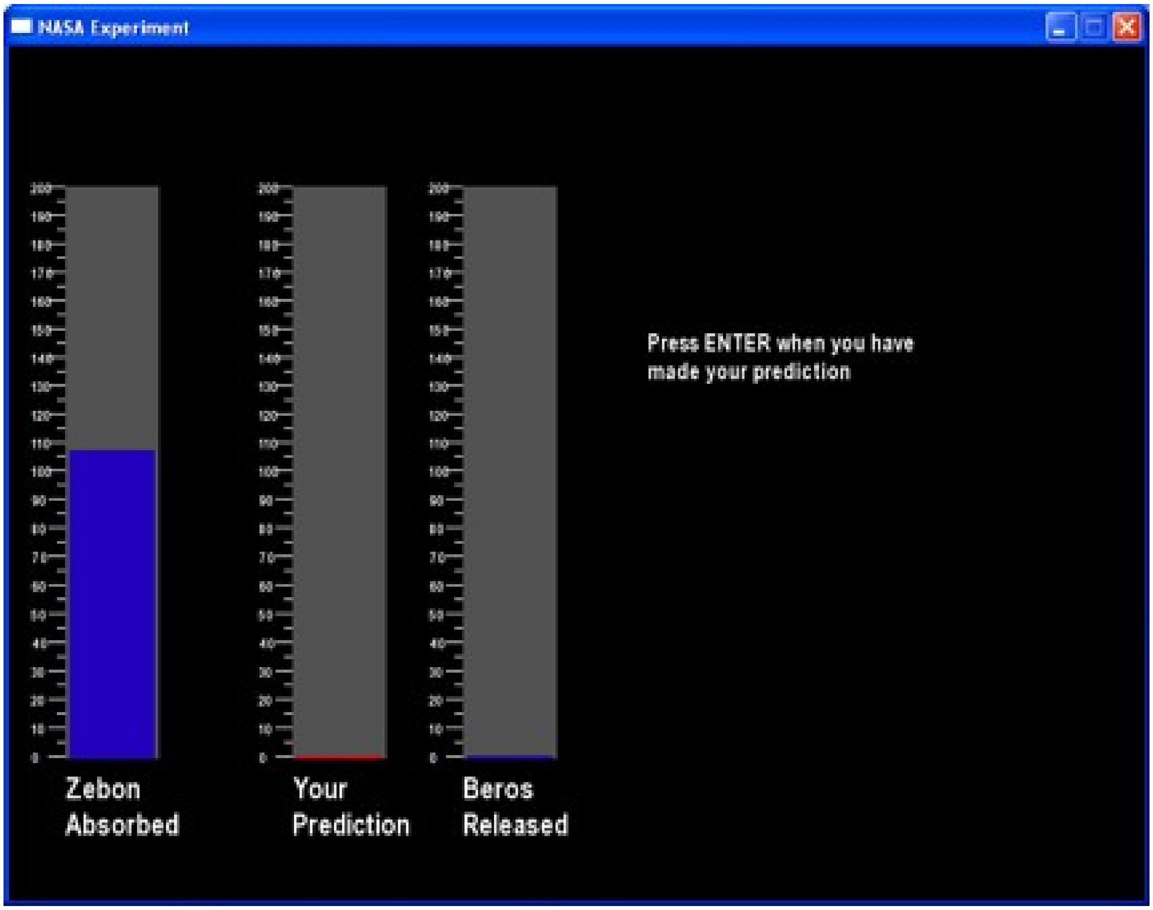

Figure 5a.

Average predicted values for extrapolation (upper panel) and interpolation (lower panel) for rule, exemplar, and sine learners in Experiment 1.

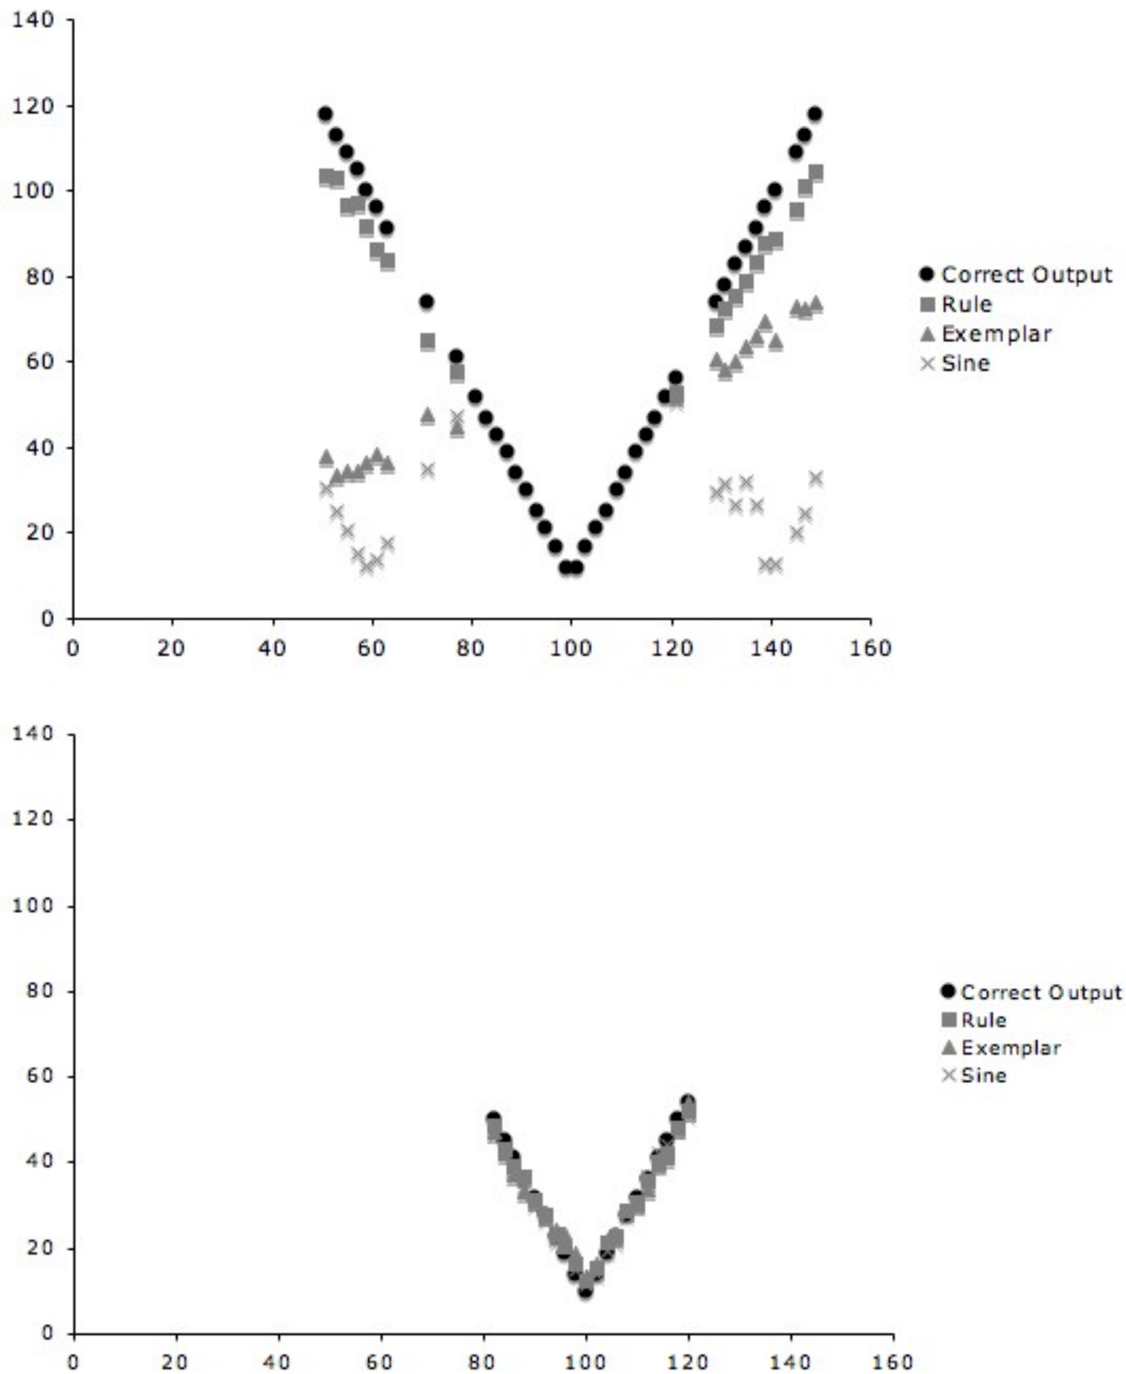

Figure 5b.

Average predicted values for extrapolation (upper panel) and interpolation (lower panel) for rule, exemplar, and sine learners in Experiment 2.

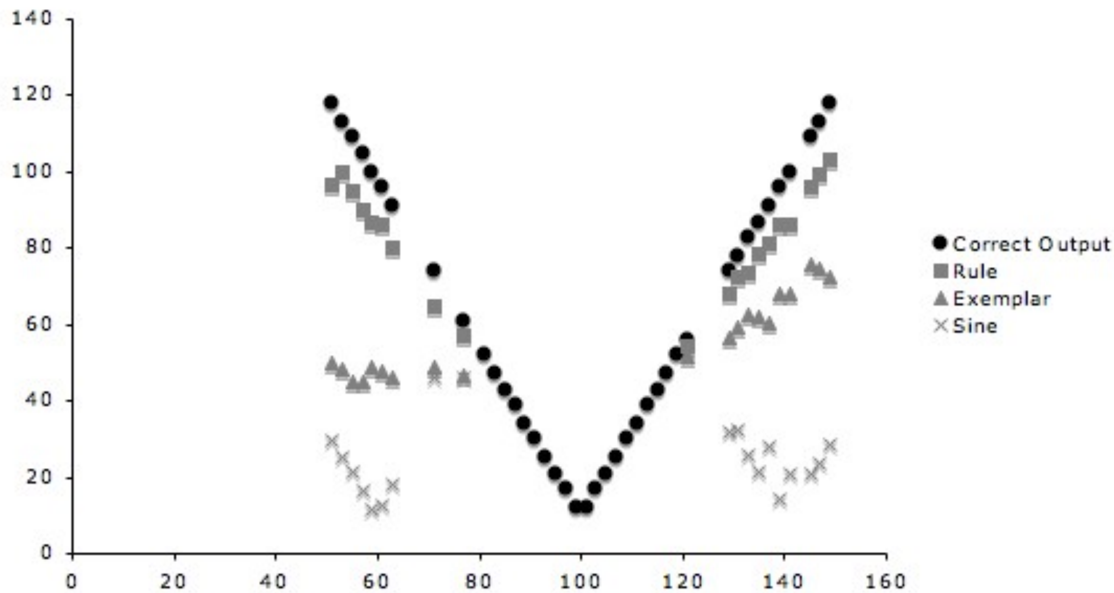

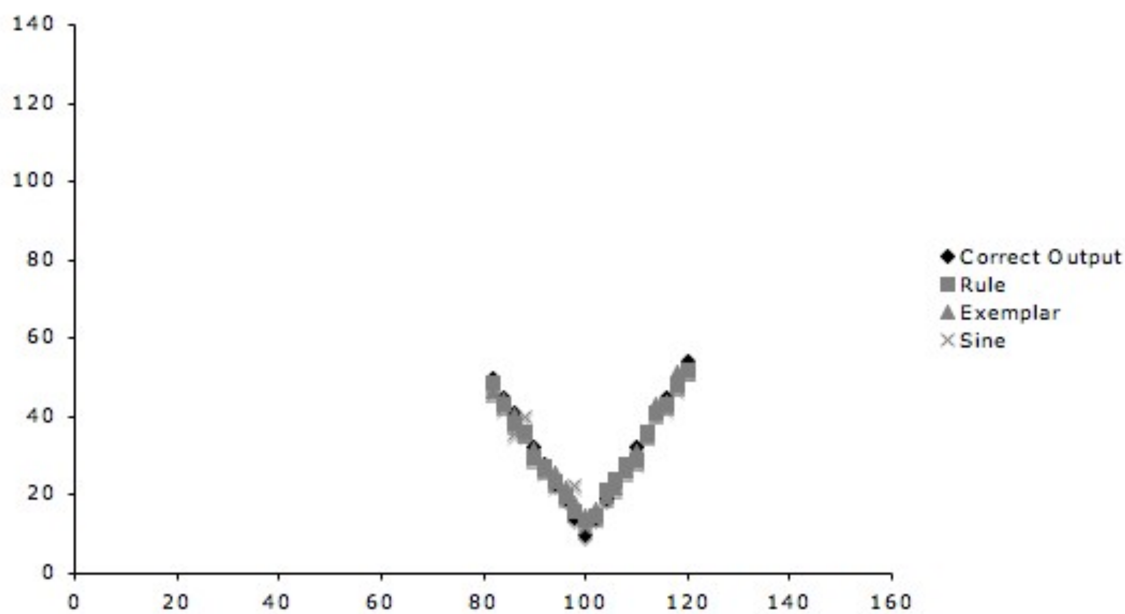

Figure 6.

Mean absolute errors in each training block as a function of condition in Experiment 1 (upper panel) and Experiment 2 (lower panel).

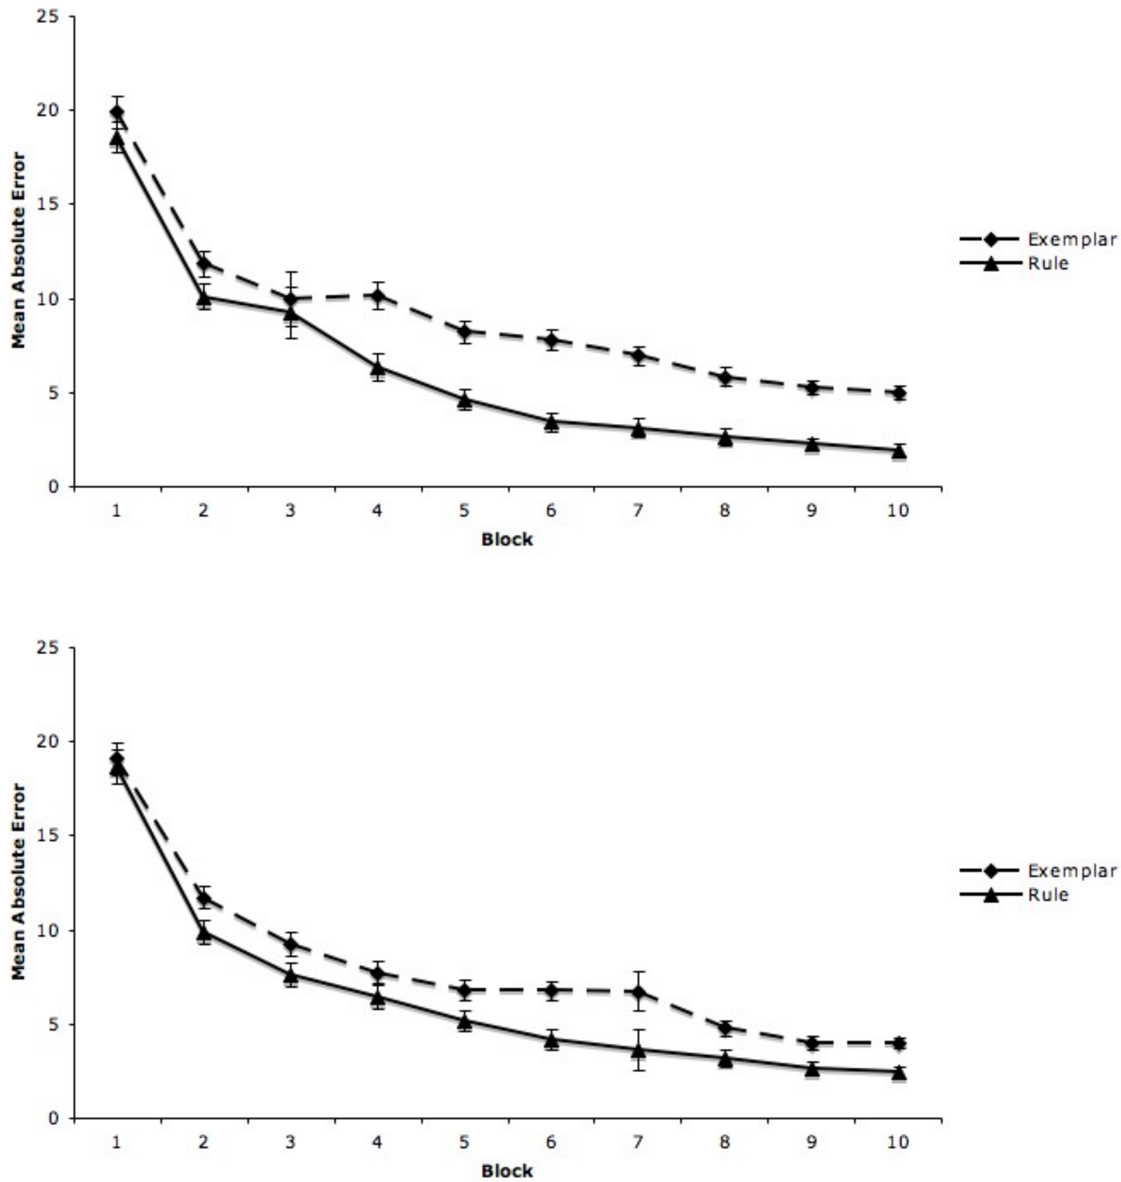

Figure 7a.

Test performance (strictly scored) as a function of item type and learning tendency in Experiment 1.

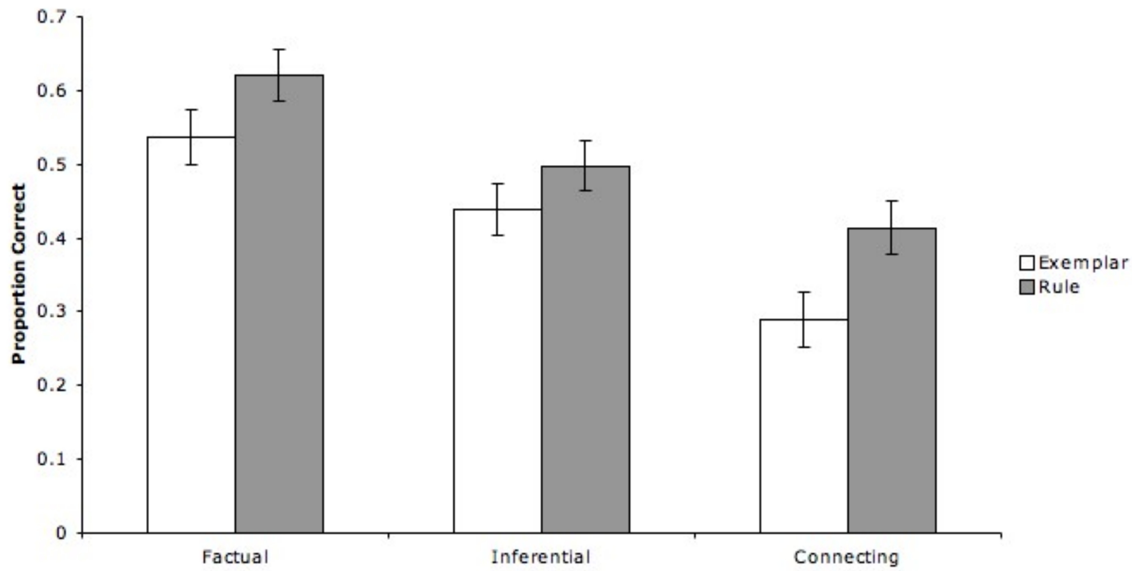

Note: Error bars represent the standard error of the mean.

Figure 7b.

Use of the correct passage as a function of item type and learning tendency in Experiment 1.

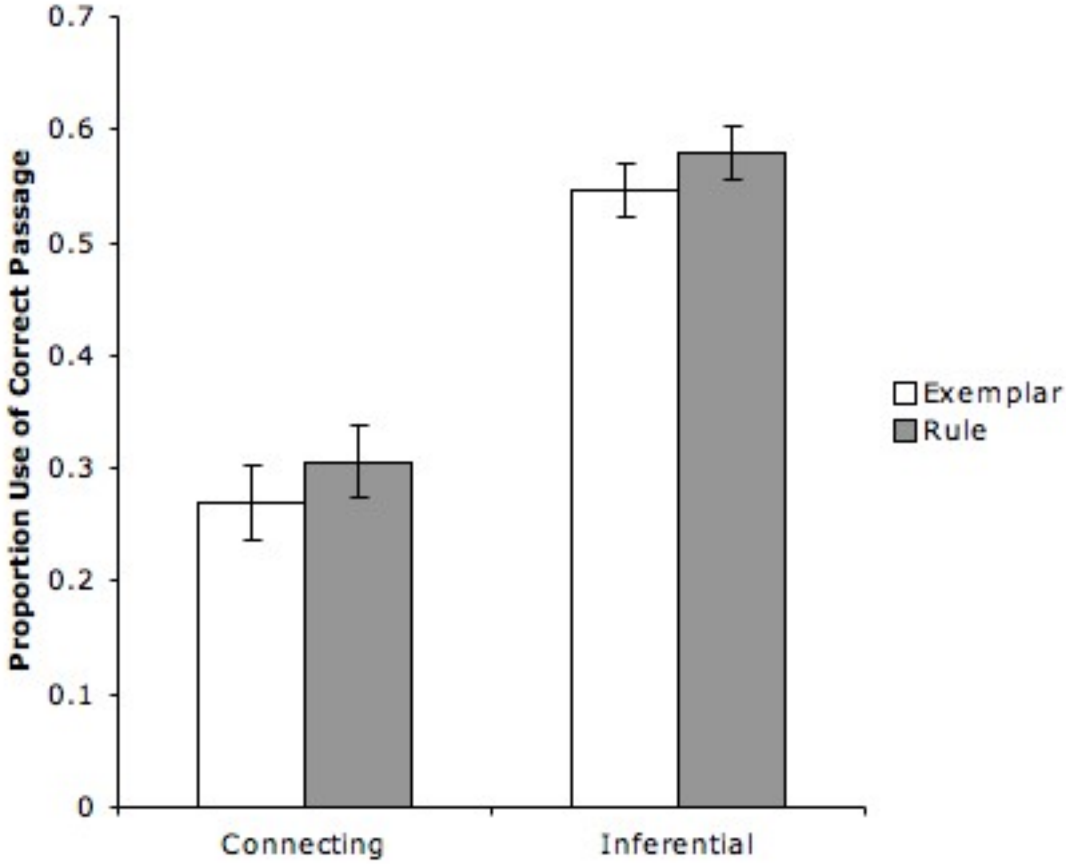

Note: Error bars represent the standard error of the mean.

Figure 8a.

Test performance (strictly scored) as a function of item type and learning tendency in Experiment 2.

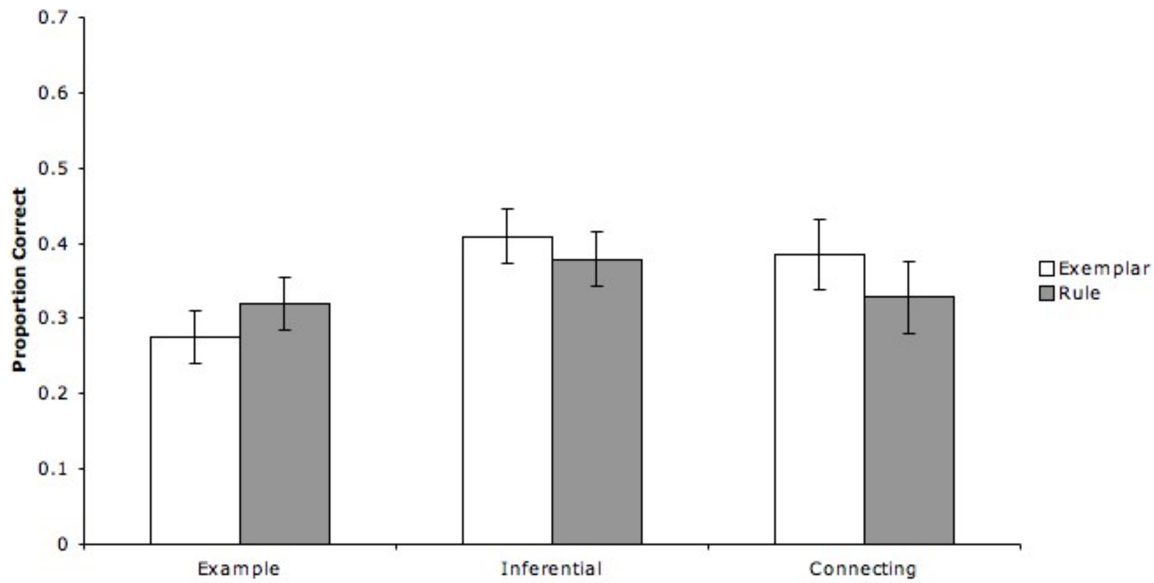

Note. Error bars represent the standard error of the mean.

Figure 8b.

Use of the correct passage as a function of item type and learning tendency in Experiment 2.

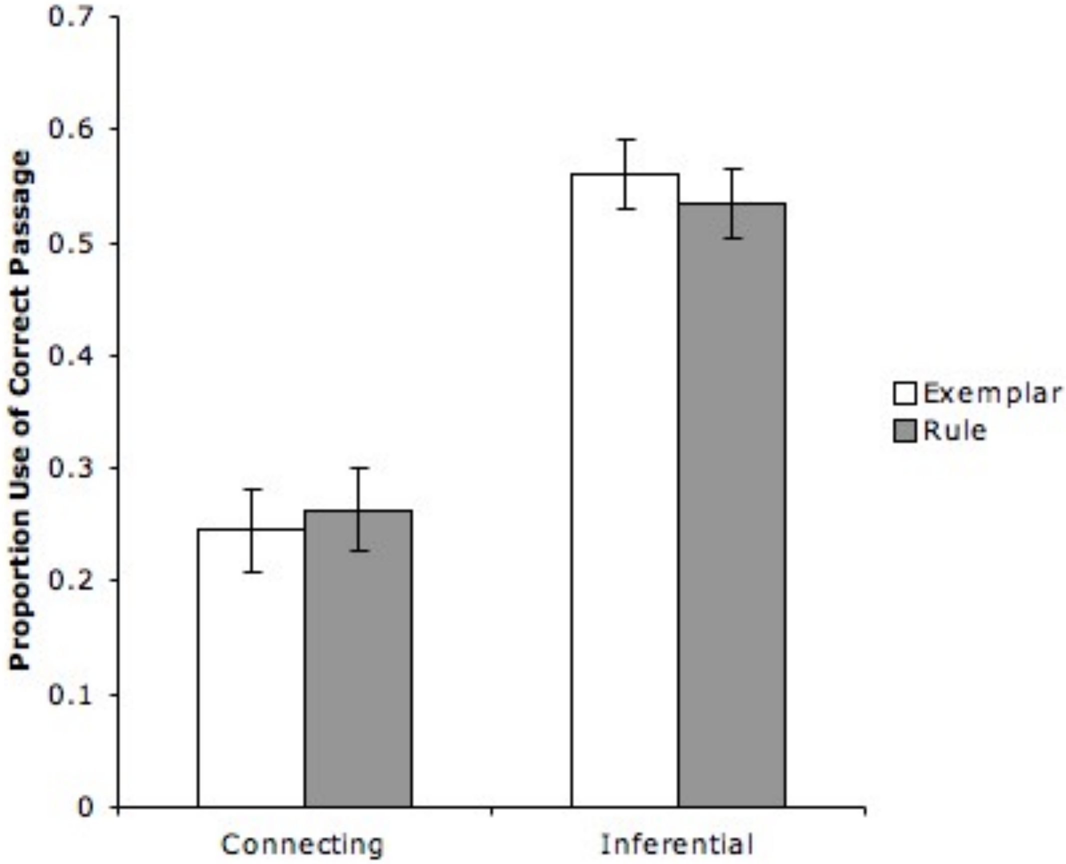

Note. Error bars represent the standard error of the mean.

Figure 9a.

Recognition and categorization performance on repeated training items on the concept-learning task.

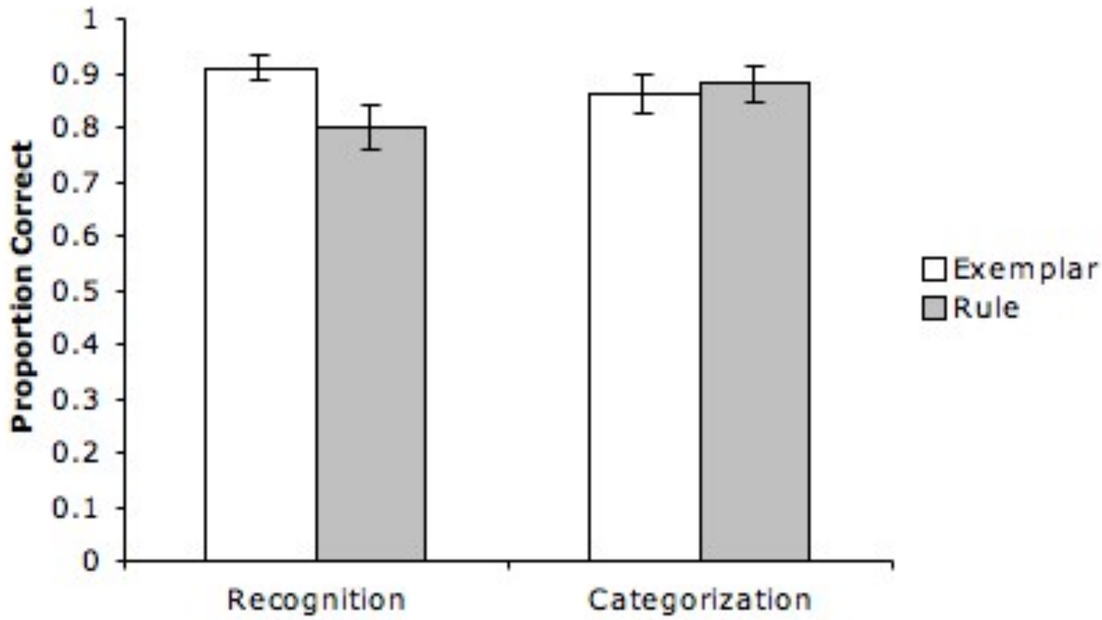

Note. Error bars represent the standard error of the mean.

Figure 9b.

Recognition performance as a function of type of lure and type of learner on the concept-learning task.

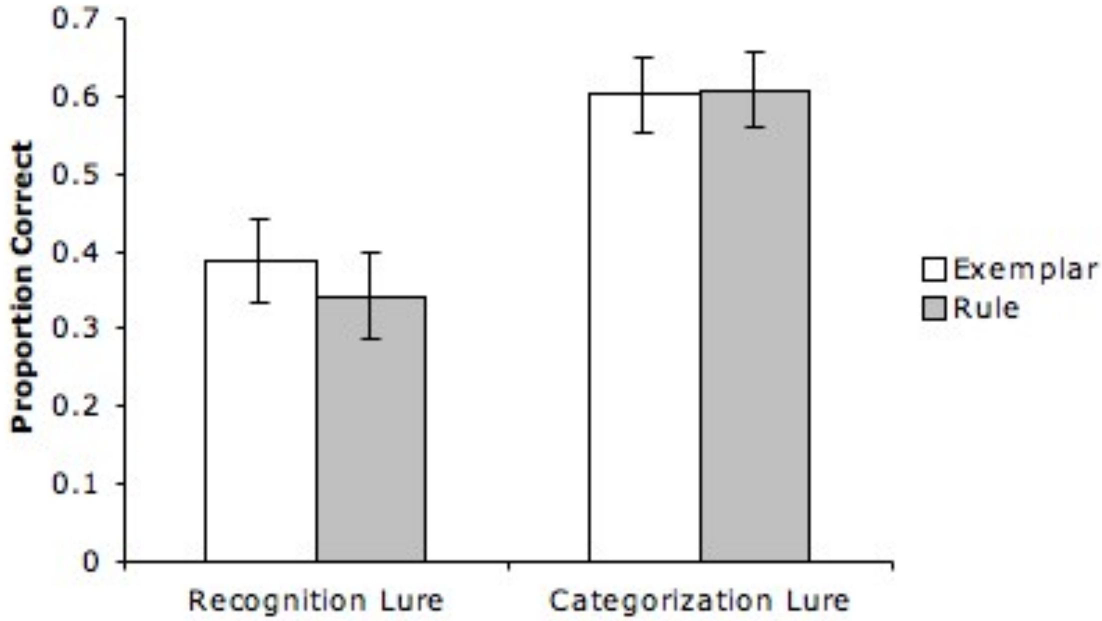

Note. Error bars represent the standard error of the mean.

Figure 9c.

Categorization performance as a function of type of lure and type of learner on the concept-learning task.

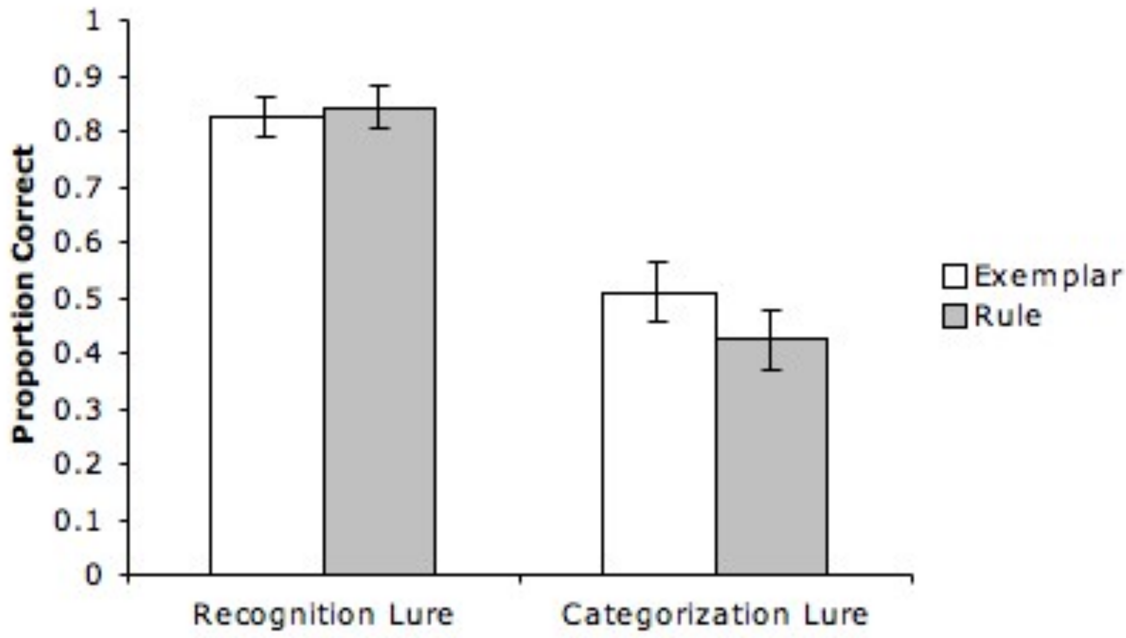

Note. Error bars represent the standard error of the mean.

Figure 10.

Proportion correct on the two-feature test as a function of item type and learner type on the abstract coherent categories task.

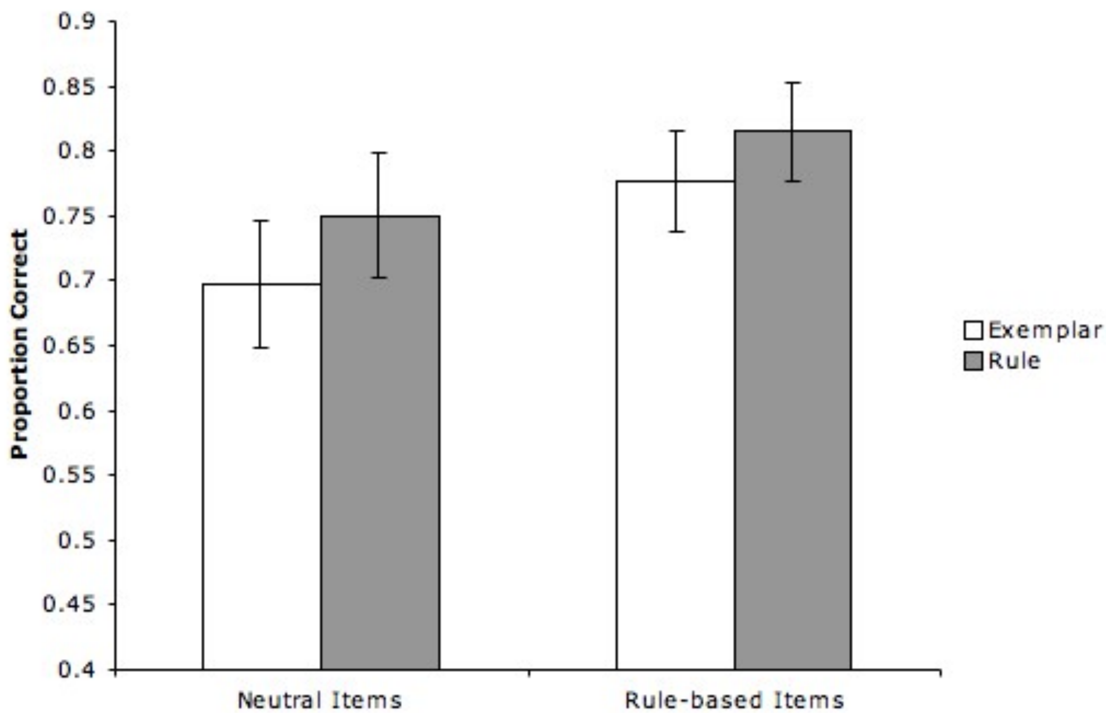

Note. Error bars represent the standard error of the mean.

Figure 11.

Ravens performance as a function of the training block on which each subject reached criterion on the function learning task combined across Experiments 1 and 2. Rule learners are depicted in the top panel; exemplar learners are depicted in the bottom panel.

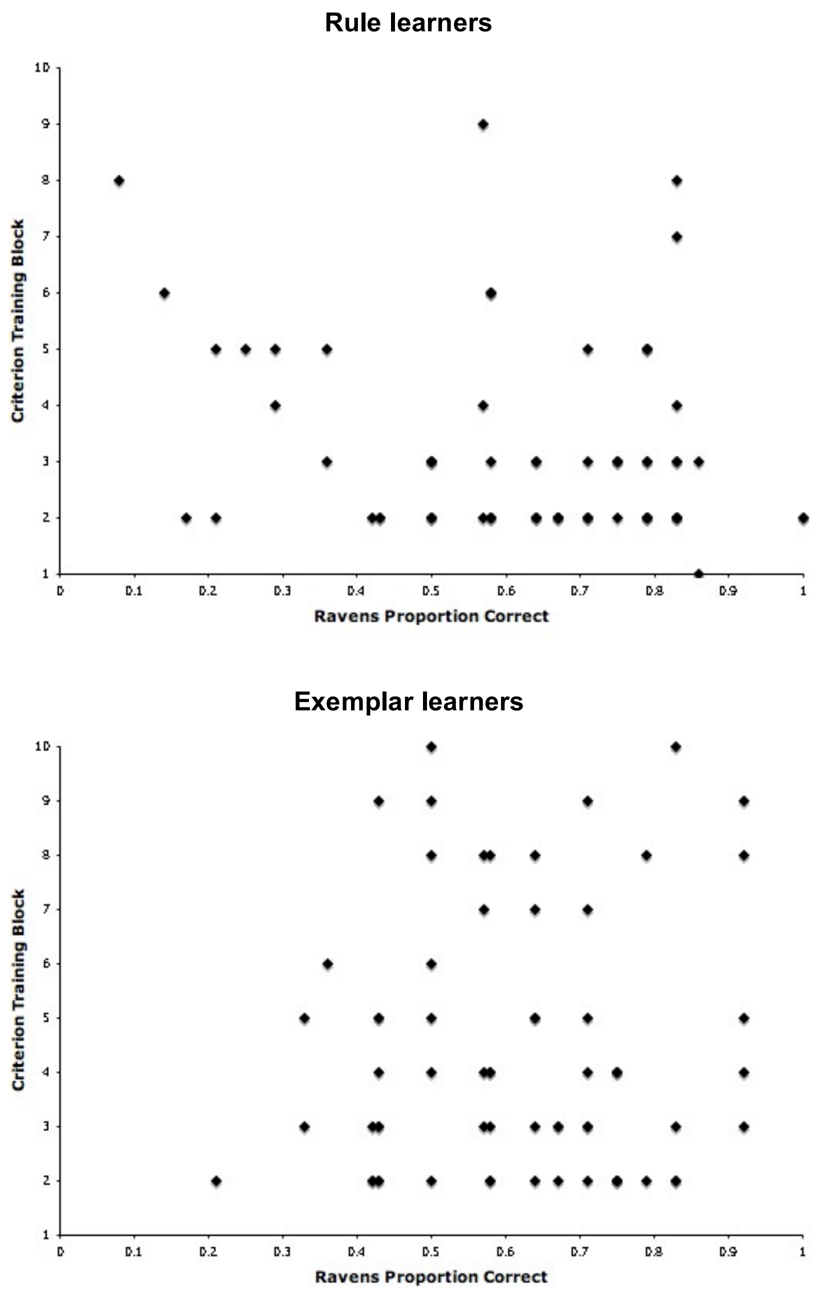

Appendix A

Convergence problems used in Experiment 1

These problems are taken from Gick and Holyoak (1983, Experiment 3). The first two problems were presented as reading comprehension stories, while the last problem was presented separately after a delay.

RED ADAIR

An oil well in Saudi Arabia exploded and caught fire. The result was a blazing inferno that consumed an enormous quantity of oil each day. After initial efforts to extinguish it failed, famed firefighter Red Adair was called in. Red knew that the fire could be put out if a huge amount of fire retardant foam could be dumped on the base of the well. There was enough foam available at the site to do the job. However, there was no hose large enough to put all the foam on the fire fast enough. The small hoses that were available could not shoot the foam quickly enough to do any good. It looked like there would have to be a costly delay before a serious attempt could be made. However, Red Adair knew just what to do. He stationed men in a circle all around the fire, with all of the available small hoses. When everyone was ready, all of the hoses were opened up and foam was directed at the fire from all directions. In this way a large amount of foam quickly struck the source of the fire. The blaze was extinguished, and the Saudis were satisfied that Red had earned his three million dollar fee.

FORTRESS

A small country was ruled from a strong fortress by a dictator. The fortress was situated in the middle of the country, surrounded by farms and villages. Many roads led to the fortress through the countryside. A rebel general vowed to capture the fortress. The general knew that an attack by his entire army would capture the fortress. He gathered his army at the head of one of the roads, ready to launch a full-scale direct attack. However, the general then learned that the dictator had planted mines on each of the roads. The mines were set so that small bodies of men could pass over them safely, since the dictator needed to move his troops and workers to and from the fortress. However, any large force would detonate the mines. Not only would this blow up the road, but it would also destroy many neighboring villages. It therefore seemed impossible to capture the fortress. However, the general devised a simple plan. He divided his armies into small groups and dispatched each group to the head of a different road. When all was ready, he gave the signal and each group marched down a different road. Each group continued down its road so that the entire army arrived together at the fortress at the same time. In this way, the general captured the fortress and overthrew the dictator.

TUMOR PROBLEM

Suppose you are a doctor faced with a patient who has a malignant tumor in his stomach. It is impossible to operate on the patient, but unless the tumor is destroyed the patient will die. There is a kind of ray that can be used to destroy the tumor. If the rays reach it all at once at a sufficiently high intensity, the tumor will be destroyed. Unfortunately, at this intensity the healthy tissue the rays pass through on the way to the tumor will also be destroyed. At lower intensities the rays are harmless to healthy tissue, but they will not affect the tumor, either. What type of procedure might be used to destroy the tumor with the rays and at the same time avoid destroying the healthy tissue?

Appendix B

Passages used in Experiment 1 and 2

The passages are arranged such that the first 6 passages were taken from Butler (2010) and the next 6 passages were written for the current study.

BATS

Although bats and birds both fly, a bat wing actually has more in common with a human arm than a bird wing. A bird's wing has fairly rigid bone structure, and the main flying muscles move the bones at the point where the wing connects to the body. In contrast, a bat has a much more flexible wing structure. It is similar to a human arm and hand, except it has a thin membrane of skin (called the *patagium*) extending between the "hand" and the body, and between each finger bone. Bats can use the wing like a hand, essentially moving through the air like a swimmer moves through water. The rigid bird wing is more efficient at providing lift, but the flexible bat wing allows for greater maneuverability.

To help them navigate and find their prey in the dark, microbat species have developed a remarkable system called echolocation. By emitting high-pitched sound waves and listening to the echoes, bats can determine with great precision the location of an object, how big it is, and the direction in which it is moving. Bats calculate the distance of the object by the amount of time it takes for the sound wave to return and the exact position of the object by comparing when the sound reaches its right ear to when the sound reaches its left ear. Similarly, a bat can tell how big an insect is based on the intensity of the echo: a smaller object will reflect less of the sound wave, and so will produce a less intense echo.

Bats have a special physiological adaptation that enables them to hang upside down. A bat's talons work like human fingers, except that humans must contract muscles to grasp an object, whereas bats must do the opposite – relax their muscles. When humans grasp an object, they contract several arm muscles, which in turn pull tendons connected to their fingers, which pull the fingers closed. To hang upside down, a bat opens its talons to grab hold of the surface, and then simply lets its body relax. The weight of the upper body pulls down on the tendons connected to the talons, causing them to clench. Since it is gravity that keeps the talons closed, instead of a contracted muscle, the bat doesn't have to exert any energy to hang upside down.

Like all mammals, bats maintain their body temperature internally. However, unlike most mammals, bats allow their body temperature to sink to the ambient temperature whenever they are not active. As their temperature drops, they enter a torpor state, in which their metabolism slows down considerably. By reducing their biological activity and not maintaining a warm body temperature, bats conserve energy. This ability is important because flying all night is hard work. When the temperature is cold for long periods during the winter months, some bats enter a deeper torpor state called hibernation. Other bat species follow a yearly migration pattern, traveling to cooler climates in the warm months and warmer climates in the cool months. This is why some regions experience “bat seasons” every year.

TROPICAL CYCLONES

Tropical cyclones often begin their lives as clusters of clouds and thunderstorms called tropical disturbances. In order to take the first step towards becoming a full-blown tropical cyclone, a disturbance must develop a pocket of low-pressure air at its center. This process, which can take anywhere from hours to days, begins with the thunderstorms in the disturbance releasing latent heat. This heat warms the air in the disturbance, causing oxygen molecules to expand and thereby lowering the density of the air. As the density of the air drops, so too does the air pressure. Once a low-pressure area exists, the first step is complete and the disturbance has the potential to take the next step in its development: beginning to rotate at high speeds.

Once rotation is initiated, a tropical cyclone builds in strength through rapidly rising air at the center of the storm. As it moves across the ocean, it sucks up warm, moist tropical air from the surface of the water and dispenses cooler air aloft. A tropical cyclone's primary energy source is the release of the heat of condensation from water vapor in this rising air. The release of heat creates a pattern of wind that circulates around a center, like water going down a drain, and brings the rotation of the tropical cyclone to high speeds. In addition to the warm air being sucked up into the center of the storm, converging winds at the surface and higher altitudes also push warm air upwards, increasing the rotation.

A tropical cyclone has two key parts. The low-pressure center of relative calm is called the eye. Weather in the eye is normally calm and free of clouds, although the sea may be extremely violent. Circular in shape, the eye may range in size from 5 to 120 miles in diameter, but most eyes are between 20 and 40 miles across. The area surrounding the eye is called the eye wall, and it consists of a dense wall of clouds and thunderstorms. The eye wall is the part of the storm where the greatest wind speeds are found, clouds reach the highest, and precipitation is the heaviest. Interestingly, the eye wall actually creates the eye by sucking out any clouds or rain in the area.

One measure of the size of a tropical cyclone is called the Radius of Outermost Closed Isobar (ROCI). The atmospheric pressure increases gradually as one moves away from the center of the storm, and the outermost closed isobar is the point at which the pressure returns to normal. ROCI is determined by measuring the radii from the center of the storm to its outermost closed isobar in each of the four quadrants surrounding the storm. The distances of the radii are then averaged to come up with a single value. If the ROCI is between 2 and 3 degrees of latitude, then the cyclone is considered “small”. A ROCI between 3 and 6 latitude degrees is considered “medium.” A “large” tropical cyclone has a ROCI of between 6 and 8 degrees.

VACCINES

A vaccine is a biological preparation that establishes or improves immunity to a particular disease. Most vaccines are prophylactic, which means that they prevent or ameliorate the effects of a future infection by any natural pathogen. The flu vaccine is an example of a prophylactic vaccine that is given annually to protect against the influenza virus. However, vaccines have also been used for therapeutic purposes, such as for alleviating the suffering of people who are already afflicted with a disease. An example of such a therapeutic use is the vaccines currently being developed for the treatment of various types of cancer. Until recently, most vaccines have been aimed at children, but the development of therapeutic vaccines has increased the number of treatments targeted at adults.

Over the following centuries, medical researchers like Edward Jenner and Louis Pasteur transformed the ancient technique of variolation into the modern day practice of inoculation with vaccines. Inoculation represented a major breakthrough because it reduced the risk of vaccination, while maintaining its effectiveness. Inoculation is the practice of deliberate infection through a skin wound. This new technique produces a smaller, more localized infection relative to variolation in which inhaled viral particles in droplets spread the infection more widely. The smaller infection works better because it is adequate to stimulate immunity to the virus, but it also keeps the virus from replicating enough to reach levels of infection likely to kill a patient.

Some vaccines are made from dead or inactivated virulent organisms that have been killed with chemicals or heat. Examples are vaccines against influenza, cholera, and hepatitis. Other vaccines contain live, attenuated virus organisms that are cultivated under conditions that disable their virulent properties. Examples include yellow fever, measles, rubella, and mumps. Aluminium-based adjuvants, such as squalene, are typically added to boost immune response. Vaccines can be monovalent or polyvalent. A monovalent vaccine is designed to immunize against a single antigen or single microorganism. A polyvalent vaccine is designed to immunize against two or more strains of the same organism, or against two or more organisms. In certain cases, a monovalent vaccine may be preferable for rapidly developing a strong immune response.

One challenge in vaccine development is economic: many of the diseases that could be eradicated with a vaccine, such as malaria, exist principally in poor countries. Although many vaccines have been highly cost effective and beneficial for public health, pharmaceutical firms and biotechnology companies have little incentive to develop vaccines for these diseases because there is little revenue potential. Even in more affluent countries, financial returns are usually minimal while the costs are great. The number of vaccines administered has actually risen dramatically in recent decades, but this rise is due to government mandates and support, rather than economic incentive. Thus, most vaccine development relies on “push” funding that is supplied by government, universities, and non-profit organizations.

BREAD

Flour provides the primary structure to bread because it contains proteins – it is the quantity of these proteins that determines the quality of the finished bread. Wheat flour contains two non-water soluble protein groups (glutenin and gliadin), which form the structure of the dough. When worked by kneading, the glutenin forms long strands of chainlike molecules while the shorter gliadin forms bridges between the strands of glutenin, resulting in a network of strands called gluten. The network of strands, or gluten, is responsible for the softness of the bread because it traps tiny air bubbles as the dough is baked. If the network of strands is more cohesive or tightly linked, the bread will be softer. Gluten development improves if the dough is allowed to rest between mixing and kneading.

The amount of flour is the most significant measurement in a bread recipe. Professional bakers use a system known as Bakers' Percentage in their recipe formulations. They measure ingredients by weight rather than by volume because it is more accurate and consistent, especially for dry ingredients. Flour is always stated as 100%, and the rest of the ingredients are a percent of that amount by weight. For example, common table bread in the U.S. uses approximately 50% water, whereas most artisan bread formulas contain anywhere from 60 to 75% water. The water (or sometimes another liquid like milk or juice) is used to form the flour into a paste or dough.

Gas-producing chemicals can also be used as a leavening agent. Whereas yeast takes two to three hours to produce its leavening action, a dry chemical leavening agent like baking powder is instantaneous. Many commerical bakeries use chemical additives to speed up mixing time and reduce necessary fermentation time, so that a batch of bread may be mixed and baked in less than 3 hours. "Quick bread" is the name that commercial bakers use for dough that does not require fermentation because of chemical additives. Often these chemicals are added to dough in the form of a prepackaged base, which also contains most or all of the dough's non-flour ingredients. Commercial bakeries also commonly add calcium propionate to retard the growth of molds.

The development of leavened bread can probably be traced to prehistoric times as well. Yeast spores occur everywhere, so any dough left to rest will become naturally leavened. For example, an uncooked dough exposed to air for some time before cooking would probably contain airborne yeasts as well as yeasts that grow on the surface of cereal grains. Thus, the most common source of leavening was early bakers retaining a piece of dough from the previous day to utilize as a form of dough starter. Although leavening is likely of prehistoric origin, the earliest archaeological evidence comes from ancient Egypt. Scientific analysis using electron microscopy has detected yeast cells in some ancient Egyptian loaves.

THE RESPIRATORY SYSTEM

When a person inhales, the diaphragm and intercostal muscles (the muscles between the ribs) contract and expand the chest cavity. This expansion lowers the pressure in the lungs below the outside air pressure. Air then flows in through the airways (from high pressure to low pressure) and inflates the lungs. The lungs are made of spongy, elastic tissue that stretches and constricts during breathing. When a person exhales, the diaphragm and intercostal muscles relax and the chest cavity gets smaller. The decrease in volume of the cavity increases the pressure in the lungs above the outside air pressure. Air from the lungs (high pressure) then flows out of the airways to the outside air (low pressure). The cycle then repeats with each breath.

Within the alveoli, gas exchange occurs through diffusion. Diffusion is the movement of particles from a region of high concentration to a region of low concentration. The oxygen concentration is high in the alveoli, so oxygen diffuses across the alveolar membrane into the pulmonary capillaries, which are small blood vessels that surround each alveolus. The hemoglobin in the red blood cells passing through the pulmonary capillaries has carbon dioxide bound to it and very little oxygen. The oxygen binds to hemoglobin and the carbon dioxide is released. Since the concentration of carbon dioxide is high in the pulmonary capillaries relative to the alveolus, carbon dioxide diffuses across the alveolar membrane in the opposite direction. The exchange of gases across the alveolar membrane occurs rapidly – usually in fractions of a second.

Several factors can trigger such an override by the autonomic nervous system. One of these factors is the concentration of oxygen in the blood. Specialized nerve cells within the aorta and carotid arteries called peripheral chemoreceptors monitor the oxygen concentration of the blood. If the oxygen concentration decreases, the chemoreceptors signal to the respiratory centers in the brain to increase the rate and depth of breathing. These peripheral chemoreceptors also monitor the carbon dioxide concentration in the blood. Another factor is chemical irritants. Nerve cells in the airways can sense the presence of unwanted substances like pollen, dust, water, or cigarette smoke. If chemical irritants are detected, these cells signal the respiratory centers to contract the respiratory muscles, and the coughing that results expels the irritant from the lungs.

Disorders of the respiratory system fall mainly into two classes. Some disorders make breathing harder, while other disorders damage the lungs' ability to exchange carbon dioxide for oxygen. Asthma is an example of a disease that influences the mechanics of breathing. During an asthma attack, the bronchioles constrict, narrowing the airways. This reduces the flow of air and makes the respiratory muscles work harder. In contrast, pulmonary edema is an example of a disease that minimizes or prevents gas exchange. Pulmonary edema occurs when fluid builds up in the area between the alveolus and pulmonary capillary, increasing the distance over which gases must exchange and slowing down the exchange. Various medical interventions are used treat disorders of the respiratory system, but coughing is the body’s main method of defense.

THE INTERNET

The story of the Internet begins with the launch of the Soviet satellite Sputnik in 1957, which spurred the United States to establish the Advanced Research Projects Agency (ARPA) in order to regain a technological lead. A project leader at ARPA, Joseph Licklider, saw great potential in universal networking and initiated a project to build a network that relied on a new technology called packet switching. Packet switching is a mode of data transmission in which data is broken into chunks, called packets, which are sent independently and then reassembled at the destination. Alternative modes of data transmission, such as circuit switching, require a fixed connection between terminals, so each circuit can handle only one user at a time. In contrast, packet switching can accommodate multiple users, optimizing network use and minimizing data transmission time.

Until the late 1980s, the networks were used for governmental and scientific research purposes only. However, this restriction on the networks came to an end when the U.S. Federal Networking Council approved the interconnection of the NSFNET to the commercial MCI Mail system in 1988. The opening of the network to commercial interests greatly accelerated the expansion of what is now called the Internet. Motivated by potential profits, commercial companies aggressively pursued the connection of existing networks and the creation of new networks. Although the Internet had existed for almost a decade, the network did not gain a public face until the 1990s. In 1991, the European Organization for Nuclear Research publicized a new project called the World Wide Web. Over the following two decades, the Internet evolved into its present-day form.

Most large communications companies that provide Internet service have their own dedicated backbones connecting various regions. In each region, the company has a Point of Presence (POP). Each POP is a place for users to access the company's network, often through a local phone number or dedicated line. Interestingly, there is no overall controlling network. Instead, several high-level networks connect to each other through a Network Access Point (NAP). Each NAP is a physical infrastructure that allows different Internet service providers to exchange traffic between their networks. Dozens of large providers interconnect at NAPs in various cities, and trillions of bytes of data flow between the networks at these points. The Internet is largely a collection of huge corporate networks that all intercommunicate at the NAPs.

What is incredible about the Internet is that a message can leave one computer and travel halfway across the world through several different networks and arrive at another computer in a fraction of a second. To accomplish this feat, all of these networks rely on routers. Routers are specialized computers that have two main functions. First, routers ensure that information makes it to the intended destination by determining where to send it along thousands of pathways. Second, routers make sure that information doesn't go where it's not needed, which is crucial for keeping large volumes of data from clogging the connections of other users. Thus, the router joins the networks so they can communicate, but also protects them from one another.

REPTILES

Most reptiles can be classified into three large groups: the turtles (order Chelonia), the snakes and lizards (order Squamata), and the alligators and crocodiles (order Crocodilia). Most reptiles share a number of general morphological features. In general, reptiles are lung-breathing vertebrates with two pairs of limbs and a horny, scaly skin. Reptiles are amniotes, which means that their large, yolky eggs have a protective layer called an amnion, which prevents them from drying out on land. Rather than laying eggs, some snakes and lizards bear their young live.

All reptiles breathe using lungs. Lung ventilation is accomplished differently in each main reptile group. In squamates, the lungs are ventilated almost exclusively by the axial musculature. This is also the same musculature that is used during locomotion. Because of this constraint, most squamates are forced to hold their breath during intense runs. Some, however, have found a way around it. Varanids, and a few other lizard species, employ buccal pumping as a complement to their normal "axial breathing." This allows the animals to completely fill their lungs during intense locomotion, and thus remain aerobically active for a long time. Crocodilians actually have a muscular diaphragm that is analogous to the mammalian diaphragm. The difference is that the muscles for the crocodilian diaphragm pull the pubis (part of the pelvis, which is movable in crocodilians) back, which brings the liver down, thus freeing space for the lungs to expand. This type of diaphragmatic setup has been referred to as the "hepatic piston."

Reptiles are cold-blooded creatures, which means that they derive their body heat from external sources (in contrast to homothermic animals that maintain a constant body temperature through internal mechanisms). Contrary to popular belief, the "cold-bloodedness" of reptiles does not mean that they maintain low body temperatures. Reptiles control their body temperature through a process of thermoregulation, and their internal temperature can fluctuate greatly according to their surroundings. Researchers have found that many reptiles exert precise control over body temperature by moving around to different areas within their surrounding habitat.

In late fall, reptiles generally begin a process called brumation, a type of dormancy similar to hibernation. However, brumation should bot be confused with hibernation; when mammals hibernate, they are actually asleep and metabolize stores of fat in order to maintain bodily functions and body temperature; when reptiles brumate, they are less active, and their metabolism slows down so they just do not need to eat as often. Reptiles can often go through the whole winter without eating. However, they do need to drink water and will often wake up to drink water and return to "sleep." The brumation period is anywhere from one to eight months depending on the air temperature and the size, age, and health of the reptile. Brumation is triggered by cold weater, lack of heat, and the decrease in the amount of hours of daylight in the winter.

LIVER

The concept that certain organs, such as the liver, brain, and heart, enjoyed a higher status than others was first proposed and accepted in the earliest days of medical thought. Indeed, the Babylonians considered the liver to be the seat and mirror of the soul and, as a consequence, this organ became the focus of divination ceremonies, in which the livers of sacrificial animals were carefully inspected by priests for signs of damage prior to being offered as gifts to the gods. The observed condition of the excised organ was taken to portend the future and, especially, to predict whether or not conditions were favorable for battle. Prayers at these solemn ceremonies were even inscribed on tablets shaped like livers, many of which were subsequently recovered from countries bordering the Mediterranean, far beyond the limits of Babylon.

A multitude of functions of the liver have already been well described, and there are many more of which relatively little is currently known. Several of these functions include detoxification, protein synthesis, and production of biochemicals necessary for digestion. One of the most important – and easily recognizable when deranged – is the metabolism of the pigment, bilirubin, a chemical predominantly derived from products released during the normal destruction of red blood cells. Yellow discoloration of the eyes and the skin (jaundice) ensues when overproduction of bilirubin exceeds the liver’s metabolic capacity or when hepatic metabolism of bilirubin is impaired.

One virus that infects the liver is Hepatitis B. Hepatitis B is an infectious illness caused by hepatitis B virus (HBV) which infects the liver of humans, and causes an inflammation call hepatitis. The acute illness causes liver inflammation, vomiting, jaundice, and rarely, death. Chronic hepatitis B may eventually cause liver cirrhosis and liver cancer – a fatal disease with very poor response to current chemotherapy. The infection is preventable by vaccination, but regardless, about a third of the world’s population, more than 2 billion people, have been infected with hepatitis B virus. This includes 350 million chronic carriers of the virus. Transmission of hepatitis B virus results from exposure to infectious blood or body fluids.

The liver is necessary for survival; there is currently no way to compensate for the absence of liver function long term. However, the human liver is one of the few glands in the body that has the ability to regenerate from as little as 25% of its tissue. This is largely due to the unipotency of hepatocytes. Resection of liver can induce the proliferation of the remained hepatocytes until the lost mass is restored, where the intensity of the liver’s response is directly proportional to the mass resected. For almost 80 years surgical resection of the liver in rodents has been a very useful model to the study of cell proliferation. In is clear that, even though ancient cultures were mistaken as to the functions of the liver, they were certainly correct in attaching so much importance to it.

Indeed, the maxim that ‘life depends on the liver’ is as pertinent today as ever before.

MCCARTHYISM

Originally coined to criticize the anti-communist pursuits of U.S. Senator Joseph McCarthy, “McCarthyism” soon took on a broader meaning, describing the excesses of similar efforts. The term is also now used more generally to describe reckless, unsubstantiated accusations, as well as demagogic attacks on the character or patriotism of political adversaries. During the McCarthy era, thousands of Americans were accused of being Communists or communist sympathizers and became the subject of aggressive investigations and questioning before government or private-industry panels, committees, and agencies. Suspicions were often given credence despite inconclusive or questionable evidence, and the level of threat posed by a person’s real or supposed leftist associations or beliefs was often greatly exaggerated. Many people suffered loss of employment, destruction of their careers, and even imprisonment.

The historical period that came to be known as the McCarthy era began well before Joseph McCarthy’s own involvement in it. Many factors contributed to McCarthyism, some of them extending back to the years of the First Red Scare, inspired by Communism’s emergence as a recognized political force. Thanks in part to its success in organizing labor unions and its early opposition to fascism, the Communist Party of the United States (CPUSA) increased its membership through the 1930s, reaching a peak of about 75,000 members in 1940-41. While the United States was engaged in World War II and allied with the Soviet Union, the issue of anti-communism was largely muted. With the end of World War II, the Cold War began almost immediately, as the Soviet Union installed repressive Communist puppet regimes across Central and Eastern Europe.

The Cold War was the continuing state of political conflict, military tension, proxy wars, and economic competition existing after World War II primarily between the Soviet Union and the United States. Although the primary participants’ military force never officially clashed directly, they expressed the conflict through military coalitions, strategic conventional force deployments, extensive aid to states deemed vulnerable, espionage, propaganda, conventional and nuclear arms races, appeals to neutral nations, rivalry at sports events, and technological competitions such as the Space Race, which began with the launch of Sputnik and culminated in the Apollo Moon landings.

Though McCarthyism might seem to be of interest only as a historical subject, the political divisions it created in the United States continue to make themselves manifest, and the politics and history of anti-Communism in the United States are still contentious. Portions of the massive security apparatus established during the McCarthy era still exist. Loyalty oaths are still required by the California Constitution for all officials and employees of the government of California. A number of observers have compared the oppression of liberals and leftists during the McCarthy period to recent actions against suspected terrorists, most of them Muslims. In *The Age of Anxiety: McCarthyism to Terrorism,* author Haynes Johnson compares the "abuses suffered by aliens thrown into high security U.S. prisons in the wake of 9/11" to the excesses of the McCarthy era. Similarly, David D. Cole has written that the Patriot Act "in effect resurrects the philosophy of McCarthyism, simply substituting 'terrorist' for 'communist.'"

VOLCANOES

A volcano is an opening in the Earth's crust that allows magma and gases from the core to escape. Volcanoes are most commonly found on the edges of tectonic plates and are caused by the gradual divergence and convergence of the plates. It is also possible for volcanoes to arise in the middle of tectonic plates by way of mantle plumes that allow magma to flow to the surface from the core. When plates diverge, magma from the core of the Earth rises to form new ocean floor. This new floor is often thin and the high pressure beneath can cause eruptions. Volcanoes caused by diverging plates are usually underneath the water and simply produce more sea floor. On the other hand, converging plates most frequently involve the subduction of an oceanic plate underneath a continental plate. This produces a large offshore trench through which magma gradually seeps. When the magma makes its way to the surface, the volcano emerges.

Volcanoes are fascinating geological features that interest many people. Thousands of visitors travel to the volcanoes of Hawaii every year to see the incredible sites. Volcanoes are also a very common choice for science fair projects. Elementary and middle school students of been making model volcanoes for many years, and it has become one of the most classic science fair projects. In order to cause the eruption of a model volcano, vinegar is usually combined with baking soda. A common mistake in making a model volcano is using baking powder instead of baking soda. Baking powder does not react with vinegar as quickly as pure baking soda, and baking powder can also start reacting on its own because it contains the acid and base needed for the production of the carbon dioxide.

The eruption of a real volcano is much more spectacular than the eruption of a model volcano. However, real volcanic eruptions are also very dangerous. Not only can the lava kill people, but volcanic ash that accompanies the lava can be hazardous. Volcanic ash consists of small tephra, which are bits of pulverized rock and glass created by volcanic eruptions, less than 2 millimeters (0.1 in) in diameter. The most devastating effect of volcanic ash comes from pyroclastic flows. These occur when a volcanic eruption creates an "avalanche" of hot ash, gases, and rocks that flow at high speed down the flanks of the volcano. These flows can be impossible to outrun.

Many cities have been wiped out by volcanoes, and their threats continue today. The ancient civilization of Pompeii was destroyed in 79 AD by Mount Vesuvius. In 1902, the city of St. Pierre in Martinique was destroyed by a pyroclastic flow which killed over 29,000 people. Mauna Loa in Hawaii has been active for at least 700,000 years. While its most recent eruption in 1984 did not cause any fatalities, the eruptions in 1926 and 1950 destroyed several villages and the city of Hilo.

FLOWERS

Flowering plants can be annual, perennial or biennial. Although flowering plants have a range of life spans and blooming periods, all flowers follow the same growing process. Flowering plants produce male pollen and have female flower parts. As the flower blooms, it produces pollen that's released into the air by rain and wind. The released pollen travels its path and seeks to fertilize the female parts of the flower. Flowers can also be fertilized by bees and other insects. When a bee lands on a flower to obtain nectar from the flower, pollen sticks to the bee and is then transferred to the next flower the bee lands on. Through this process, a flower in one garden could pollinate a flower a mile away!

Flowers, like all plants need sunlight, water, and nutrients to grow. Sunlight is essential for photosynthesis. Photosynthesis is a process that converts carbon dioxide into organic compounds, especially sugars, using the energy from sunlight. Photosynthesis occurs in two stages. In the first stage, light-dependent reactions capture the energy of light and use it to make the energy-storage molecules ATP and NADPH. During the second stage, the light-independent reactions use these products to capture and reduce carbon dioxide.

Water is necessary for all life, but different flowers need different amounts of water. For example, perennial flowers need less water than potted flower. It is important to know how much water flowers need before planting them so the flowers do not get over watered and die. Flowers absorb water through their roots by the process of diffusion. In a garden, the roots of flowers will actually grow in search for water supplies. The proper nutrients for flowers can come from the surrounding environment, but when planting flowers in a garden, it is often helpful to fertilize. Fertilization helps flowers bloom bigger and last longer. Though naturally growing flowers may not have ample fertilization, it's good practice to fertilize flowers at least once or twice each year. Fertilization feeds flowers with the nutrients that soil might not provide. These nutrients include nitrogen, phosphorus and potassium. Nitrogen promotes the growth of foliage and other green structures of the plant. Phosphorus promotes strong root development and flower strength. Potassium promotes the overall health and strength of the entire plant and its flowers. Fertilizer replenishes the surrounding soil and balances the pH levels to complement the flower's acidic requirements.

Pruning- along with proper fertilization, watering and sunlight--promotes vigorous flower growth. Pruning is the process of removing stems, branches and flowers strategically from the plant. When completed successfully, the plant blooms with a plentiful amount of flowers that are of a greater quantity and quality than the previous blooms. While pruning involves removing dead or wilted branches and flowers, new growth can also be eliminated to make room for additional growth.

BALLOONS

A balloon is an inflatable flexible bag filled with a gas, such as helium, hydrogen, nitrous oxide, oxygen, or air. Modern balloons can be made from materials such as rubber, latex, polychloroprene, or a nylon fabric, while some early balloons were made of dried animal bladders. Some balloons are used for practical purposes such as meteorology, medical treatment, military defense, or transportation. Today however, most balloons are used for decorative purposes. Decorative balloons come is all shapes, colors, and sizes. Balloons can be found at almost any celebration, but this decoration necessity would not exist with out the work of Michael Faraday who invented the rubber balloon 1824.

While the most common use for balloons today is decoration, balloons were used for other purposes for many years. The first record of balloons is from 220-80 AD when Zhuge Liang of the Shu Han kingdom used airborne lanterns for military signaling. There is also speculation that the Nazca culture of Peru began using hot air balloon 1500-2000 years ago to design their famous ground lines and figures, the largest of which is 660 feet across. Surprisingly, the balloon is the oldest successful human-carrying flight technology. On November 21, 1783, Jean Francois Pilatre de Rozier and Francois Laurent d'Arlandes made the first hot air balloon trip.

A hot air balloon consists of a bag called the envelope with an opening at the bottom called the mouth or throat. This envelope is capable of containing heated air and is usually made out of light-weight, but strong, synthetic fabric. The fabric is often coated with silicone or polyurethane to make it impermeable to air. Suspended beneath the envelope is a gondola or wicker basket, which carries passengers and a source of heat, in most cases an open flame. At the top of the balloon, there is a vent that enables the pilot to allow hot air to escape through the top of the balloon in order to control the rate of decent. Like any aircraft, it is important for the pilot to make a smooth landing in order to ensure the safety of those on board the aircraft.

Today, hot air balloons are very popular. There are many hot air balloon festivals around the world that millions of people attend each year. At these festivals, balloons of all shapes can be found. The art of hot air balloon design has become increasingly complex and designs now range from the traditional round balloon to hotdogs, flowers, cows and insects. Hundreds of hot air balloons fill the sky at these festivals and many also include balloon races. The first hot air balloon races was the Gordon Bennett Balloon Race, which took place in 1906 in Paris France and was won by Americans Frank Lahm, and Henry Hersey. Lahm went on to become the first Army's first certified pilot in 1909.

Appendix C

Questions and grading criteria used in Experiments 1 and 2

The passages are broken down by item type. Factual, Inferential, and Connecting questions were used in Experiment 1. Example, Inferential, and Connecting questions were used in Experiment 2. Factual and Inferential questions were adapted from Butler (2010), while Example and Connecting questions were written for this study.

NOTE: For the Factual and Example questions, any answer was considered as having come from the correct passage because the correct passage was listed in the question. Strict scoring indicates the criterion for a correct score. Lenient scoring indicates the criterion for partial credit.

FACTUAL

Bats have specially adapted talons that enable them to hang upside down. How do these talons function?

> Answer: A bat must relax its muscles to grip an object, which is the opposite of how human fingers work. The weight of the upper body pulls down on the tendons connected to the talons, causing them to clench and gravity keeps the talons closed.
>
> Strict scoring: Relax muscles; hang upside down, gravity/body weight closes the talons.
>
> Lenient scoring: Mentioning only relaxing muscles, or only gravity- not both

When bats sleep during the day, they enter a torpor state. What happens to bats physiologically when in a torpor state?

> Answer: Bats allow their body temperature to sink to the ambient temperature whenever they are inactive. As their body temperature drops, they enter a torpor state. When in a torpor state, a bat's metabolism slows down, reducing biological activity and conserving energy.
>
> Strict scoring: Body temperature sinks and metabolism slows conserving energy
>
> Lenient scoring: only mentioning one of the above

Vaccines are biological preparations that commonly used in modern medicine. What are the two main ways in which vaccines are used today?

Answer: Most vaccines are used for prophylactic purposes, which means that they prevent or ameliorate the effects of a future infection by any natural pathogen. However, vaccines have also been used for therapeutic purposes, such as alleviating the suffering of people who are already afflicted with a disease.

Strict scoring: To prevent disease and for therapeutic purposes

Lenient scoring: if only give one use

Vaccines vary in terms of their valence. What does the valence of a vaccine refer to?

Answer: The valence of the vaccine refers the number of different antigens contained in the vaccine. A monovalent vaccine is designed to immunize against a single antigen or single microorganism. A polyvalent vaccine is designed to immunize against two or more strains of the same organism, or against two or more organisms.

Strict scoring: Valence= the number of diseases the vaccine treats

Lenient scoring: no partial credit given for this item

Communications companies that provide Internet service to individuals depend on Points of Presence (POPs) and Network Access Points (NAPs). What is the difference between POPs and NAPs?

Answer: A Point of Presence is a place for users to access an Internet service provider's network, often through a local phone number or dedicated line. In contrast, a Network Access Point is a physical infrastructure that allows different Internet service providers to exchange traffic between their networks.

Strict scoring: POP= place for users to access Internet. NAP= place where Internet provides exchange traffic between networks

Lenient scoring: if they get one correct

Routers are crucial to the workings of the Internet. What two main functions do they serve?

Answer: Routers are specialized computers that have two main functions. First, routers ensure that information makes it to the intended destination by determining where to send it along thousands of pathways. Second, routers make sure that information doesn't go where it's not needed, which is crucial for keeping large volumes of data from clogging the connections of other users.

Strict scoring: Make sure information goes to the right place and make sure info does not go where it is not needed

Lenient scoring: if they get one of the two

The part of a tropical cyclone surrounding the eye is called the eye wall. What are the conditions in the eye wall like?

Answer: The area surrounding the eye is called the eye wall, and it consists of a dense wall of clouds and thunderstorms. The eye wall is the part of the storm where the greatest wind speeds are found, clouds reach the highest, and precipitation is the heaviest.

Strict scoring: Dense wall of clouds, highest winds, most precipitation, most intense part of the storm

Lenient scoring: if only mention one thing. Ex: high winds. Need at least 2.

The Radius of Outermost Closed Isobar (ROCI) is a measure of the size of a tropical cyclone. How is ROCI determined?

Answer: The Radius of Outermost Closed Isobar (ROCI) is determined by measuring the radii from the center of the storm to its outermost closed isobar in the four quadrants surrounding the storm. The outermost closed isobar is the point at which the atmospheric pressure returns to normal as it gradually increases from the storm center. The distances of the radii are averaged to come up with a single value.

Strict scoring: Measure radius from eye to where the pressure returns to normal-in 4 quadrants, average.

Lenient scoring: radius from eye to eye wall/ or where pressure returns to normal. (Have to write average the 4 quadrants for full credit)

Flour contains proteins. How do these proteins contribute to the consistency or texture of bread?

Answer: When worked by kneading, the non-water soluble proteins in flour form a network of strands called gluten, which is responsible for the softness of the bread because it traps tiny air bubbles as the dough is baked. If the network of strands is more cohesive or tightly linked, the bread will be softer.

Strict scoring: The proteins create stands that trap air bubbles and create the softness of the bread

Lenient scoring: only mentioning the strands or just saying that they make to bread softer- need both parts to get full credit

Professional bread makers use a system called Bakers' Percentage. How does this system work?

Answer: Bakers' Percentage is a system in which ingredients are measured by weight instead of by volume. Measurement by weight is more accurate and consistent, especially for dry ingredients. Flour is always stated as 100%, and the rest of the ingredients are a percent of that amount by weight.

Strict scoring: Measured by weight instead of volume AND Flour is 100% (everything else is a percentage of the flour)

Lenient scoring: one of the two

In the human respiratory system, a low concentration of oxygen in blood can trigger breathing automatically. How does this occur?

Answer: Low concentration of oxygen in the blood will trigger an override by the autonomic nervous system. Specialized nerve cells within the aorta and carotid arteries called peripheral chemoreceptors monitor the oxygen concentration. If the oxygen concentration decreases, the chemoreceptors signal the respiratory centers in the brain to increase the rate and depth of breathing.

Strict scoring: Low concentration of oxygen triggers special nerve cells (must mention the cells) that tell the brain to increase breathing

Lenient scoring: anything about the diaphragm contracting and expanding, or concentration gradients

There are two main classes of breathing disorders that can affect the human respiratory system. How does each class of disorder affect the respiratory system?

Answer: Disorders of the respiratory system fall mainly into two classes. Some disorders make breathing harder, while other disorders damage the lungs' ability to exchange carbon dioxide for oxygen.

Strict scoring: Making breathing harder and preventing lungs to exchange carbon dioxide for oxygen (gas exchange)

Lenient scoring: one of the two

INFERENTIAL

The U.S. Military is looking for inspiration in developing a new type of aircraft that promotes increased maneuverability. How would this new type of aircraft differ from traditional aircrafts like fighter jets?

Answer: Traditional aircrafts are modeled after bird wings, which are rigid and good for providing lift. Bat wings are more flexible, and thus an aircraft modeled on bat wings would have greater maneuverability.

Strict scoring: More flexible/maneuverable wings (don't have to say bats- this information was clearly from the bats passage)

Lenient scoring: .5= things other than flexibility that still make sense

Correct passage scoring: Bats- if they get the flexible wings, it had to come from Bats

Submarines use sound waves (SONAR) to navigate underwater. Using SONAR, how does a submarine determine that an object is moving towards it (i.e. rather than away from it)?

Answer: The submarine can tell the direction that an object is moving by calculating whether the time it takes for the sound waves to return changes over time. If the object is moving towards the submarine, the time it takes the sound wave to return will get steadily shorter. Also, the intensity of the sound wave will increase because object will reflect more of the sound wave as it gets closer.

Strict scoring: Calculate the time it takes for sound waves to return, if the sound waves return faster, object is moving toward

Lenient scoring: 0.5= if they don't mention how to tell if the object is moving towards them (waves returning faster)

Correct passage scoring: Bats- if mention frequency of echoes, comes from bat passage.

Controlled burning involves setting small fires as a forest management technique. How might this method be used to prevent wildfires?

Answer: Controlled burning involves setting small fires under controlled conditions that eliminate the dry brush that fuels wildfires and limits the risk of the fire spreading out of control.

Strict scoring: Small controlled fires will prevent large uncontrollable fires, pruning

Lenient scoring: .5= destroy brush, create space between potential fires

Correct passage scoring: Mention vaccines or pruning

Research in some fields, such as renewable energy, is not commercially profitable. Where might funding come from to encourage companies to conduct research on the development of things like renewable energy?

Answer: Like vaccine development, research on renewable energy technology relies on "push" funding that is supplied by government, universities, and non-profit organizations.

Strict scoring: The government, universities, nonprofit groups –must mention two or something about the idea of push funding.

Lenient scoring: .5= only mention one (e.g. “government”)

Correct passage scoring: If mention any of these, then came from correct passage.

When engineers move historic buildings from one location to another, it is a challenge to move such a massive object. How do engineers accomplish this daunting task?

Answer: Packet switching is a mode of data transmission in which data is broken into chunks, called packets, which are sent independently and then reassembled at the destination. Engineers use a similar method in which they take apart the building, move the pieces of the building to the new location, and then reassemble them.

Strict scoring: Break the building down into small pieces, move the pieces, and put it back together

Lenient scoring: 0.5= just “piece by piece” or incomplete description

Correct passage scoring: If mention breaking down, get credit

In 1983, an old “radio telephone” patent expired, allowing more companies access to this technology. Why would the expiration of the “radio telephone” patent affect the mobile phone industry?

Answer: The interconnection of the NSFNET to the commercial MCI Mail system in 1988 signaled the opening of the network to commercial interests, which greatly accelerated the expansion of the Internet. Likewise, the expiration of the “radio telephone” patent opened the mobile phone industry to commercial interests and led to its expansion.

Strict scoring: Expiration of the patent opens the industry to commercial interest (competition), which leads to improvement; must mention both commercial interest and improvements

Lenient scoring: 05. = just saying “competition”

Correct passage scoring: Anything about commercial interests

How would hot summer temperatures affect the confined air in car tires?

Answer: Heat in the summer causes oxygen molecules in car tires to expand. However, since there is nowhere for the air to expand, the air pressure increases.

Strict scoring: The heat causes air molecules to expand, increasing the air pressure,

Lenient scoring: 0.5= just saying air expands. Or says that the air pressure would increase but for the wrong reason. Or tire gets larger

Correct passage scoring: Must mention air pressure

In order for a car to run properly, the pistons inside the engine require energy to spin. What might be the process that is responsible for spinning the engine components of a car?

Answer: In a car engine, gasoline is burned inside the cylinders, giving rise to a tremendous amount of heat, and this heat does the work of spinning the engine components.

Strict scoring: Gasoline is burned (or combustion) which creates heat, the heat spins the engine

Lenient scoring: 0.5= just combustion (need to have that the energy to spin the pistons comes from heat)

Correct passage scoring: Must mention heat or energy and spinning

Many products that can be traced to ancient times have been enhanced recently by the use of chemical additives. Paint is one product that has been updated – what function might the chemical additives in paint serve?

> Answer: Quick bread is the name that commercial bakers use for dough that does not require fermentation because of chemical additives, which speeds up mixing time. Similarly, chemical additives are used in paint to speed up the drying time.
>
> Strict scoring: Chemical additives make the paint dry quicker
>
> Lenient scoring: 0.5= Any of the following- last longer, prevent weathering
>
> Correct passage scoring: Have to get that is makes it dry faster just like chemical additives in bread make it rise faster or any other blatant connection

Cladosporium is a type of mold found in the air that can induce asthmatic symptoms in people. While eliminating cladosporium would help to reduce asthma attacks, why would this task be difficult to achieve?

> Answer: Yeast spores naturally occur everywhere, including in the air and many surfaces. Likewise, cladosporium occurs in the air and thus eliminating it would be very difficult.
>
> Strict scoring: Yeast is in the air, it is hard to eliminate because it is everywhere
>
> Lenient scoring: 0.5= it is hard to get rid of things in the air (need to say WHY to get full credit)
>
> Correct passage scoring: Must mention connection to yeast

A bellows is a compressible container with an outlet nozzle that allows a metal worker to manipulate air pressure in order to deliver air in iron smelting. How might a bellows work?

> Answer: Breathing in humans depends on air pressure. Similar to the lungs, when a bellows is expanded, it fills with air (high to low pressure). When a bellows is compressed, it increases the pressure in the bellows above the outside air pressure and the air flows out.
>
> Strict scoring: Changes in air pressure cause the air to go in and out. Low pressure inside cause air to come it, compressing it increases pressure so air goes out.

Lenient scoring: 0.5= if they just say compress the container. Need to say WHY the air moves (difference in air pressure) to get full credit

Correct passage scoring: Anything about air pressure counts

When a cube of sugar is placed into hot tea, the particles dissolve and spread throughout the cup. Why does a sugar cube dissolve in hot tea?

Answer: Within the alveoli, gas exchange occurs through diffusion. Diffusion is the movement of particles from a region of high concentration to a region of low concentration. When a high concentration of sugar (the cube) is placed in hot tea, the sugar molecules with diffuse throughout the water because the concentration of sugar is lower.

Strict scoring: Diffusion- the sugar cube is a high concentration the tea is low concentration of sugar, dissolving makes the concentrations equal

Lenient scoring: 0.5= evening out temperature or breaking down molecules

Correct passage scoring: Have to get diffusion to get correct passage or concentrations

CONNECTING

How do bats survive during winter months?

Answer: Bats enter a state of hibernation by lowering their body temperature. They survive during this time by slowly using up the fat that they built up before the winter months in order to maintain slowed bodily functions.

Strict scoring: Hibernation or torpor state. Must mention at least two things or migration and something that happens during hibernation

Lenient scoring: 0.5= if just say hibernation or migration

Correct passage scoring: Must mention that they live off stores of fat

What is a specific prevention of damage to the liver and how is that prevention designed?

Answer: Hepatitis B is a virus that attacks the liver that can be prevented with a vaccine that is designed through dead or inactivated virulent Hep B organisms that were killed with chemicals or heat.

Strict scoring: Hepatitis B vaccination

Lenient scoring: 0.5= Just mention Hep B or just vaccine

Correct passage scoring: Must mention both hepatitis B and how vaccines are designed

What major global events led to the invention of the internet?

Answer: Throughout the Cold War, there was a Space Race between the Soviet Union and the United States, which eventually led the United States to establish the Advanced Research Projects Agency (ARPA) in order to gain a technological lead.

Strict scoring: The launch of Sputnik, Cold war, Space Race; must mention 2 of these

Lenient scoring: 0.5= deprivitization of the internet; just Sputnik or just Cold War

Correct passage scoring: Must mention the Cold War or Space Race to get credit. Just Sputnik doesn't count

What causes hot air balloons to rise?

Answer: The burner, or heat source, heats the air with in the balloon. The air molecules then expand which lowers the density of the air within the balloon. When the air within the balloon is lower than the air surrounding the balloon, the balloon with rise.

Strict scoring: Air inside balloon is heated by heat source (open flame). Less dense air inside; decreased air pressure

Lenient scoring: 0.5= hot air, or expansion of air

Correct passage scoring: Must mention why hot air rises- molecules expand/less dense; air pressure

How do chemical leavening agents make bread rise?

Answer: Baking powder, the most common chemical used to leaven bread contains both acid and base that react to produce carbon dioxide.

Strict scoring: Chemical leaving agents (baking powder) release gas

Lenient scoring: 0.5= form air pockets, works instantly

Correct passage scoring: Releases carbon dioxide specifically

How do plants get water?

Answer: Plants absorb water through their roots by the process of diffusion. In diffusion, molecules pass through membranes from higher concentration to water concentration. When a plant is watered, there is a higher concentration of water molecules around the roots than inside the roots, so water goes into the root and can then be used throughout the plant.

Strict scoring: Through their roots, by diffusion- high concentration to low concentration

Lenient scoring: 0.5= rain, by watering, just through roots

Correct passage scoring: Diffusion

EXAMPLE

How large is the largest of the Nazca ground lines?

Answer: 660 feet across

Strict scoring: 600 or 660 feet across

Lenient scoring: switch to miles or something close to 660 (e.g. 6600)

What is the name of the thin membrane of skin found on a bat's wing?

Answer: Patagium

Strict scoring: Very close to patagium in spelling and/or pronunciation

Lenient scoring: Words that start with "p" but do not look much like patagium

What is the percent of water in most artisan bread?

Answer: 60-75%

Strict scoring: 60 or 60% to 70 or 75%

Lenient scoring: must be close to range (down to 50, up to 80); only mentioning one end of range

What energy-storage molecules are created during the first stage of photosynthesis?

Answer: ATP and NADPH

Strict scoring: close to correct on both

Lenient scoring: only one correct

What was the name of the project leader at ARPA who first explored the use of packet switching?

Answer: Joseph Licklider

Strict scoring: Something that is close to Licklider in spelling or pronunciation

Lenient scoring: Joseph only or “starts with L”

What percentage of liver tissue is made up of hepatocytes?

Answer: 70-80%

Strict scoring: 70 or 75% to 80%

Lenient scoring: one number within range

Who wrote The Age of Anxiety: McCarthyism to Terrorism?

Answer: Haynes Johnson

Strict scoring: Either Haynes or Johnson

Lenient scoring: anything that looks close to Haynes or “starts with a J”

What is the name of the protective layer, which covers the eggs of reptiles?

Answer: Amnion

Strict scoring: Amnion (misspellings allowed)

Lenient scoring: words that start with “a” and are somewhat close to amnion

What is an example of a disorder that minimizes or prevents gas exchange in the lungs?

Answer: Pulmonary edema

Strict scoring: Pulmonary edema (misspellings allowed)

Lenient scoring: close to or just pulmonary; not asthma

How large is the eye of most tropical cyclones?

Answer: 20-40 miles across

Strict scoring: 20-40 miles, meters in diameter/radius

Lenient scoring: only one number right or listed, close to range but wrong numbers

What is an example of an aluminum-based adjuvant?

Answer: squalene

Strict scoring: squalene or something close to it

Lenient scoring: other words that start with “s” or a definition instead of an example

How many people in Martinique were killed in a 1902 volcanic eruption?

Answer: over 29,000

Strict scoring: anything between 25,000 and 30,000

Lenient scoring: 20,000 or a rearrangement of numbers, e.g. 2090

Appendix D

Function learning classification guide used in Experiments 1 and 2

This guide was used to re-analyze the data from Experiments 1 and 2 according to an extreme groups approach.

Overall Guidelines

1. Before beginning classification, be sure the graphs have been randomized and that you have something to cover the middle section (training range) of each graph. This is to ensure that you are not thrown off by patterns that might emerge between training and extrapolation. We are interested in extrapolation performance only.
2. On each graph, cover the training section (80-120) and examine the patterns of both left and right extrapolation. When you feel confident about your response, sort the graphs into three piles: exemplar, rule, or ambiguous (details below).

Rule learners

1. Rule learners are typically defined as anyone who has a clear negative slope on the left and positive slope on the right side of the function. The slopes may range from exactly on the V-shaped function to much lower slopes, so long as it is clear that both sides are sloping in the correct direction according to the underlying function.

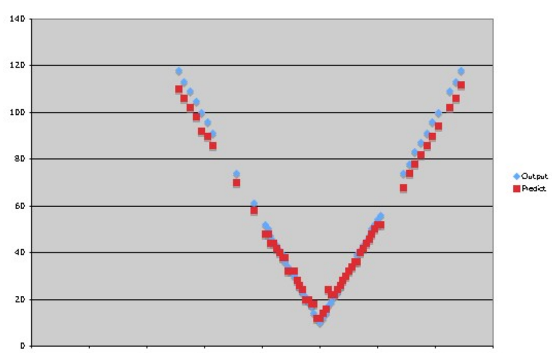

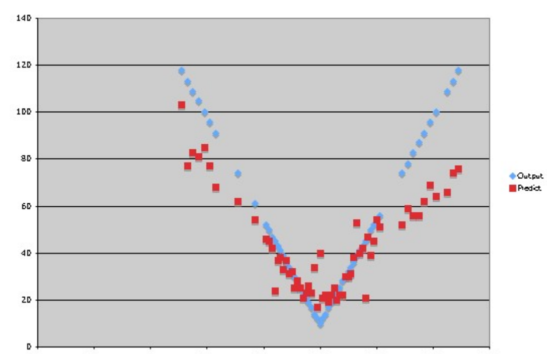

2. If the pattern *clearly* follows a sine function (oscillating with very little spread/scatter, showing distinct lower vertices), this is also classified as a rule learner.

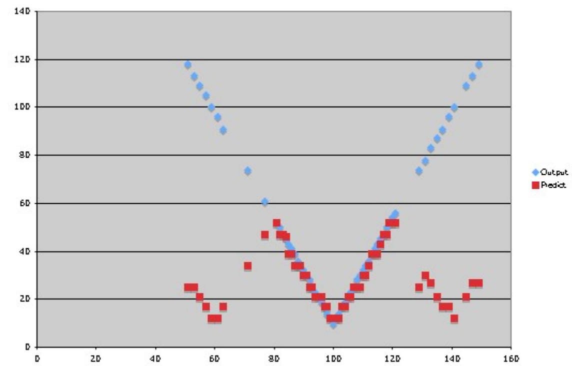

3. If there are *very distinct* downward slopes on both sides (with very little spread, not scattered, very linear), this is also classified as a rule learner.

Exemplar learners

1. Exemplar learners are typically defined as those individuals showing flat extrapolation (i.e. near a zero slope on both sides).

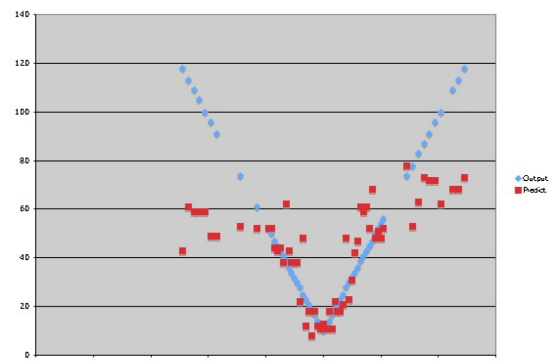

2. Exemplar learners may also show extrapolation that has a large spread or looks scattered (i.e. no pattern to the data), so long as there is not a clear upward or downward linear slope.

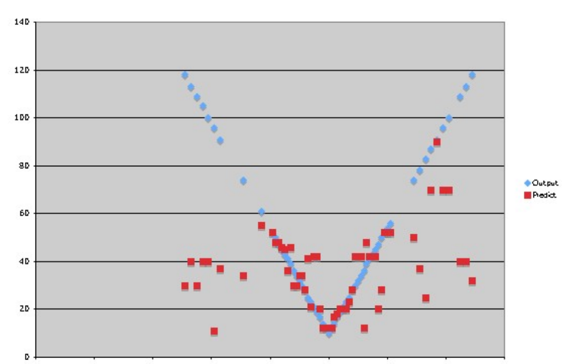

3. A small amount of positive slope is allowed on the *right side only*. If the slope is close to that of the function, it does not count as a small positive slope.

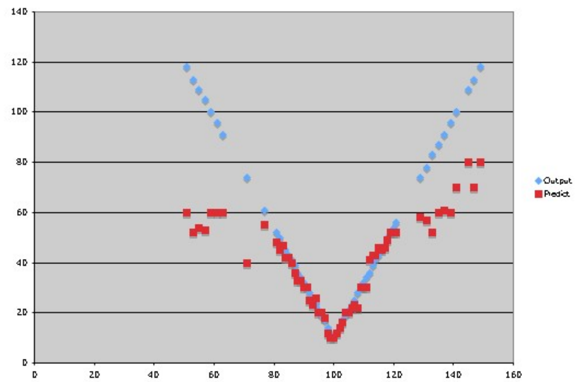

Ambiguous

1. If the data does not fall into one of the distinct categories above it is considered ambiguous. There are many possible patterns of data that could be considered ambiguous. A few of the more common include:
   a. If the slope on both the left and right sides are clearly positive.

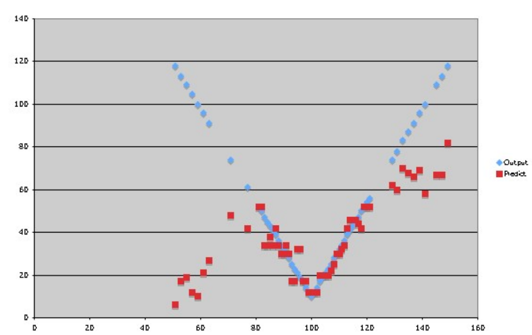

   b. If it is difficult to tell if the slope is flat or negative, flat or positive, scattered or linear. When in doubt, call it ambiguous.
   c. If one side closely approximately the slope of the function and the other is flat/scattered.
2. Less common examples include:
   a. If there are multiple discrepant points: most of the points show flat extrapolation, but a few closely approximate the function or vice versa.

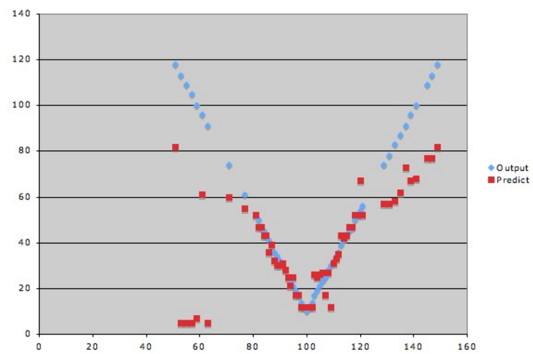

b. If the extrapolation patterns are parallel (both positive, but in the same range)
c. Any other “weird” pattern to the extrapolation data.

Appendix E

Concept learning stimuli from one of the counterbalances used in Experiment 2.

Training Items Recognition Lures

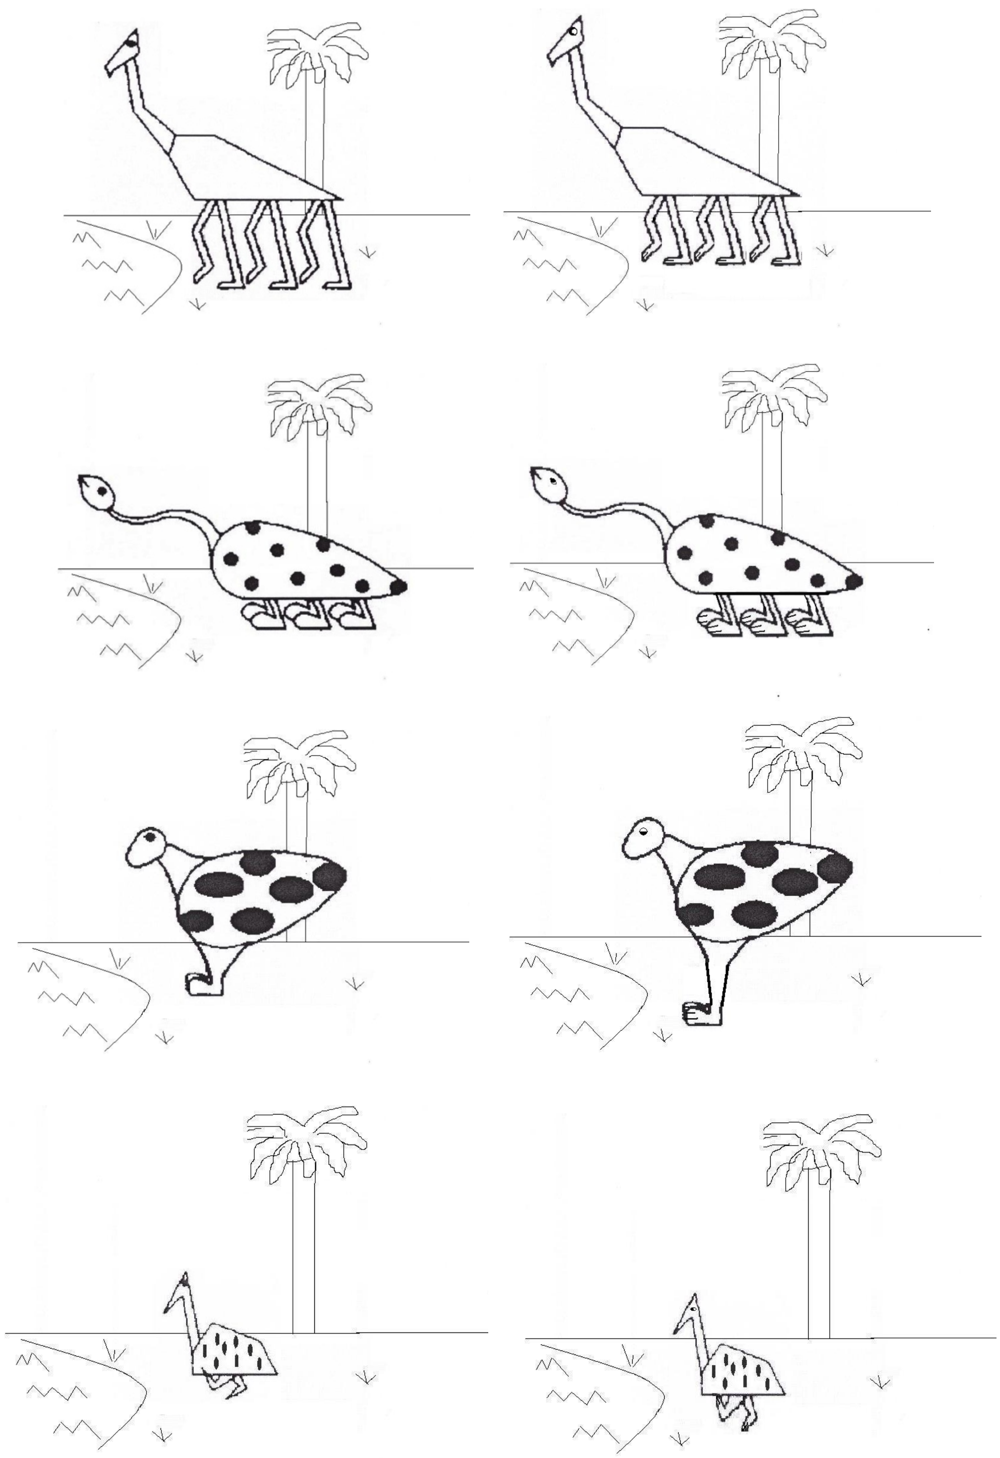

Training items Categorization Lures