Published in final edited form as:

Acta Psychol (Amst). 2013 November ; 144(3): . doi:10.1016/j.actpsy.2013.09.005.

Differential impact of relevant and irrelevant dimension primes on rule-based and information-integration category learning

Lisa R. Grimm^{a,1} and W. Todd Maddox^b

^aDepartment of Psychology, The College of New Jersey

^bDepartment of Psychology, University of Texas at Austin (Maddox)

Abstract

Research has identified multiple category-learning systems with each being “tuned” for learning categories with different task demands and each governed by different neurobiological systems. Rule-based (RB) classification involves testing verbalizable rules for category membership while the information-integration (II) classification requires the implicit learning of stimulus-response mappings. In the first study to directly test rule priming with RB and II category learning, we investigated the influence of the availability of information presented at the beginning of the task. Participants viewed lines that varied in length, orientation, and position on the screen, and were primed to focus on stimulus dimensions that were relevant or irrelevant to the correct classification rule. In Experiment 1, we used an RB category structure, and in Experiment 2, we used an II category structure. Accuracy and model-based analyses suggested that a focus on relevant dimensions improves RB task performance later in learning while a focus on an irrelevant dimension improves II task performance early in learning.

Keywords

categorization; memory systems; dimension relevance; priming; category learning

1. Introduction

Categorization is ubiquitous in human thinking, and as such, has been extensively studied by psychologists interested in a wide-variety of processes from language to object recognition to reasoning and decision-making. Some early work on basic categorization focused on the presence of a single categorization system (Nosofsky & Johansen, 2000) while other work demonstrated the existence of multiple memory systems (Ashby & Ell, 2001; Ashby & O'Brien, 2005; Ashby & Maddox, 2005). Specifically, researchers argued that there was an explicit-hypothesis testing system that was recruited to process rule-based (RB) classification and another implicit-hypothesis testing system that was recruited to process information-integration (II) classification. RB classification tasks are constructed to involve rules for category membership that are verbalizable (Ashby, Alfonso-Reese, Turken, & Waldron, 1998). Because the classification rule is verbalizable, the optimal strategy is to test and discard rules until the correct rule is discovered and classification accuracy improves. For example, if short, shallow-oriented lines are in category A and all others are in category

¹This research was supported by NIMH grant R01 MH077708 and NIDA grant DA032457 to WTM. We thank Benjamin Lewis, Erin Haughee, Jesse Taylor, Daniel Carlin, Janika Berridge, Kristin Martin, Kristen Duke, Micajah Spoden and Seth Koslov for assistance with data collection. We also benefited from helpful conversations with Jonathan Rein.

Address correspondence to Lisa R. Grimm, Department of Psychology, The College of New Jersey, P.O. Box 7718, Ewing, NJ 08628-0718. Send electronic mail to grimm@tcnj.edu.

B then this classification task might be solved by first trying a rule on length, and ultimately discarding this rule for a rule that makes a decision on length and a decision on orientation and then combines these decisions to generate the correct categorization response. In contrast, II classification tasks involve the predecisional integration of information across dimensions (Ashby & Gott, 1988) and therefore the optimal classification rules are not verbalizable (Ashby et al., 1998). For example, a possible II task rule could require that participants place lines that are longer than they are steep into a category. In this case the participant would rely on the implicit-learning system to slowly learn the association between the stimulus and the appropriate response.

There is a huge literature on the dissociations between RB and II classification learning (e.g., Ashby, Queller, & Berretty, 1999; Ashby, Maddox, & Bohil, 2002; Ashby, Ell, & Waldron, 2003; Maddox, Ashby, & Bohil, 2003; Maddox & Ing, 2005) and the different neurobiological systems that may be recruited for explicit and implicit hypothesis testing (e.g., Rao, Bobholz, Hammeke, Rosen, Woodley, Cummingham, et al., 1997; Maddox & Filoteo, 2001; Seger & Cincotta, 2002; Filoteo, Maddox, Salmon, & Song, 2005; Filoteo, Maddox, Simmons, Ing, Cagigas, Matthews, & Paulus, 2005). Most relevant to the current studies is the literature that focuses specifically on the impact of working memory. RB classification is thought to recruit the explicit-hypothesis testing system and actively engage working memory to generate and apply candidate classification rules (Ashby et al., 1998). In contrast, II classification does not rely on working memory because learning occurs below conscious awareness and is the result of stimulus-response association learning (Maddox, Ashby, Ing, Pickering, 2004; Zeithamova & Maddox, 2007). Maddox et al. (2004) and Zeithamova and Maddox (2006; 2007) demonstrated that adding a verbal or visuospatial working memory load decreases RB but not II learning. In Maddox et al., participants viewed a classification stimulus, made a classification judgment, and received corrective feedback. Next, they completed a sequential verbal working memory task in which a digit memory set (set size 4) was presented and followed by a memory probe (i.e., “Was this item in the memory set?”). Zeithamova and Maddox (2007) adapted this working memory task by requiring that participants remember object locations instead of digits, which transformed the task from a verbal working memory task to a visuospatial task.

While tempting to assume that working memory loads will always show this dissociation between RB and II learning, there is an interesting set of predictions generated by the COmpetition between Verbal and Implicit Systems (COVIS) model (Ashby et al., 1998; Ashby & Maddox, 2011) argued to underlie RB and II learning. COVIS postulates that II tasks recruit visual cortical areas and the tail of the caudate nucleus, while RB tasks use the prefrontal cortex, the head of the caudate nucleus, and the anterior cingulate. It is believed that both the explicit and implicit systems are initially recruited to solve a classification problem (Ashby & Maddox, 2005; Zeithamova & Maddox, 2006). Each system generates a response and response selection is governed by the weight of the responses, determined by the past success of responses from that system. Moreover, participants are initially inclined to favor responses from the explicit system. Filoteo, Lauritzen, and Maddox (2010) tested the interesting prediction that adding a working memory task would engage working memory and make participants less likely to rely on rules which would also require working memory. This interference would result in participants disengaging from the explicit system, and allowing the implicit system to learn the classification task and operate without competition. Consistent with their prediction, they found that the working memory task impaired RB learning but improved II learning. This result is consistent with work by DeCaro, Thomas, and Beilock (2008) that examined the impact of working memory capacity. They found that individuals with a lower working memory capacity performed better on II learning tasks but worse on RB tasks relative to those with a higher working memory capacity.

All of the prior work examining working memory capacity has used traditional working memory dual tasks. Our studies ask a different question: Does the content of the verbal working memory store matter? Instead of relying on working memory tasks that are unrelated to the classification task, we used a method that allows us to directly compare content that is relevant and irrelevant to task performance. Our experiments are the first to consider the influence of dimension priming on the learning of RB and II category structures. In Experiment 1, we use a perceptual RB task, and in Experiment 2, we use a perceptual II task. For each task, participants viewed lines that varied in length, orientation, and position on the screen. A conjunctive RB rule (i.e., lines that are long and steep are in one category) could be used to perfectly classify the stimuli in Experiment 1, while in Experiment 2, an information-integration rule (i.e., lines that are longer than steep are in one category) could be used to perfectly classify the stimuli.

To examine the influence of dimension primes on each system, we varied the information available at the start of the experiment. Participants were told that a focus on the position, length, or orientation of the lines led to good performance in prior participants; control participants were told nothing. This simple manipulation allowed us to examine whether RB and II tasks would be differentially impacted when participants are told to focus on dimensions relevant or irrelevant to the classification rule.

Given the prior research on working memory, the predictions for Experiment 1 are fairly straightforward. We predicted that RB learning would benefit from a focus on relevant dimensions (i.e., length and orientation) but not from a focus on the irrelevant dimension (i.e., position). A focus on an irrelevant dimension is analogous to the prior work using unrelated working memory loads and therefore should hurt RB task performance. The participants should fill the verbal working memory store generating and testing rules that are not helpful to performance. In contrast, a focus on relevant task dimensions should help performance because the verbal store is being recruited for activities that directly benefit task performance.

There are two possible outcomes for Experiment 2 based on the prior literature. One possibility is that II learning will be unaffected by providing participants with a focus on relevant or irrelevant task dimensions. Simply, because II learning is thought to proceed without relying on working memory resources, the addition of content to working memory should not impact performance. This is consistent with several studies (Maddox et al., 2004; Zeithamova & Maddox, 2006; 2007) demonstrating that adding a verbal working memory load did not affect II learning. In contrast, the work by Filoteo et al. (2010) suggests that adding irrelevant content to working memory improves II learning and DeCaro et al. (2008) demonstrated that a larger working memory capacity resulted in worse II learning. If one assumes that individuals with a larger working memory capacity focused on content relevant to task performance, we should find a corresponding decrease in II performance when we ask participants to focus on relevant task dimensions. In contrast, consistent with Filoteo et al., asking participants to focus on the irrelevant task dimension of position would fill the verbal working memory store and allow for a disengagement of the explicit system, thereby permitting the II system to learn and generate the correct stimulus-response mappings that involve the task-relevant length and orientation dimensions. As such, we predicted that the opposite pattern would be true for II learning when compared to RB learning. II learning should benefit from a focus on an irrelevant dimension. This prediction is consistent with findings in Ashby and Crossley (2010) that using an RB explicit strategy may limit the use of the II system.

2. Experiment 1

2.1. Material and Methods

2.1.1. Participants and Design—One hundred forty-two undergraduate students at The College of New Jersey participated for course credit. Participants were randomly assigned to one of four dimension-prime conditions (Position, Length, Orientation, or Control) yielding a 4 Condition $\times$ 12 Block mixed-factorial design with Condition between participants and block within participants.

2.1.2. Materials—Dimensions were primed by manipulating the instructions participants read prior to completing the category learning task. All groups received basic instructions explaining the nature of the task (e.g., viewing lines that vary in length, orientation, or position) and the goal of the task (e.g., to learn how to correctly classify the stimuli into two categories). The Control group received no additional instructions. For the other 3 prime groups, participants received an additional hint to focus on a specific dimension. For example, for the Length group, participants were told that prior participants found that creating rules using the length of the line led to good task performance. Corresponding hints were presented to the Orientation and Position groups.

2.1.3. Stimuli and Stimulus Presentation—Participants viewed stimuli on a computer screen and were asked to classify a set of items into one of two categories. The stimuli to be categorized were lines that varied across items in their length, orientation, and position within a box on the screen. For Category A, there were 24 stimuli sampled from each of 12 bivariate normal distributions on length and orientation resulting in 288 stimuli. For Category B, there were 72 stimuli sampled from 4 bivariate normal distributions on length and orientation resulting in 288 stimuli. The position dimension was sampled independently of length and orientation for each category: Category A used a univariate normal distribution with a mean of 253 pixels and a standard deviation of 75 and Category B used a univariate normal distribution with a mean of 397 pixels and a standard deviation of 75. The stimulus structure is shown in Figure 1. The lines were presented inside of a black 650 $\times$ 650 pixel box, centered vertically (see Figure 2), and were randomly ordered for each participant in each block. There were 48 trials in each block and 12 blocks.

The stimuli were generated such that using the orientation of the line or the length of the line to classify the stimuli will result in 83% accuracy for a block of trials, while using position would result in 50% accuracy for a block of trials. These unidimensional rules are fairly easy to verbalize and are salient to participants (Maddox, Baldwin, & Markman, 2006). However, there is an optimal decision bound for this task that, if used, will yield 100% accuracy on the task. This decision criterion requires a rule that takes into account both length and orientation. This rule is: If the length is long and the orientation is steep, then respond Category A; otherwise, respond Category B (see Figure 1 for a graphical representation of this rule). In order for participants to perform well in the task, they need to abandon the use of unidimensional rules in favor of the more complex conjunctive one.

2.1.4. Procedure—Participants were tested in individual cubicles. Participants first read the basic task instructions and the hint (if in either the Length, Orientation, or Position group). Each participant completed 12 blocks of trials with 48 trials. For each trial, the stimulus was displayed until the participant responded “A” or “B”. After the response, the stimulus disappeared for 500 ms and then the participants received corrective feedback (Correct or Incorrect with the correct category identified). Following feedback, the screen turned black for 250ms for the inter-trial-interval.

2.2. Results

First, we present analyses from our category learning task that uses proportion correct as the dependent variable. For all analyses, we use an alpha level of .05 to determine statistical significance and two-tailed *p*-values are reported unless otherwise noted. Our post hoc tests use Bonferroni corrections. Second, we present model-based analyses where we report the results of computational modeling to characterize the rules used by participants on a block-by-block basis. To foreshadow, these results are consistent with the proportion correct data and support the effectiveness of our dimension primes.

2.2.1. Task-based Analyses—The data were analyzed using an Analysis of Variance (ANOVA) with Condition (Position, Length, Orientation, and Control) as a between-participants' factor and Block (12) as a within-participants' factor. The dependent measure was the proportion of correct responses in each block of trials. This analysis revealed that the accuracy of classification improved over time (main effect of Block, $F(11,1518) = 26.85$, $p < .001$, partial $\eta^2 = .16$). Also, Length priming resulted in higher accuracy than Position priming; there was a marginally significant main effect of Condition, $F(3,138) = 2.20$, $p = .091$, partial $\eta^2 = .05$. Post hoc tests revealed that only Position and Length were reliably different, $p = .013$ (Position ($M = .80$, $SD = .07$); Length ($M = .85$, $SD = .05$); Orientation ($M = .82$, $SD = .11$); Control ($M = .83$, $SD = .05$)). Lastly, there was a two-way interaction between Condition and Block, $F(33,1518) = 1.57$, $p = .022$, partial $\eta^2 = .03$.

To examine this interaction, we considered whether our effect emerged at different stages of learning by looking at the effect of Condition within Blocks. First, we examined the effect of Condition within every Block and then considered Condition differences within learning stages. The position group performed worse than the length group in 5 blocks and also worse than the orientation group in the last block of trials (all $p < .05$) (the main effect of Condition appeared in Blocks 5, 6, 7, 11, and 12, all $p < .05$). To identify learning stages, we both considered prior work using a similar task (Grimm, Markman, Maddox, & Baldwin, 2009) and our current pattern of data. Using both prior and current patterns of data, we assume that the first 4 blocks of trials represents early learning because it is clear that the pattern shifts after the first 4 blocks (see Figure 3). Therefore, we analyzed the first 4 blocks of trials and then we analyzed the last 8 blocks of trials.

For the first 4 blocks, the data were analyzed using an ANOVA with Condition (Position, Length, Orientation, and Control) as a between-participants' factor and Block (4) as a within-participants' factor. The dependent measure was the proportion of correct responses in each block of trials. Participants improved over the first 4 blocks (a main effect of Block, $F(3,414) = 41.81$, $p < .001$, partial $\eta^2 = .23$). In fact, post hoc tests revealed that all blocks were different from each other, all $p < .001$, with the exception of blocks 3 and 4 which were not different from each other, $p = .143$. There was neither a significant main effect of Condition, $F(3,138) = 0.54$, $p = .658$, partial $\eta^2 = .01$, nor a two-way interaction between Condition and Block, $F(9,414) = 0.27$, $p = .265$, partial $\eta^2 = .03$.

For the last 8 blocks, the data were analyzed using an ANOVA with Condition (Position, Length, Orientation, and Control) as a between-participants' factor and Block (8) as a within-participants' factor. Participants improved only slightly over the last 8 blocks (marginally significant main effect of Block, $F(7,966) = 1.91$, $p = .065$, partial $\eta^2 = .01$). The Position group performed worse than the Length group and the Control group (main effect of Condition, $F(3,138) = 3.01$, $p = .033$, partial $\eta^2 = .06$) with post hoc tests revealing that Position was reliably different from Length, $p = .004$, and marginally different from Control, $p = .054$ [Position ($M = .81$, $SD = .07$); Length ($M = .87$, $SD = .06$); Orientation ($M = .84$, $SD = .12$); Control ($M = .85$, $SD = .06$)]. Lastly, there was not a two-way interaction between Condition and Block, $F(21,966) = 1.19$, $p = .252$, partial $\eta^2 = .03$.

2.2.2. Model-based Analyses—An advantage of using this classification task is that we have computational models that allow us to characterize participants' responses on a block-by-block basis. Models allow us to determine the types of strategies used by participants during classification learning as multiple different strategies can yield the same accuracy rate. We predict that participants will start our task by generating data consistent with a rule that matches the hint they were provided. For the Control group, following Maddox, Baldwin, & Markman (2006), we hypothesize that participants will start with a simple unidimensional rule on position to classify the stimuli, as this has been demonstrated to be the most salient dimension. Moreover, because we are using a verbalizable rule-based task, we predict that those primed with Length or Orientation as beneficial will be more likely to have their data be consistent with the more complex and optimal conjunctive-rule on length and orientation by the end of the task.

To test these hypothesis, we fit a series of decision-bound models to the data for each participant for each block (Ashby & Maddox, 1993; Maddox & Ashby, 1993). The unidimensional model on position assumes that the participant used a criterion on position and put all of the lines to the left in one category and all of the lines to the right in the other category. The unidimensional model on orientation assumes that the participant's criterion involved one response for shallow lines and another response for steep lines. The unidimensional model on length assumes one response for short lines and another response for long lines. Each of these unidimensional models uses two free parameters: one decision criterion and one noise parameter. The conjunctive model assumes that the participant used length and orientation. We fit two different conjunctive models. First, we fit an optimal model which assumes the participant used the optimal criterion on both length and orientation. This model only has one free noise parameter. Second, we fit a suboptimal model which assumes that the participant used criteria on both length and orientation but these criteria were not optimal. Therefore, this model has three free parameters: one for the length criterion, one for the orientation criterion, and one noise parameter.

The model parameters were estimated using maximum likelihood (Ashby, 1992). We found the best fitting model using: $AIC = 2r - 2\ln L$ (Akaike, 1974; Takane & Shibayama, 1992) where r is the number of parameters in the model and $\ln L$ is the log likelihood of the model given the data. This criterion allows us to assess the goodness-of-fit of models that differ in the number of free parameters, and select the model that provides the most parsimonious account of the data (i.e., the model with the smallest AIC value).

First, to assess the extent to which participant data initially corresponded to one of the rules identified in the hints, we considered the proportion of participants whose first block of trials was consistent with unidimensional rules on position, length, or orientation. As can be seen in Table 1, participants' data in the Position and Control groups was consistent with early position rule use, while the length rule was most often consistent with participant data in the Length group and the orientation rule was most often consistent with participant data in the Orientation group.

Interestingly, many participants were fit by more complicated conjunctive rules even in the first block of trials. That said, these conjunctive rules still contained the dimension initially primed. For the Position group, 39% were best fit by a conjunctive model using position (length and position or orientation and position), and thus 78% of participants initially used a rule on position. For the Length group, 62% were best fit by a conjunctive model using length (length and position or length and orientation), and thus 84% of participants initially used length to create a rule. For the Orientation group, 21% was best fit by a model on orientation and position or orientation and length, and thus 63% initially relied on orientation. Lastly, for the Control, 54% were best fit by a conjunctive model.

Of importance to our initial hypothesis, we considered the discovery and use of the optimal classification rule on length and orientation by participants across our 4 groups and across trials. First, we considered the proportion of participant data accounted for by the optimal classification rule. Averaging across all blocks, the optimal rule accounted for 88% of the responses for participants in the Length group, 86% of the responses for the Orientation group, 84% of the responses for the Control group, and 82% of the responses for the Position group. As with the accuracy data, this pattern is representative of the last 8 blocks of trials and is even more pronounced when only the last block is considered (90% for the Length group, 89% for the Orientation group, 86% for the Control group, 83% for the Position group). Second, using AIC fits, we analyzed which model was selected as the model best fitting participant data for each participant for each block of trials. As shown in Figure 4, participants, initially primed with length and orientation, were most likely to have their data increasingly be consistent with the optimal rule on length and orientation. In contrast, participants primed with position were unlikely to find and use the correct strategy. Comparing the Position group to each of the other groups using sign tests, the Position group was less likely to be fit by the optimal model (less on all 12 blocks of trials as compared to each of the other groups), all $ps < .05$. Comparing the Control group to the Length and Orientation groups separately using sign tests, the control group was also less likely to be fit by the optimal model (less on 11 of 12 blocks), both $ps < .05$.

2.3. Discussion

Participants learned to classify simple stimuli into two categories. The predetermined correct rule used length and orientation dimensions (i.e., lines that were both long and steep belonged to a category) and therefore the position dimension was irrelevant to the task. We hypothesized that in a rule-based classification task participants would benefit from receiving a focus on dimensions relevant to the classification rule and would not benefit from receiving a focus on the irrelevant dimension. Moreover, we anticipated that these effects would appear later in learning because the participant needs to test and discard unidimensional rules in favor of the correct conjunctive rule which will need to be fine-tuned over time to optimize performance. Our predictions are supported by the data. We did not find effects of prime early in learning, but did find effects later in learning; that is, participants in the Position group did not perform as well as the other groups. Moreover, modeling shows that participants' data reflected the optimal model most frequently for the Length and Orientation groups and least frequently for the Position group.

3. Experiment 2

Experiment 2 uses an information-integration category structure (e.g., Maddox & Ashby, 2004), and a classification task that requires that participants learn to classify lines into two categories. As shown in Figure 5, the stimulus dimensions are the same as those in Experiment 1, but in this case a plane cuts the stimulus space on the diagonal creating a rule that is difficult to verbalize using length and orientation (e.g., a stimulus goes in a category because it is longer than it is steep). This rule is hard to verbalize because the dimensions are measured in different units. As demonstrated previously by Maddox, Baldwin, and Markman (2006), this structure requires that participants abandon the explicit rule-based system in favor of the implicit learning system to achieve a high level of performance.

In this experiment, we examine the implications of initially focusing participants on dimensions that are relevant or irrelevant to the classification rule. Unlike Experiment 1, we predict that a focus on length and orientation will be detrimental to performance. This prediction is consistent with the finding from DeCaro et al. (2008) that participants, likely using their larger working memory capacity to focus on task-relevant information, performed worse on an II classification task. Moreover, we predict that a focus on position

will improve performance, which is consistent with Filoteo et al. (2010). Filoteo et al. demonstrated that creating a working memory load, which was irrelevant to task performance, improved II learning. Therefore, we predict the opposite pattern of data from Experiment 1.

3.1. Material and Methods

3.1.1. Participants and Design—One hundred eleven undergraduate students at the University of Texas at Austin participated for course credit. As in Experiment 1, participants were randomly assigned to one of four dimension-prime conditions (Position, Length, Orientation, or Control) yielding a 4 Condition $\times$ 12 Block mixed-factorial design with Condition between participants and block within participants.

3.1.2. Stimuli and Stimulus Presentation—Other than the change in stimuli to form categories consistent with an information-integration rule, stimuli selection and presentation were identical to Experiment 1. The accuracy of rules on the various dimensions was held constant across experiments.

3.1.3. Materials and Procedure—The Materials and Procedure were identical to Experiment 1.

3.2. Results

As in Experiment 1, we first examine the proportion of correct responses and then present model-based analyses.

3.2.1. Task-based Analyses—The data were analyzed using an ANOVA with Condition (Position, Length, Orientation, and Control) as a between-participants' factor and Block (12) as a within-participants' factor. The dependent measure was the proportion of correct responses in each block of trials. Participants improved over the course of the experiment (main effect of Block, $F(11,1177) = 29.48$, $p < .001$, partial $\eta^2 = .22$). Also, the Position group performed better than the Orientation group (marginally significant main effect of Condition, $F(3,107) = 2.57$, $p = .058$, partial $\eta^2 = .07$). Post hoc tests revealed that Position and Orientation were reliably different, $p = .007$ (Position ($M = .83$, $SD = .05$); Length ($M = .81$, $SD = .06$); Orientation ($M = .78$, $SD = .11$); Control ($M = .80$, $SD = .07$)). Lastly, as in Experiment 1, there was a two-way interaction between Condition and Block, $F(33,1177) = 1.74$, $p = .006$, partial $\eta^2 = .05$.

To examine this interaction consistent with Experiment 1, we examined the effect of Condition within every Block and then considered whether our effect occurred at different stages of learning. The Position group performed better than all other groups in Block 1, better than the Orientation group in Blocks 2 and 3, better than the Control and Length groups in Block 4, and better than the Orientation group in Block 6 (all $p < .05$ for main effects tests for Blocks 1, 2, 3, 4, and 6, and post hoc tests, except for main effect in Block 4, $p = .186$, and Position vs. Length post hoc in Block 4, $p = .065$). Next, we examined learning stages and analyzed the first 4 blocks of trials and the last 8 blocks of trials separately. As in Experiment 1, our assumption that the first 4 blocks of trials represents early learning is nicely represented in the data as it is clear that the pattern shifts after the first 4 blocks (see Figure 6).

For the first 4 blocks, the data were analyzed using an ANOVA with Condition (Position, Length, Orientation, and Control) as a between-participants' factor and Block (4) as a within-participants' factor. Participants improved over the first 4 blocks of trials (main effect of Block, $F(3,321) = 44.90$, $p < .001$, partial $\eta^2 = .30$). In fact, post hoc tests revealed that all

blocks were different from each other, all $p < .01$. There was also a significant main effect of Condition, $F(3,107) = 5.66$, $p = .001$, partial $\eta^2 = .14$. Post hoc tests revealed that the Position group performed reliably better than participants from all of the other conditions, all $p < .02$, and the differences between Orientation and Control ($p = .084$) and Orientation and Length ($p = .090$) were marginally significant (Position ($M = .81$, $SD = .03$); Length ($M = .76$, $SD = .08$); Orientation ($M = .72$, $SD = .11$); Control ($M = .76$, $SD = .08$)). Lastly, there was a two-way interaction between Condition and Block, $F(9,321) = 2.84$, $p = .003$, partial $\eta^2 = .07$. To examine the interaction, the effect of Condition was considered within each block separately. In each case, the Position group had the best performance. There was a main effect of Condition in block 1 ($F(3,107) = 8.26$, $p < .001$, partial $\eta^2 = .19$), block 2 ($F(3,107) = 5.67$, $p = .001$, partial $\eta^2 = .14$), and a marginally significant main effect in block 3 ($F(3,107) = 2.44$, $p = .068$, partial $\eta^2 = .04$).

For the last 8 blocks, the data were analyzed using an ANOVA with Condition (Position, Length, Orientation, and Control) as a between-participants' factor and Block (8) as a within-participants' factor. This analysis revealed a main effect of Block, $F(7,749) = 2.94$, $p = .005$, partial $\eta^2 = .03$. Post hoc tests revealed that block 5 was different from all other blocks, all $p < .025$, but that none of the other blocks were different from each other. There was not a significant main effect of Condition, $F(3,107) = 1.06$, $p = .371$, partial $\eta^2 = .03$. Lastly, there was not a two-way interaction between Condition and Block, $F(21,749) = 0.84$, $p = .667$, partial $\eta^2 = .02$.

3.2.2. Model-based Analyses—Again, as in Experiment 1, we considered the proportion of participants who had their first block of trials consistent with unidimensional rules on position, length, or orientation that corresponded to the rules provided in the hints to participants. As can be seen in Table 2, replicating Experiment 1, participants' data in the Position and Control groups was consistent with early position rule use. Unlike Experiment 1, we did not find that participants' data in the Length and Orientation groups was consistent with length or orientation rule use, respectively. In fact, very few or none of the participants' data was best fit by these rules.

That said, if we examine the use of more complicated conjunctive rules in the first block of trials, for the Position group, 66% were best fit by a conjunctive model using position (length and position or orientation and position), and thus 93% of participants initially used a rule on position. For the Length group, 22% were best fit by a conjunctive model using length (length and position), and thus 26% of participants initially used length to create a rule. For the Orientation group, 59% was best fit by a model on orientation and position. Therefore, the overwhelming majority of Position participants and the majority of Orientation participants at least begin by using strategies consistent with the dimension suggested in their prime.

To further examine whether the task-performance data corresponded to the use of the optimal classification rule by participants, as in Experiment 1, we calculated the proportion of participants' data best fit by the optimal classifier in each block of trials and analyzed which model was selected as the best fitting model using AIC fits. First, as can be seen in Figure 7, averaging across participants, we find the same pattern as in the accuracy data. Averaging across the first 4 blocks, the optimal rule accounted for 81% of the responses for participants in the Length group, 75% of the responses for the Orientation group, 80% of the responses for the Control group, and 84% of the responses for the Position group.

Second, using AIC fits, we analyzed which model was selected as the model best fitting participant data for each participant for each block of trials. Approximately a quarter of participants, across all groups, had data best fit by the optimal classifier or the suboptimal

classifier in the last block of trials (Position: 24%, Length: 30%, Orientation: 26%, Control: 25%). Interestingly, the average percent best fit across the first 4 blocks (where we found accuracy differences across our groups) does not correspond to the accuracy data (Position: 15%, Length: 36%, Orientation: 16%, Control: 13%). The open question is why our Position group demonstrated better task accuracy in the first 4 blocks of trials because this groups' data was not best fit by the optimal or suboptimal classifier. We believe the pattern in accuracy can be explained by two different factors: the percent of data accounted for by the optimal model as noted above and modeling random performance. Participants in the Position group did not have their data best fit by the random responding model in any of the first 4 blocks (i.e., 0% in each block). In contrast, using the average % best fit in the first 4 blocks, 14% the Orientation group data, 7% of the Control group data, and 4% of the Length group data was best fit by the random response model. This suggests that, although the optimal model was not selected using AIC fit as the best fitting model, participant responses in the Position group were most closely aligned with the optimal classifier as compared to the other groups.

3.3. Discussion

As in Experiment 1, participants learned to classify simple stimuli into two categories with the position dimension being irrelevant to the task, but in this experiment the predetermined correct rule was an information-integration rule on length and orientation. We hypothesized the opposite pattern of data as found in Experiment 1. Specifically, we believed that participants receiving hints to use the rule-relevant dimensions (i.e., length and orientation) would perform worse than participants told to focus on the rule-irrelevant dimension (i.e., position). Lastly, we anticipated that we would see this benefit early in learning before participants in all groups would shift to relying on the implicit system. Our predictions are supported by the data. The Position group performed significantly better than the other groups during early learning and their data was best accounted for by the optimal classifier.

4. General Discussion

In two experiments, we examined the impact of a working memory load on RB and II classification learning by varying the information available at the start of the task. Participants viewed lines that varied in length, orientation, and position on the screen and learned to classify them into two categories following an RB rule in Experiment 1 and an II rule in Experiment 2. We hypothesized that the RB and II tasks would be differentially affected by having participants focus on dimensions that were relevant (i.e., length and orientation) or irrelevant (i.e., position) to the task. Specifically, the RB task would benefit from a focus on the relevant dimensions, while the II task would benefit from a focus on the irrelevant dimension. Our data support these predictions.²

There are two main limitations of the current work. First, in Experiment 1, our orientation-primed group performed very similarly to the control group. While it is not ideal that our orientation group did not perform better than the control, we do have evidence that the orientation-group data was better fit by the optimal classifier. This suggests that the orientation-primed group was better able to find the correct classification rule but was not applying the rule with enough precision. Second, our hints to use specific dimensions are not as apparent in the modeling results for Experiment 2 as in Experiment 1, in that the length-primed group was less likely to initially use a rule on the length dimension. While

²We combined the data from both experiments and conducted an ANOVA with Experiment (Experiment 1, Experiment 2) and Condition as between-participants factors and Block as a within-participants factor to verify the interaction of Experiment and Condition. There was a two-way interaction between Experiment and Condition, $F(3,245) = 3.16$, $p = .025$, partial $\eta^2 = .04$, and a main effect of Experiment, $F(1,245) = 5.84$, $p = .016$, partial $\eta^2 = .02$.

suboptimal, this result is completely consistent with the accuracy data, such that the length-primed and control participants have similar accuracy patterns.

4.1. Conclusions

Our work highlights the importance of considering factors that differentially impact RB and II learning. For the RB task in Experiment 1, we demonstrate that participants primed with a focus on the irrelevant position dimension performed reliably worse than the other groups. Moreover, interestingly, this effect is most pronounced in later learning. Model-based analyses reveal that participants primed with the relevant dimensions of length and orientation generated data most consistent with the optimal classification rule. In contrast, for the II task in Experiment 2, we find that the position group performed reliably better than the other groups and that this effect was largest in early learning, where the position group performed significantly better than all other groups. Furthermore, the model-based analyses find that the position group generated data most consistent with the optimal classification rule.

Interestingly, we find that a focus on relevant and irrelevant dimensions not only impacts task performance overall but impacts task performance at different stages of learning. The RB task requires that participants start testing simple rules and then switch to testing more complex rules, and therefore a focus on relevant dimensions helps in this more complicated process of finding and fine-tuning complex rules. In contrast, because II tasks require that participants implicitly learn the correct stimulus-response mappings, we find that an initial explicit focus on an irrelevant dimension allows the II system to better learn mappings without early interference from the RB system. This qualifies the finding of Ashby and Crossley (2010) that using an RB explicit strategy may limit the use of the II system. We demonstrate that this limitation may exist if the RB system is actively engaged with dimensions that are required for good task performance, but not on irrelevant dimensions.

Our work suggests that it is critical to understand the circumstances that improve RB and II learning. All of the prior work demonstrates that adding a working memory load that is irrelevant to the task hurts RB learning, and we further demonstrated that adding a relevant working memory load improves RB learning. For II learning, the past literature has been more mixed. Maddox et al. and Zeithamova and Maddox demonstrated that working memory load manipulations only affected RB learning, while Filoteo et al. and DeCaro et al. found that working memory impacted II learning. We believe that this difference was likely caused by differences in task difficulty—a suggestion proposed by Filoteo et al.—with more difficult tasks demonstrating the improvement in II learning. Future research should examine the interaction between working memory load and task difficulty.

References

- Akaike H. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*. 1974; 19:716–723.
- Ashby, FG., editor. *Multidimensional models of perception and cognition*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.; 1992.
- Ashby FG, Alfonso-Reese LA, Turken AU, Waldron EM. A neuropsychological theory of multiple systems in category learning. *Psychological Review*. 1998; 105:442–481. [PubMed: 9697427]
- Ashby FG, Crossley MJ. Interactions between declarative and procedural-learning categorization systems. *Neurobiology of Learning and Memory*. 2010; 94:1–12. [PubMed: 20304078]
- Ashby FG, Ell SW. The neurobiology of human category learning. *Trends in Cognitive Sciences*. 2001; 5:204–210. [PubMed: 11323265]
- Ashby FG, Ell SW, Waldron EM. Procedural learning in perceptual categorization. *Memory & Cognition*. 2003; 31:1114–1125. [PubMed: 14704026]

- Ashby FG, Gott RE. Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1988; 14:33–53.
- Ashby FG, Maddox WT. Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology*. 1993; 37:372–400.
- Ashby FG, Maddox WT. Human category learning. *Annual Review of Psychology*. 2005; 56:149–178.
- Ashby FG, Maddox WT. Human category learning 2.0. *Annals of the New York Academy of Sciences*. 2011; 1224:147–161. [PubMed: 21182535]
- Ashby FG, Maddox WT, Bohil CJ. Observational versus feedback training in rule-based and information-integration category learning. *Memory & Cognition*. 2002; 30:666–677. [PubMed: 12219884]
- Ashby FG, O'Brien JB. Category learning and multiple memory systems. *Trends in Cognitive Science*. 2005; 2:83–89.
- Ashby FG, Queller S, Berretty PM. On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics*. 1999; 61:1178–1199. [PubMed: 10497436]
- DeCaro MS, Thomas RD, Beilock SL. Individual differences in category learning: Sometimes less working memory capacity is better than more. *Cognition*. 2008; 107:284–294. [PubMed: 17707363]
- Filoteo JV, Lauritzen JS, Maddox WT. Removing the frontal lobes: The effects of engaging executive functions on perceptual category learning. *Psychological Science*. 2010; 21:415–423. [PubMed: 20424079]
- Filoteo JV, Maddox WT, Salmon DP, Song DD. Information-integration category learning in patients with striatal dysfunction. *Neuropsychology*. 2005; 19:212–222. [PubMed: 15769205]
- Filoteo JV, Maddox WT, Simmons AN, Ing AD, Cagigas XE, Matthews S, Paulus MP. Cortical and subcortical brain regions involved in rule-based category learning. *NeuroReport*. 2005; 16:111–115. [PubMed: 15671857]
- Grimm LR, Markman AB, Maddox WT, Baldwin GC. Stereotype threat reinterpreted as a regulatory mismatch. *Journal of Personality and Social Psychology*. 2009; 96:288–304. [PubMed: 19159133]
- Maddox WT, Ashby FG. Comparing decision bound and exemplar models of categorization. *Perception and Psychophysics*. 1993; 53:49–70. [PubMed: 8433906]
- Maddox WT, Ashby FG. Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioural Processes*. 2004; 66:309–332. [PubMed: 15157979]
- Maddox WT, Ashby FG, Bohil CJ. Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2003; 29:650–662.
- Maddox WT, Ashby FG, Ing AD, Pickering AD. Disrupting feedback processing interferes with rule-based but not information-integration category learning. *Memory & Cognition*. 2004; 32:582–591. [PubMed: 15478752]
- Maddox WT, Baldwin, G.C., & Markman AB. A test of the regulatory fit hypothesis in perceptual classification learning. *Memory & Cognition*. 2006; 34:1377–1397. [PubMed: 17263064]
- Maddox WT, Filoteo JV. Striatal contributions to category learning: Quantitative modeling of simple linear and complex nonlinear rule learning in patients with Parkinson's disease. *Journal of the International Neuropsychological Society*. 2001; 7:710–727. [PubMed: 11575593]
- Maddox WT, Ing AD. Delayed feedback disrupts the procedural-learning system but not the hypothesis testing system in perceptual category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2005; 31:100–107.
- Nosofsky RM, Johansen MK. Exemplar-based accounts of multiple-system phenomena in perceptual categorization. *Psychonomic Bulletin & Review*. 2000; 7:375–402. [PubMed: 11082849]
- Rao SM, Bobholz JA, Hammeke TA, Rosen AC, Woodley SJ, Cunningham JM, Cox RW, Stein EA, Binder JR. Functional MRI evidence for subcortical participation in conceptual reasoning skills. *Neuroreport*. 1997; 8:1987–1993. [PubMed: 9223090]
- Seger CA, Cincotta CM. Striatal activation in concept learning. *Cognitive, Affective, & Behavioral Neuroscience*. 2002; 2:149–161.
- Takane, Y.; Shibayama, T. *Structure in stimulus identification data*. Hillsdale, NJ: Earlbaum; 1992.

- Zeithamova D, Maddox WT. Dual-task interference in perceptual category learning. *Memory & Cognition*. 2006; 34(2):387–398. [PubMed: 16752602]
- Zeithamova D, Maddox WT. The role of visuo-spatial and verbal working memory in perceptual category learning. *Memory & Cognition*. 2007; 35(6):1380–1398. [PubMed: 18035635]

Figure 1.
Stimulus space used in Experiment 1 with correct conjunctive rule on length and orientation dimensions represented.

Figure 2.
Example of stimulus display used in both experiments.

Figure 3. Proportion correct in each block for participants in the Control, Position, Length, and Orientation-prime conditions in Experiment 1.

Figure 4.

The proportion of participants' data best fit by the optimal conjunctive model on length and orientation dimensions for the 4 experimental groups (Position, Length, Orientation, and Control) in Experiment 1.

Figure 5.
Stimulus space used in Experiment 2 with optimal classifier on length and orientation represented.

Figure 6. Proportion correct in each block for participants in the Control, Position, Length, and Orientation-prime conditions in Experiment 2.

Figure 7.

The average proportion of participants' data best accounted for by the optimal model on length and orientation dimensions for the 4 experimental groups (Position, Length, Orientation, and Control) in Experiment 2.

Table I

Proportion of participants in the experimental groups best fit by simple models corresponding to primed dimensions in Block 1 (Experiment 1)

	Position	Length	Orientation	Control
Position Rule	.39	.06	.03	.27
Length Rule	.06	.22	.08	.00
Orientation Rule	.04	.03	.42	.14

Table II

Proportion of participants in the experimental groups best fit by simple models corresponding to primed dimensions in Block 1 (Experiment 2)

	Position	Length	Orientation	Control
Position Rule	.27	.11	.00	.15
Length Rule	.00	.04	.07	.07
Orientation Rule	.00	.00	.00	.00