

## **UC Merced**

# **Proceedings of the Annual Meeting of the Cognitive Science Society**

### **Title**

Differential Cue Saliency, Blocking and Learned Inattention

### **Permalink**

<https://escholarship.org/uc/item/5h38z6sf>

### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 28(28)

### **ISSN**

1069-7977

### **Authors**

Heckler, Andrew F.  
Kaminski, Jennifer A.  
Sloutsky, Vladimir M.

### **Publication Date**

2006

Peer reviewed

# Differential Cue Salience, Blocking and Learned Inattention

**Andrew F. Heckler (heckler.6@osu.edu)**  
Department of Physics, Ohio State University  
191 W. Woodruff Avenue, Columbus, OH 43210 USA

**Jennifer A. Kaminski (kaminski.16@osu.edu)**  
Center for Cognitive Science, Ohio State University  
210A Ohio Stadium East, 1961 Tuttle Park Place, Columbus, OH 43210, USA

**Vladimir M. Sloutsky (sloutsky.1@osu.edu)**  
Center for Cognitive Science, Ohio State University  
208C Ohio Stadium East, 1961 Tuttle Park Place, Columbus, OH 43210, USA

## Abstract

A classic blocking design is modified by varying the relative salience of the blocked and the blocking cue. It is found that, in accordance with the learned inattention theory of blocking, the amount of blocking is diminished when the blocked cue is more salient than the blocking cue. In addition, when studying the patterns of answers about the control cues, it is found that low salience cues were preferred over high salience cues in the case when the blocking cue was also of low salience. This novel result may be explained by the simultaneous blocking of two separate dimensional values of the blocked cue, and this results in category blocking.

## Introduction

When two cues are simultaneously paired with an outcome, the learned associative strengths of each cue with the outcome often differ. An excellent example of this is the phenomenon known as blocking. In blocking, cues A and B are paired with an outcome, and the association of cue B with the same outcome is significantly weakened if it is also learned that A alone is associated with the outcome. For the last several decades the phenomenon of blocking has played a major role in many studies of associative learning. It is a robust phenomenon first observed in rats (Kamin, 1969) and since then has been observed in a variety of situations and species, including humans (e.g., Mackintosh & Turner, 1971; Kruschke & Blair, 2000; Crookes & Moran, 2003; Arcediano, Escobar & Miller, 2004).

There are two principle theories explaining the phenomenon of blocking. One explanation uses the influential Rescorla-Wagner model which claims that the blocked cue B is only weakly associated with the outcome because it is extraneous information: it does not add to the accuracy of predicting the outcome, since A has already been learned to be a perfect predictor (Rescorla & Wagner, 1972). Thus in this model B is simply not learned. However, Mackintosh and Turner (1971) found in rats and Kruschke and Blair (2000) found in humans that there is a subsequent attenuation of learning of the

blocked cue in a new learning situation. These results cannot be explained by the Rescorla-Wagner model or its other variations (e.g., Van Hamme & Wasserman, 1994), since they predict that nothing (or very little) was learned about the blocked cue. To accommodate these findings of attenuation of subsequent learning, a second explanation of blocking was developed (Sutherland & Mackintosh, 1971; Kruschke & Blair, 2000). In short, this explanation involves the modeling of attentional shifts and learned attention or inattention to specific cues. In this model, the cue B is blocked not because it accrues no associative strength, rather it was blocked because it was learned to be ignored. Later, when trying to associate B with another outcome, the learner has difficulty because B was previously learned to be ignored (rather than not learned at all). Thus learned inattention naturally explains both classical blocking results and subsequent attenuation of learning.

In the above-mentioned discussion of blocking, it was noticed early on that the relative salience of the cues play significant role in the magnitude of blocking in rats and pigeons (Kamin, 1969; Hall et al., 1977). However, the vast majority of the studies above involve blocking with cue of approximate equal salience, and to our knowledge there are no published studies of the effect of differential cue salience on blocking in humans. The issue of differential salience in blocking is important for two reasons. The first is practical: cues often vary significantly in salience in real world situations thus understanding the any possible effects of differential cue salience is important for applications such as classroom learning. Second, the issue of blocking and differential cue salience is especially germane when considering the paradigm that blocking is due to attentional shifts and learned inattention. More specifically, if blocking is one way to shift attention to favor one cue over another, and differential salience is another mechanism that also shifts attention, then one would expect an interaction between blocking and differential salience. For example, if cue B is blocked because the subject has learned to not attend to it and to attend to another predictive cue A, then if the

salience of B is increased relative to A, the more attention will inherently be paid to it and the more difficult to it may be to block. Likewise, if the salience of A is increase relative to B, then one would expect an enhancement of the blocking effect. Therefore, in this study we will examine the effect of differential cue salience on blocking in humans.

When adding in differential cue salience to blocking, there is another issue to consider which may help shed light on the extent to which attentional shifting plays a role in learning. This stems from the fact that when adding salience, one is necessarily adding another variable dimension to the stimulus. With two perceptual dimensions to be learned, in certain conditions, there may arise the possibility that both dimensions will be blocked, either in tandem or separately. Significantly, in this experiment, the added dimensional value can be constant across stimuli, thus creating categories, such as the category “colorful”. In this study we find evidence that multiple dimensions can be blocked, including the blocking of categories, and this adds a new challenge to an explanation of blocking using the Rescorla-Wagner model or the attentional shifting model.

## Experiment

### Method

Experiment 1, design shown in Table 1, follows a typical blocking design. However the actual learning task involves the learning of the association of shapes and colors with an outcome, in contrast to many such experiments which involve the association of words (symptoms) with an outcome (disease). In particular, in Experiment 1, participants were told that they were learning about an appliance factory and they learned to associate various cues (diagrams of computer chip components that have different shapes and/or colors) with specific outcomes (appliances). On each trial, they were presented with a diagram representation of a computer chip with either one or two “components” on it, placed randomly in one of four positions on the chip. The components were simple geometric shapes such as a triangle or an oval. The participants were told that the presented chip was installed inside a particular appliance (represented as a picture) to help it operate, and they had to answer which component (or combination of components) went with each appliance. After each trial, they were given feedback providing the correct answer.

Participants were randomly placed in one of six conditions: 3 salience conditions  $\times$  2 shape conditions. The experiment aims to study the effect of perceptual salience on blocking, and the shape conditions, which are two random permutations of the particular shapes with the appliances, where added to counterbalance any effects of particular shapes. The cues (computer chip components) were geometric shapes placed on a simple diagram representing a circuit board. The cues were classified as

salient or non-salient. The salient cues were colored (e.g., red, blue, green) and somewhat larger than the non-salient cues which were all light gray. To demonstrate that the colorful cues were indeed more salient than the grey cues, nine participants (separate from the those in experiment 1) were shown blank grids and asked to indicate as quickly as possible whether or not they saw a shape on the grid (i.e. cue on the chip). They saw 35 chips (14 highly salient shapes, 14 less salient shapes, and 7 with no shape). The response time was lower for the highly salient shapes than for less salient shapes, independent samples t-test  $t(16) = 2.76, p < .014$ .

Table 1. Design of Experiment

Phase	Condition		
	A salient	AB salient	B salient
Training	A*→1	A*→1	A→1
Phase 1	G*→2	G*→2	G→2
Training	A*B→1	A*B*→1	AB*→1
Phase 2	C*D→3	C*D*→3	CD*→3
	E*F→4	E*F*→4	EF*→4
Test for blocking	e.g. B, D, BD	e.g. B*, D*, B*D*	e.g. B*, D*, B*D*
Inverse questions	1,2,3,4	1,2,3,4	1,2,3,4

Note: Letters denote cue (computer chip component) and numbers denote outcome (appliance). An asterisk (e.g. A\*) denotes a high salience cue.

The three salience conditions are shown in Table 1. In the “A salient” condition the A cue is more salient than the blocked cue B. In the “AB salient” condition, both cues were of equivalent (high) salience, and in the “B salient” condition, the blocked cue B was more salient than cue A. In each condition, the cue A is first associated alone with a particular outcome (e.g. blender), then in phase two both cues A and B are associated with the same outcome. In addition to A and B, other control cues such as C and D are associated with another outcome in phase two, with no training on these cues in the first phase. The learned inattention theory predicts that since cue A was at first solely associated with the blender, when both A and B are subsequently associated with the blender, the learner will learn to not attend to B and only attend to cue A, since knowing B does not reduce the error rate in predicting the outcome. However, cues C and D were learned only simultaneously, and both will be learned equally well (at least in the same salience condition). During the test phase, two kinds of questions are used to test for blocking: single-cue prompts and two-cue prompts. In each case, the prompt is novel. For example, the learners are shown a chip with component B (cue B) only and are asked to chose in which appliance it would be installed. Likewise they are also asked about a chip with both components B and D. The standard blocking results

predict that at least for the equal salience condition, the learner will not choose the previously associated outcome as much for B as for D. Likewise, when both cues are presented together, the learner will disregard B and choose the outcome associated with D. Both of these results indicate blocking occurred, but since they may probe different learning or reasoning mechanisms, both were included to determine the consistency and robustness of the results. In the differential salience conditions, the learned in attention theory predicts that blocking will be enhanced when A is more salient and blocking will be reduced or eliminated when B is more salient.

**Participants** 156 undergraduate students (92 females and 64 males) from Ohio State University participated in the experiment and received partial credit for an introductory psychology course. Fifty two students were assigned to each of three salience conditions that specified which cues were salient.

**Materials and Design** Table 1 shows the abstract design. Training Phase 1 had 20 blocks of 2 trials (one for each cue), Training Phase 2 had 20 blocks of three trials. The first testing phase had one block of six trials with one-cue questions (3 B and 3 D), and 2 blocks of nine trials of both one and two-cue questions. Each trial consisted of one multiple choice question. Participants were shown a computer chip with a combination of components (cues) and asked which of six appliances (outcomes) uses the chip. The training phases included corrective feedback, and the testing phases did not. In the final testing phase, one block of four “inverse” questions were presented, asking participants to match a particular outcome (appliance) with a choice of cues. This is a total of 128 trials.

**Procedure** All training and testing was presented to individual participants on a computer screen in a quiet room. They proceeded through training and testing at their own pace; and their responses were recorded.

## Results and Discussion

The participants successfully learned during the training sessions, with an average score of 87% correct in Phases 1 and 2 (chance was 16%). This excludes 6 participants who scored 2 standard deviations below average on at least three of the training cue types. A one-way ANOVA revealed that the scores on the training phases were independent of condition ( $p > 0.2$ ).

The rich results of the testing phase provide a measurement of blocking in three different ways. In the first and most straightforward way, Figure 1 displays the response percentages for the relevant single-cue questions. (Full details on response percentages in the testing phase are in Table 2.) The standard blocking result is replicated in the AB-salient condition when both cues

have equal salience. In this case, the learners scored lower (chose the previously associated appliance less) with the blocked cue than with the control cue [paired t-test  $t(48) = 5.73, p < 0.001$ ], and an effect size of 0.82. However, in the B-salient condition, while the blocked cue score is lower than the control cue score, [paired t-test  $t(50) = 2.2, p = 0.032$ ], the difference, and thus the amount of blocking is significantly reduced with an effect size of 0.31. Overall, the scores on the blocked cue were dependent on the condition, [ $F(2,149) = 23.6, p < 0.001$ ].

In the case of the two-cue questions (Table 2), there are two separate measures of blocking. The results are much the same as the single-cue questions, only there is stronger evidence that the blocking was greatly suppressed or eliminated in the B-salient condition. When the two-cue test stimuli include the blocked cue and a similar-salience cue (e.g. B\*D\*), the standard blocking result is replicated in the AB-salient condition: the learners chose the outcome associated with the control cue over the outcome associated with the blocked cue 35.7% to 4.1% [ $\chi^2(df=1, N=98)/2 = 12.32, p < 0.001$ ], where, as Kruschke and Blair (2000) have done, the  $\chi^2$  value is divided by 2, as a conservative estimate of the of the possible lack of independence between the two repetitions of the questions (Wickens, 1989). However, in the B-salient condition, there was no significant preference between the blocked and control cue, with the respective percentages of 14.9% and 10.9% [ $\chi^2(df=1, N=100)/2 = 0.31, n.s.$ ], thus the two-cue questions indicate that the blocking was greatly suppressed in the B salient condition, and this suppression is larger than for the one-cue questions.

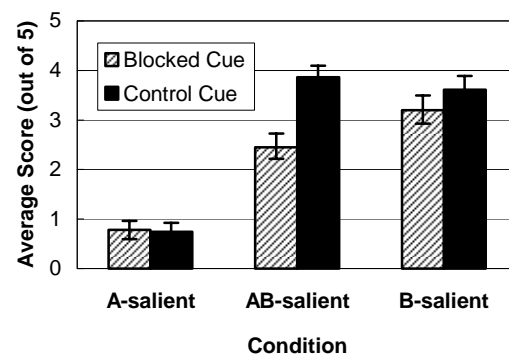


Figure 1. Total Average Scores for B and D single cue questions. The AB-salient condition shows significant blocking, while A-salient and B-salient do not. Error bars represent standard error of mean.

Examination of table 2 reveals similar results for the two cue question E\*B\* vs. E\*D\* in the AB-salience condition, with a significant preference for choosing the outcome associated with D\* over B\* [paired t-test  $t(48) = 3.8, p < 0.001$ ], indicating blocking. Furthermore comparing EB\* vs. ED\*, there was a marginal but no significant preference for the outcome associated with the

similar-salience control cue over the blocked cue [paired t-test  $t(50)=1.7, p=0.1$ ].

The results for the A-salient condition in Figure 1 show that the blocked cue B had very low scores. The average one-cue score of 17.1% of the blocked cue is at chance (16%), and it lower than either of the blocked cue scores in the other conditions [post-hoc Tukey test,  $ps<0.001$ ]. Therefore either B was either completely blocked or somehow not noticed and not learned. The control cue D, which also had low salience compared to the cue it was paired with in the training phase, had virtually identical results as the blocked cue B in the A-salient condition. Therefore D was also either somehow completely blocked or somehow not learned. This results suggests that the participants may have learned to ignore all low salience (grey) cues in this condition, perhaps due to overshadowing or some kind of categorical blocking, as is discussed in the final section..

Table 2 reveals an unexpected pattern in responses to the two cues question involving the control cues D and E (or equivalently C and F), which also indicates that something other than overshadowing, and perhaps categorical blocking is taking place. In the A-salient condition, when prompted with cues ED, the learners overwhelmingly choose the outcome associated with E over D (83.7% to 2.9%). This is consistent with the single-cue questions, show that D was either completely overshadowed or somehow blocked. In the AB-salient

condition, when all of the cues have equal salience, there is a no significant preference for the outcome associated with D over E (32.7% to 18.4%) [ $\chi^2(1)/2=1.96, p=0.16$ ], which is to be expected, since D and E have equal salience. In the B-salient condition, the cue E is less salient than cue D, so one might expect an overwhelming preference for cue D, analogous to the A-Salient condition in which there was a preference for cue E when it was more salient than D. However, in the B-salient condition, the there was no significant preference for either the outcome associated with the low salience cue E or the outcome associated with the higher salience cue D (52.9% to 33.7% respectively) [ $\chi^2(1)/2=1.84, p=0.17$ ]—if anything there was a slight (but insignificant) preference for the *low* salience cue. A similar pattern occurred in the responses to “inverse” questions, when the learners were asked to choose which cue goes with a specific outcome. When asked which cue went with appliance 3, the participants in the A-salient condition preferred the high salient (C\*) over the low salient cue (D) (94.2 to 0%); in the AB-salient they had no preference for the equal salience cues (48.1% to 42.3%) [ $\chi^2(1)=0.09, p=0.76$ ], and a slight but insignificant preference for the low salient cue (C) over the high salient cue (D\*) (57.7% to 40.4%) [ $\chi^2(1)=1.28, p=0.26$ ].

Table 2. Percentage choice response. Letters A-E denote cues (computer chip components) and numbers 1-6 denote chosen outcomes (appliances) during the testing phase. The “correct” choices given in training are shown in Table 1. Data in bold font are discussed in the text.

Condition	Cue	Outcome						
		Blank	1	2	3	4	5	6 (don't know)
A-Salient	B	0.4	<b>17.1</b>	9.7	14.0	13.2	12.4	33.3
	D	1.2	9.2	<b>15.4</b>	6.5	9.6	21.5	36.5
	A*	1.9	89.4	3.8	2.9	1.9	0.0	0.0
	E*	0.0	3.8	1.0	1.9	89.4	1.0	2.9
	BD	1.9	<b>8.7</b>	<b>9.6</b>	11.5	10.6	11.5	46.2
	E*B	0.0	<b>3.8</b>	1.9	7.7	80.8	3.8	1.9
	E*D	0.0	3.8	<b>2.9</b>	2.9	<b>83.7</b>	0.0	6.7
AB-Salient	B*	0.8	<b>46.5</b>	5.4	7.7	15.8	8.5	15.4
	D*	1.2	6.2	<b>74.6</b>	1.2	3.8	4.6	8.5
	A*	1.0	88.5	1.0	3.8	3.8	0.0	1.9
	E*	1.0	4.8	8.7	4.8	61.5	4.8	14.4
	B*D*	1.0	<b>4.1</b>	<b>35.7</b>	1.9	1.9	9.6	46.2
	E*B*	1.0	<b>10.6</b>	7.7	5.8	30.8	10.6	33.7
	E*D*	1.0	1.0	<b>32.7</b>	3.3	<b>18.4</b>	7.1	36.5
B-Salient	B*	0.4	<b>64.0</b>	3.9	3.1	3.5	5.0	20.2
	D*	0.4	2.3	<b>71.2</b>	0.8	4.6	2.7	18.1
	A	0.0	88.5	1.0	5.8	1.9	0.0	2.9
	E	0.0	7.7	5.8	5.8	73.1	1.9	5.8
	B*D*	1.0	<b>10.9</b>	<b>14.9</b>	2.9	0.0	5.8	65.4
	EB*	0.0	<b>28.8</b>	1.0	1.0	52.9	2.9	13.5
	ED*	0.0	1.0	<b>33.7</b>	1.0	<b>52.9</b>	1.0	10.6

## General Discussion

The results of this experiment provides strong evidence that differential cue salience significantly effects blocking. When initially learned cue A and the blocked cue B have similar salience, the classic blocking results were replicated. When the initially learned cue was more salient than the blocked cue (the A-salient condition) blocking was enhanced. Most importantly, in the B-salient condition when the blocked cue B was more salient than the initially learned cue A, blocking was significantly reduced or eliminated. Other researchers such as Hall et al. (1977) and Arcediano, Escobar, and Miller (2004) have found this same kind of attenuation of blocking in rats and pigeons, and this study is an extension of the results to humans.

These results are consistent with the model of attentional shifting and learned inattention as an explanation of blocking. In this model of error reduction, the learner shifts attention away from the blocked cue B in order to devote more attentional resources, which are limited, to the cue A, which is always predictive (whereas B is sometimes absent, thus not useful). The experiment in this study supports the idea that attention may also be shifted by differential cue salience. When cue B is more salient than cue A, more attention is shifted to B, thus blocking is attenuated. In the opposite case when A is more salient than B, even more attention is shifted to A, besides that due to blocking and blocking is enhanced.

In modeling attentional shifting in learning, Kruschke (2001) has included cue salience as a parameter in his EXIT model, and has successfully predicted classic blocking results with equal cue salience. However, we know of no experiments that have tested this model. Nonetheless, a similar effect of cue salience on associative learning was studied by Bohil, Markman, and Maddox (2005), where they found that differential cue salience can interact with and even mimic the inverse base rate effect in categorization tasks. They propose that their results support models of attentional shifting and that “any stimulus element that can cause attention to shift appropriately can lead to a learning effect analogous to the inverse base rate effect”. The results of our experiment are consistent with Bohil et al. (2005) in that differential cue salience does effect a specific associative learning phenomenon. It is worth noting that the cues they used were labels (disease symptoms), and the differences in feature salience were conceptual (sore muscles vs. paralysis). This is in contrast to our experiment in which the cues are perceptual (colored shapes) and the difference in salience is also perceptual (colored and large vs. grey and small).

It is important to consider that the attenuation or enhancement of blocking may be at least partially due to a known salience effect that is separate from blocking. It is known that when two cues are presented, the more salient one will be utilized more— this phenomenon is known as

overshadowing (Edgell et al., 1992; Edgell et al., 1996; Krushke & Johansen, 1999). This is not at odds with our conclusion that this experiment supports the attentional theory of blocking, as overshadowing has also been explained in terms of attentional shifting (Krushke & Johansen, 1999). Nonetheless, when comparing choice percentages for various cues (which all are 100% predictive), it is to be expected that the more salient cue will be favored. Overshadowing alone could explain the results of the A-salient condition, where the high salient cues are strongly preferred and the low salient cues are at or below chance. However there is reason to believe that it may be only partially responsible for the signal. Evidence for this comes the choice patterns for control cues on the two-cue questions. In contrast to the A-salient condition, in the B-salient condition there was certainly not a strong preference for the salient cue, rather if anything a slight preference for the *low* salient control cue, even though the type and number of training trials for these control cues was the same for both conditions. This cannot be explained by overshadowing, but the results of both conditions can be understood by one explanation, which we will call “category blocking”. Notice that each cue has two feature dimensions: shape and color. More specifically, we can categorize the colors into two categories: “colored” or “grey”. If we denote cue A as colored shape one, or “Colored( $s_1$ )” and B as “grey shape two” as Grey( $s_2$ ), etc, then the relevant training for the condition can be abbreviated as in Table 3.

Table 3. Blocking model of the category Colored or Grey.  $s_n$  denote specific cue shapes within the category.

A-Salient Condition:	B-Salient Condition:
Colored( $s_1$ ) $\rightarrow$ 1	Grey( $s_1$ ) $\rightarrow$ 1
Colored( $s_1$ ), Grey( $s_2$ ) $\rightarrow$ 1	Grey( $s_1$ ), Colored( $s_2$ ) $\rightarrow$ 1
Colored( $s_3$ ), Grey( $s_4$ ) $\rightarrow$ 2	Grey( $s_3$ ), Colored( $s_4$ ) $\rightarrow$ 2
Colored( $s_5$ ), Grey( $s_6$ ) $\rightarrow$ 3	Grey( $s_5$ ), Colored( $s_6$ ) $\rightarrow$ 3

Note that, ignoring the shapes  $s_n$  for the moment, in the A-salient condition, the category “grey” should be blocked and in the B-salient condition “colored” should be blocked. Therefore, we propose that since the stimuli have two feature dimensions, both dimensions are experiencing blocking effects. In this case one of the dimensions fits within a category, namely Colored or Grey, thus for this dimension the whole category may be blocked. The choice percentages are consistent with this explanation. For example, in the A-salient condition, both  $s_2$  (ie. cue B) and Grey are blocked. This shows up in both the choice percentages and scores of B, which are very low. In effect, the cue B has four effects working against it: the cue is “double blocked” (shape and grey are blocked), the blocking is enhanced because of a difference in salience, and there is overshadowing. In the case of the low-salience control cue  $s_4$  (cue D), which is Grey, one would expect it to be both “category blocked”

(Grey is blocked) and overshadowed, and this is consistent with the very low scores and choice percentages of D compared to the colored cue  $s_5$  (cue E). Since all of these effects work to attenuate the choice percentages, it is difficult to know which is dominant.

The effect is more dramatic in the B-salient condition. In this case  $s_2$  (cue B) and colored shapes are blocked. However, the blocked dimension of color is highly salient, so its blocking should be diminished and it would also tend to overshadow non-salient cues, thus there are competing factors to decrease the blocking of B, which is observed in the choice and score data. The strong signal of category blocking comes from the two-cue choice scores of control cues  $s_5$ (cue E) and  $s_4$ (cue D). Normally, one would not expect a different signal than from A-salient condition, since the training of the control cues is identical in both conditions and ostensibly independent of the first training phase. If anything overshadowing predicts that there should always be a preference for the more salient cue when prompted with the control cues (Edgell et al, 1992). This is overwhelmingly true in the A-salient condition, but not true in the B-salient condition. This might be explained by the fact that in the B-salient condition, the category Colored is blocked, thus there would be a preference for Grey, or the low salient cue, and this blocking of the high salient category overcomes the opposite effect of overshadowing.

As mentioned earlier, there is an inherent confound in the design: when varying relative salience, one necessarily varies the distinguishability of the stimuli (hence varying two dimensions simultaneously). This introduces a possible confound for our conclusions about category blocking since overshadowing and blocking both affect associative strengths. We have argued that these two effects behave differently in different conditions, allowing us to isolate their effects. Nonetheless, this suggests another design: although varying the differential salience necessarily implies varying the distinguishability of the cues, it is possible to vary the distinguishability without varying the differential salience. Thus, while this experiment provides evidence for the novel result that more than one dimension can be blocked at a time and that this can result in category blocking, in order to confirm this result it would be worthwhile to design an experiment which more carefully controls each dimension to explicitly test for category blocking and remove the confound of salience.

### Acknowledgments

This research is supported by a grant from the Institute of Educational Sciences of the U.S. Department of Education (#R305H050125) to Andrew F. Heckler and Vladimir M. Sloutsky.

### References

Arcediano, F., Escobar, M., & Miller M. M. (2004). Is stimulus competition an acquisition deficit of a

- performance deficit? *Psychonomic Bulletin and Review*, 11(6), 1105-1110.
- Bohil, C. J., Markman A. B., & Maddox, W. T. (2005). A feature-salience analogue of the inverse base-rate effect. *Korean Journal of Thinking and Problem Solving*, in press.
- Crookes, A. E., & Moran, P. M. (2003). An investigation into age and gender differences in human Kamin blocking, using a computerized task. *Developmental Neuropsychology*, 24(1), 461-477.
- Edgell, S. E., Bright R. D., Ng P. C., Noonan T. K., & Ford L. A. (1992). The effects of representation on the processing of probabilistic information. In B. Burns (Ed.), *Precepts, Concepts and Categories*. New York: Elsevier Science.
- Edgell, S. E., Castellan, N. J., Roe, R. M., Barnes, J. M., Ng P. C., Bright R. D., & Ford L. A. (1996). Irrelevant information in probabilistic categorization. *Journal of Experimental Psychology: Learning Memory and Cognition* 22), 1463-1481.
- Hall, G., Mackintosh, N. J., Goodall, G., & Dal Martello, M. (1977). Loss of control by a less valid or less salient stimulus compounded with a better predictor of reinforcement. *Learning and Motivation*, 8, 145-158.
- Kamin L. J. (1969). Predictability, surprise, attention, and conditioning. In B.A Campbell & R. M. Church (Eds.), *Punishment*. New York: Appleton-Century-Crofts.
- Kruschke, J. K., (2001) The inverse base rate effect is not explained by eliminative inference. *Journal of Experimental Psychology: Learning Memory and Cognition* 27, 1385-1400.
- Kruschke J. K., & Johansen, M. K. (1999). A model of probabilistic category learning. *Journal of Experimental Psychology: Learning Memory and Cognition* 25 (5), 1083-1119.
- Kruschke, J. K., & Blair, N. J. (2000). Blocking and backward blocking involve learned inattention. *Psychonomic Bulletin and Review*, 7 (4), 636-645.
- Mackintosh, N. J., & Turner, C. (1971). Blocking as a function of novelty of CS and predictability of UCS. *Quarterly Journal of Experimental Psychology*, 23, 359-356.
- Rescorla, R. A., & Wagner, A. R., (1972). A theory of Pavlovian conditioning: Variation in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning: II. Current research and theory*. New York: Appleton-Century-Crofts.
- Sutherland, N. S., & Mackintosh, N. J., (1971). *Mechanisms of Animal Discrimination Learning*. New York: Academic Press.
- Van Hamme, L. J., & Wasserman E. A. (1994). Cue competition and causality judgments: the role of nonpresentation of compound stimulus events. *Learning and Motivation* 25, 127-151.
- Wickens, T. D. (1989). *Multiway Contingency Tables Analysis for the Social Sciences*, Hillsdale, NJ: Erlbaum.