Sub-sampling of cues in associative learning

Omar D. Perez[1, 3], Edgar H. Vogel[4], Sanjay Narasiwodeyar[2] & Fabian A. Soto[2]

[1] Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, California

[2] Department of Psychology, Florida International University, Florida, USA

[3] Nuffield College CESS-Santiago, Faculty of Business and Economics, University of Santiago, Santiago, Chile

[4] Faculty of Psychology, University of Talca, Talca, Chile

Author Note

Correspondence concerning this article should be addressed to Omar D. Perez, 1200 E. California Blvd., Pasadena, California, CA91125. E-mail: odperez@caltech.edu

Abstract

Theories of learning distinguish between elemental and configural stimulus processing depending on whether stimuli are processed independently or as whole configurations. Evidence for elemental processing comes from findings of summation in animals where a compound of two dissimilar stimuli is deemed to be more predictive than each stimulus alone, whereas configural processing is supported by experiments employing similar stimuli that fail to find summation. However, in humans the summation effect is robust and impervious to similarity manipulations.In three experiments in human causal learning, we show that summation can be obliterated when partially reinforced, common cues are added to the compound in training and test (AXY+, BXY+, CXY-, DXY-; test with ABX and ABY). We also show that this effect holds only when the partially reinforced cues are similar to the reinforced cues. We propose that these findings could be captured by a model where subjects sample one stimulus in each trial according to its predictive value and salience and update the predictions of the sampled stimulus by reward prediction error. This model captures the present results and other phenomena from the learning literature which are usually interpreted in favor of configural processing of stimuli.

Keywords: generalization, Rescorla-Wagner, configural, summation, elemental

Word count: 10,106

Sub-sampling of cues in associative learning

Introduction

What do organisms do when confronted with a situation where two cues that provide independent information about the same amount of reward are present? How much reward do they predict? In animal studies, a typical procedure to study this type of generalization problem is the summation design, where two cues A and B are separately paired with the same amount of reinforcement during a first stage of training and responding to a compound of the two stimuli, AB, is assessed in a final test. A summation effect is obtained when subjects respond more to AB than to each of A or B alone. In humans, an analogous effect is obtained in human predictive learning experiments when the AB compound is deemed more predictive of an outcome than each of A or B alone (Aydin and Pearce, 1997; Kehoe *et al.*, 1994; Soto *et al.*, 2009; Thein *et al.*, 2008).

Two types of computational models in the Pavlovian conditioning literature have been pitted against each other using this summation design. In these models, the internal representation of stimuli is taken to be either elemental (meaning that an agent represents and processes each stimulus independently) or configural (if compounds of stimuli are represented and processed as whole configurations that enter into a single association with the outcome). The cardinal example of elemental models is the influential Rescorla-Wagner model (1972), which assumes that subjects represent A and B independently, and that the presence of the compound AB simply makes subjects sum their individual predictions. What the Rescorla-Wagner model proposes is that the conjoint presentation of stimuli A and B will produce a total prediction given by $v_{total} = v_{AB} = v_A + v_B$. Using this rule, the model can directly explain summation with the presentation of a compound of cues that have been previously trained with the same outcome value, and makes efficient use of the available evidence under the assumption that the two cues are independent causes of such outcome (Pérez *et al.*, 2018).

Configural models, by contrast, take a sharply different approach. Under this view, subjects process and associate whole configurations with the outcomes that follow. Total responding to a compound AB depends on the similarity between the training configurations A and B, and the testing configuration AB. In the basic formulation of configural theory proposed by Pearce (1987; 1994), responding to the compound AB in a simple summation design will depend on the similarity between the training configurations A and B, and the testing configuration AB. More formally, the total predictive value of AB is given by $v_{AB} = E_{AB} + \sum_j e_{j,AB}$, where $e_{j,AB}$ is the predictive value of configuration j that generalizes to configuration AB (including stimuli A and B). Generalization, in turn, depends upon the similarity between the configurations, which is a direct function of the number of common elements they share with one another. In Pearce's formulation of configural theory, similarity between A and AB is given by $S(A, AB) = \frac{N_c}{N_A} \times \frac{N_c}{N_{AB}}$, where N_c is the number of common elements between the configurations (1 element, A) and N_A, N_{AB} is the total number of elements in configurations A and AB (2 and 1 elements, respectively). Under this rule, the outcome prediction from AB when both A and B have converged to the same predictive value during training is equal to $v_{AB} = 0.5$, which is simply the average of the prediction from each of its components A and B, anticipating an absence of summation.

Empirically, studies in animals have shown that summation is modulated by a number of factors, among which the similarity between the stimuli forming the compound appears critical. Indeed, summation is usually stronger when the components A and B are from different sensory modalities (e.g., auditory and visual) (Kehoe *et al.*, 1994; Redhead, 2007; Thein *et al.*, 2008) but not with stimuli from the same modality (e.g., two visual stimuli) (Aydin and Pearce, 1995; Aydin and Pearce, 1997). Apart from modality, spatial and temporal contiguity can also impact summation (Glautier, 2002; Glautier *et al.*, 2010). More recent developments of elemental and configural models have incorporated similarity as a critical factor determining the level of summation obtained (Soto *et al.*, 2014b;

Thorwart *et al.*, 2012; Wagner, 2008). Under these flexible models, dissimilar stimuli (such as those coming from different modalities) will prompt elemental processing and a level of summation closer to the one predicted by elemental models. Configural processing, on the other hand, will be prompted by similar stimuli (such as those coming from the same modality) and summation should be weaker and closer to the average of the predictive value of each of the components alone (Harris, 2006; McLaren and Mackintosh, 2002; Melchers *et al.*, 2008; Pérez *et al.*, 2018; Soto *et al.*, 2014a; Thorwart *et al.*, 2012; Wagner, 2008). More recent models predict that, apart from perceptual similarity, higher spatial or temporal contiguity should encourage configural processing and reduce the summation effect (Soto *et al.*, 2014b; see also Thorwart *et al.*, 2012).

In spite of being a widely-accepted factor in modulating summation in these models, obtaining an effect of similarity on summation has been an extremely difficult task in humans. In a recent series of studies, we tested this prediction and found no evidence for similarity affecting summation across several manipulations of similarity, including color, shape and spatial contiguity (Pérez *et al.*, 2018). Across six studies, we consistently observed our participants showing strong summation, disregarding similarity and spatial contiguity when making their predictions. In that work, we concluded that the properties of the visual stimuli where such that participants could easily identify them on the screen during the test regardless of manipulations of similarity. Under those conditions, participants can simply rely on a strategy where they identify the number of stimuli on the screen, assume them as independent causes of the outcome, and invariably sum the individual predictions of each stimulus in the compound (Pérez *et al.*, 2018).

The goal of Experiment 1 was to avoid that strategy. To this end, we designed a task where the number of stimuli on the screen during both training and testing was the same. In addition, we increased the complexity of the compound stimuli to make it more difficult

for participants to distinguish between them. We found that summation can be weakened and modulated by similarity in this new design. In Experiments 2 and 3 we found a potential mechanism for this modulation: limited sampling and processing of stimuli within compounds. We formalize this limited sampling idea in the last section of the paper and demonstrate how this approach can provide an additional interpretation for a wide range of data on generalization across different species.

Results

Experiment 1

In the three experiments reported in this paper participants were presented with a task in which they were asked to play the role of an allergist who had to predict the levels of allergy caused by different drugs in a hypothetical patient, Mr. X (Pérez *et al.*, 2018; Soto *et al.*, 2009). During training, a configuration comprising one of the target and the two non-target cues was presented at the center of the screen (see Figure 1). Participants were required to give an assessment of the level of allergy that each drug would cause in a fictitious patient, Mr. X, on a scale of 0 to 35. After clicking with the mouse on the rating scale to make their predictions, participants had to confirm them by clicking on a button below the rating scale. Cues were presented 30 times each, in randomized order. Trials were self-paced.

During the training stage, and after entering their prediction and confirming it, participants were presented with a feedback screen with the message “Correct!” if participants’ estimation was correct, or the message “Incorrect!” if the estimation was incorrect (see Figure 2). After, they were presented with the outcome message “Allergic reaction: 10 points in the allergy scale” for cues that predicted allergy or “No allergic reaction: 0 points in the allergy scale” for cues that did not predict allergy. To encourage learning and attention to the task, each incorrect assessment was followed by a 5-second message while only a 1-second message was employed for correct assessments. The testing

stage was similar, with two key differences. First, participants were presented not only with the same stimuli they had seen during training (AXY, BXY, CXY, DXY), but also, to test the summation effect, with stimuli including A and B in the same configuration (ABX or ABY). Second, no feedback was provided during the test stage. Each stimulus was presented two times during the test stage.

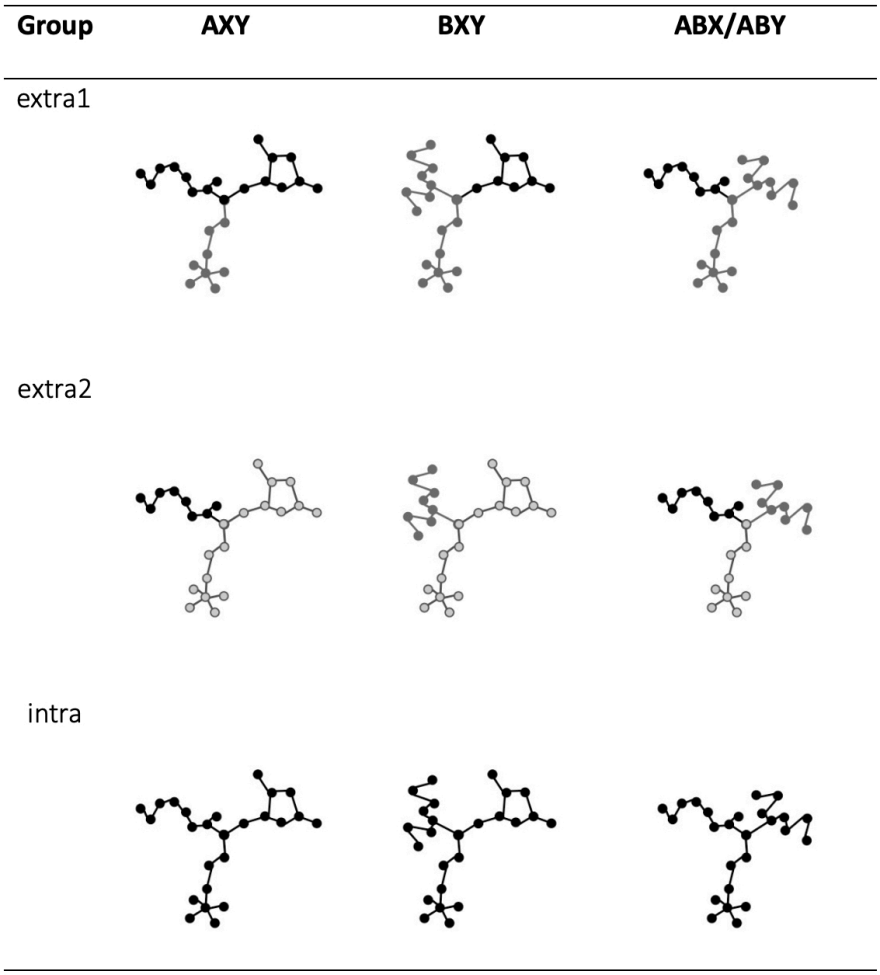

Figure 1. **Stimuli used in the experiments reported in this paper.** Each stimulus was formed by a central point which branched out into three different components. Target cues A and B were accompanied by two other non-target cues, X and Y. The compound AB was accompanied by one of the non-target stimuli X or Y. All stimuli were presented in three different planar orientations, so that each component was equally likely to appear in one of three positions. Cues A, B, C and D were the target cues. Cues X and Y were contextual or non-target cues that accompanied target cues in each configuration. Stimuli in group *intra* varied only in shape. Stimuli in groups *extra1* and *extra2* varied in both shape and color.

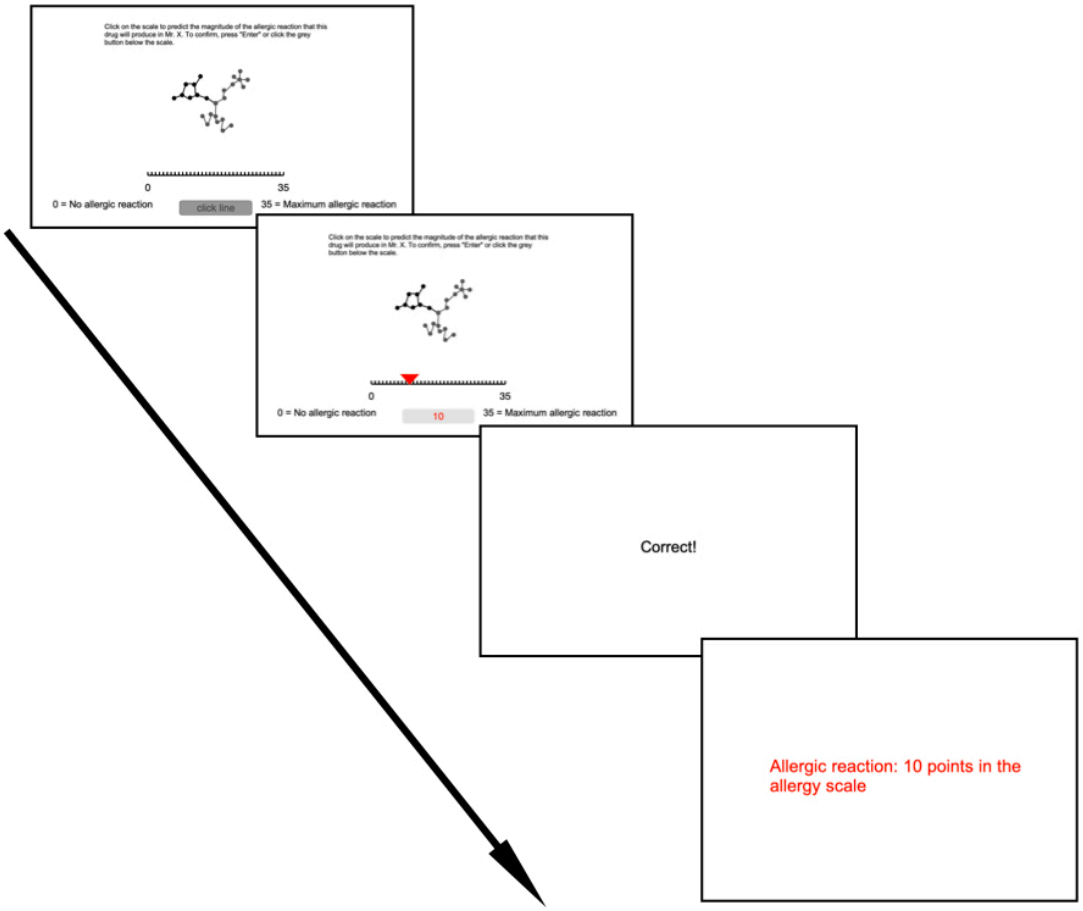

Figure 2. **Trial design for all experiments.** In each training trial, participants observed one target cue (A,B, C or D) and two non-target cues (X and Y) in the same configuration (AXY, BXY, CXY, DXY) and were asked to rate the level of allergy they thought would be produced by the stimulus on a scale from 0 to 35. After an inter-stimulus interval, feedback was presented ("Correct!" if the prediction was correct, "Incorrect!" if the prediction was incorrect), followed by the message "Allergic reaction: 10 points in the allergy scale" for cues that predicted allergy or "No allergic reaction: 0 points in the allergy scale" for cues that did not predicted allergy. During the test phase, the cues were presented in the same way as in training, but now the compound AB, together with one non-target cue (X or Y) was presented in some trials (i.e., ABX, ABY). The predictions were assessed in the same way as in training but no feedback was given during the test stage.

In Experiment 1, we hypothesized that the strong and robust summation found in previous human studies (Pérez *et al.*, 2018) resulted from the use of cues that could be easily parsed by participants. Thanks to properties like their reliable spatial position and the absence of non-target cues, a simple counting strategy brings about reliable and strong summation independently of similarity. Therefore, a task where stimuli are part of the same configuration consisting of a fixed number of components should impair participants' deployment of a counting strategy.

To test this prediction, in all trials we presented participants with a configuration of three cues, so that each trial included one target cue (A, B, C or D) and two non-target cues (X and Y). As shown in Figure 1, all stimuli consisted of the three cues joined at a central point. In training trials, one of the target cues (A or B) was presented together with the two non-target cues X and Y. To test for summation, during the test stage in some trials both A and B were presented in the same configuration, but with only one of the non-target cues (either X or Y, counterbalanced) so that again the stimuli consisted of three components. The two additional target cues, C and D, were associated with no outcome (0 points of allergy). These cues were included to check which participants correctly learned the contingencies between the cues and their respective outcomes.

Apart from the change in the number of stimuli between training and testing (1 in training, 2 in testing), in the studies reported in Perez et al. (2018) the two components A and B were each presented in a specific spatial position, which may have also facilitated parsing them into independent cues, prompting a counting strategy and a strong summation effect. To avoid such strategy, we rotated the stimuli across trials so that each cue was presented in each of the three possible spatial positions within the configuration. This ensured that the number of stimuli on each configuration was the same for both training and test phases, and that the spatial position of the elements was irrelevant for solving the task. Under this design, participants should not be able to easily deploy a

counting strategy to predict the level of outcome of the compound ABX/ABY during the test.

We manipulated the similarity between cues within a configuration in three different groups. Group *intra* included cues that varied only in shape (i.e., intra-dimensional differences, see Figure 1). For the two *extra* groups, cues varied both in shape and in color (i.e., extra-dimensional differences). In group *extra1*, A and B had different color (black vs gray), but the colors were shared with the non-target cues X and Y (one of them black and the other gray, see top-left stimulus in Figure 1). This means that although A and B can be easily distinguished, they cannot be easily distinguishable from the non-target cues X and Y. Finally, in group *extra2*, A and B had different color (black vs. gray), and they also differed in color with the non-target cues X and Y (light gray with black contours). Provided that participants were using a counting strategy in all previous studies, we expected that summation will be weaker under these conditions. In particular, we expected no summation for group *intra* and positive summation in groups *extra1* than *extra2*.

Figure 3 presents the mean causal ratings to the trained (AXY and BXY) and novel (mean of ABX and ABY) cues in the three groups of Experiment 1. The figure indicates that there was a noticeable summation effect in the form of higher ratings to the novel cue than to trained cues in group *extra2*, but no evident summation in similar comparison in groups *extra1* and *intra*. A 3 (group) X 3 (cue) mixed design ANOVA confirmed the reliability of these observations in showing a reliable group x cue interaction ($F(2, 55.5) = 4.00, p = .02, \eta_p^2 = 0.13, 90\%CI\ [.04, .17]$). Subsequent simple effects of cue in each group revealed that the mean ratings to ABX/ABY reliably differed from ratings to the trained AXY and to BXY in group *extra2* ($t(26) = 4.820, p < .001, D = 0.773, 95\%CI\ [2.813, 6.916]$ and $t(26) = 4.815, p < .001, D = 0.773, 95\%CI\ [2.832, 6.797]$, respectively) but did not in groups *extra1* ($t(12) = 0.77, p = .46, BF_{01} = 2.79$ in both cases) and *intra*

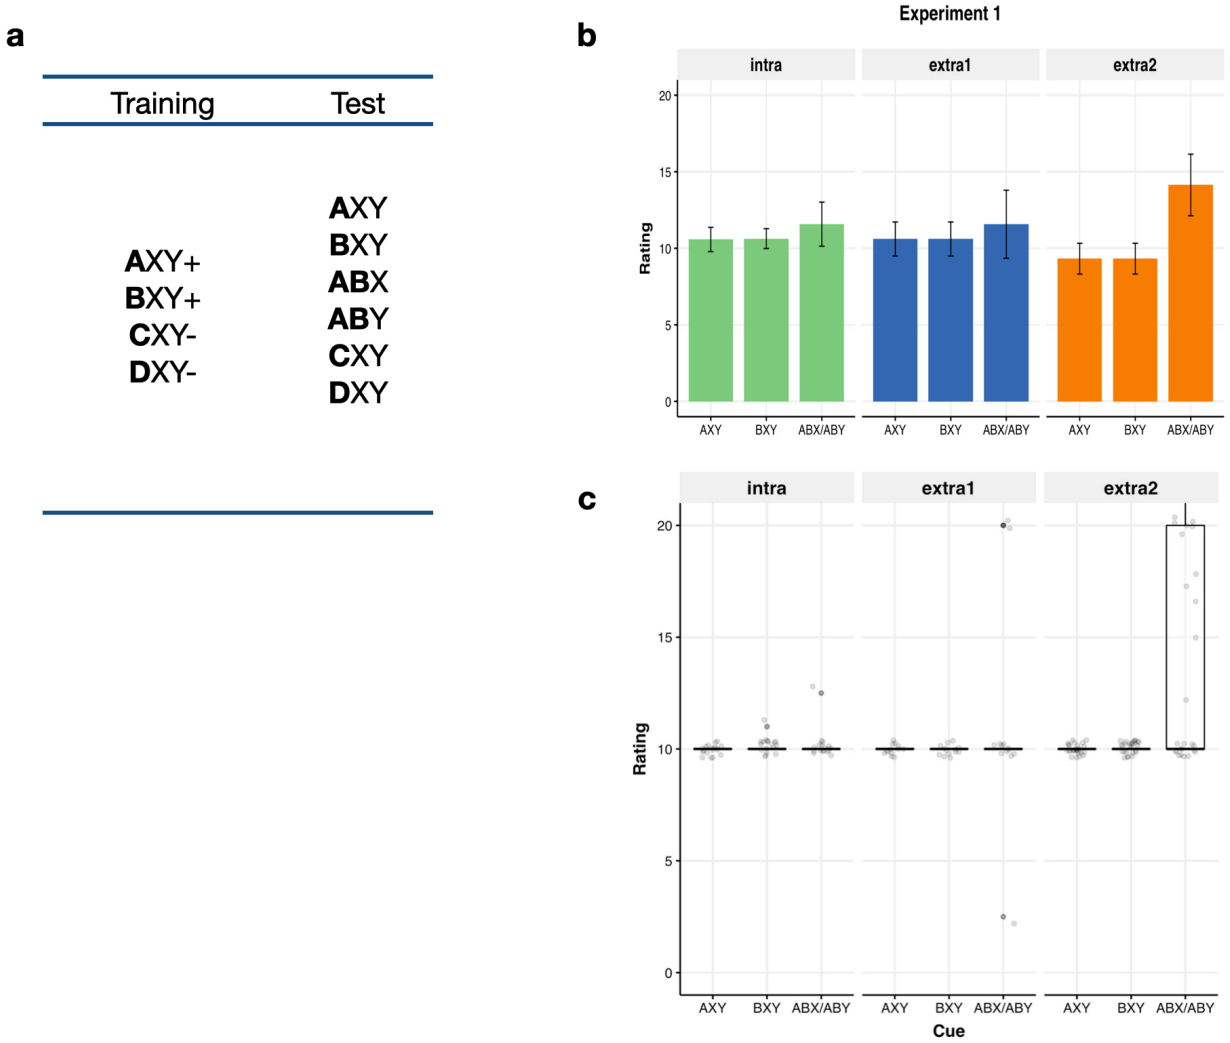

Figure 3. **Design and results of Experiment 1**. **(a)** Letters denote cues, represented by different chemical shapes that can cause different levels of allergy to a fictitious patient. Each cue is represented by a capital letter. Cues A, B, C and D were the target cues (shown in bold), whereas X and Y were non-target cues. In Experiment 1, all cues were followed (+) or not followed (-) by the same level of allergy (10 points of allergy out of a total of 35). **(b)** Average ratings given to each cue during the test. Error bars represent 95% within-subject confidence intervals. **(c)** Boxplots and individual ratings (shown in jittered points) given to each cue during the test.

$(t(17) = 1.16, p = .26, BF_{01} = 2.30$, and $t(17) = 1.17, p = .26, BF_{01} = 2.28$, respectively).

The main effects of cue and group were also reliable $(F(1, 55.05) = 14.97, p < .001, \eta_p^2 = 0.21, 90\%CI\ [.10, .31]$ and $F(2, 55) = 3.80, p = .03, \eta_p^2 = 0.12, 90\%CI\ [.01, .25]$, respectively), but they are mainly accounted by the summation effect of group *extra2* described above. Moreover, the summation and absence of summation effects observed in this experiment can be further illustrated through a summation index computed as the mean causal rating for novel compound (ABX and ABY) minus the mean causal ratings for the trained compounds (AXY and BXY). The outcome of this computation is summarized in Table 1, which indicates that the summation effect was not significantly different from zero in group *extra1* $(t(12) = 0.80, p = .46, BF_{01} = 2.80)$ and group *intra* $(t(17) = 1.16, p = .26, BF_{01} = 2.30)$. By contrast, a significant summation effect was found for group *extra2* $(t(26) = 4.02, p < .001, BF_{01} = 0.01)$ The effect of similarity between target and non-target cues on summation was evident in a detectable difference between groups *extra2* and *extra1* $(t(31.77) = 2.22, p = .03, D = 0.67, 95\%\ CI[-0, 01, 1.35])$, and between *extra2* and *intra* $(t(42.50) = 4.54, p = .01, D = 0.71, 95\%CI\ [0.09, 1.32])$. There was no detectable difference in summation between groups *extra1* and *intra* $(t(22.55) = -0.03, p = .98, D = -0.01, 95\%CI\ [-0.75, 0.74])$

Group	N	Summation	ci_low	ci_high
extra1	13	1.57	-0.42	3.56
extra2	27	4.13	2.33	5.93
intra	18	1.58	0.27	2.89

Table 1

Mean and 95% confidence intervals for the summation effect for all groups of Experiment 1. Confidence intervals were calculated following the method proposed by Morey (2008).

To our knowledge, this is the first demonstration of humans showing little to no summation in a predictive learning task in absence of manipulations aimed to prompt rational rules that neutralize prior assumptions of independence between the cues (group *intra*). In addition, we found an effect of similarity on summation, indicated by a higher summation effect in group *extra2* than in group *extra1*. The results of this experiment agree with our hypothesis that stimulus similarity affects the summation effect when participants cannot deploy a counting strategy to assess their predictions of the compound ABX/ABY.

Contemporary learning models with flexible representations assume that the key variable controlling the summation effect is the similarity between the target cues A and B. However, in contrast to this assumption, the main driver of the summation effect in Experiment 1 was the similarity between the targets A and B and non-targets cues X and Y, rather than the similarity between the two targets A and B. In group *extra1*, these two cues were as dissimilar as in group *extra2*. If the effect of cue similarity acted through the mechanisms proposed by these learning models, then we should have obtained a similar summation effect in these two groups. On the other hand, it is also true that the design we have employed here differs considerably from a regular summation design, in that we are employing contextual or common cues during the course of training. These common cues were partially reinforced during training, and thus it is expected that they develop some association with the outcome. Moreover, any model accepting some way of coding for stimulus conjunctions predicts that after AXY+ and BXY+ there should be some ”external inhibition”, or negative summation, when testing ABX and ABY, due to the replacement of those components that code for the conjoint occurrence of A, X and Y, and of B, X and Y, by new elements coding for the conjoint occurrence of A, B and X, and of A, B and Y. This loss of association will counteract to some extent the gain of association provided by the addition of A and B. It is nevertheless still conceivable that learning models incorporating flexible representations based on the similarity between the components within a configuration could describe our data for some combination of values

in their parametric space.

To explore this possibility, we simulated the predictions of the most influential of these models: the Replaced Element Model (REM) proposed by Wagner and his colleagues (Brandon *et al.*, 2000; Vogel *et al.*, 2017; Wagner, 2008). In short, REM assumes that stimuli are formed by a finite but large number of elements, each of which accrues learning by reward prediction-error. Unlike classic elemental models such as Rescorla-Wagner (Rescorla and Wagner, 1972), REM assumes that the presentation of a given cue activates two types of elements: Context-independent and context-dependent elements. The context independent elements are activated whenever the target cue is presented independently on whether other cues are present or absent. In turn, context-dependent elements, apart from the cue they represent, are activated only when other specific cues are present or absent. In a typical summation experiment, for example, cue A would be represented by a certain number of context-independent elements, denoted as a_i, and a certain number of context-dependent elements denoted as $a_{\neg b}$ and, likewise, cue B would be represented by b_i and $b_{\neg a}$. In the summation test, context-dependent elements of each cue will be replaced so that AB will be now represented by a_i, b_i, a_b and b_a, where a_b and b_a are the elements replacing the context-dependent elements $a_{\neg b}$ and $b_{\neg a}$. By this strategy, REM can generate outcomes that would look either "elemental" (i.e., AB>A, B) or "configural" (i.e., AB=A, B) depending on proportion of context-dependent elements over the total number of elements representing cues A and B, which is given by the free parameter r ($0 < r < 1$). Wagner and his collaborators have suggested that r might represent the degree of similarity between a pair of cues and recommended the use of r values ranging between 0.2 to 0.3 for simulating standard phenomena of Pavlovian conditioning (Brandon *et al.*, 2000; Vogel *et al.*, 2017; Wagner, 2008).

These predictions, however, do not necessarily follow from more complex designs such as the one employed in our experiment (where target cues were consistently reinforced/not

reinforced with contextual or non-target cues) as an assumption is required regarding the similarity between pairs of components in each configuration. Figure 4a depicts how we treated in our simulations the elements in a compound in each group. In group *intra*, all cues were of the same color. By contrast, cues A, C and X differed in color with B, D and Y in group *extra1*. In group *extra2*, X and Y were of the same color, which in turn differed from cues A, C and cues B, D. Figure 4b shows in matrix form the assumptions made in terms of the replacement parameter in each group. First, in order to give the model a chance to produce summation, we made the very conservative assumption of $r = 0$ for any pair of cues of different color. With this assumption, the groups would differ only in the number of cues having the same color, in which case r was varied across a broad parametric range from 0 (no replacement) to 1 (full replacement) in steps of 0.02.

Figure 4c presents the predictions of REM for the three groups reported in Experiment 1. In general, the model predicts a weaker summation effect under the common-cue design compared to a simple summation design. As expected, summation decreases monotonically for all groups as the number of elements replaced increases. Note that when $r = 0$ (that is, assuming no perceptual interaction between the cues forming the compounds), REM predicts equal summation in all groups. Furthermore, negative summation (i.e., external inhibition; Pavlov, 1927) is predicted by REM for values of $r >= 0.1$ in group *intra*, whereas positive summation is predicted for the group *extra1* across the same range of r values. When $r = 0.1$, the model predicts no summation in group *intra* but positive summation in both *extra* groups. Thus, although REM replicates the order of responding of Experiment 1, at no point in the parametric space does it predict zero summation for groups *intra* and *extra1* and, when the latter group approaches zero, it does it in the range of negative summation for *intra*.

The results of these simulations prompted us to search for alternative explanations for the data obtained in this experiment. To this end, we took a closer look at the

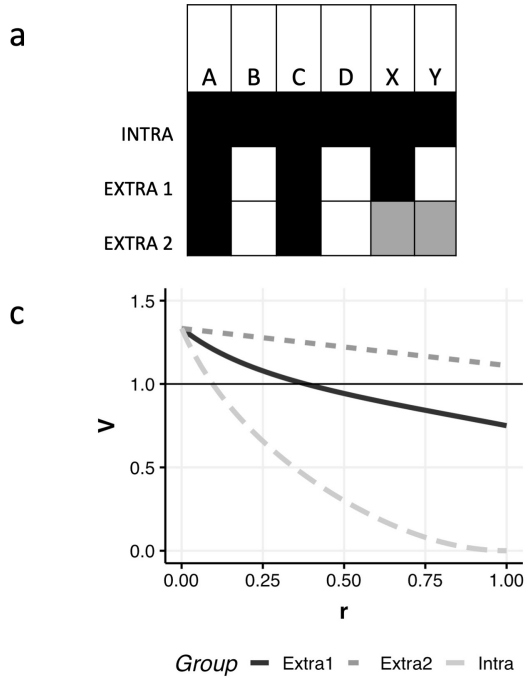

b

Intra

	B	C	D	X	Y
A	r	r	r	r	r
B		r	r	r	r
C			r	r	r
D				r	r
X					r

Extra 1

	B	C	D	X	Y
A	0	r	0	0	0
B		0	r	0	0
C			0	r	0
D				0	r
X					0

Extra 2

	B	C	D	X	Y
A	0	r	0	0	0
B		0	r	0	0
C			0	0	0
D				0	0
X					r

Figure 4. **Simulations of the REM model for Experiment 1. (a)** Schematic representation of the cues that have the same or different color in the three groups of Experiment 1. Notice that the black, white and grays colors were used for illustrative purposes only (see method section for details). **(b)** r values for different pairs of cues in the three groups. It was assumed that the level of replacement (r) of different-color cues was equal to zero. **(c)** Predicted asymptotic associative strength for a summation with common cues experiment (AXY+BXY+CXY-DXY-, test with AXY, BXY, ABX and ABY) as a function of the level of replacement of same-color cues according to the REM model. The simulations were conducted with the ALTSim simulator (Thorwart *et al.*, 2009) using the learning rule suggested by Schulteis et al. (2008). The parameters employed were $\alpha = 0.1, \beta^{+} = 0.2, \beta^{-} = 0.04$.

distribution of scores in each group and found a very interesting pattern of results. This is summarized in Figure 3c, which depicts the boxplots and individual scores for each group of Experiment 1. As can be appreciated, although all groups exhibited the same median of 10 points of allergy for all cues, they differed in the distribution of scores to ABX/ABY. Consider, for instance, the very thin boxes for the ratings to all cues in groups *intra* and *extra1*, which indicate that about 75% of the ratings are contained within a very small range around 10 points of allergy. This pattern differs noticeably from the wider box of group *extra2*, where only half of the participants rated ABX/ABY with 10 points of allergy. Of course, the higher summation observed in this group is explained by a significant

proportion of participants rating ABX/ABY higher than the median. Accordingly, a Friedman's test showed a significant difference in the distribution of the scores across cues in group *extra2* ($X^2(2) = 26.00, p < .001$) but not in groups *intra* ($X^2(2) = 3.71, p = .16$) and *extra2* ($X^2(2) = 0.67, p = .72$). Dunn-Bonferroni post-hoc tests were carried out in group *extra2* showing significantly higher scores to ABX/ABY than to AXY ($p = 0.02$) and to BXY ($p = 0.02$) after Bonferroni adjustments.

However, what is most noticeable in the boxplots of Figure 3c is the fact that a non-negligible number of participants rated ABX/ABY with scores around 20 points of allergy in all groups. Thus, it seems that participants rated ABX/ABY with either 10 or 20 points resulting in bimodal distributions (Hardigan's Dip Test: $D = 0.10, p < .001$) These descriptive observations about bimodality were reinforced by a mixture model which yielded a generative model with two Gaussians for the ratings of AB across all groups (Gaussian 1: $\mu_1 = 9.75, \sigma = 2.65$; Gaussian 2: $\mu_2 = 22.10, \sigma = 2.65$). In sum, it appears that the summation effect was determined by two distinguishable sets of participants: the majority, who gave ratings of exactly 10—the rating predicted by AXY or BXY alone—and the rest of participants, who gave ratings of exactly 20, which is the outcome that should be given to ABX/ABY if they were processing both A and B independently, just as in our previous reports of summation in humans (Perez *et al.*, 2018).

These observations led us to hypothesize that participants' lack of summation in groups *intra* and *extra1* may have been a consequence of limited processing of cues, in that they might be sampling one of the cues A or B during the test phase when presented with the compound ABX/ABY, ignoring the rest of the components of the stimulus. This, of course, should be a direct consequence of the experimental design, which was aimed to make it harder for participants to parse the stimuli during the task, thereby avoiding the counting strategy we thought was behind the consistent summation observed in our previous experiments (Perez *et al.*, 2018). By contrast, the higher dissimilarity between

cues in the compound ABX/ABY in group *extra2* may have allowed participants to process and sample all the cues in the compound. The following two experiments were designed to explore this possibility.

Experiment 2

To test the sub-sampling hypothesis, in Experiment 2 we used the same stimuli as in group *intra* of Experiment 1 but assigned different outcome values (i.e., allergy levels) to cues AXY and BXY in different groups (see Figure 5a). Group *intra* was a replication of the same group from the previous experiment in which the outcome associated to both AXY and BXY was 10 points of allergy. We expected the majority of participants to rate the compounds ABX and ABY in group *intra* as producing 10 points of allergy, replicating the lack of summation of Experiment 1. By contrast, group *intra2* involved two different outcome values: 10 points of allergy for cue AXY and 8 points of allergy for cue BXY. To the extent that our stimulus manipulation prompts limited sampling and processing of a single cue, the majority of participants in group *intra2* should rate the compound ABX/ABY as producing either 8 or 10 points of allergy, indicating that they have responded in accord with the value of only one of the two components.

The predictions of this sub-sampling hypothesis differ from those of contemporary associative theories in two respects. First, with their usual parametric assumptions, associative approaches like Wagner (2008)'s REM and Pearce (1994)'s configural approach predict lower response to ABX/ABY than to both AXY and BYX in both groups (*intra* and *intra1*) (see Figure 4). By contrast, to the extent that participants engage in a sub-sampling process, we should expect the mean of ABX and ABY to be approximately equal to the mean of AXY and BXY in both groups. Second, and perhaps most importantly, associative theories do not make a clear prediction with respect to the relationship between the ratings to ABX/ABY and each of AXY and BXY, which are reinforced towards different asymptotes in group *intra2* (AXY to 10 and BXY to 8).

According to the sub-sampling hypothesis, however, provided that ratings to AXY differ from BXY in group *intra2*, it is expected that ABX/ABY should be rated equally to one of the reinforced cues and differently from the other.

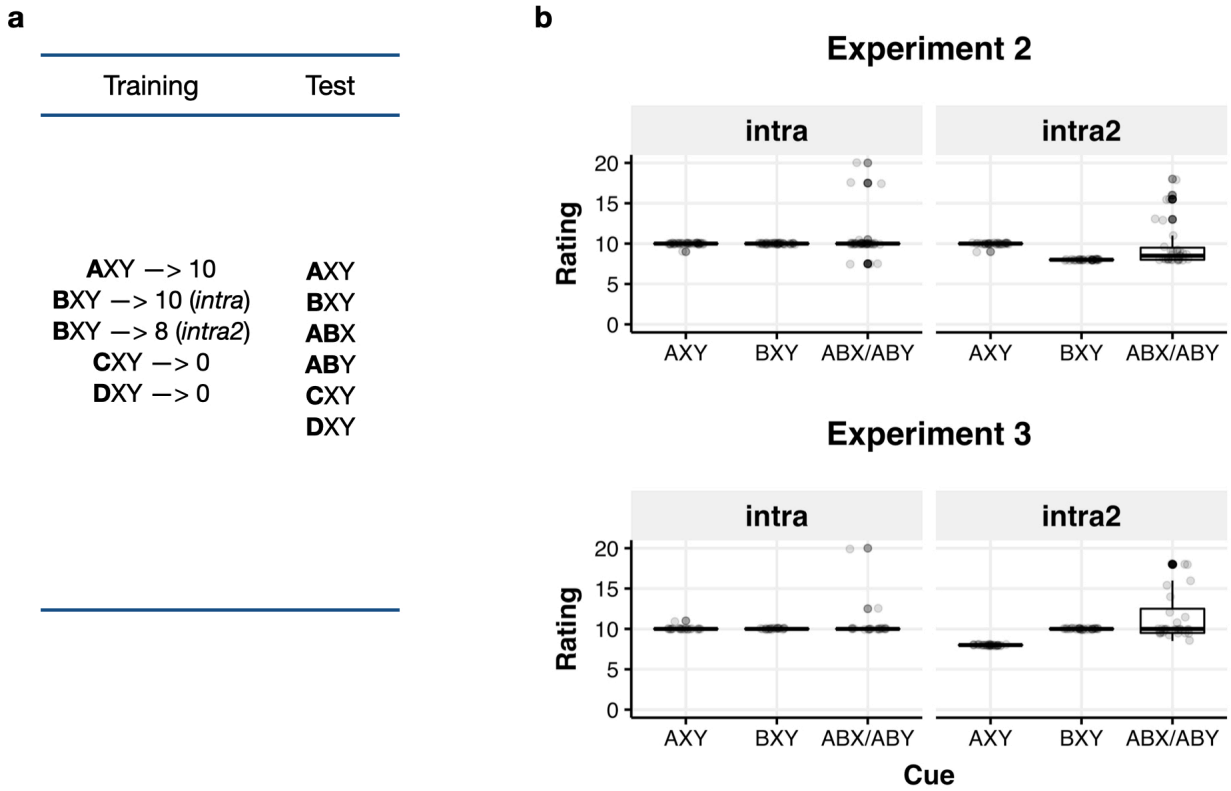

Figure 5. **Design and results of Experiments 2 and 3**. **(a)** In Experiments 2 and 3, cues were followed by different levels of allergy, which are represented by the numbers shown next to each of them. The only difference between Experiments 2 and 3 was the assignment of allergy levels to cues AXY and BXY in group *intra2*: in Experiment 3 the outcomes of AXY and BXY were swapped, so that they predicted 8 and 10 points of allergy, respectively. **(b)** Boxplots and individual ratings (shown in jittered points) given to each cue during the test.

Since the main interest of this experiment was to test whether participants rate the ABX/ABY compound similarly to either AXY or BXY, we focus our analysis on the distribution of scores and the non-parametric Friedman's test. The top plot of Figure 5b presents the boxplot and individual scores for the trained (AXY and BXY) and novel (mean of ABX and ABY) cues in the two groups of Experiment 2. As expected, most participants rated AXY and BXY in accordance with their trained outcomes (10 for AXY

in both groups, and 8 and 10 for BXY in groups *intra* and *intra1*, respectively). In group *intra*, median scores were virtually identical for cues AXY, BXY and ABX/ABY with very little variation around 10 points of allergy, which, of course, led to no reliable differences in the three distributions ($X^2(2) = 0.75, p = .16$). As seen in the figure, reliable differences were apparent in group *intra2* ($X^2(2) = 44.69, p < .001$), where the scores to ABX/ABY (med=8.5) were closer to BXY (med=8) than to AXY (med=10). Dunn-Bonferroni post-hoc tests indicated that the distribution of ABX/ABY differed significantly from both AXY ($p = .01$) and BXY ($p =< .01$) after Bonferroni adjustments. These results are in line with the sub-sampling hypothesis, in that the predictive value that participants assigned to the novel compound was based on the predictive value of just one of its constituent elements rather than on a combination of the two.

While these data are in agreement with our sub-sampling hypothesis in that the majority of participants gave ABX/ABY a rating equal to that of BXY, a question arises as to why this might have been the case. One simple possibility is that stimulus B was somehow more salient than stimulus A, attracting more attention during the task. Experiment 3 was designed to confirm this possibility.

Experiment 3

The aim of Experiment 3 was to further test the sub-sampling approach based on the hypothesis that B was a more salient stimulus for participants in group *intra2* of Experiment 2. To this end, we swapped the outcome values produced by AXY and BXY in group *intra2*, so that AXY predicted 8 points of allergy and BXY 10 points of allergy. To the extent that B is truly a more salient stimulus in attracting participant's attention, the majority of participants should give a higher rating to the compound ABX/ABY (closer to 10 rather than 8) in response to the compound presentation during the test.

The results of Experiment 3 are summarized in the bottom plot of Figure 5c. As can

be appreciated, the outcome of this experiment is basically the mirror image of Experiment 2: In the case of group *intra*, we replicated for the third time the absence of any reliable difference among the cues ($X^2(2) = 2.46, p = .30$) which varied in a very small range around 10 points of allergy. In contrast, in group *intra2* the Friedman's test indicate reliable differences among the cues ($X^2(2) = 55.00, p < .001$), where the ratings to ABX/ABY (med=10) were, again, very similar to BXY (med=10) and different from AXY (med=8) ($p = 1.00$ and $p < .001$, respectively, according to Dunn-Bonferroni pairwise comparisons after Bonferroni adjustment). Thus, ratings to the novel compounds ABX and ABY seem to be determined by the predictive value of BXY in this experiment too. This change in the distribution of ratings from Experiment 2 to Experiment 3 is in agreement with our hypothesis that the compound ABX/ABY is rated by the majority of participants as equal to the outcome predicted by BXY due to this cue being more salient and attracting attention during a sampling process.

Taken together, the studies reported here provide good evidence for the hypothesis that sub-sampling of cues can bring about weak to no summation in humans. We speculated that this process might explain several results in the literature usually attributed to configural processing of stimuli, including those in which no summation is found. In the next section we present a formal model describing how sub-sampling might occur in a variety of tasks in which configural processing seemed to have occurred.

A model of limited cue sampling

In this section we formalize a computational model based on the idea that both human and animals could be sub-sampling cues in generalization experiments where similarity or any other factor makes it hard for them to parse and identify the cues individually, such as in the case of groups *intra* reported here. As we have discussed, we believe this might be the main reason why summation was not obtained in those participants. We show in addition how many of the previous results on generalization

which are consistent with configural processing can also be captured by a sub-sampling approach [1].

The model we propose learns according to reinforcement prediction-error (Bush and Mosteller, 1951). In this model, which we call the *sub-sampling model*, the predictive value or *associative strength* of stimulus i in trial n, v_i^n, is updated in accord with Equation 1

$$v_i^{n+1} = v_i^n + \alpha_i \beta(\lambda^n - v_i^n) \tag{1}$$

, where α and β are learning rate parameters ($\alpha \in [0, 1], \beta \in [0, 1]$) which represent how much the subject updates the value of cue i in trial n and λ^n is an indicator function taking a value of one if a reinforcer is presented in trial n and zero if it is not. This algorithm assumes that the change in the associative or predictive value of stimulus i is determined by the difference between the observed outcome and the current outcome expected from that stimulus.

The learning rule specifies how the predictive value of each cue should be updated according to the specific arrangement of contingencies between each cue and the reinforcing outcome in the experiment, but it does not commit to any form of representation for instances when the cues are presented in conjunction or in isolation. The assumption that stimuli presented by the experimenter are represented independently and that compounds of stimuli can be fully parsed by the agent is theoretically problematic.

We follow instead an approach offered by Wagner and Rescorla. As in REM, under this modified Rescorla-Wagner model any component can activate its own elemental

[1] In what follows, we will refer to both REM with $r = .5$ and Pearce's model as "configural", recognizing that we are referring to the predictions made by these models and not the type of stimulus representation assumed in their formulations.

representation and acquire its own predictive value in accord with Equation 1, but the critical difference concerns how the agent represents cues in isolation and in compound. When a stimulus is presented in compound with other stimulus, Wagner and Rescorla (1972) proposed that an additional element can enter into an association with the outcome. We assume that this element, which they refer to as a *unique-cue* element, follows the same learning algorithm as in Equation 1.

It is worth noting that the unique cue is usually interpreted as an internal representation of a compound, but this interpretation is not strictly necessary. In fact, in the present model we assume that the unique cue is added not because there is an additional internal representation; we simply recognize that stimuli in a compound hold visual properties that are not present when stimuli are presented in isolation. For example, as long as visual cues are presented in close proximity during compound trials (i.e., either overlapping or close to one another), it is possible that the subject will sample visual information in some areas of the display that are unique to compound trials (e.g., line intersections, corners, etc.), which makes the compound trial different to a single cue trial by the addition of those features in the configuration.

To model the agent's sampling process, in addition to the learning algorithm that updates the value of a cue according to Equation 1 we assume a system which samples in the configuration for a single stimulus to process. We assume that the probability of stimulus i being processed is given by a softmax function incorporating salience and predictive value of each stimulus (Sanderson and Bannerman, 2011; Sutton and Barto, 1998):

$$p_i = \frac{e^{(\eta \alpha_i v_i)}}{\sum_j e^{(\eta \alpha_j v_j)}} \tag{2}$$

, where η is a decisiveness or *temperature* parameter that determines the extent to which

the agent is biased to sample cues with low salience or predictive value ($j = [1, 2, ..., i, ...k]$) [2]. Finally, in each trial the actor samples from the stimulus array $S = [S_1, S_2, ..., S_k]$ for a single stimulus to process. We assume that the sampled cue S_i follows a categorical distribution:

$$Prob[S = S_i|\mathbf{p}] = Categorical(\mathbf{p}) \tag{3}$$

where $\mathbf{p} = [p_1, p_2, ..., p_k]$ is the vector of probabilities given by Equation 2 and k is the number of stimuli presented in a given trial during training and testing. To illustrate how this model operates, see the schematic representation of Figure 6. In this example, we assume that the compound AB is presented to the subject. In this case, the configuration produces an additional unique cue X for this particular combination of cues. Once the agent samples cue A (shown in red), its value is then updated by the critic according to the difference between the current value or prediction (v_i) and the outcome observed (λ) in that trial.

For all the simulations in this section, we assume that the salience of the unique cues is equal to the salience of the target cues; that is, $\alpha_{\text{target}} = \alpha_{\text{nonTarget}}$, and that the order of presentations of different trial types in an experiment is random. The former assumption—that the salience of the cues presented and the unique cues represented by the subject in compound trials are equal—is the assumption generally held by Pearce (1994; 2002) in his modelling of configural theory and is the one that, unless otherwise noted, we follow in all the subsequent simulations in this paper.

[2] Formally, the model should include the absolute value of v, since cues with high inhibitory strength (negative v) should command more attention than other cues in a given array (Parkhurst *et al.*, 2002). However, such implementation is only relevant in a model that allows for negative associative values for cues, which is not the focus of our model.

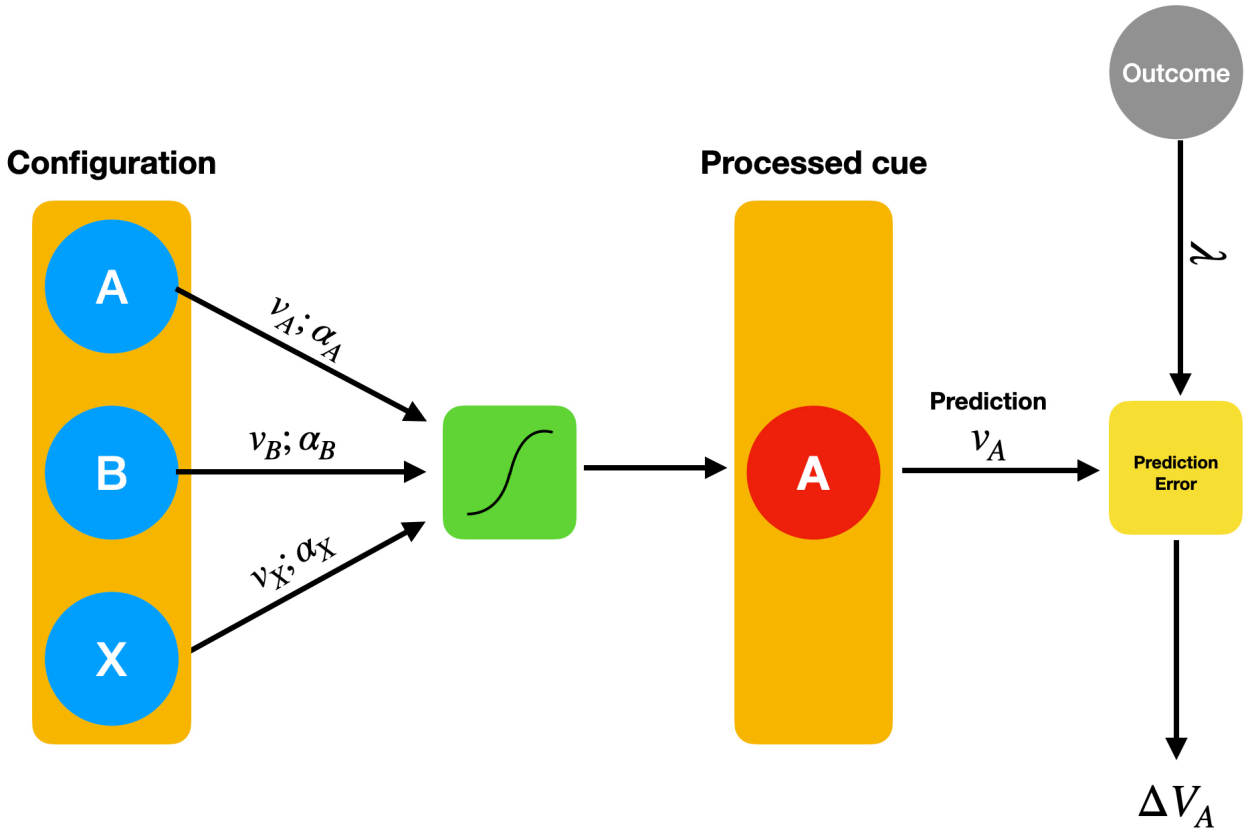

Figure 6. **A sub-sampling model.** On each trial where the agent is presented with a compound of cues that predict an outcome, the agent samples one of the cues according to their current value and salience. In this example, the compound AB is presented together with the unique cue X, which is only present when A and B are presented in compound. For illustrative purposes, we assume that cue A has been sampled in this trial (colored in red). Once the agent samples the cue, it updates its value according to a reward prediction error rule. Only the value of the sampled cue (A, in this example) is updated for the following trial.

To allow the agent to sample cues with low predictive value we set the temperature parameter η to 30 and the initial predictive value of each cue, both target and unique cues, v_i, $i \in [1, ..., k]$, to 0.05. Unless otherwise noted, the value of λ was set to 1 for reinforced trials and to zero for non-reinforced trials. Finally, we assume that the agent perceives an additional unique cue for each one of the possible combinations of cues (for example, if the compound ABC is presented, we assume that there is an additional unique cue represented for each possible pair—AB, AC, BC—and the compound of three cues—ABC. The compound ABC is thus represented as ABCXYZV, where X, Y, Z, and V are the additional unique cues). We ran 80 simulations for each experimental design. The values shown in the following figures are the average values across all simulations.

In this first version of the model, we did not attempt provide a definite solution to the elemental-configural controversy, but rather to highlight some aspects in which a sub-sampling process might be the best explanation of the data. Therefore, to facilitate the exposition, we focused on the simple summation task in which many studies have found absence of summation in animals, rather than the more complex summation design with common cues employed in this paper. We will also test the model against classic paradigms in the animal literature which are usually taken as evidence of configural processing.

Summation

As previously noted, support for configural theory comes from studies in pigeon auto-shaping (Aydin and Pearce, 1995; Aydin and Pearce, 1997). In these studies, a summation effect is usually not found when visual stimuli are presented in compounds. The purpose of this section is to simulate those summation procedures and the results obtained in Experiments 2 and 3 reported here.

We start by replicating the conditions of autoshaping in pigeons. To this end, we assume for simplicity that both components predict the same outcome value and have

equal saliencies ($\alpha_A = \alpha_B = .4$). The results of this simulation are shown in the right panel of Figure 7a. As shown in the figure, the sub-sampling model correctly predicts the failure to find a summation effect in these experiments. The left hand panel of Figure 7a shows a similar design reported by Rescorla and Coldwell using pigeon autoshaping where the authors failed to find a summation effect (1995). Although their results are compatible with a configural approach, the sub-sampling model is also able to predict this result because, during the test with AB, the system randomly samples the unique cue that is only perceived in the testing phase (ABX, where X is the unique cue), which tends to bring down responding to AB (ABX) compared to the elements A and B.

In a second set of simulations, we investigated the predictions of the sub-sampling model when cue B predicts a lower outcome value than A, trying to match the conditions of Experiment 2. As suggested by the data, participants' behavior was driven by the higher salience of B, which led to the majority of them to sample this cue during the test stage. To account for the fact that B is more salient, we set the value of α_A to .4 and the value of α_B to .5. A value of .4 was also set to the unique cue X ($\alpha_X = .4$). To account for different outcome values predicted by each cue, we set $\lambda_A = 1 > \lambda_B = .95$. The model correctly predicts the pattern of results of Experiment 2, in that responding to A, the most salient cue, is higher than to B (see Figure 7b). The model also replicates our finding that responding to AB would be closer to B than to A. Lastly, we tried to match the conditions of Experiment 3 by reversing the roles of A and B, so that B predicts a higher outcome value than A ($\lambda_A = .95, \lambda_B = 1$). Again, the sub-sampling model correctly captures the pattern of behavioral results in this experiment, in that responding to AB should be closer to the outcome predicted by B, and higher than in Experiment 2 (see Figure 7c).

Differential summation

In another experiment in pigeon auto-shaping, Pearce and colleagues (**pearce1997configuqral**) found that responding at test for the compound of three cues,

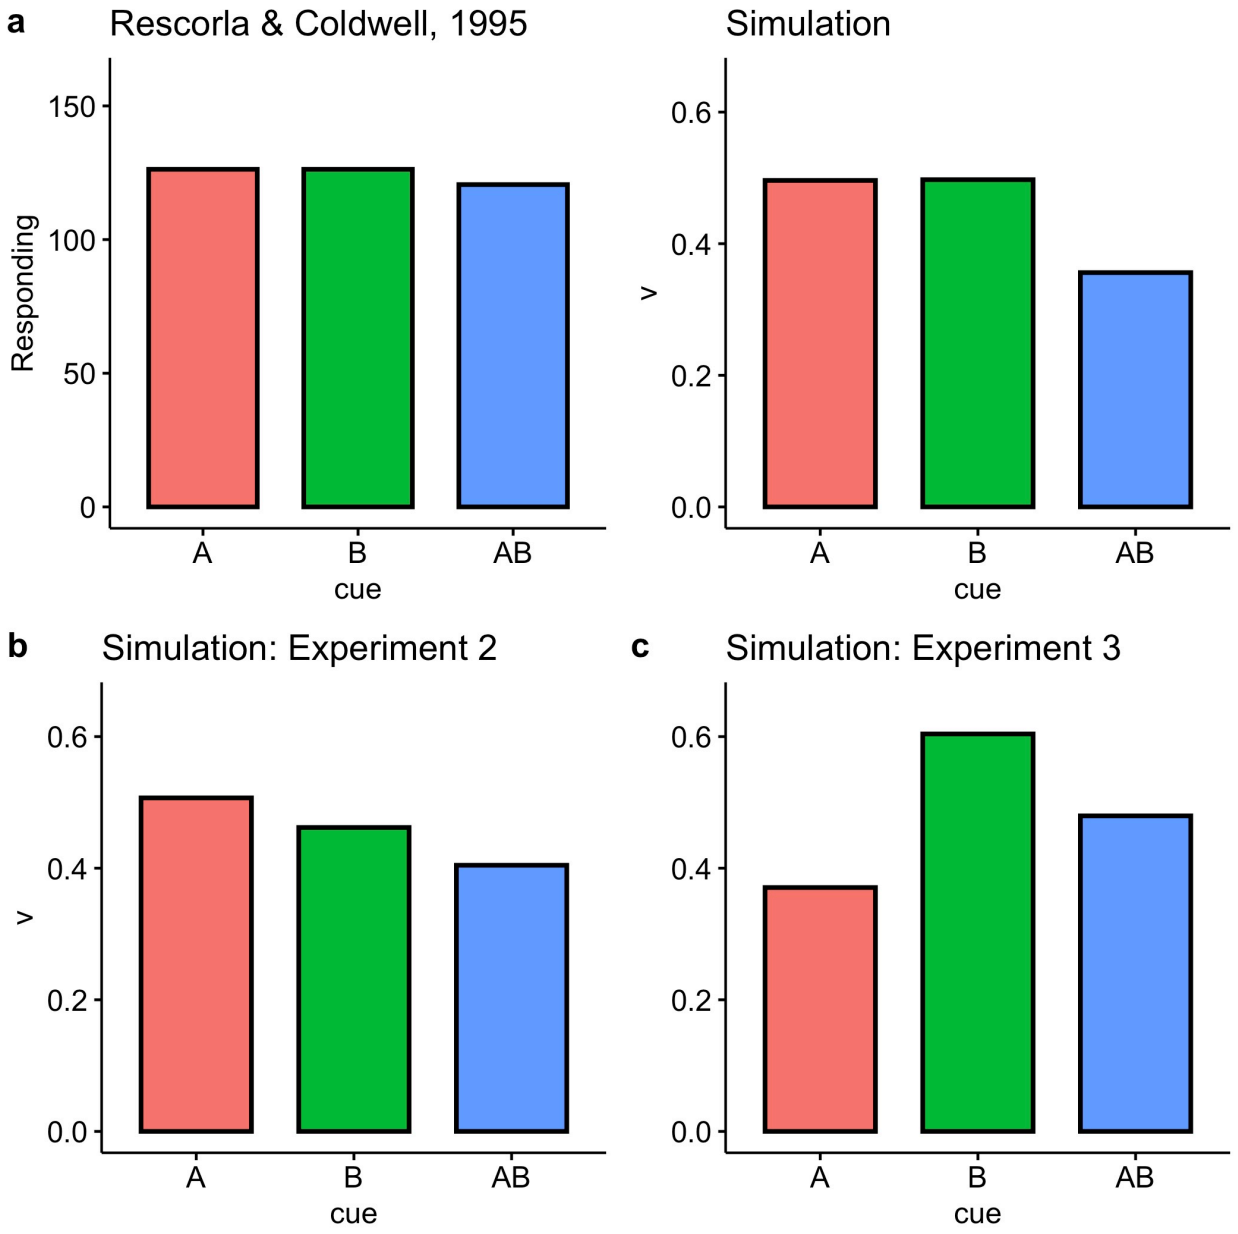

Figure 7. (a) Simulations of the sub-sampling model for different summation designs. The left panel shows the results obtained by Rescorla and Coldwell (1995) using pigeon autoshaping. The right panel shows the simulations of the model assuming that the saliencies of A and B are equal. (b) Results of a simulation of Experiment 2, where the salience of B is assumed to be higher than that of A, but B predicts a lower outcome value. (c) Results of a simulation of Experiment 3, where the salience of B was higher than that of A but the outcome value predicted by B was higher than the value predicted by A.

ABC, was weaker when the three cues were separately paired with a reinforcer (A+, B+, C+) than when the cues were paired with the same reinforcer, but in compounds (AB+, AC+, BC+) (see Figure 8, left panel). Elemental models make the opposite prediction, because compound training makes each cue in the compound acquire only half of the predictive strength as in the case of single training. Pearce's configural model, in contrast, can readily account for these results under the assumption that the similarity between the compounds comprising two cues (AB, AC, BC) and the tested compound ABC is higher than that between each A, B and C and the compound ABC.

Figure 8 (right panel) depicts the simulations of this design in the sub-sampling model, which correctly captures these data. The reason why the sub-sampling model correctly predicts this result concerns the addition of a cue for the compound ABC during testing. In the case of single-cue training, at the end of training ABC is comprised of four unique cues with very low predictive values (.05 each, by assumption), whereas the same ABC compound is comprised of the same four unique cues, some of which have been experienced by the agent during training and have therefore acquired higher predictive values than in the case of single-cue training. During the presentation of ABC at test, the agent sometimes sample these higher-valued unique cues. As a consequence, the agent tends to respond more to the compound ABC after compound training than after single-cue training, in agreement with the results observed by Pearce et al. (1997). In other words, our model proposes that it is not the similarity between two-stimuli compounds like AB and the test compound ABC what increases summation in this design, but rather the availability of configural cues that can be sampled during both training and testing.

Reversing a conditioned inhibitor

Pearce and Wilson ran an experiment which included a negative patterning design of the form A+, AB- in a first phase, and reinforced presentations of B alone in a second phase (B+). In the learning literature, the first phase turns cue B to what is referred to as

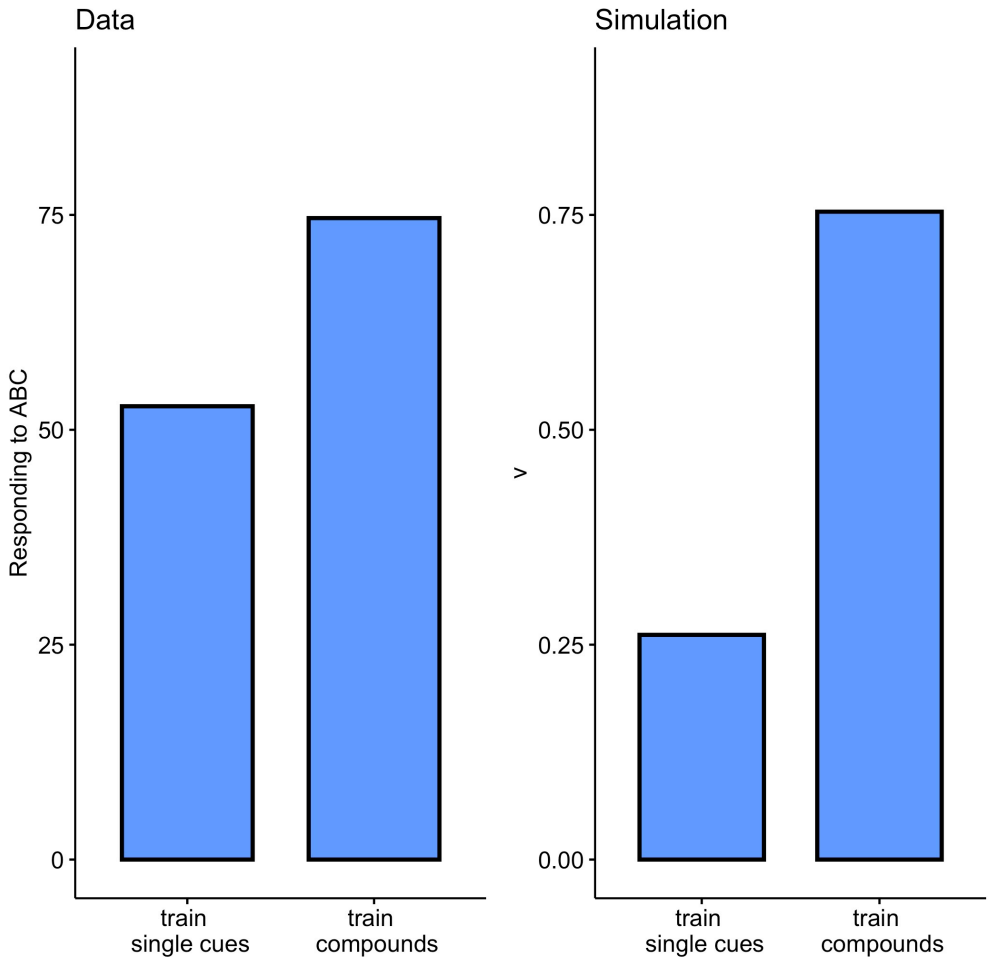

Figure 8. **Simulations of a sub-sampling model for a differential summation design.** Left panel. Results obtained by Pearce and colleagues (1997) in pigeon autoshaping. The bars show the responding to the compound ABC after training with the single cues A, B and C or after training with the compounds AB, BC, and AC. Right panel. Simulations of the sub-sampling model for the same design.

a *conditioned inhibitor*, meaning that B signals the absence of an otherwise present reinforcer. In contrast to an elemental theory which would predict B to recover its predictive value so that responding to AB should be higher than A or B alone at the end of the second phase, Pearce and Wilson observed instead lower responding to AB than A or B alone (see Figure 9, left panel). This result is also anticipated by the sub-sampling model (see Figure 9, right panel). During the first stage, A acquires more value than both B and the unique cue represented for AB, call it X. During the second stage, B acquires value independently of the sub-sampling process. During the final test with AB, however, the agent still samples the unique cue X, whose value has not been modified during the second phase, staying at a low level. Responding to AB is therefore lower than to either A or B.

Negative patterning

In another experiment in pigeons, Pearce and Redhead (1993, Exp. 1, cited in Pearce, 1994) found that subjects mastered a negative patterning discrimination (A+ AB-) more easily than the same discrimination with an added redundant stimulus C (AC+ ABC-). Figure 10a shows the results of this experiment. The left panel shows the results observed after training with A+ and AB- and the right panel shows the results observed after training with AC+ and ABC-. Given that the similarity between A and AB is lower than that of AC and ABC, an elemental approach anticipates that it will be easier for agents to master the discrimination in the group that is trained with A and AB separately than in the group trained with AC and ABC as compounds. Pearce's configural model, by contrast, correctly predicts this result. The similarity between the training set AC and the compound ABC is higher than that between A and AB and, given that similarity determines how much the compound ABC is processed in each trial; it follows that agents should learn to respond less to the compound AB than to the compound ABC during training (see Pearce, 1994; Pearce, 2002)).

It is relatively clear that when a subject's tendency is to explore constantly solving

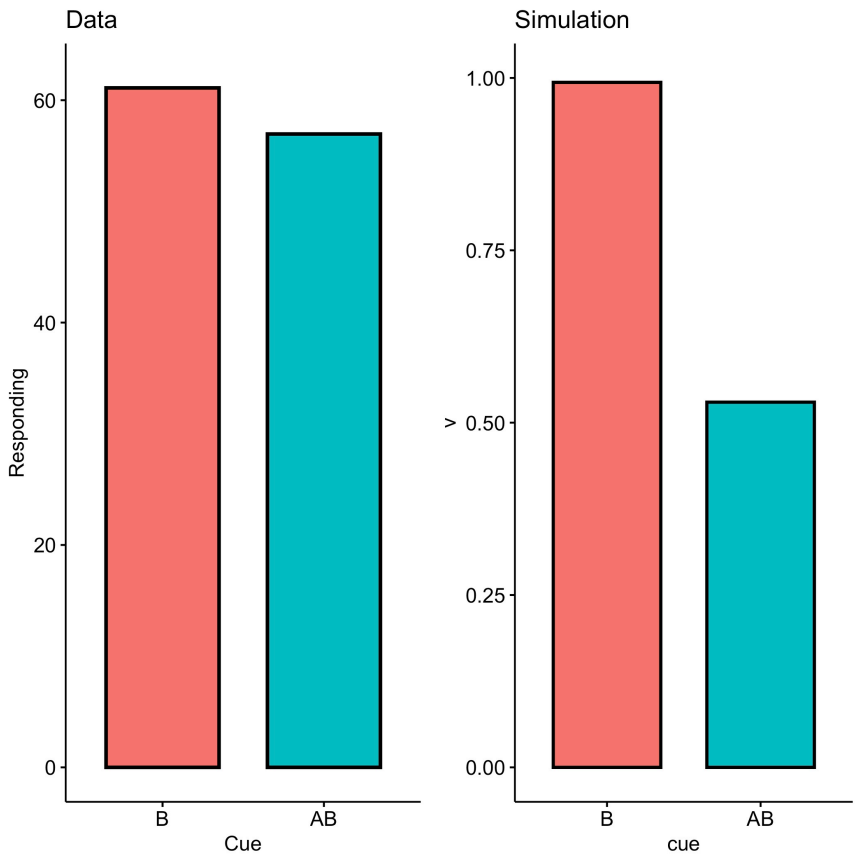

Figure 9. **Simulations of a sub-sampling model for Experiment 3 of Pearce and Wilson (1991).** The experiment comprises a first phase where A is reinforced and the compound AB is not (A+, AB-). In the second phase B is reinforced in isolation. The left panel shows the original data obtained by these authors. Simulations of the sub-sampling model are shown in the right panel.

this problem should be extremely difficult, as there needs to be less sampling of unique cues that have a low value during training. We therefore tested the model assuming a mildly exploratory agent ($\eta = 5$) for which the unique cue for the compound ABC, that we call X, is more salient than the other cues, that is, $\alpha_A = \alpha_B = .4 < \alpha_X = .7$, where X is the unique cue for the compound ABC. Figure 10 shows the results of the simulations for these two negative patterning designs. The left panel shows the predictive value of A and AB during training with a A+, AB- discrimination; the right panel shows a similar simulation for a AC-, ABC- design. The results of this simulation are reasonably in keeping with those obtained by Pearce and Redhead (1993): the difference in value between A and

AB is higher than that between AC and ABC, and the order of values also follows their empirical observations.

This set of simulations provides computational evidence for the hypothesis that sub-sampling of cues could mimic results that are usually attributed to configural learning, including those that fail to find a summation effect, and goes beyond the empirical evidence presented in Experiments 1-3 by showing how more complex designs can be captured by a reinforcement learning rule and a sampling mechanism that processes a single stimulus in a compound.

Discussion

This paper reports the first empirical evidence of similarity affecting summation in humans. In Experiment 1, when stimulus properties were designed to make it difficult for subjects to parse them and deploy counting strategies, the difference in similarity between target and non-target cues within a compound was reflected in different ratings for the compound ABX/ABY during the test. In addition, although previous studies had demonstrated that summation could disappear when participants assumptions are affected by prompting rational rules concerning the independence of cues in a compound (Pérez *et al.*, 2018), in Experiment 1 we observed null summation in the absence of such manipulations. We replicated this absence of summation across three different experiments.

By simulation, we found that these data cannot be easily accommodated by current associative learning models. Indeed, for the type of complex summation design employed in our experiments, flexible models such as REM predict negative summation for a wide range of values in its parametric space. These results follow directly from the assumptions of these models. In the case of REM, when the novel cues ABX and ABY are presented during the test, REM assumes an additive contribution from all the representations involved, each of which is acquired independently. For example, when ABX is presented,

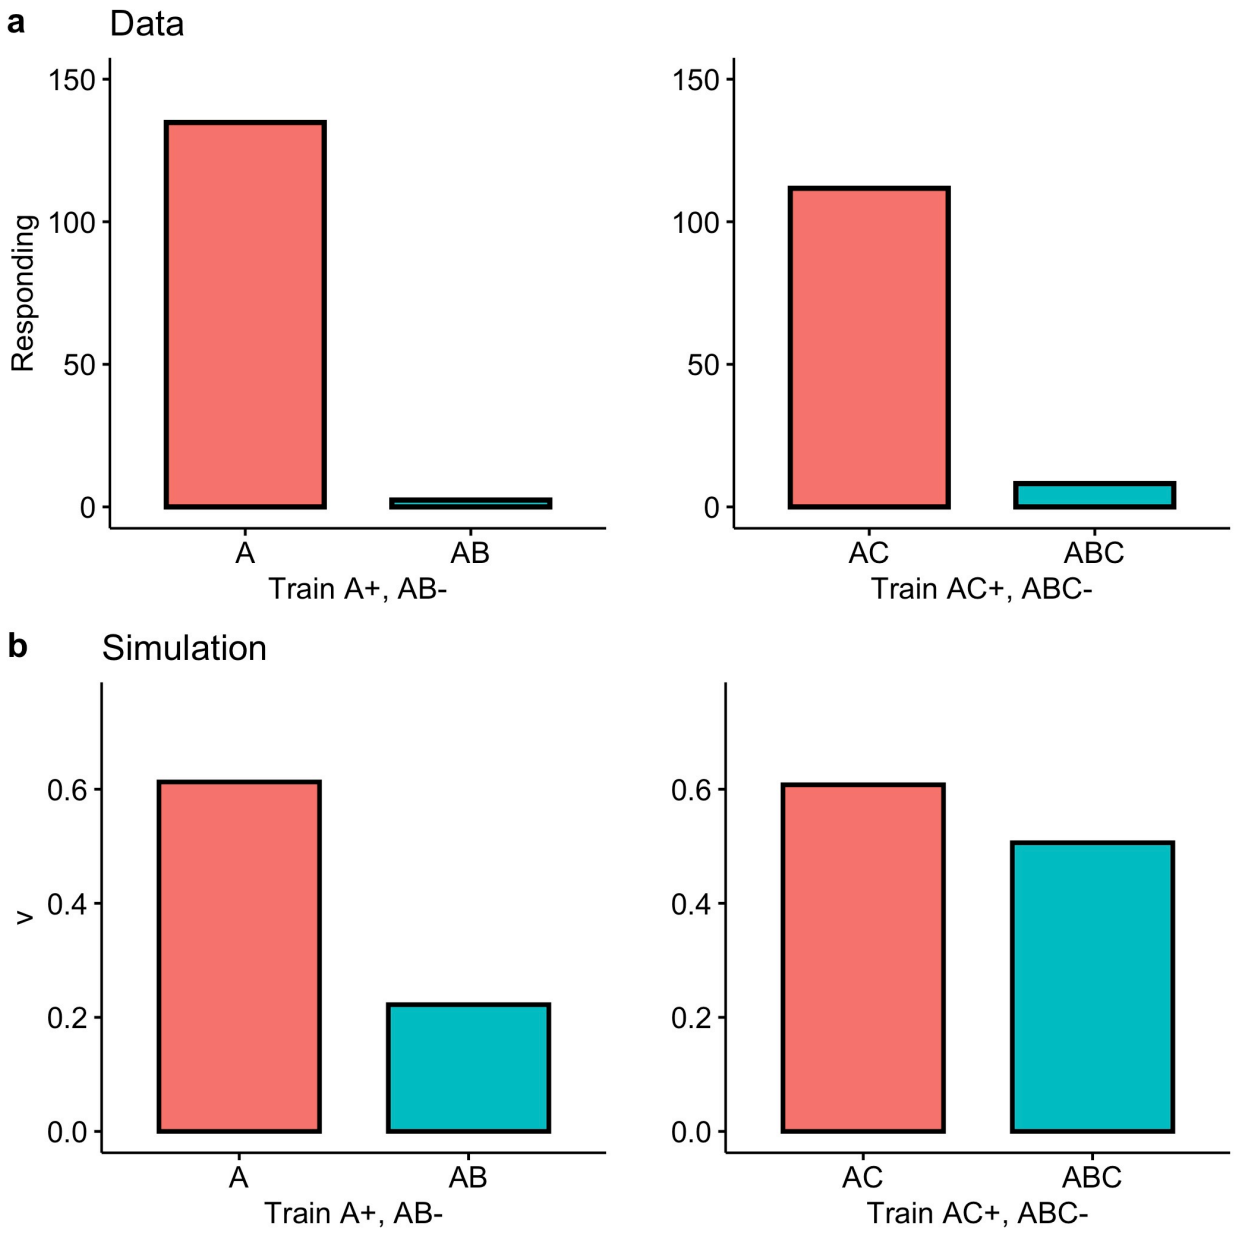

Figure 10. **Simulations of a sub-sampling model for a negative patterning design. (a)** Data obtained by Pearce and Redhead (1993) in pigeons. The left panel shows responding to the compound to A and AB at the end of training after a discrimination with reinforced trials of A and non-reinforced trials with AB (A+, AB-). The right hand panel shows responding to the compounds AC and ABC at the end of training after a discrimination with reinforced trials of AC and non-reinforced trials with ABC (AC+, ABC-). **(b)** Simulations of the sub-sampling model for a group trained on the A+, AB- discrimination (left panel) and for a group trained on the AC+, ABC- discrimination.

context independent elements a_i, b_i, and x_i and context dependent elements a_x and b_x contribute to the prediction. At the same time, however, some predictive value is also lost when ABX is presented due to the replacement of several excitatory elements (e.g., $y_i, x_y, a_y, b_y, a_{xy}$,and b_{xy}) by associatively neutral elements (e.g., a_b, b_a, a_{bx}, a_{by}). The overall result, of course, will depend on the assumptions made about the replacement between components, but the predominant result in REM (and Pearce's model) is negative summation (see Figure 4). These predictions are not consistent with the results obtained in Experiment 1.

We hypothesized an alternative view for these results. We thought that the data was more readily interpreted as coming from a sub-sampling process where participants sample one of the stimulus in a compound and learned only about that sampled stimulus in any given trial. This, we believed, was a consequence of a design with compounds that are not easily parsed into different components. In Experiments 2 and 3 we obtained evidence for this hypothesis by showing that the majority of participants rated a compound ABX/ABY equal to the value of the most salient cue in the compound. Again, associative models cannot explain these results as they predict lower responding to ABX/ABY in groups *intra2* than in group *intra*—in the former group the asympotic values for AXY and BXY are 8 and 10, respectively, whereas these values are both 10 for group *intra*. Moreover, when we swapped the outcome values predicted by AXY and BXY between experiments 2 and 3, participants followed the value of the most salient stimulus, BXY.

We presented a computational implementation of this hypothesis. By simulations we showed how a sub-sampling model can capture not only the patterns of results of summation designs presented here, but also a wider range of previous phenomena that are usually attributed to configural processing of stimuli in the animal literature. In contrast with these models, however, we obtained these results by taking a view that relies on objective properties of stimuli—some of which are absent when stimuli are presented in

isolation—rather than in internal representations, as it is usually assumed in associative learning models.

In a separate set of simulations (see Supplemental material online) we tested this sub-sampling approach against classic designs in the learning literature and observed that the model could predict several of them. One such phenomenon is external inhibition (Pavlov, 1927) in which less responding is obtained to a compound AB during testing after training with one of the elements than responding to the single cue A alone. Pearce's configural model can readily explain this result because the activation of the configural unit AB depends on the similarity between AB and A, which is only .5. In the sub-sampling model, by contrast, external inhibition is anticipated because after A has acquired significant predictive value during training, the presentation of AB during the test is equivalent to ABX, where X represents those properties of the stimuli that are added when it is presented in compound. Subjects then sample in this space and process B and X according to the softmax activation function. Since B and X have both a low predictive value, and are sometimes sampled, the average responding to AB will be lower than to A [3].

A number of additional phenomena, usually known as cue competition phenomena, are also anticipated by a sub-sampling approach. In these type of procedures it is usually found that the predictive value of a cue can impede the value otherwise normally accrued to other target cues. Blocking is perhaps the best-known example of cue competition (Kamin, 1969; Mackintosh and Turner, 1971). In a blocking procedure, after having trained A to predict an outcome, A is trained in compound with another cue B during a second stage. Responding to B is then tested against a group that did not undergo the first stage with A and it is usually found that the value of this latter cue B is higher than that of the

[3] A similar mechanism for external inhibition was suggested by Pavlov (1927), who reasoned that the presentation of the added cue B during the test stage would make the subject "orient" (i.e., attend) to it, so that the previous cue A would not be processed (i.e., sampled) as it would otherwise be, explaining why responding to the compound AB during the test is usually lower than responding to A alone.

cue B trained in compound. In terms of the present model, the second stage in which A and B are trained in compound makes subjects rarely sample the added cue B as a consequence of the higher predictive value of A at the start of the second stage. Compared to a group that has not undergone the first stage of training, responding to B will be lower, and blocking follows. In addition, when compounds are presented during training the sampling mechanism of the sub-sampling model makes it so that, for the same amount of training, each cue is sampled only half the time it would have been processed during training with either of the single cues alone; the sub-sampling process will then explain why training compounds brings about less responding to each cue than training the cues in isolation (overshadowing; Mackintosh, 1976).

It is also possible to entertain further interpretations regarding the mechanisms involved in a sub-sampling process. Indeed, the great majority of data consistent with configural processing (including the absence of a summation effect) come from auto-shaping experiments in pigeons using visual stimuli (Aydin and Pearce, 1995; Aydin and Pearce, 1997). In an auto-shaping procedure, hungry pigeons receive pairings of visual stimuli and food, which results in the animals approaching the stimuli and pecking at them. Two notable features of autoshaping experiments may explain why summation is not obtained. First, the proximity of the pigeon to the screen in which stimuli are displayed may limit its ability to sample the large and complex stimuli usually displayed in compound generalization experiments. In addition, just before the final ballistic movement that produces a peck, the peck's target is centered on the area dorsalis of the pigeon's retina, which is considered a "second fovea" due to its high density of cells. That is, pigeons specifically sample visual information from the pecked area, and they tend to peck on a circumscribed area of the screen displaying visual features predictive of reinforcement (Soto *et al.*, 2012).

What these two factors suggest is that pigeons may be using a similar process to the

one proposed in this paper. In fact, it is possible that pigeons sample and process only the visual features of the stimuli that they happened to be pecking during a particular trial, similar to what humans did in the present experiments. For example, in a simple summation design (A+, B+; test AB), the pigeon may be sampling only part of one stimulus and predict the outcome level associated with that specific stimulus part only.

This interpretation makes the auto-shaping procedure seem very similar to an overt visual search task (Wolfe, 2018) in which the pigeon samples information from the display by pecking and foveating a circumscribed area until it encounters an area previously associated with reinforcement, or *target*. At this point, the pigeon might continue to sample information from the target while relevant information from the rest of the display is never sampled. In the search literature this effect is known as *satisfaction of search*, and it leads to a sub-optimal sampling of useful information in the stimulus, missing any other targets that may be present. This type of search, known as *serial*, is considered to be inefficient, taking longer as the display size increases and being prone to satisfaction of search. By contrast, people in predictive learning experiments usually see the stimuli from a distance and can process all the cues at the same time, just as they did in our previous experiments on summation (Perez *et al.*, 2018). When there is only a few of such cues that are easily distinguishable, they can be detected in parallel , which is an efficient sampling strategy (Wolfe, 2018). In that case, all the relevant cues are sampled from the stimulus array. We believe this inefficient search is a plausible interpretation for these previous data and a process which might be underlying the results of the experiments reported here. The hypothesis is reinforced by the fact that similarity is thought to be one of the critical factors determining if search is parallel or serial (Duncan and Humphreys, 1989). In particular, increasing the similarity between the target and non-target cues makes search appear more serial, producing less summation. In contrast, more summation should obtain if similarity is decreased so that search appears more parallel. This is consistent with the effect of similarity we observed in Experiment 1.

Although largely speculative at this juncture, the merits of a visual search approach as the general mechanism underlying sub-sampling are not difficult to test in future studies. For example, a natural follow-up experiment should include a manipulation where participants are given limited time during the test with ABX/ABY in group *extra2* of Experiment 1. Under time pressure, participants' attentional capacities should be constrained and search stop as soon as they find one of the cues in the compound, providing evidence for a sequential search process under increasing difficulty. Likewise, parallel processing should be prompted with cues that are more easily parsed and with less time constrains, so that all cues should be processed under those conditions, and summation obtained.

Theoretically, a natural extension of the model would be to include an explicit "difficulty" rule that could allow the agent to process more than one cue at a given trial depending on how easy is to parse the cues and process them independently, including the case where all the components are processed and learned about. A first approximation could be to formulate the mechanism in a one-or-none fashion. Under this view, the agent should arbitrate between models following a rule such as

$$M = \begin{cases} m_{sub} & \text{if } d > D \\ m_{full} & \text{if } d \leq D. \end{cases} \tag{4}$$

where d is a parameter that determines how difficult is to parse components—which should be a direct function of the similarity between them—and D is a threshold that varies across individuals and which will determine whether sub-sampling or full sampling is deployed.

In sum, the empirical and computational evidence presented in this paper suggests that many results in the literature can be explained by an alternative approach where in

each trial subjects sub-sample cues within a compound and update their predictions for the sampled cue by the same principles of reward prediction error embodied in current models of associative and reinforcement learning (Kaelbling *et al.*, 1996; Vogel *et al.*, 2004). This approach can explain the empirical evidence in humans presented in this paper better than current models of learning, and a simple formalization of its principles captures a wide range of additional phenomena from the animal literature which are largely consistent with those predicted by configural models. We believe this interpretation is testable, important, and should be considered as a plausible alternative hypothesis in the literature on learning and generalization.

Materials and Methods

Participants were tested in desktop Windows (c) computers running Psychopy (version 1.75 for Experiment 1; version 1.82.4 for Experiments 2 and 3; Peirce, 2007). Responses were recorded from standard PC keyboards.

Statistical analysis

Statistical analyses were performed using RStudio and SPSS 27.0. For all the pre-planned comparisons we calculated a Welsh t-test and included Cohen's D, along with a 95% confidence interval on this estimate, as a measure of effect size. When reporting interactions between factors, we computed η_p^2. Following Steiger (2004), we report a 90% confidence interval on this estimate. The reliability of the results was contrasted against the usual criterion of $\alpha = .05$. When a null effect was expected, a Bayes Factor in favor of the null (BF_{01}) is reported.

Experiment 1

Participants. Eighty-six undergraduate students from Florida International University participated in Experiment 1. Participants did not have previous experience with the experimental procedure and were tested simultaneously and in the same room.

The Institutional Review Board of Florida International University approved all the studies in this paper (IRB-15-0460). Written informed consent was obtained from all participants. They were given course credit for their participation.

Participants were randomly assigned to one of three groups: *intra* ($n_{intra} = 27$), *extra1* ($n_{extra1} = 28$) and *extra2* ($n_{extra2} = 31$). In all of the studies in this paper, we determined our exclusion criteria before data collection. Following the same exclusion criteria of our previous study on summation (Pérez *et al.*, 2018), participants that did not give on average a rating between 7 and 13 points of allergy to cues that predicted allergy and between 0 and 3 points of allergy to cues that did not predict allergy, were left out from the analysis. The final number of participants in Experiment 1 was $n_{intra} = 18$, $n_{extra1} = 13$, $n_{extra2} = 27$.

Procedure. Groups differed in the similarity between cues in the display (see Figure 1). Each stimulus was created from three different cues that "branched out" from a central point. Among these branches, only one of them represented the target cue associated with either allergy or no allergy during training (A, B, C or D). The other two branches were non-target cues that could not predict the presence or absence of allergy (X and Y). During the test, the compound AB was comprised by two target branches together with an additional non-target cue (ABX or ABY). In group *intra*, all these "branches" were of the same color (black), but differed in shape. In group *extra1*, A and B differed in color (grey and black), but they shared color with the non-target cues (X and Y, one grey and one black). In group *extra2*, the target cues were the same as in group *extra1*, but now the non-target cues had a distinctive color as well. In all groups, A and B, which predicted allergy, shared color with cues C and D, which predicted no allergy. Thus, all participants, regardless of group, had to attend to shape; color was irrelevant to solve the discrimination.

Experiment 2

Participants

75 undergraduate students from Florida International University were randomly assigned to one of two groups ($n_{intra} = 40, n_{intra2} = 35$) and were compensated with course credit for their participation. The exclusion criteria were the same as those of Experiment 1, except that the admittance interval for the mean rating to A and B in group intra2 was set to [5-11]. The final number of participants per group was $n_{intra} = 39$, $n_{intra2} = 33$.

Procedure

The procedure was the same as described for group *intra* of Experiment 1, with only one exception: In group *intra2*, stimulus BXY was associated with 8 points of allergy during training (see Figure 5a).

Experiment 3

Participants

80 undergraduate students from Florida International University were randomly assigned to one of two groups ($n_{intra} = 42, n_{intra2} = 38$) and were tested under the same conditions of Experiment 2. The final number of participants per group was $n_{intra} = 24$, $n_{intra2} = 33$.

Procedure

The procedure was the same as in Experiment 2, except that the outcome values for A and B were interchanged in group *intra2*. In this experiment, the outcome assigned to BXY was 10 while a value of 8 was assigned to AXY.

Acknowledgements

We thank Jaron Colas and Tony Dickinson for their valuable comments on a previous version of this manuscript. S.D.G.

Open Practices Statement

The data and materials for all experiments are available at https://osf.io/xqnpk/

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

Aydin, A. and Pearce, J. M. (1995). Summation in autoshaping with short-and long-duration stimuli. *The Quarterly Journal of Experimental Psychology*, **48**: 215–234.

Aydin, A. and Pearce, J. M. (1997). Some determinants of response summation. *Animal Learning & Behavior*, **25**: 108–121.

Brandon, S. E., Vogel, E. H., and Wagner, A. (2000). A componential view of configural cues in generalization and discrimination in Pavlovian conditioning. *Behavioural brain research*, **110**: 67–72.

Bush, R. R. and Mosteller, F (1951). A model for stimulus generalization and discrimination. *Psychological review*, **58**: 413–423.

Duncan, J. and Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, **96**: 433–458.

Glautier, S (2002). Spatial separation of target and competitor cues enhances blocking of human causality judgements. *Q J Exp Psychol B*, **55**: 121–135.

Glautier, S, Redhead, E., Thorwart, A., and Lachnit, H. (2010). Reduced summation with common features in causal judgments. *Experimental psychology*, **57**: 252–259.

Harris, J. a. (2006). Elemental representations of stimuli in associative learning. *Psychological review*, **113**: 584–605.

Kaelbling, L. P., Littman, M. L., and Moore, A. W. (1996). Reinforcement Learning : A Survey. *Journal of Artificial Intelligence Research*, **4**: 237–285.

Kamin, L. (1969). Predictability, surprise, attention, and conditioning. *Punishment and aversive behavior*.

Kehoe, E. J., Horne, A. J., Horne, P. S., and Macrae, M. (1994). Summation and configuration between and within sensory modalities in classical conditioning of the rabbit. *Animal Learning {&} Behavior*, **22**: 19–26.

Mackintosh, N. J. (1976). Overshadowing and stimulus intensity. *Animal learning {&} behavior*, **4**: 186–192.

Mackintosh, N. J. and Turner, C. (1971). Blocking as a function of novelty of CS and predictability of UCS. *The Quarterly Journal of Experimental Psychology*, **23**: 359–366.

McLaren, I. P. L. and Mackintosh, N. J. (2002). Associative learning and elemental representation: II. Generalization and discrimination. *Animal learning & behavior*, **30**: 177–200.

Melchers, K. G., Shanks, D. R., and Lachnit, H. (2008). Stimulus coding in human associative learning: flexible representations of parts and wholes. *Behavioural processes*, **77**: 413.

Morey, R. D. *et al.* (2008). Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorial in Quantitative Methods for Psychology*, **4**: 61–64.

Parkhurst, D., Law, K., and Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision research*, **42**: 107–123.

Pavlov, I. P. (1927). No Title. *Conditioned Reflexes. An Investigation of the Physiological Activity of the Cerebral Cortex.* London.

Pearce, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological review*, **94**: 61–73.

Pearce, J. M. (1994). Similarity and discrimination: a selective review and a connectionist model. *Psychological review*, **101**: 587–607.

Pearce, J. M. (2002). Evaluation and development of a connectionist theory of configural learning. *Animal learning & behavior*, **30**: 73–95.

Pearce, J. M., Aydin, A., and Redhead, E. S. (1997). Configural analysis of summation in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, **23**: 84.

Peirce, J. W. (2007). PsychoPy—psychophysics software in Python. *Journal of neuroscience methods*, **162**: 8–13.

Perez, O. D., Milton, A. L., Aitken, M. R. F., and Dickinson, A (2018). A re-examination of responding on ratio and regulated-probability interval schedules. *bioRxiv*.

Pérez, O. D., San Martín, R., and Soto, F. A. (2018). Exploring the Effect of Stimulus Similarity on the Summation Effect in Causal Learning. *Experimental Psychology*, **65**: 183–200.

Redhead, E. (2007). Multimodal discrimination learning in humans: Evidence for configural theory. *Quarterly Journal of Experimental Psychology*, **60**: 1477–1495.

Rescorla, R. and Coldwell, S. E. (1995). Summation in autoshaping. *Animal Learning & Behavior*, **23**: 314–326.

Rescorla, R. and Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical conditioning II: Current research and theory*, **2**: 64–99.

Sanderson, D. J. and Bannerman, D. M. (2011). Competitive short-term and long-term memory processes in spatial habituation. *Journal of experimental psychology. Animal behavior processes*, **37**: 189–199.

Schultheis, H., Thorwart, A., and Lachnit, H. (2008). Rapid-REM: A MATLAB simulator of the replaced-elements model. *Behavior Research Methods*, **40**: 435–441.

Soto, F. A., Gershman, S., and Niv, Y (2014a). Toward a General Theory of Generalization: Rational Principles of Dimensional Generalization Can Explain Compound Generalization Phenomena. *Psychological review*: 1–79.

Soto, F. A., Vogel, E. H., Castillo, R. D., and Wagner, A. (2009). Generality of the summation effect in human causal learning. *Quarterly journal of experimental psychology (2006)*, **62**: 877–889.

Soto, F. A., Gershman, S. J., and Niv, Y. (2014b). Explaining compound generalization in associative and causal learning through rational principles of dimensional generalization. *Psychological review*, **121**: 526.

Soto, F. A., Siow, J. Y., and Wasserman, E. A. (2012). View-invariance learning in object recognition by pigeons depends on error-driven associative learning processes. *Vision Research*, **62**: 148–161.

Steiger, J. H. (2004). Beyond the F test: Effect size confidence intervals and tests of close fit in the analysis of variance and contrast analysis. *Psychological methods*, **9**: 164–182.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement learning: An introduction*. Vol. 1. 1. MIT press Cambridge.

Thein, T., Westbrook, R. F., and Harris, J. A. (2008). How the associative strengths of stimuli combine in compound: Summation and overshadowing. *Journal of Experimental Psychology: Animal Behavior Processes*, **34**: 155–166.

Thorwart, A., Livesey, E. J., and Harris, J. A. (2012). Normalization between stimulus elements in a model of Pavlovian conditioning: Showjumping on an elemental horse. *Learning & behavior*, **40**: 334–346.

Thorwart, A., Schultheis, H., König, S., and Lachnit, H. (2009). ALTSim: A MATLAB simulator for current associative learning theories. *Behavior Research Methods*, **41**: 29–34.

Vogel, E. H., Castro, M. E., and Saavedra, M. A. (2004). Quantitative models of Pavlovian conditioning. *Brain research bulletin*, **63**: 173–202.

Vogel, E. H., Ponce, F. P., and Wagner, A. R. (2017). A theoretical analysis of transfer of occasion setting: SOP with replaced elements. *Behavioural processes*, **137**: 19–32.

Wagner, A. (2008). Evolution of an elemental theory of Pavlovian conditioning. *Learning & behavior*, **36**: 253–65.

Wolfe, J. M. (2018). Visual Search. en. *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience*. Ed. by J. T. Serences. Vol. 2. Hoboken, NJ: John Wiley & Sons: 1–55.