Peak Shift and Rules in Human Generalization

Article *in* Journal of Experimental Psychology Learning Memory and Cognition · January 2018

DOI: 10.1037/xlm0000558

Jessica C Lee
UNSW Sydney

Journal of Experimental Psychology:
Learning, Memory, and Cognition
2018, Vol. 44, No. 12, 1955–1970

0278-7393/18/$12.00 http://dx.doi.org/10.1037/xlm0000558

Peak Shift and Rules in Human Generalization

Jessica C. Lee, Brett K. Hayes, and Peter F. Lovibond
University of New South Wales

Two experiments tested whether a peak-shifted generalization gradient could be explained by the averaging of distinct gradients displayed in subgroups reporting different generalization rules. Across experiments using a causal judgment task (Experiment 1) and a fear conditioning paradigm (Experiment 2), we found a close concordance between self-reported rules and generalization gradients using a continuous stimulus dimension (hue). Both experiments also showed an overall peak-shifted gradient after differential conditioning, but not after single cue conditioning. Importantly, the peak shift could be decomposed into linear and peaked gradients when participants were divided into rule subgroups. Our results highlight the need to consider individual differences in the rules that participants derive in human generalization studies and suggest that in some situations, peak shift may be a consequence of averaging across diverse rule subgroups.

Keywords: generalization, peak shift, rule, associative learning, discrimination

Supplemental materials: http://dx.doi.org/10.1037/xlm0000558.supp

Generalization concerns the transfer of learning from past instances to novel instances and is an adaptive ability that allows us to anticipate events and engage appropriately with the world. Investigating the theoretical processes underlying generalization is paramount to understanding behavior, since new stimuli that we encounter in the world are never exactly the same as previously experienced instances (Shepard, 1987). There are also important clinical implications in studying generalization, as maladaptive generalization has been implicated in clinical disorders, such as panic disorder (Lissek et al., 2010), posttraumatic stress disorder (Grillon & Morgan, 1999; Morey et al., 2015), and generalized anxiety disorder (Lissek et al., 2014).

Empirical investigation of generalization originated in the animal conditioning literature, where generalization along a stimulus dimension was assessed following different types of conditioning procedures. In single cue conditioning, an animal is trained to respond in the presence of a single stimulus (S+) by rewarding responses in the presence of that stimulus in an operant conditioning procedure, or a single stimulus (conditionaled stimulus [CS]+) is followed by an unconditioned stimulus (US) such as food or shock, in a Pavlovian procedure. In differential conditioning, another stimulus (S−) is presented during which responses are never rewarded, or in the Pavlovian case, another stimulus (CS−) is presented that is not followed by the US. Learning is evident when the animal responds (either an instrumental response or a conditioned response) in the presence of the S+ or CS+, and in the case of differential conditioning, suppresses responding to the S− or CS−. After training, generalization testing is typically assessed by varying the dimension of interest and measuring the degree to which the animal responds, producing a generalization gradient.

In the first study to examine generalization gradients empirically, Guttman and Kalish (1956) tested pigeons using stimuli varying along a visual wavelength (hue) dimension. Following single cue training with keylights of different wavelength, they found peaked, symmetrical gradients with the highest rates of responding at the value of the S+ (see Figure 1 for an example). Such peaked gradients are usually taken to indicate generalization on the basis of similarity to the physical features of the S+. If, however, animals are given differential training whereby they learn to discriminate between a S+ and a S− lying on the same dimension (i.e., "intradimensional" or within-dimension discrimination), the peak of the gradient can shift beyond the S+ in the direction away from the S− (Figure 1). This "peak shift" effect was first reported by Hanson (1957; see Purtle, 1973, for a review) and is reliably found in a variety of animal species and stimulus dimensions (given appropriate parameters on the training and test procedures; see Purtle, 1973). A related and more commonly found feature of the postdiscrimination generalization gradient is that it is asymmetrical, with more responding to stimuli on the opposite side of the S+ to the S− (an "area shift"; see Honig & Urcuioli, 1981; Figure 1).

Peak shift has received a large amount of interest in the animal conditioning and ethology literature due to its demonstration of a seemingly adaptive behavior involving greater responding for a novel stimulus compared with a trained stimulus with which the animal has had prior experience. Peak shift has been proposed as an explanation for a wide range of phenomena, including why caricatures of faces are better recognized than actual faces (Lewis & Johnston, 1999), and why animals come to prefer slightly exaggerated features in mating partners (ten Cate, Verzijden, &

This article was published Online First March 19, 2018.

Jessica C. Lee, Brett K. Hayes, Peter F. Lovibond, School of Psychology, University of New South Wales.

This study was funded by an Australian Research Council Discovery Grant (DP160101907) awarded to Peter Lovibond and Brett Hayes. We thank Nik Raic for assistance with data collection.

Correspondence concerning this article should be addressed to Jessica Lee, School of Psychology, University of New South Wales, UNSW Sydney, NSW 2052, Australia. E-mail: jessica.lee@unsw.edu.au

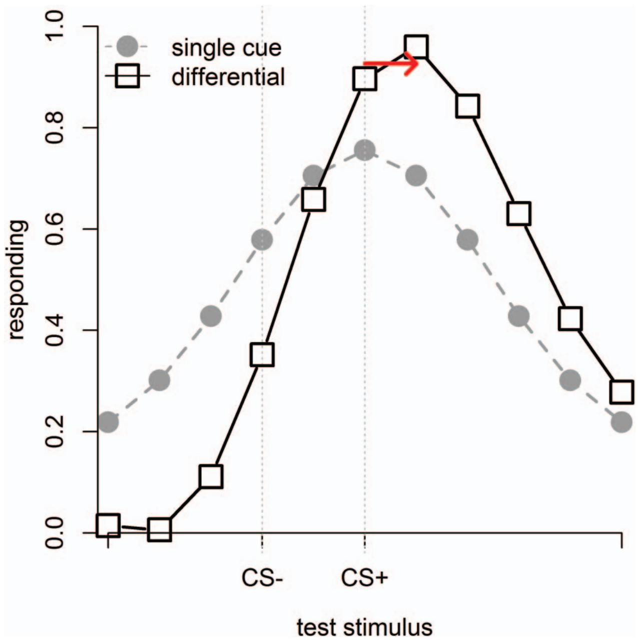

Figure 1. Example of typical generalization gradients following single cue and differential Pavlovian conditioning. CS+ (S+ in instrumental conditioning) is the conditioned stimulus paired with an outcome, and CS− (S− in instrumental conditioning) is the conditioned stimulus not paired with an outcome. See the online article for the color version of this figure.

Etman, 2006). As detailed in the following sections, peak shift has traditionally been explained in associative terms—as a consequence of interactions between conditioned excitation and inhibition. In this article, however, we argue that understanding peak shift benefits from taking a broader theoretical perspective, which conceives of human learning and generalization as a process of rule generation and testing, where the participant is actively seeking to determine the class or category of stimuli that reliably predict an outcome (e.g., an aversive event like electric shock). We see our approach as broadly consistent with other work that has argued that more complex conceptual processes play a key role in human conditioning (e.g., Dunsmoor & Murphy, 2015).

Associative Accounts of Peak Shift in Infrahuman Animals and Humans

An influential account of peak shift was provided by Spence (1937). Spence's account postulated that during differential conditioning, an excitatory gradient is established around the S+ and an inhibitory gradient around the S−. On test, these gradients sum algebraically to determine the net associative strength of the generalization stimuli, and accordingly, responses to these stimuli. Assuming that the excitatory gradient was narrower and taller than the inverse of the inhibitory gradient (later confirmed in studies by Honig, Boneau, Burstein, & Pennypacker, 1963; see Honig & Urcuioli, 1981, for a review), subtracting the inhibitory gradient from the excitatory gradient produces a peak-shifted gradient. In single cue conditioning, the absence of an inhibitory gradient means that the excitatory gradient alone determines generalization, producing a peaked gradient with the highest responding at the S+ (Figure 1).

In a similar fashion, more recent associative theories (Blough, 1975; Ghirlanda & Enquist, 1998; McLaren & Mackintosh, 2002) explain peak shift by postulating elemental representation of stimulus dimensions and graded (i.e., Gaussian) overlapping activation across these elements by the S+ and S−. Similar to Spence's account, their central tenet is that activation of the elements on the dimension by the S+ and S− interact such that elements that are activated maximally by the S+ do not accrue the most associative strength (unlike in single cue training) because of their concurrent inhibitory activation by the S−. Rather, the stimulus that produces the greatest *discrepancy* in activation between the S+ and S− will produce the highest level of conditioned responding during generalization. Following differential training with a perceptually similar S−, this will typically be a stimulus that is slightly displaced from the S+ in the direction away from the S− (Figure 1). Associative accounts therefore provide an adequate and precise explanation of the shape of generalization gradients following different training procedures in nonhumans (see Mackintosh, 1974, for a review).

In humans, however, peak shift is relatively elusive in comparable single-outcome conditioning designs (but see Dunsmoor & LaBar, 2013; Struyf, Iberico, & Vervliet, 2014) and appears in two-choice discrimination tasks only under certain conditions (see Livesey & McLaren, 2009; Wills & Mackintosh, 1998). In contrast to the animal literature, a variety of generalization gradients have been found following both single cue and differential conditioning, suggesting that there are other mechanisms that influence generalization. In particular, in addition to peaked gradients, participants can also exhibit a monotonically increasing (e.g., linear) gradient with the highest level of responding at the extreme end of the dimension (Dunsmoor, Mitroff, & LaBar, 2009; Laberge, 1961; see also Livesey & McLaren, 2009, and Wills & Mackintosh, 1998, for examples in two-choice discrimination learning). Such linear gradients are suggestive of the use of relational rules (e.g., "the greener the stimulus, the more likely the outcome") derived from participants noticing the relation between the S+ and S−.

Peak Shift and the Role of Rules

There is a growing interest in investigating the role of cognitive processes such as rule formation, categorization, and inductive reasoning in the learning and generalization of conditioned responses, and recognition that these processes exert powerful influences on generalization over and above perceptual similarity (Dunsmoor & Murphy, 2015; Dymond, Dunsmoor, Vervliet, Roche, & Hermans, 2015). For example, Dunsmoor and Murphy (2014) have shown that participants are more willing to generalize from typical to atypical exemplars than from atypical to typical exemplars in fear conditioning. This result is significant because the physical similarity between the CS+ and generalization stimuli was identical across conditions. The result suggests that generalization was influenced by participants' knowledge of category structure, in line with results found in studies of inductive reasoning with verbal materials (e.g., see Hayes & Heit, in press, for a review). Other studies have shown that instructional manipulations influence the rules that participants form, and hence their subsequent generalization (Ahmed & Lovibond, 2015; Boddez, Bennett, van Esch, & Beckers, 2017; Vervliet, Kindt, Vansteenwegen, & Hermans, 2010), and Dunsmoor, Martin, and LaBar (2012) have demonstrated that generalization of conditioned fear occurs across conceptually related stimuli. Such results suggest that perceptual similarity is just one of many potential stimulus relations that can influence generalization.

Recognition of individual variability in learning strategies is also evident in the category learning literature. For example, some participants categorize on the basis of exemplar similarity, while others use dimensional or conjunction rules (Little & McDaniel, 2015; McDaniel, Cahill, Robbins, & Wiener, 2014; Nosofsky, Palmeri, & McKinley, 1994). The tendency to search for rules has been suggested to be a stable psychological trait that influences learning strategy across a wide range of cognitive tasks (Don, Goldwater, Otto, & Livesey, 2016; McDaniel et al., 2014) and may be related to working memory capacity or fluid intelligence (McDaniel et al., 2014, but see Little & McDaniel, 2015). Alternatively, the likelihood of deriving a relational rule may also depend on mastery of the training material (DeLosh, Busemeyer, & McDaniel, 1997; McDaniel & Busemeyer, 2005; Shanks & Darby, 1998; but see McDaniel et al., 2014), or if the stimulus or task is complex, whether participants happen to attend to a particular stimulus feature. Acknowledging the existence of individual differences in generalization leads to the conclusion that analyzing group-level data alone can be misleading (see Estes, 1956, for a similar view and Maddox, 1999, for a simulation in category learning). Furthermore, analyzing aggregate data can obscure interesting results when subgroups of participants are generalizing in different ways (see Livesey & McLaren, 2009, Experiment 2). If we accept that group-level data can be decomposed into distinct gradients exhibited by subgroups of participants using different strategies, an alternative explanation for peak shift may be derived.

A peak-shifted gradient obtained at the group-level is explainable if a certain subgroup of participants generalize according to a linear relational rule (e.g., "the bluer the stimulus, the more likely the outcome is"; Figure 2) and another subgroup generalize according to symmetrical stimulus similarity (e.g., "the closer the stimulus is to the bluey-green stimulus that led to shock, the more likely the outcome is"). The former rule will produce a linear gradient, while the latter rule will produce a peaked gradient around the CS+, with decreasing responding as similarity to the CS+ decreases; Figure 2).[1] Crucially, if the "similarity" subgroup produces the highest level of responding at the CS+, and the gradient exhibited by the "linear" subgroup is sufficiently steep, the peak of responding of the averaged gradient should be not at the CS+, but at a stimulus slightly removed from the CS+ in the direction away from the CS− (in Figure 2, the peak is shifted to the right of the CS+).

Thus, the *averaging* of generalization gradients from these subgroups would produce a peak-shifted gradient, despite no *individual* displaying a peak shift. This explanation is equally applicable to the phenomenon of area shift since a linear gradient obviously contains an area shift (Figure 2). By contrast, the same peak-shifted gradient would not be predicted to occur following single-cue training because linear rules could be derived in either direction along the dimension (e.g., if the dimension ranges from green to blue, "greener/bluer stimuli lead to the outcome"), making the resultant gradient symmetrical and flat. Such a result would suggest that peak shift can arise from averaging data over groups, and therefore would not uniquely support an associative account.

The Current Studies

The aim of the current studies was to test whether an overall peak shift could be obtained from the combination of "linear" and "similarity" rule subgroups. Hue (restricted to blue-green) was chosen as the generalization dimension to emulate the early animal generalization literature, which primarily used key lights of different hues (e.g., Hanson, 1957). We used a causal judgment (Experiment 1), as well as a fear conditioning paradigm (Experiment 2). This was important to test whether results obtained in causal judgment replicated when using a biologically significant outcome (shock). The general procedure was similar to that of Ahmed and Lovibond (2016, 2017) and Wong and Lovibond (2017) in that participants were divided into subgroups based on their reported rules in a questionnaire. Verbal report has been used successfully in a variety of tasks to distinguish between different strategies and rules (e.g., Gluck, Shohamy, & Myers, 2002; Little & McDaniel, 2015; Regehr & Brooks, 1995; Smith & Sloman, 1994) and has produced distinct gradients between rule subgroups in previous studies (Ahmed and Lovibond, 2016, 2017; Wong & Lovibond, 2017).

Our study advances beyond these previous studies in a number of ways. First and foremost, none of these earlier studies specifically aimed to examine peak shift and none reported a significant peak shift in the generalization test. Second, the stimuli used in the previous experiments were not optimal for identifying linear and similarity-based rules used by different participations. Ahmed and Lovibond, 2016 used circle size, which is an "intensity" dimension. Intensity dimensions have been shown to produce linear generalization gradients even in animals (e.g., Grice & Saltz, 1950; Hull, 1949; Razran, 1949). This complicates the interpretation of any linear gradients obtained. Ahmed and Lovibond, 2017 and Wong and Lovibond (2017) used a dimension with a clearly defined midpoint (a dot in a square box that varied its horizontal location within the box) that was therefore noncontinuous. The middle location (the CS+) was highly salient and may have facilitated grouping of the dimension into three categories: left of center, center, and right of center. Thus, the stimuli employed in these previous studies may have promoted the formation of explicit generalization rules and masked evidence of associative learning processes. In other words, the high degree of concordance between generalization gradients and verbalizable rules may have been a product of the stimuli used in these studies.

In addition to a group that underwent differential conditioning (Differential group), we also included a control group that underwent single cue conditioning (Single Cue group) with a single CS+. Comparing the differences in gradients between groups allowed us to test how adding a CS− affects generalization through changing the distribution of subgroups using different rules, as well as the direction of the linear rules. Specifically, linear rules might be produced in either direction in single cue conditioning, whereas differential conditioning should bias linear rules in the direction consistent with the training contingencies. The novel prediction was that the bias in the linear rules would produce

[1] In these experiments, it is typical to use a CS+ that lies somewhere in the middle of the dimension.

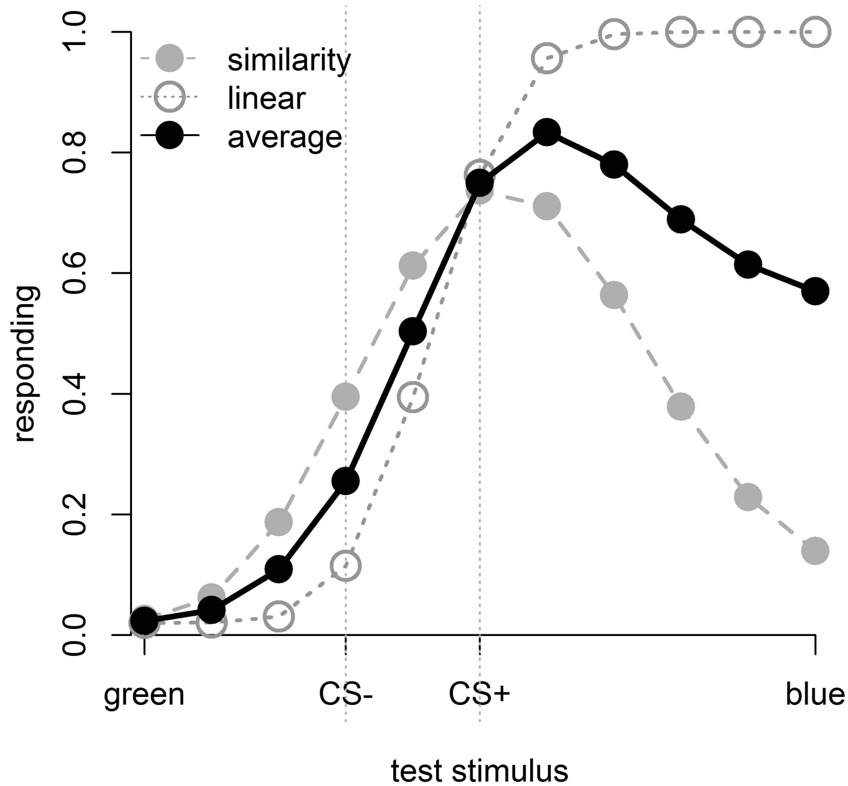

Figure 2. Example generalization gradients from Linear and Similarity subgroups following differential training and resulting average group-level gradient showing a peak shift.

peak shift when combined with the peaked gradients displayed in the Similarity subgroup.

Experiment 1

Experiment 1 utilized a causal judgment paradigm. Causal judgment paradigms have traditionally been interpreted as measuring associative learning despite the explicit nature of the learning conditions and verbalizable nature of outcome measures (see Shanks, 2007). Thus, despite being quite different from the conditioning paradigms used in animal studies, causal ratings in this task can be seen as an index of associative strength between stimuli (CSs) and outcomes (USs). Participants underwent conditioning with either a single CS+ (Single Cue group) or a CS+ and CS− (Differential group), with the CS+ followed by the outcome 75% of the time, and the CS −never followed by the outcome. Partial reinforcement was used to avoid participants reaching ceiling levels of expectancy at the end of training which would restrict the opportunity to observe responding beyond the CS+ that is needed to observe both linear gradients and peak shift. For the Differential group, the CS+ and CS− lay on the same dimension (hue) and were perceptually similar, approximating the procedures typically used in the peak shift literature.

Hue was chosen as the continuous stimulus dimension to measure generalization to provide consistency with the animal conditioning literature. It was hypothesized that for both Single Cue and Differential groups, a linear generalization gradient would be exhibited in the subgroups reporting a linear rule, a peaked gradient with the highest ratings at the CS+ would be evident in the subgroups reporting a similarity rule, and a flat generalization gradient would be shown in the subgroups who either reported miscellaneous rules ignoring hue or did not report any rule. In the Differential group, the presence of the CS− should promote formation of a linear rule in the direction consistent with the relational difference between the CS+ and CS− (e.g., if the CS+ was greener than the CS−, then participants should show the highest expectancy of the outcome at the greenest stimulus). By contrast, linear rules reported by Single Cue group participants were not expected to be in a consistent direction since participants would not have experience with a CS− to promote formation of a linear rule in a particular direction. When combined with the peaked gradients in the Similarity subgroup, the Differential group may show an overall peak-shifted gradient that is composed of two distinct gradients, while the Single Cue group should show a peaked gradient centered at the CS+.

Method

Participants. A total of 181 first-year psychology students (111 women, $M_{\text{age}} = 19.09$, $SD = 2.11$) at the University of New South Wales participated in exchange for partial course credit. All participants were recruited via an internal website. Participants were randomly allocated to the Single Cue ($n = 91$) or Differential group ($n = 90$). Participants were excluded if they indicated that they were colorblind (1 participant), or if they failed the training criterion (a further 61 participants). The criterion was an average causal rating >50 for the CS+ in the final block (last 4 trials) of training for both groups, and for the Differential group, an additional requirement that ratings for the CS− were <20. After exclusions, a total of 119 participants remained (46 in the Differential group and 73 in the Single Cue group). The high rate of exclusions based on training performance was most likely because of the perceptual difficulty of the discrimination (Figure 3), the small number of training trials, as well as partial reinforcement creating uncertainty in participants' prediction of the outcome.

Apparatus. The experiment was programmed using Psychtoolbox (Brainard, 1997; Pelli, 1997) and run using Matlab on standard PC computers connected to a 23-inch monitor. Participants made responses using a standard PC keyboard and mouse in individual cubicles.

Stimuli. There were 11 circle stimuli (200 pixels in diameter) presented in the experiment that varied in hue along the green-blue dimension (Figure 3). The stimuli were created on the HSB (hue, saturation, brightness) scale by varying hue (H), keeping saturation (S) and brightness (B) levels constant at 100% and 75% respectively. The minimum and maximum hue values were .403 and .555, with equal spacing between stimuli along the dimension (see online supplemental materials for exact hue values).

Procedure. Experiment 1 was approved by the University of New South Wales Human Research Ethics Advisory Panel. The experiment was composed of a training phase, test phase, and post-test questionnaire. All instructions and stimuli were presented on a light gray background that had RGB (red, green, blue) values of (200, 200, 200).

Training. Participants in the Single Cue group received training with a single aqua (blueish-green) stimulus at the midpoint on the tested dimension (Stimulus 6, S6), while participants in the Differential group received discrimination training with a CS+ (S6) and a CS− (S4 on the dimension; Figure 3). Because the direction of the dimension was essentially arbitrary for the Single Cue group (i.e., S1 and S11 were interchangeable), a direction had to be specified so that the generalization gradients could be compared meaningfully between groups. Therefore, for both Single Cue and Differential groups, the direction of the hue dimension (going from S1 to S11) was counterbalanced (green to blue or blue to green) such that for the Differential group, the CS− was either greener or bluer than the CS+ (which was always the same stimulus, S6). This meant that for every participant in the Differential group who received training with CS− (S4), another participant in the Single Cue group was put in the same counterbalancing condition such that the physical stimulus representing S4 was identical across pairs of participants.

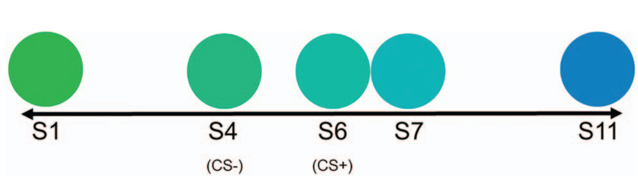

Figure 3. Stimuli comprising the hue dimension in Experiment 1. The direction of the dimension (going from S1 to S11) was counterbalanced between participants (either green-blue, as above, or blue-green). S6 was always the CS+ and S4 was always the CS− in the Differential group. S7 was the stimulus predicted to show peak shift. See the online article for the color version of this figure.

Training consisted of 12 presentations of the CS+ for the Single Cue Group, and 12 presentations of both the CS+ and CS− for the Differential group, with the CS+ reinforced at a rate of 75%. Trials were randomized in blocks of four trials for the Single Cue group and eight trials for the Differential group, with the constraint that the first CS+ trial in each block had to be reinforced, and no more than two presentations of the same stimulus could occur in a row in the Differential group.

Participants were asked to make predictions about whether a vending machine would dispense a snack (the outcome) based on the symbol (CS) displayed on the vending machine. On each trial, participants were presented with the symbol displayed on the vending machine, and made a prediction about the likelihood of snack delivery. Each symbol was a colored circle (200 pixels in diameter) presented in a black 300 × 300 pixel square outline. The symbol was presented first and after 1-s delay, a visual analogue rating scale appeared underneath along with the question "How likely is it that your snack will be delivered?". The scale ranged from "Certain NO snack" to "Certain snack," and the midpoint and endpoints were marked with a tick. There were no numerical anchors. Participants responded by using the mouse to click any point on the scale and pressed the spacebar once they were finished with their rating. The trials were not timed, and participants could change their rating as many times as they wanted. The next trial began after a blank screen intertrial interval (ITI) of 2 s. The instructions given to participants prior to training did not explicitly refer to the color of the circle and did not encourage participants to form rules. Rather, they stressed that the aim of the task was to make predictions about whether a snack would be delivered based on the symbol presented.

Generalization test. Prior to the test phase, both groups were instructed that they would no longer receive feedback about whether a snack would be delivered, but to continue making ratings about the likelihood of snack delivery regardless. The test phase consisted of two presentations of each of the 11 test stimuli, with the order of presentation randomized in each block of 11 trials. All other task features were the same as the training phase.

Questionnaire. After the causal judgment task was finished, participants completed a written questionnaire, administered on two separate pages. Participants first answered a two-alternative forced-choice (2AFC) question asking whether they thought there was a relationship between the symbols and snack delivery (yes or no). If participants answered "yes" they were asked to describe the relationship, giving as much detail as possible, and also indicate whether they figured out this relationship during the first phase of the experiment (when there was feedback), or during the second phase of the experiment (when there was no feedback). If participants indicated that they did not think there was a relationship between the symbols and snack delivery, they were not required to fill out the remaining questions on that page, and progressed to the second page of the questionnaire.

The first 2AFC question on the second page asked participants how many different colored circles they saw during the training phase (1 or 2), and the second 5AFC question asked participants to select the option (i.e., rule) they thought best described the relationship between the symbols and snack delivery (see online supplemental materials). Participants were asked to choose the best option from two linear functions (with separate options for each direction along the hue continuum, i.e., greener or bluer stimuli were more likely to lead to the outcome), a similarity option stating that only a single stimulus led to shock and that all others did not, an option stating that there was no relationship between the stimuli, or another rule which they were asked to describe.[2] Note that this was the second opportunity for participants to either claim that there was no relationship between the symbol and the outcome or describe an alternative rule. The final question asked participants to indicate whether they were colorblind.

Data analysis. To test for the presence of a linear gradient, a linear trend analysis was conducted across the 11 test stimuli. To test for peaked gradients, a quadratic trend analysis was conducted. However, a quadratic trend is difficult to interpret in isolation because the form of the gradient can vary considerably and still show a significant quadratic trend (e.g., a gradient that increases and then flattens). Thus, planned paired *t* tests between the CS+ and each endpoint (i.e., S1 vs. CS+, CS+ vs. S11) were also included as a more stringent test of whether a gradient was peaked.

Associative theories and empirical demonstrations of peak shift predict the peak of responding to be only slightly removed from the S+ when the S+ and S− are highly similar and the breadth of generalization is narrow. Thus, to test for the presence of peak shift, paired *t* tests were conducted using the stimulus immediately next to the CS+ in the direction away from the CS− (S7). A significant peak shift would be demonstrated if there were a significant rise (from CS+ vs. S7) and fall (from S7 to S11) in expectancy at this peak. Because multiple analyses were done within each subgroup and overall, Holm-Bonferroni-correction to the critical alpha value was used to control the family wise Type I error rate. Additional Bayes Factor analyses were carried out with the "BayesFactor" package (Rouder, Speckman, Sun, Morey, & Iverson, 2009) in R (R Core Team, 2015).

Questionnaire coding. Reported rules in the open-ended question were classified into four categories: linear, similarity, no relationship, and other.[2] The reported rules were then classified by a second rater, blind to the coding of the first rater. There was substantial agreement (Landis & Koch, 1977) between the two raters using Cohen's kappa, $k = .80$, $p < .001$. Disagreements between raters were resolved via discussion. A two-step procedure was then employed to determine the final rule subgroups. Participants who reported an unambiguous linear or similarity generalization rule in the open-ended question were simply put into those subgroups ($n = 69$). If participants indicated that they did not think there was a relationship between the symbols and the outcome, or reported other rules (e.g., vague rules, rules referring to sequences of trials, rules describing the training contingencies but not referring to the dimension), they were assigned to whatever rule they had endorsed in the 5AFC question. The final four rule subgroups were: linear, similarity, no relationship, and other.

Results and Discussion

Training. The training data were analyzed in an analysis of variance (ANOVA) with presentation order as a within-subjects factor for both groups, and trial type (CS+ vs. CS−) as an additional within-subjects factor for the Differential group. In the Single Cue group, there was a marginal linear trend in causal judgments to the CS+ over the training trials, $F(1, 72) = 4.12$, $p = .046$, $\eta_p^2 = .054$ (Figure 4). In the Differential group, there was a significant overall difference between causal judgments for the

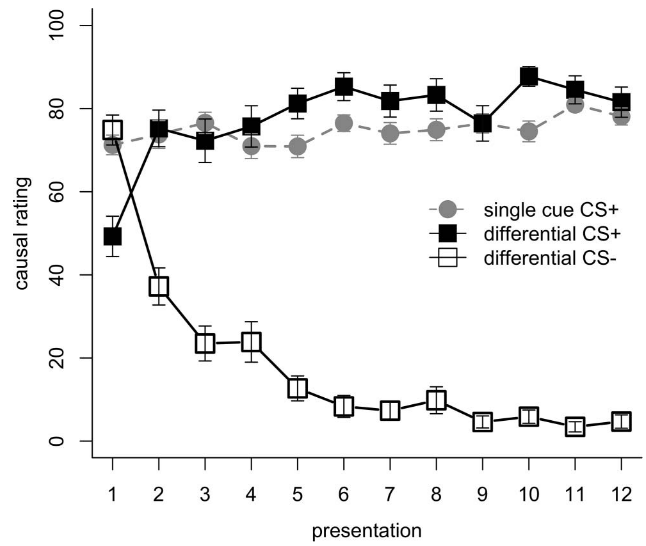

Figure 4. Causal ratings over stimulus presentations in training for Experiment 1. The CS+ was reinforced at a rate of 75%, and the CS− was never reinforced. Within-subject error bars calculated according to the method proposed by Cousineau (2005) with the correction by Morey (2008).

CS+ and CS−, $F(1, 45) = 533.6$, $p < .001$, $\eta_p^2 = .922$, and a significant interaction between the linear trend for the CS+ and CS−, $F(1, 45) = 113.8$, $p < .001$, $\eta_p^2 = .717$, indicating that causal judgments diverged over the training trials (Figure 4). Thus, there was evidence of learning in both groups.

Generalization test. The overall generalization gradients were first analyzed to look for peak shift in the Differential group, and then planned contrast analyses were conducted for each rule subgroup separately. The overall generalization gradients are shown in Figure 5. Figure 5 shows that adding the CS− in the Differential Group changed the shape of the generalization gradients. There was a significant overall linear, $F(1, 117) = 91.3$, $p < .001$, $\eta_p^2 = .438$, and quadratic trend, $F(1, 117) = 70.6$, $p < .001$, $\eta_p^2 = .376$. Both linear, $F(1, 117) = 61.5$, $p < .001$, $\eta_p^2 = .345$, and quadratic, $F(1, 117) = 19.0$, $p < .001$, $\eta_p^2 = .140$, trends interacted with group. The gradient for the Single Cue group appears to be flat, while the gradient for the Differential group is more linear. This was confirmed statistically in that there was no significant linear trend in the Single Cue group, $F(1, 72) = 1.61$, $p = .208$,

[2] Note that in the original classification there was an additional "step-function" group designed to capture a rule whereby participants divided the dimension into two categories, and grouped stimuli on one side of the boundary (e.g., "bluer than the CS+) together in causing shock. The treatment of these "step" rules in the questionnaire was exactly the same as the linear rules, meaning that participants who reported an unambiguous step rule in the free-report question were classified immediately into the Step subgroup. There were also two additional options in the forced-choice question, phrased in terms of either "green or greener stimuli" or "blue or bluer" stimuli leading to the outcome (see online supplemental materials). The final rule subgroups were Linear, Step, Similarity, No Relationship, and Other. Because of a high degree of similarity between the gradients in this Step subgroup with the Linear subgroup (see online supplementary materials), we collapsed the Linear and Step subgroups into a final Linear subgroup for the analyses.

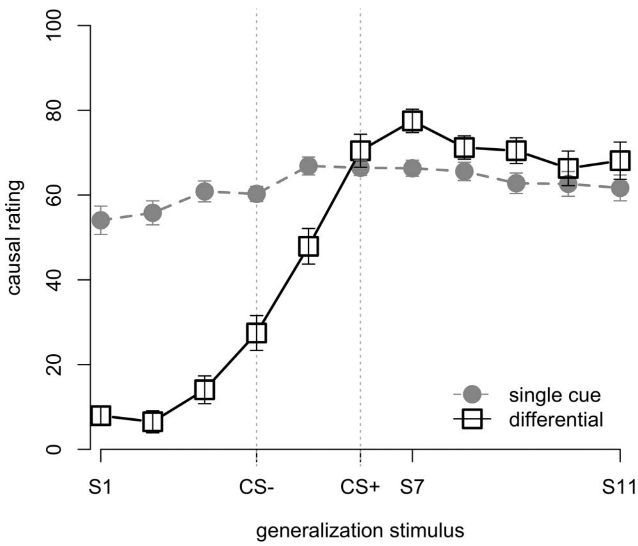

Figure 5. Overall generalization gradients for both groups in Experiment 1. The CS+ (a bluey-green circle) was S6 on the dimension, and S1 and S11 represented the extreme hue values on the dimension (blue or green, depending on counterbalancing). Within-subject error bars calculated according to the method proposed by Cousineau (2005) with the correction by Morey (2008).

η_p^2 = .022, but there was in the Differential group, $F(1, 45) = 171.9$, $p < .001$, $\eta_p^2 = .793$. The quadratic trend was significant in both Single Cue, $F(1, 72) = 15.8$, $p < .001$, $\eta_p^2 = .180$, and Differential groups, $F(1, 45) = 43.4$, $p < .001$, $\eta_p^2 = .491$. Thus, in the Single Cue group, the gradient is best described as slightly peaked.

Figure 5 also shows that for the Differential group, the peak of responding was slightly removed from the S+ in the direction away from the S− at S7, the stimulus predicted to show a peak shift. However, planned *t* tests comparing ratings for S7 against the CS+ (S6) and against the endpoint of the dimension (S11) showed that both the rise, $t(45) = 2.04$, $p = .048$, $SEM = 3.46$, and fall in ratings, $t(45) = 1.64$, $p = .109$, $SEM = 5.73$ were not significant after Holm-Bonferroni corrections. Thus, although the peak of the gradient was not at the CS+, the evidence for this overall peak shift was weak.

Subgroup analysis. Table 1 shows the number of participants categorized in each rule subgroup. It is apparent that although many participants in both groups reported a linear rule, a larger proportion of those in the Differential group reported a linear rule than in the Single Cue group. Unsurprisingly, a large proportion of participants in the Single Cue group also claimed that there was no relationship between the stimulus dimension and the outcome.

Linear subgroups. The number of participants deriving a rule in each direction (consistent/inconsistent with the training contingencies, greener/bluer stimuli more likely to lead to the outcome) is displayed in Table 2. It is clear that differential training led to a strong tendency to derive a linear rule consistent with the direction of the training contingencies. There was also a large bias in the Single Cue group in the tendency to derive a linear rule favoring greener stimuli causing the outcome rather than bluer stimuli.[3]

The gradients for the Linear subgroups are presented in Figure 6. The Linear gradients are collapsed over reported direction (bluer vs. greener more likely to lead to the outcome) so that S11 was always the stimulus that should be given the highest probability of leading to the outcome according to the training contingencies in the Differential group or the reported rule in the Single Cue group (see online supplemental materials for a more thorough breakdown of the linear gradients as well as the gradients from subgroups with small *n*s). There was a significant linear trend in both the Single Cue, $F(1, 42) = 49.7$, $p < .001$, $\eta_p^2 = .542$, and Differential groups, $F(1, 33) = 110.0$, $p < .001$, $\eta_p^2 = .769$. However, the gradient in the Differential group was significantly steeper than in the Single Cue group, $F(1, 75) = 14.6$, $p < .001$, $\eta_p^2 = .163$. This suggests that while the pattern of generalization was consistent across Single Cue and Differential groups, the presence of a CS− can be seen to strengthen belief or confidence in a linear rule because those participants had experience with an additional point along the dimension.

Similarity subgroups. For the Similarity subgroups, there was an overall quadratic trend, $F(1, 17) = 131.0$, $p < .001$, $\eta_p^2 = .885$, which interacted with conditioning group, $F(1, 17) = 10.6$, $p = .005$, $\eta_p^2 = .385$, reflecting the fact that the quadratic trend for the Differential Group was stronger than for the Single Cue group. To show that the gradient was indeed peaked, planned *t* tests were carried out for both groups. There was a significant increase in ratings from S1 to the CS+ for both the Single Cue, $t(7) = 2.90$, $p = .023$, $SEM = 12.53$, and Differential Group, $t(10) = 7.53$, $p < .001$, $SEM = 10.00$, and a significant decrease in ratings from the CS+ to S11 for both the Single Cue, $t(7) = 8.12$, $p < .001$, $SEM = 5.73$, and Differential Group, $t(10) = 4.49$, $p = .001$, $SEM = 11.03$. Thus, it is clear that peaked gradients were exhibited in both Similarity subgroups after single cue and differential training, but that the addition of the CS− in the Differential group sharpened the gradient.

It is interesting to note that in the Similarity subgroup of the Differential group, the peak of the causal ratings was actually not at the CS+, but rather at S7, the stimulus predicted to have the highest ratings in a peak-shifted gradient. If it is accepted that participants are able to learn associatively and consciously report using similarity as the basis of their generalization then, at first glance, this result may be interpreted as evidence of associative learning in a small subset of participants. However, planned *t* tests revealed that although the fall in accuracy from S7 to S11 was significant, $t(10) = 4.85$, $p = .001$, $SEM = 10.91$, the increase in accuracy from S6 to S7 was not significant, $t(10) = .339$, $p = .742$, $SEM = 9.82$. Furthermore, only one participant from this subgroup gave a numerically higher rating to S7 over S6 (see online supplemental materials for the gradients for each individual). Thus, there was no statistical evidence for a classic peak shift (significant rise and fall in causal ratings) in the Similarity subgroup who underwent differential training.

To check whether the failure to find a significant peak shift in the Differential-Similarity subgroup was a genuine null effect or because of limited power in the small subgroup, we conducted additional Bayes Factor (BF) *t* tests using the BayesFactor package

[3] Based on responses in the questionnaire, one potential reason for this imbalance is that a large number of participants associated the color green with "go," and thus arbitrarily chose linear rules in this direction in the Single Cue group.

Table 1
Number of Participants in Each Rule Subgroup in Each Experiment

Experiment	Group	Linear	Similarity	No relationship	Other
1	Single Cue	43 (58.9%)	8 (11.0%)	18 (24.7%)	4 (5.5%)
	Differential	34 (73.9%)	11 (23.9%)	1 (2.2%)	0 (0%)
2	Single Cue	8 (21.1%)	14 (36.8%)	14 (36.8%)	2 (5.3%)
	Differential	14 (42.4%)	18 (54.5%)	1 (3.0%)	0 (0%)

Note. Percentages expressed according to the number of participants in each group in each experiment (i.e., row totals).

in R using the default scale parameters. These *t* tests assume a Cauchy prior distribution over possible effect sizes concerning the difference between two samples, allocating highest prior probability to small effect sizes (see Rouder et al., 2009). A $BF_{10} < .33$ is usually taken as moderate evidence in favor of the null hypothesis (no difference between samples), a $BF_{10} > 3$ taken as moderate evidence in favor of the alternative hypothesis (a difference between samples), and a BF_{10} between .33 and 3 indicates indeterminate evidence that does not clearly favor either the null or the alternative hypothesis. Using this Bayesian *t* test, we obtained a BF_{10} of 1515.1 (strong evidence in favor of the alternative hypothesis) for the fall in causal ratings from S7 to S11, but a BF_{10} of .240 (moderate evidence in favor of the null hypothesis) for the rise in causal ratings from the CS+ to S7. Thus, although the size of the Differential-Similarity subgroup was small, there was evidence in favor of the null hypothesis (i.e., no peak shift) rather than indeterminate evidence.

No relationship subgroup. There were a large number of participants in the Single Cue group who reported no relationship between the stimuli and the outcome ($n = 18$) compared with the Differential group ($n = 1$). For this subgroup, the linear trend was not significant, $F < 1$, and the quadratic trend was marginally nonsignificant, $F(1, 17) = 4.30$, $p = .054$, $\eta_p^2 = .202$, confirming that the gradient in the No Relationship subgroup was relatively flat.

Summary

There was generally good correspondence between reported rules and generalization gradients along the hue dimension, and the shape of the overall gradients could be decomposed into distinct linear and peaked gradients exhibited in the Linear and Similarity subgroups respectively. The numerical peak in the overall gradient in the Differential group was not at the CS+ but at the test stimulus slightly displaced from CS+ in the direction away from CS−. The gradient did not decline at the extreme end of the dimension as expected in a classic peak shift, but this result can be attributed to the relatively small number of participants generalizing on the basis of similarity as compared with those using a linear rule. It follows that a more classic peak shift pattern may have been found had the ratio of Similarity to Linear subgroups been higher.

Table 2
Number of Participants Who Reported a Linear Rule in Each Direction

Experiment	Group	Consistent	Inconsistent	Greener	Bluer
1	Single cue			35	8
	Differential	32	2	17	17
2	Single cue			5	3
	Differential	14	0	8	6

Note. Rules are expressed as being either consistent or inconsistent with the training contingencies and whether they reported that a greener or bluer stimulus was more likely to lead to the outcome. Note that the consistent/inconsistent coding is orthogonal to the greener/bluer coding of the linear rules.

Experiment 2

The aim of Experiment 2 was to investigate whether peak shift could be mediated through a mixture of generalization rules used by different subgroups in a fear conditioning paradigm. Fear conditioning more closely approximates traditional animal generalization studies in using a biologically significant outcome, and hence may provide a better chance of observing participants who generalize on the basis of similarity. Moreover, in a fear conditioning study with humans, Wong and Lovibond (2017) found an approximately equal number of participants using linear and similarity rules in generalization, albeit along a different (noncontinuous) stimulus dimension to that studied here. According to our mixture-of-rules account, such a balance of linear and similarity rules should lead to an overall peak shift in generalization. We measured outcome expectancy, widely regarded as a valid index of fear conditioning (Boddez et al., 2013), as well as skin conductance.

Like Experiment 1, this study used stimuli varying in blue-green hue for training and generalization. However, pilot testing showed that a large number of participants in the Differential group failed to acquire the discrimination. This was most likely because of the longer ITI needed to allow skin conductance to return to baseline on each trial. Thus, the hue dimension was extended slightly to increase the perceptual discriminability of the CS+ and CS−.

Method

Participants. 110 University of New South Wales students (69 women, $M_{age} = 20.5$, $SD = 3.79$) participated in exchange for partial course credit or payment (AUD$15). Participants were randomly allocated to the Single Cue group ($n = 54$) or the Differential group ($n = 56$). The same exclusion criteria were used as in Experiment 1. Four participants were excluded because they indicated that they were colorblind, and a further 36 participants failed the training criterion. After exclusions, 71 participants remained (33 in the Differential group and 38 in the Single Cue group).

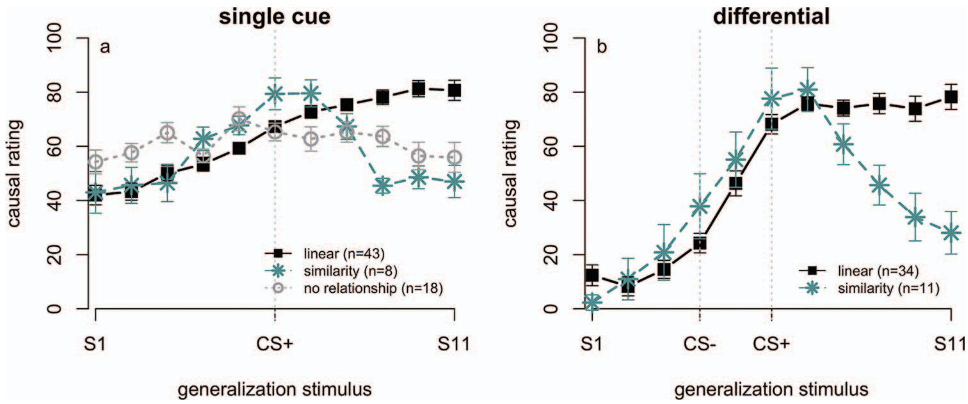

Figure 6. Generalization gradients for each rule subgroup in Experiment 1. The "linear" subgroup reported generalizing on the basis of a stimulus relation (i.e., greener or bluer stimuli led to a higher likelihood of the outcome), the "similarity" subgroup identified a single stimulus (i.e., the conditioned stimulus [CS]+) leading to the outcome with all other stimuli having a lower or zero likelihood of leading to the outcome, and the "no relationship" subgroup reported that there was no relationship between the stimuli and outcomes. Note that for the Single Cue group, linear gradients are plotted in the reported direction of the rule. Within-subject error bars calculated according to the method proposed by Cousineau (2005) with the correction by Morey (2008). See the online article for the color version of this figure.

Apparatus. The shock was delivered through stainless steel electrodes attached to the distal and middle segments of the participants' index finger of their nondominant hand. The electrodes measuring skin conductance level (SCL) were attached to the distal and proximal segments of the ring finger of the same hand. Isotonic gel was applied to the skin conductance electrodes unless the initial reading was very high, and all electrodes were secured with Leukoplast tape.

A semicircular dial with a rotary pointer was clamped to the table in front of participants in front of their dominant hand. A label was placed on the dial such that (approximately) 0 degrees represented 0% expectancy of shock (accompanied with the words "CERTAIN NO SHOCK") and 180 degrees represented 100% expectancy of shock (accompanied with the words "CERTAIN SHOCK"), and just under 0 degrees was labeled "Off." Tick marks were placed at intervals of 10. The presentation of stimuli, instructions, and triggering of shocks was controlled by Matlab and programmed using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Skin conductance was recorded throughout the whole experiment using AD instruments hardware and LabChart software.

Procedure. Experiment 2 was approved by the University of New South Wales Human Research Ethics Committee. The experiment was composed of a training phase (where the shock electrodes were detached for the first two blocks and attached for the last block), an expectancy rating test phase with the shock electrodes detached, a skin conductance test phase where the shock electrodes were reattached, and a written questionnaire (see online supplemental materials for additional details). The reason for administering both training and test in two separate phases, one with the electrodes attached and one without, was to minimize habituation to the shock as much as possible (as in Wong & Lovibond, 2017). Participants were screened for heart conditions, and underwent a shock work-up procedure to select an appropriate level of shock prior to beginning the experiment. Throughout the experiment, the participant was seated in a dimly lit cubicle with the door closed while the experimenter monitored the experiment from outside. As in Experiment 1, participants were randomly allocated to either the Single Cue or Differential group and dimension counterbalancing (green to blue or blue to green) was matched between groups.

Phase 1 (Training—electrodes detached). Similar to Experiment 1, the training phase consisted of 12 presentations of each stimulus (12 trials in total for the Single Cue group and 24 trials in total for the Differential group), with the CS+ reinforced at 75% and the CS− never reinforced. Trials were randomized in the same way as Experiment 1, ensuring that the first CS+ trial of each block was reinforced. Each stimulus consisted of a colored circle with a radius of 200 pixels presented in the middle of the screen. The hue values for the dimension were expanded by 10%, keeping the same hue values for the CS+ (S6, the midpoint, see online supplemental materials). The saturation, brightness values, and background color were the same as Experiment 1.

The trial structure consisted of a 10-s baseline period, a 10-s stimulus presentation period, a 2-s period where feedback was either presented or not presented, a 2-s period where participants were presented with the message "Please turn the expectancy dial back to the 'Off' position," and a variable ITI period (ranging between 5 and 15 s in Phases 1 and 3, and between 15 and 25 s in phases 2 and 4).

Phase 2 (Training—electrodes connected). After eight trials in the Single Cue group, or 16 trials in the Differential group had been completed, the program froze for 30 s, and instructions came on screen asking participants to wait for the experimenter. The experimenter went into the cubicle to reattach the shock electrodes and verbally reiterated that they would now be receiving an actual shock in addition to the picture of shock, and that they should use what they had learned in the first phase to continue making expectancy ratings in Phase 2. Phase 2 (the last four or eight trials in the Single Cue and Differential group, respectively) ran exactly the same way as Phase 1, except that a shock of 0.5-s duration was delivered via the electrodes at the same time as the onset of the shock feedback picture.

Phase 3 (Expectancy Generalization Test—electrodes disconnected). After Phase 2 had finished, participants were told that they would receive a break from shocks and the electrodes were disconnected again. The experimenter explained that in the following phase they would no longer receive feedback or shocks, but that they should continue to make "hypothetical" expectancy ratings as if the electrodes were still connected and shock were still possible. After checking for understanding, the experimenter left the cubicle. The expectancy test consisted of 11 trials—one presentation of each of the generalization stimuli in randomized order.

Phase 4 (Skin Conductance Generalization Test—electrodes connected). After Phase 3, participants were told that now it was possible to receive shock again, and the experimenter reconnected the shock electrodes. The skin conductance test consisted of five stimuli only—S1, S4 (CS−), S6 (CS+), S7, and S11. Each stimulus was presented once in randomized order, and no shocks were presented in this phase.

Questionnaire. At the end of the experiment, all electrodes were detached and the participant was led outside the cubicle to complete a questionnaire. The questionnaire was similar to that used in Experiment 1, except that the rules and questions were phrased in terms of predicting shock. A total of 25 participants were classified into a subgroup based on their reported rule, while the forced-choice question was used for the remaining participants.

Results and Discussion

There were no significant results in the skin conductance data in generalization testing, and in general, the data were highly variable. As such, they will not be presented here (see online supplemental materials for analysis of the skin conductance data).

Training. In expectancy ratings, there was a significant linear increase in ratings for the CS+ in the Single Cue group, $F(1, 37) = 20.1$, $p < .001$, $\eta_p^2 = .352$. In the Differential group, ratings for the CS+ were significantly higher than ratings for the CS− overall, $F(1, 32) = 553.1$, $p < .001$, $\eta_p^2 = .945$. Again, there was a significant interaction between the linear trend in ratings for the CS+ and CS−, $F(1, 32) = 174.3$, $p < .001$, $\eta_p^2 = .845$, indicating that expectancy ratings for the CSs diverged over training and that participants learned the discrimination (Figure 7).

Generalization Test. Figure 8 displays the overall generalization gradients in each group. There were some similarities to Experiment 1, but also some important differences. The gradient for the Single Cue group appears to be symmetrical and more peaked than in Experiment 1. In the Differential group, the gradient now appears distinctly peak-shifted, with a clear decrease on the extreme right of the dimension that was not present in Experiment 1. Similar to Experiment 1, there was an overall linear, $F(1, 69) = 47.2$, $p < .001$, $\eta_p^2 = .406$, and quadratic trend, $F(1, 69) = 34.1$, $p < .001$, $\eta_p^2 = .330$. The linear trend interacted with group, $F(1, 69) = 26.3$, $p < .001$, $\eta_p^2 = .276$, while the quadratic trend did not, $F < 1$. Again, we can conclude that adding the CS− in the Differential Group changed the shape of the generalization gradient.

As in Experiment 1, there was a significant linear trend in the Differential group, $F(1, 32) = 63.7$, $p < .001$, $\eta_p^2 = .666$, but not

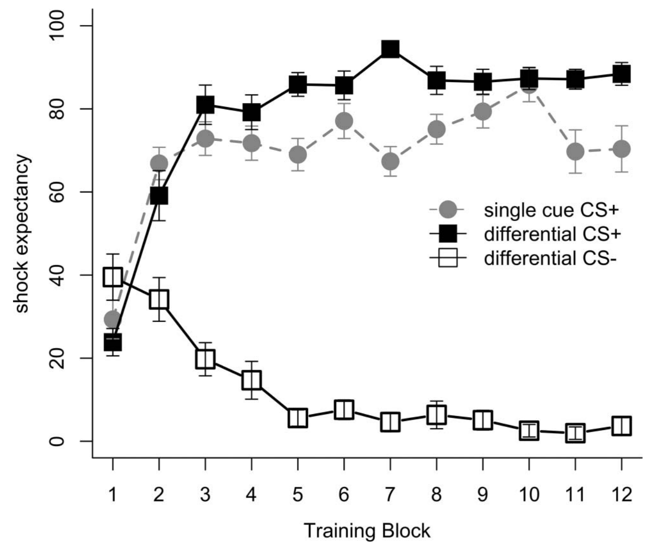

Figure 7. Shock expectancy ratings over stimulus presentations in training for Experiment 2. The CS+ was reinforced at a rate of 75% and the CS− was never reinforced. Within-subject error bars calculated according to the method proposed by Cousineau (2005) with the correction by Morey (2008).

in the Single Cue group, $F(1, 37) = 1.72$, $p = .198$, $\eta_p^2 = .044$. Both the Single Cue, $F(1, 37) = 14.9$, $p < .001$, $\eta_p^2 = .287$, and Differential groups, $F(1, 32) = 20.6$, $p < .001$, $\eta_p^2 = .392$, showed a significant quadratic trend in their overall gradients. To further test whether the Single Cue group's gradient was peaked, planned t tests were conducted. The comparison between S1 and the CS+ was significant, $t(37) = 3.91$, $p < .001$, $SEM = 6.93$, but the comparison between the CS+ and S11 was marginally nonsignificant, $t(37) = 1.89$, $p = .066$, $SEM = 8.32$.

To test whether a peak shift occurred in the Differential group, the peak (S7) was compared with S6 (CS+) and S11 (the endpoint). There was a significant rise between the CS+ and S7, $t(32) = 2.12$, $p = .042$, $SEM = 6.27$, and a significant fall from S7 to S11, $t(32) = 3.44$, $p = .002$, $SEM = 6.72$. Thus, the Differential group showed a significantly peak-shifted gradient.

Subgroup analysis. Figure 9 shows the gradients in expectancy ratings obtained for each rule subgroup in the Single Cue (left panels) and Differential (right panels) groups. The number of participants assigned to each rule subgroup is shown in Table 1 (see online supplemental materials for the gradients from subgroups with small n). The table shows the ratio of the Similarity to Linear subgroup participants is more balanced in this experiment than in Experiment 1.

Linear subgroups. For participants who reported using a linear rule (see Figure 9), there was a significant linear trend, $F(1, 20) = 45.8$, $p < .001$, $\eta_p^2 = .696$, which, unlike Experiment 1, did not interact with group, $F(1, 20) = 1.16$, $p = .295$, $\eta_p^2 = .055$. Unsurprisingly, the significant linear trend was present in both the Single Cue, $F(1, 7) = 17.5$, $p = .004$, $\eta_p^2 = .715$, and Differential groups, $F(1, 13) = 36.9$, $p < .001$, $\eta_p^2 = .739$.

Similarity subgroups. Like Experiment 1, there was a significant overall quadratic trend, $F(1, 30) = 32.2$, $p < .001$, $\eta_p^2 = .517$, which interacted with group, $F(1, 30) = 13.3$, $p = .001$, $\eta_p^2 = .307$,

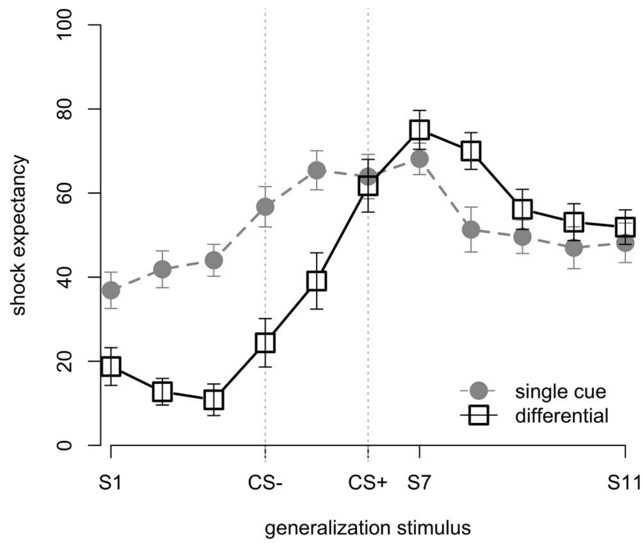

Figure 8. Overall generalization gradients for both groups in Experiment 2. The CS+ (a bluey-green circle) was S6 on the dimension, and S1 and S11 represented the extreme hue values on the dimension (blue or green, depending on counterbalancing). Within-subject error bars calculated according to the method proposed by Cousineau (2005) with the correction by Morey (2008).

again reflecting the sharper gradient seen in the Differential group. Planned *t* tests comparing S1 to S6 and S6 to S11 were both significant, for both groups, *t*s(13) > = 3.24, $p = .006$, $SEM = 10.41$, indicating peaked gradients. Again, it is interesting to note that the gradient in the Differential group appears to be peak-shifted, with the highest ratings given to S7, as in Experiment 1. However, while the decrease in ratings from the peak (S7) to S11 was significant, $t(17) = 4.63$, $p < .001$, $SEM = 8.14$, the increase from CS+ to S7 was not, $t < 1$ (see online supplemental materials for the individual gradients from the Differential-Similarity subgroup). We conducted the same Bayesian *t* tests as in Experiment 1 to verify this null effect. Like Experiment 1, there was strong evidence for the fall in ratings from S7 to S11, $BF_{10} = 130.2$, but there was moderate evidence in favor of the null hypothesis when comparing ratings from the CS+ to S7, $BF_{10} = .290$. Thus, the Bayesian *t* tests support the hypothesis that participants gave the same ratings for CS+ and S7.

No relationship subgroup. Despite the appearance of a peaked gradient in the No Relationship subgroup in the Single Cue group, there were no significant linear, $F < 1$, or quadratic trends, $F(1, 13) = 2.88$, $p = .114$, $\eta_p^2 = .181$. There was also no significant difference between ratings of S1 and of the CS+, $t(13) = 1.97$, $p = .070$, $SEM = 12.13$, nor between the CS+ and S11, $t(13) = 1.67$, $p = .120$, $SEM = 14.00$.

Summary

Experiment 2 showed similar results to Experiment 1 in that the empirical generalization gradients corresponded to participants' reported rules. In this experiment, and consistent with previous animal conditioning studies, an overall peak-shifted gradient was found in the Differential group, and an overall peaked gradient was found in the Single Cue group. As we predicted, these overall gradients can be explained through a combination of the distribution of rule subgroups, and the different gradients exhibited by each subgroup. While the proportion of participants generalizing on the basis of similarity did not differ greatly between conditioning groups (Table 1), participants in the Differential group primarily derived linear rules in the direction consistent with their training contingencies, while participants in the Single Cue group did not derive linear rules in a consistent direction. This shift in the direction of linear rules as result of differential training led to an overall peak-shifted gradient in the Differential group that was not present in the Single Cue group.

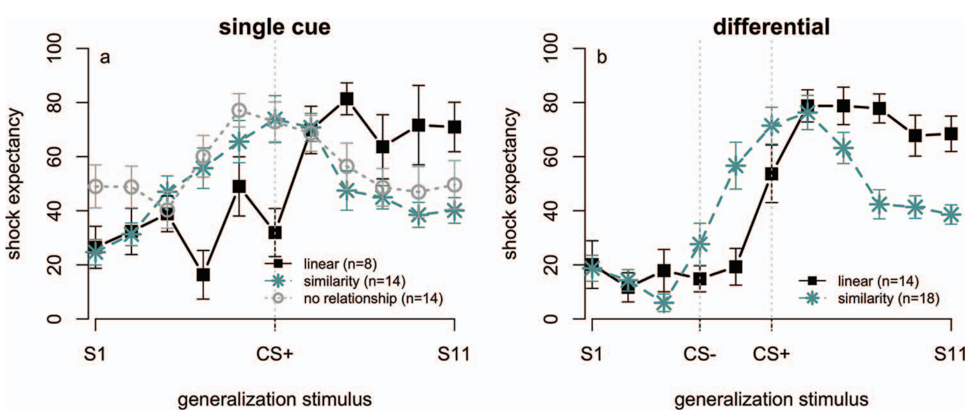

Figure 9. Generalization gradients for each rule subgroup in Experiment 2. The "linear" subgroup reported generalizing on the basis of a stimulus relation (i.e., greener or bluer stimuli led to a higher likelihood of the outcome), the "similarity" subgroup identified a single stimulus (i.e., the CS+) leading to the outcome with all other stimuli having a lower or zero likelihood of leading to the outcome, and the "no relationship" subgroup reported that there was no relationship between the stimuli and outcomes. Note that for the Single Cue group, linear gradients are plotted in the reported direction of the rule. Within-subject error bars calculated according to the method proposed by Cousineau (2005) with the correction by Morey (2008). See the online article for the color version of this figure.

Overall, the generalization gradients obtained within each rule subgroup were broadly consistent with Experiment 1, despite the introduction of a biologically relevant shock outcome, the slight extension of the hue dimension, and the increased difficulty of learning introduced by lengthening the ITI. However, the Experiment 2 results extend those of Experiment 1 in an important way by finding clear evidence of overall peak shift, which was based on a mixture of generalization rules across different subgroups.

General Discussion

In the current studies, we demonstrated that the shape of empirical generalization gradients conformed to participants' reported rules, and further, that a peak-shifted gradient could be found at the group-level by averaging subgroups of participants who reported different generalization rules. Notably, we showed that in fear conditioning, an overall peak shift in generalization of outcome expectancy along a perceptual dimension was the result of a mixture of gradients from participants who reported generalizing on the basis of a relational rule, and from those who reported generalizing on the basis of similarity to the trained CS+. These results demonstrate the importance of considering individual differences in generalization rules, and show that peak shift in humans can result from averaging distinct gradients arising from separate rule subgroups.

In a causal judgment paradigm (Experiment 1) and a fear conditioning paradigm (Experiment 2), we found generalization gradients in causal judgments (Experiment 1) and shock outcome expectancy (Experiment 2) that corresponded to participants' verbalizable rules. A linear generalization gradient was exhibited by a subgroup who reported generalizing on the basis of a relational (linear) rule, and a peaked gradient was found in the subgroup who reported generalizing according to stimulus similarity to the CS+. Participants reporting no relationship between the stimuli and the outcomes exhibited noisy, but mostly flat gradients. It is worth highlighting that although the gradients displayed in the Linear subgroups of the Differential group are consistent with a relational rule (i.e., increasing along the hue continuum in the direction of the CS+ and away from the CS−), the gradients appeared to be more like a step function than a linear function (Figures 6 and 9). In both experiments (especially Experiment 2), the gradient between the CS+ and CS− appears to be steep but the gradient beyond the CS+ is quite flat. This pattern of results might be explained if participants notice the difference between the CS+ and the CS− and classify one stimulus as "green" and the other as "blue" during differential training. The use of these labels might facilitate sorting of the generalization stimuli into these two categories on test, effectively treating stimuli on the extreme ends of the dimension as equivalent and producing flat gradients at each end. Despite this, the overall pattern of results broadly replicate the findings of Ahmed and Lovibond (2016, 2017), and Wong and Lovibond (2017) but with a continuous stimulus dimension (hue), which overcomes issues with the stimuli used in these previous studies (e.g., the use of intensity or noncontinuous dimensions).

A Rule-Based Explanation of Peak Shift

The key novel finding from our study is that peak shift can be a result of averaging group data, and that it can be decomposed into distinct linear and peaked gradients by asking participants to report their generalization rules. The finding of a significant overall peak shift was restricted to Experiment 2 and appeared to be mediated by the relative proportion of participants reporting linear as opposed to similarity-based generalization rules. Therefore, an overall peak shift might only be obtained by averaging subgroups of participants who primarily report generalizing on the basis of either a linear or similarity rule when the numbers of participants in each subgroup are roughly equal. Evidence for overall peak shift was weaker in Experiment 1, when there was a higher proportion of linear rules (and fewer similarity rules) reported by those in the Differential condition.

Given the many differences between the paradigms used in the two studies, we can only speculate about the reasons for differences in the distribution of rule subgroups. One likely possibility is that the higher proportion of linear rules in Experiment 1 was because of the considerably shorter ITI used in that study compared with Experiment 2. The short ITI means that "extreme" values on the test dimension were more likely to appear close together in time. This contrast may have highlighted the hue relations (i.e., greener and bluer), producing a higher proportion of linear rules. The fact that many participants in the Single Cue group reported linear rules supports this idea, because they only had experience with a single point on the dimension and yet reported a linear rule based on this ambiguous information.

One limitation of Experiment 2 is that we did not obtain an overall peak shift in the skin conductance measure. While there was evidence of acquisition in the skin conductance data (see online supplemental materials), the generalization gradients obtained were highly variable, and generally flat. This may be attributed to our procedure (testing expectancy first, and skin conductance separately) being insufficient in minimizing habituation to shock. Indeed, during training, the highest skin conductance levels were recorded on the first reinforced CS+ trial in both Single Cue and Differential groups, and declined over the next three CS+ presentations (see online supplemental materials). Fear learning is known to extinguish rapidly and has been noted to be highly variable when using a single test trial for each generalization stimulus (Vervoort, Vervliet, Bennett, & Baeyens, 2014), meaning that power is limited when examining skin conductance, especially after dividing participants into rule subgroups. However, when significant effects with skin conductance have been found using paradigms broadly similar to those used here (Wong & Lovibond, 2017), they have generally aligned with expectancy ratings. This suggests that reported rules do predict generalization patterns on physiological as well as self-report measures.

The mechanism through which peak shift occurs in our study is quite different to that posited by associative models. Instead of the CS− serving only as a source of inhibition, the CS− in our study served to promote formation of linear rules in the direction consistent with the experienced training contingencies. It is thus easy to see how an overall area-shift could emerge by combining gradients from the Linear and Similarity subgroups, while the conditions needed for a peak-shift are considerably stricter and dependent on the parameters of the linear and similarity gradients themselves. This aligns with the fact that peak shift is a relatively more elusive phenomenon in the human generalization literature than area shift (see Honig & Urcuioli, 1981). It also provides a possible explanation for why Wong and Lovibond (2017) did not find an overall peak shift in their study, since the salient nature of their CS+ (a dot in the center of a rectangular frame) meant that the highest ratings were given to the CS+ in both Linear and Similarity subgroups. The current study is thus consistent with the empirical results found in animals, but suggests that in humans, peak shift can be mediated through the distribution of rule subgroups.

Note that our results do not allow us to determine whether these differences in reported rules are because of differences in prior experience, or reflect stable individual differences in ability or rule use.[4] There is evidence from related literature that working memory predicts whether individuals derive an abstract rule or rely on exemplar memory in a function learning task (McDaniel et al., 2014). To speculate, in our task, similarity-based responding may be due to a more perceptual process that relies on episodic memory and less on working memory. In contrast, the extraction of an explicit linear rule might be seen as more abstract reasoning that is dependent on an individual's working memory capacity. In any case, our experiments show that despite potential stable individual differences in generalization rules, the distribution of rule subgroups is also clearly influenced by the nature of the task (causal judgment vs. fear conditioning), as well as the choice of stimuli (e.g., hue in our experiments vs. a noncontinuous dimension in Wong & Lovibond, 2017).

Methodological Considerations

It is worth noting that generalization gradients can vary greatly with the choice of testing procedure. Gradients are generally steeper with within-subjects testing (the procedure used here) than with between-subjects testing (e.g., see Vervliet, Iberico, Vervoort, & Baeyens, 2011), implying that participants continue to learn about the range of potential stimuli during the test phase (associatively or through formation of new rules). In the context of our experiments, it is reasonable to ask whether participants generalize on the basis of rules formed during training on the very first test trial, or whether they derive the rules throughout the test phase while reflecting back on the training contingencies. Despite the fact that we withheld feedback during the generalization test, learning about the "stimulus space" might change or reinforce rules that participants derived during training (see Livesey & McLaren, 2009). Future studies that measure generalization between-subjects is needed to confirm whether this occurs.

An additional consideration concerning our testing procedure is that our choice of the CS− as S4 and CS+ as S6 on the dimension may have resulted in slight "adaptation level" effects (see Capehart, Tempone, & Hebert, 1969; Thomas, 1993). According to this adaption level account, during differential training, participants would store a representation of the average stimulus value seen overall (i.e., S5, the average of S4 and S6), and learn to respond by comparing each training stimulus to this adaptation value. Over the course of the randomized test trials, the adaptation level would shift from S5 (the average of the training stimuli) to a value close to S6 (the average of the test stimuli). If participants continue to use a similar response rule as in training (e.g., "give high ratings to stimuli greener than the adaptation level"), then this may result in a peak shift. Although there is some evidence supporting this account (e.g., Thomas, Mood, Morrison, & Wiertelak, 1991; Thomas & Jones, 1962), many of these experiments have used quite extreme asymmetrical testing to produce peak shifts (e.g., Thomas et al., 1991), while in the current experiments the adaptation level during training (S5) was quite close to the adaptation level during test (S6). In addition, most of the previous demonstrations of adaptation effects used absolute identification tasks, where participants are instructed to remember and respond to a single stimulus (the target stimulus) throughout the task. Arguably, this task is rather different to that of a conditioning task.

Another issue is whether to treat generalization as an "active" process that generates similar behavior to a novel (but different) stimulus, or as a failure to discriminate between a generalization stimulus (GS) and the CS+ (Struyf, Zaman, Vervliet, & Van Diest, 2015; see also the early debate between Hull, 1943; and Lashley & Wade, 1946). Like the bulk of the generalization literature, we have assumed that participants view each stimulus along the dimension as a distinct stimulus and can discriminate between them and therefore generalize actively. In other words, we have assumed that generalization does not mean lack of discrimination. While our stimuli were difficult to discriminate (as seen in our high exclusion rates in the Differential group), the fact that we imposed a training criterion, and the sensible generalization gradients obtained suggest that participants could distinguish between adjacent stimuli. Still, it would be interesting to explore whether perceptual discrimination can predict the types of rules participants derive, or examine the relationship between peak shift and ability to discriminate between test stimuli in future studies.

Associative and Rule-Based Generalization

From an associative view, it may be argued that only the Similarity subgroups are learning and generalizing associatively. The peak-shifted shape of the gradients in both experiments is certainly consistent with an associative account, although our tests for this peak shift were not significant. Given many more participants, these gradients may have shown a significant peak shift, providing evidence of associative processes. While this is possible, the individual gradients for each participant (see online supplemental materials) do not support this view. While many participants in the Differential-Similarity subgroups showed an area shift in their gradients, very few (1/13 in Experiment 1 and 1/18 in Experiment 2) showed a clear peak-shifted gradient. This stands in contrast to the animal literature, where there is often a remarkable uniformity in the shape of the gradients when peak shift is found (e.g., see the individual gradients in Blough, 1973).

Another consideration is that our division of participants into rule subgroups was based on self-report, and therefore may not have been completely reliable. Thus, participants in the Differential-Similarity subgroup may still have entertained some belief in a linear rule. Therefore, if a significant peak shift in this subgroup were obtained, it could still be explained as the result of combining linear and similarity-based generalization. This idea that participants may derive and express learning consistent with multiple rules receives support from the categorization literature where individual participants have been observed to display a mixture of rule-based and exemplar-based (i.e., similarity-based) responding (e.g., Erickson & Kruschke, 1998; Hahn, Prat-Sala,

[4] Note that we did not find evidence interactions between rule subgroup and training performance (see online supplemental materials).

Pothos, & Brumby, 2010). A more comprehensive assessment of rules taking into account belief in multiple rules would be needed in order to judge whether peak shift (at an individual- or group-level) is due to associative processes. Another possibility is to test whether participants who display a peak shift can identify the relational difference between the CS+ and CS−, since an associative model can predict peak shift in the absence of this information. In summary, although the gradients exhibited in the Differential-Similarity subgroups appear to be peak-shifted in both experiments, there is still considerable variability at the subject-level, and it is only when these gradients are combined with the linear gradients that an overall peak shift emerges.

Conclusion

The present study adds to the literature which shows that there are large individual differences in the strategies and rules that participants use when confronted with learning and generalization tasks (e.g., Erickson & Kruschke, 1998; Gluck et al., 2002; Little & McDaniel, 2015; Smith & Sloman, 1994), and that verbal report is a useful way of separating out these strategies. We extended the results of Ahmed and Lovibond, 2016, 2017 and Wong and Lovibond (2017) using a continuous stimulus dimension, showing that generalization gradients differ according to reported rules. Crucially, we showed how an overall peak-shifted gradient in fear conditioning could be found by combining gradients from subgroups of participants who primarily reported generalizing on the basis of a linear rule, or stimulus similarity. Our study shows that peak shift can emerge from distinct rule subgroups and therefore does not uniquely support traditional associative accounts based on perceptual similarity. These results highlight the need for researchers in the field to look beyond group generalization gradients and carefully examine qualitative differences between individuals in their generalization strategies. Generalization in conditioning does not always involve associative processes and, in some cases, is better conceived of as a process of rule generation and application.

References

Ahmed, O., & Lovibond, P. F. (2015). The impact of instructions on generalisation of conditioned fear in humans. *Behavior Therapy, 46,* 597–603. http://dx.doi.org/10.1016/j.beth.2014.12.007

Ahmed, O., & Lovibond, P. F. (2016). Rule-based processes in generalisation and peak shift in human fear conditioning. Manuscript submitted for publication.

Ahmed, O., & Lovibond, P. F. (2017). The influence of inferred rules on generalization in human predictive learning. Manuscript submitted for publication.

Blough, D. S. (1973). Two-way generalization peak shift after two-key training in the pigeon. *Animal Learning & Behavior, 1,* 171–174. http://dx.doi.org/10.3758/BF03199068

Blough, D. S. (1975). Steady state data and a quantitative model of operant generalization and discrimination. *Journal of Experimental Psychology, 104,* 3–21.

Boddez, Y., Baeyens, F., Luyten, L., Vansteenwegen, D., Hermans, D., & Beckers, T. (2013). Rating data are underrated: Validity of U.S. expectancy in human fear conditioning. *Journal of Behavior Therapy and Experimental Psychiatry, 44,* 201–206. http://dx.doi.org/10.1016/j.jbtep.2012.08.003

Boddez, Y., Bennett, M. P., van Esch, S., & Beckers, T. (2017). Bending rules: The shape of the perceptual generalisation gradient is sensitive to inference rules. *Cognition and Emotion, 31,* 1444–1452.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10,* 433–436. http://dx.doi.org/10.1163/156856897X00357

Capehart, J., Tempone, V. J., & Hebert, J. (1969). A theory of stimulus equivalence. *Psychological Review, 76,* 405–418. http://dx.doi.org/10.1037/h0027418

Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorials in Quantitative Methods for Psychology, 1,* 42–45. http://dx.doi.org/10.20982/tqmp.01.1.p042

DeLosh, E. L., Busemeyer, J. R., & McDaniel, M. A. (1997). Extrapolation: The sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 968–986. http://dx.doi.org/10.1037/0278-7393.23.4.968

Don, H. J., Goldwater, M. B., Otto, A. R., & Livesey, E. J. (2016). Rule abstraction, model-based choice, and cognitive reflection. *Psychonomic Bulletin & Review, 23,* 1615–1623. http://dx.doi.org/10.3758/s13423-016-1012-y

Dunsmoor, J. E., & LaBar, K. S. (2013). Effects of discrimination training on fear generalization gradients and perceptual classification in humans. *Behavioral Neuroscience, 127,* 350–356. http://dx.doi.org/10.1037/a0031933

Dunsmoor, J. E., Martin, A., & LaBar, K. S. (2012). Role of conceptual knowledge in learning and retention of conditioned fear. *Biological Psychology, 89,* 300–305. http://dx.doi.org/10.1016/j.biopsycho.2011.11.002

Dunsmoor, J. E., Mitroff, S. R., & LaBar, K. S. (2009). Generalization of conditioned fear along a dimension of increasing fear intensity. *Learning & Memory, 16,* 460–469. http://dx.doi.org/10.1101/lm.1431609

Dunsmoor, J. E., & Murphy, G. L. (2014). Stimulus typicality determines how broadly fear is generalized. *Psychological Science, 25,* 1816–1821. http://dx.doi.org/10.1177/0956797614535401

Dunsmoor, J. E., & Murphy, G. L. (2015). Categories, concepts, and conditioning: How humans generalize fear. *Trends in Cognitive Sciences, 19,* 73–77. http://dx.doi.org/10.1016/j.tics.2014.12.003

Dymond, S., Dunsmoor, J. E., Vervliet, B., Roche, B., & Hermans, D. (2015). Fear generalization in humans: Systematic review and implications for anxiety disorder research. *Behavior Therapy, 46,* 561–582. http://dx.doi.org/10.1016/j.beth.2014.10.001

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General, 127,* 107–140. http://dx.doi.org/10.1037/0096-3445.127.2.107

Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin, 53,* 134–140. http://dx.doi.org/10.1037/h0045156

Ghirlanda, S., & Enquist, M. (1998). Artificial neural networks as models of stimulus control. *Animal Behaviour, 56,* 1383–1389. http://dx.doi.org/10.1006/anbe.1998.0903

Gluck, M. A., Shohamy, D., & Myers, C. (2002). How do people solve the "weather prediction" task?: Individual variability in strategies for probabilistic category learning. *Learning & Memory, 9,* 408–418. http://dx.doi.org/10.1101/lm.45202

Grice, G. R., & Saltz, E. (1950). The generalization of an instrumental response to stimuli varying in the size dimension. *Journal of Experimental Psychology, 40,* 702–708. http://dx.doi.org/10.1037/h0054435

Grillon, C., & Morgan, C. A., III. (1999). Fear-potentiated startle conditioning to explicit and contextual cues in Gulf War veterans with posttraumatic stress disorder. *Journal of Abnormal Psychology, 108,* 134–142. http://dx.doi.org/10.1037/0021-843X.108.1.134

Guttman, N., & Kalish, H. I. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology, 51,* 79–88. http://dx.doi.org/10.1037/h0046219

Hahn, U., Prat-Sala, M., Pothos, E. M., & Brumby, D. P. (2010). Exemplar similarity and rule application. *Cognition, 114,* 1–18. http://dx.doi.org/10.1016/j.cognition.2009.08.011

Hanson, H. M. (1957). Discrimination training effect on stimulus generalization gradient for spectrum stimuli. *Science, 125,* 888–889. http://dx.doi.org/10.1126/science.125.3253.888

Hayes, B. K., & Heit, E. (in press). Induction 2.0. *Wiley Interdisciplinary Reviews: Cognitive Science*.

Honig, W. K., Boneau, C. A., Burstein, K. R., & Pennypacker, H. S. (1963). Positive and negative generalization gradients obtained after equivalent training conditions. *Journal of Comparative and Physiological Psychology, 56,* 111–116. http://dx.doi.org/10.1037/h0048683

Honig, W. K., & Urcuioli, P. J. (1981). The legacy of Guttman and Kalish (1956): Twenty-five years of research on stimulus generalization. *Journal of the Experimental Analysis of Behavior, 36,* 405–445. http://dx.doi.org/10.1901/jeab.1981.36-405

Hull, C. L. (1943). *Principles of behavior: An introduction to behavior theory*. Oxford, England: Appleton-Century.

Hull, C. L. (1949). Stimulus intensity dynamism (*V*) and stimulus generalization. *Psychological Review, 56,* 67–76. http://dx.doi.org/10.1037/h0058051

Laberge, D. (1961). Generalization gradients in a discrimination situation. *Journal of Experimental Psychology, 62,* 88–94. http://dx.doi.org/10.1037/h0043035

Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics, 33,* 159–174. http://dx.doi.org/10.2307/2529310

Lashley, K. S., & Wade, M. (1946). The Pavlovian theory of generalization. *Psychological Review, 53,* 72–87. http://dx.doi.org/10.1037/h0059999

Lewis, M. B., & Johnston, R. A. (1999). Are caricatures special? Evidence of peak shift in face recognition. *European Journal of Cognitive Psychology, 11,* 105–117. http://dx.doi.org/10.1080/713752302

Lissek, S., Kaczkurkin, A. N., Rabin, S., Geraci, M., Pine, D. S., & Grillon, C. (2014). Generalized anxiety disorder is associated with overgeneralization of classically conditioned fear. *Biological Psychiatry, 75,* 909–915. http://dx.doi.org/10.1016/j.biopsych.2013.07.025

Lissek, S., Rabin, S., Heller, R. E., Lukenbaugh, D., Geraci, M., Pine, D. S., & Grillon, C. (2010). Overgeneralization of conditioned fear as a pathogenic marker of panic disorder. *The American Journal of Psychiatry, 167,* 47–55. http://dx.doi.org/10.1176/appi.ajp.2009.09030410

Little, J. L., & McDaniel, M. A. (2015). Individual differences in category learning: Memorization versus rule abstraction. *Memory & Cognition, 43,* 283–297. http://dx.doi.org/10.3758/s13421-014-0475-1

Livesey, E. J., & McLaren, I. P. L. (2009). Discrimination and generalization along a simple dimension: Peak shift and rule-governed responding. *Journal of Experimental Psychology: Animal Behavior Processes, 35,* 554–565. http://dx.doi.org/10.1037/a0015524

Mackintosh, N. J. (1974). *The psychology of animal learning*. London, England: Academic Press.

Maddox, W. T. (1999). On the dangers of averaging across observers when comparing decision bound models and generalized context models of categorization. *Perception & Psychophysics, 61,* 354–374. http://dx.doi.org/10.3758/BF03206893

McDaniel, M. A., & Busemeyer, J. R. (2005). The conceptual basis of function learning and extrapolation: Comparison of rule-based and associative-based models. *Psychonomic Bulletin & Review, 12,* 24–42. http://dx.doi.org/10.3758/BF03196347

McDaniel, M. A., Cahill, M. J., Robbins, M., & Wiener, C. (2014). Individual differences in learning and transfer: Stable tendencies for learning exemplars versus abstracting rules. *Journal of Experimental Psychology: General, 143,* 668–693. http://dx.doi.org/10.1037/a0032963

McLaren, I. P. L., & Mackintosh, N. J. (2002). Associative learning and elemental representation: II. Generalization and discrimination. *Animal Learning & Behavior, 30,* 177–200. http://dx.doi.org/10.3758/BF03192828

Morey, R. D. (2008). Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology, 4,* 61–64. http://dx.doi.org/10.20982/tqmp.04.2.p061

Morey, R. A., Dunsmoor, J. E., Haswell, C. C., Brown, V. M., Vora, A., Weiner, J., . . . the VA Mid-Atlantic MIRECC Workgroup. (2015). Fear learning circuitry is biased toward generalization of fear associations in posttraumatic stress disorder. *Translational Psychiatry, 5,* e700. http://dx.doi.org/10.1038/tp.2015.196

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review, 101,* 53–79. http://dx.doi.org/10.1037/0033-295X.101.1.53

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442. http://dx.doi.org/10.1163/156856897X00366

Purtle, R. B. (1973). Peak shift: A review. *Psychological Bulletin, 80,* 408–421. http://dx.doi.org/10.1037/h0035233

Razran, G. (1949). Stimulus generalization of conditioned responses. *Psychological Bulletin, 46,* 337–365. http://dx.doi.org/10.1037/h0060507

R Core Team. (2015). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from https://www.R-project.org/

Regehr, G., & Brooks, L. R. (1995). Category organization in free classification: The organizing effect of an array of stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 347–363. http://dx.doi.org/10.1037/0278-7393.21.2.347

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian *t* tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review, 16,* 225–237. http://dx.doi.org/10.3758/PBR.16.2.225

Shanks, D. R. (2007). Associationism and cognition: Human contingency learning at 25. *Quarterly Journal of Experimental Psychology, 60,* 291–309. http://dx.doi.org/10.1080/17470210601000581

Shanks, D. R., & Darby, R. J. (1998). Feature- and rule-based generalization in human associative learning. *Journal of Experimental Psychology: Animal Behavior Processes, 24,* 405–415. http://dx.doi.org/10.1037/0097-7403.24.4.405

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science, 237,* 1317–1323. http://dx.doi.org/10.1126/science.3629243

Smith, E. E., & Sloman, S. A. (1994). Similarity- versus rule-based categorization. *Memory & Cognition, 22,* 377–386. http://dx.doi.org/10.3758/BF03200864

Spence, K. W. (1937). The differential response in animals to stimuli varying within a single dimension. *Psychological Review, 44,* 430–444. http://dx.doi.org/10.1037/h0062885

Struyf, D., Iberico, C., & Vervliet, B. (2014). Increasing predictive estimations without further learning: The peak-shift effect. *Experimental Psychology, 61,* 134–141. http://dx.doi.org/10.1027/1618-3169/a000233

Struyf, D., Zaman, J., Vervliet, B., & Van Diest, I. (2015). Perceptual discrimination in fear generalization: Mechanistic and clinical implications. *Neuroscience and Biobehavioral Reviews, 59,* 201–207. http://dx.doi.org/10.1016/j.neubiorev.2015.11.004

ten Cate, C., Verzijden, M. N., & Etman, E. (2006). Sexual imprinting can induce sexual preferences for exaggerated parental traits. *Current Biology, 16,* 1128–1132. http://dx.doi.org/10.1016/j.cub.2006.03.068

Thomas, D. R. (1993). A model for adaptation-level effects on stimulus generalization. *Psychological Review, 100,* 658–673. http://dx.doi.org/10.1037/0033-295X.100.4.658

Thomas, D. R., & Jones, C. G. (1962). Stimulus generalization as a function of the frame of reference. *Journal of Experimental Psychology, 64,* 77–80. http://dx.doi.org/10.1037/h0043304

Thomas, D. R., Mood, K., Morrison, S., & Wiertelak, E. (1991). Peak shift revisited: A test of alternative interpretations. *Journal of Experimental*

Psychology: Animal Behavior Processes, 17, 130–140. http://dx.doi.org/10.1037/0097-7403.17.2.130

Vervliet, B., Iberico, C., Vervoort, E., & Baeyens, F. (2011). Generalization gradients in human predictive learning: Effects of discrimination training and within-subjects testing. *Learning and Motivation, 42,* 210–220. http://dx.doi.org/10.1016/j.lmot.2011.03.004

Vervliet, B., Kindt, M., Vansteenwegen, D., & Hermans, D. (2010). Fear generalization in humans: Impact of verbal instructions. *Behaviour Research and Therapy, 48,* 38–43. http://dx.doi.org/10.1016/j.brat.2009.09.005

Vervoort, E., Vervliet, B., Bennett, M., & Baeyens, F. (2014). Generalization of human fear acquisition and extinction within a novel arbitrary stimulus category. *PLoS ONE, 9,* e96569. http://dx.doi.org/10.1371/journal.pone.0096569

Wills, S., & Mackintosh, N. J. (1998). Peak shift on an artificial dimension. *The Quarterly Journal of Experimental Psychology, 51B,* 1–31.

Wong, A. H. K., & Lovibond, P. F. (2017). Rule-based generalisation in single-cue and differential fear conditioning in humans. *Biological Psychology, 129,* 111–120. http://dx.doi.org/10.1016/j.biopsycho.2017.08.056

Received May 14, 2017
Revision received October 27, 2017
Accepted January 4, 2018 ■