Title:

To perceive or not to perceive? Individual differences in perception predict rule induction and response generalization

Authors:

Jonas Zaman^{a,b} and Jessica C. Lee^c

a Health Psychology, Faculty of Psychology and Educational Sciences, KU Leuven, Tiensestraat 102, box 3726, 3000 Leuven, Belgium

b Centre for learning and Experimental Psychopathology, KU Leuven, Tiensestraat 102, box 3726, 3000 Leuven, Belgium

c University of New South Wales, Sydney, NSW 2052, Australia

Manuscript information:

Number of figures: 5

Number of Tables: 2

Word count: 7060

Supplemental information: No

Corresponding author:

Jonas Zaman, PhD

Health psychology

Tiensestraat 102 - bus 3726

3000 Leuven

Tel. +32 16 37 31 97

Fax +32 16 3 26144

E-mail: Jonas.zaman@kuleuven.be

Institutional URL : <https://ppw.kuleuven.be/ogp>

Abstract

When novel stimuli trigger a previously learned response, this can be due to failure to perceive the novel stimulus as different from the trained stimulus (perception), or active extrapolation of learned properties from the trained stimulus (induction). To date, there has been little investigation of how individual differences in perceptual ability relate to differences in induction. In this paper, we perform cluster analysis in six datasets (four published datasets and two unpublished datasets, N = 992 total) to examine the relationship between individual differences in perception and induction, as well as the utility of perception in predicting generalization gradients. The datasets were obtained from predictive learning tasks where participants learned associations between different colored cues and the presence or absence of a hypothetical outcome. In these datasets, stimulus perception and response generalization (expectancy ratings) were assessed in separate phases. Using cluster analyses, we identified similar subgroups of *good* and *bad perceivers* in all six datasets, with distinct patterns of response generalization between these subgroups. Based on the differences in stimulus perception, we could predict where across the stimulus range generalized responses would differ between subgroups as well as the direction of the difference. Furthermore, participants classified as *good perceivers* were more likely to report a similarity generalization rule than a relational or linear rule, providing evidence that individual differences in perception predict differences in induction. These findings suggest that greater consideration should be given to inter-individual variability in perception and induction and their relationship in explaining response generalization.

Keywords: generalization, perception, rules, conditioning, individual differences

1. Introduction

How we respond to novel stimuli is the outcome of a process that is luckily not based solely on randomness but relies on different sources of information to generate an *adaptive* response (Zaman, Chalkia, et al., 2020). This ability to generalize from previous learning is highly functional as it makes optimal use of acquired knowledge in order to avoid unnecessary (re)learning and costly mistakes. The field of associative learning has a long tradition and well-established procedures to measure response generalization in the laboratory (Guttman & Kalish, 1956; Honig & Urcuioli, 1981). First, a response is trained using conditioning procedures. In simple conditioning, a cue (conditional stimulus, CS+, e.g., a tone) is paired with an outcome (unconditional stimulus, US; e.g., a shock). After repeated pairings, the CS+ comes to elicit the conditioned response (e.g., fear) as a result of its association with the US. In a subsequent generalization phase, responses to a range of test stimuli are recorded, with the stimulus set most often comprising minor variations of a specific stimulus feature (e.g., the frequency of a tone or the hue of a key light). When aggregated across subjects, responses tend to form a bell-shaped gradient across the stimulus dimension that peaks at the location of the CS+. If instead of simple conditioning, differential conditioning is conducted with an additional similar cue (e.g., a slightly higher or lower tone) that predicts US absence (CS-), the gradient often becomes asymmetrical and sigmoidal in shape (Ghirlanda & Enquist, 2003; Mednick & Freedman, 1960), and the peak of the gradient can sometimes shift (Lee, Hayes, & Lovibond, 2018; Purtle, 1973; Zaman, Struyf, Ceulemans, Vervliet, & Beckers, 2020).

Despite abundant examples of response generalization within the laboratory and real-life (Enquist & Ghirlanda, 2005; Ghirlanda & Enquist, 2003; Mednick & Freedman, 1960), insight into the underlying mechanisms remains unclear for several reasons (Struyf, Zaman, Vervliet, & Van Diest, 2015; Zaman, Chalkia, et al., 2020). The first issue is that current evidence suggests that there are multiple mechanisms through which generalized response can occur, including a failure to discriminate or inductive extrapolation (Lovibond, Lee, & Hayes, 2020; Zaman, Chalkia, et al., 2020). In other words, response generalization may be an interaction between perceptual and cognitive processes. The identification and differentiation regarding these and other mechanisms based solely on patterns of generalized responding, however, remains difficult as different mechanisms can yield the same behavioral phenotype (Struyf, Zaman, Vervliet, & Van Diest, 2015). The second complexity arises from the fact that response generalization in humans displays a high amount of variability (Lee, Mills, Hayes, & Livesey, 2020; Zaman, Chalkia, et al., 2020), with large individual differences in inductive rules (Lee et al., 2018; Lovibond et al., 2020; Wong & Lovibond, 2017) and perceptual ability (Struyf, Zaman, Hermans, & Vervliet, 2017; Struyf et al., 2015; Zaman, Ceulemans, Hermans, & Beckers, 2019; Zaman, Struyf, Ceulemans, Beckers, & Vervliet, 2019; Zaman et al., 2020; Zaman, Chalkia, et al., 2020; Zaman,

Struyf, Ceulemans, Vervliet, & Beckers, 2020; Zenses, Lee, & Zaman, 2020). Accounting for these differences in both cognitive and perceptual factors is therefore critical to explain human response generalization.

A series of studies recently demonstrated that perceptual errors influence generalization patterns in various ways and determine generalized responding to a large extent (Struyf, Zaman, Hermans, & Vervliet, 2017; Struyf et al., 2015; Zaman, Ceulemans, Hermans, & Beckers, 2019; Zaman, Struyf, Ceulemans, Beckers, & Vervliet, 2019; Zaman et al., 2020; Zaman, Chalkia, et al., 2020; Zaman, Struyf, Ceulemans, Vervliet, & Beckers, 2020; Zenses, Lee, & Zaman, 2020). Responding to a test stimulus was attenuated on trials when it was falsely perceived as the CS+. Similarly, responding was lower on trials where the stimulus was perceived as different from the CS+ (even during CS+ trials). In all studies, we found many perceptual errors during the test phase for all stimuli with distinct patterns of generalized responding when accounting for these perceptual errors. Furthermore, we identified subgroups of *good* and *bad perceivers* using cluster analysis that were associated with divergent generalization gradients, with a much broader gradient in the latter compared to the former subgroup (Zaman, Ceulemans, et al., 2019; Zaman, Struyf, et al., 2019; Zaman et al., 2020). Thus, we were able to account to a large extent for interindividual differences in generalized responding through accounting for perception.

As already mentioned, in addition to perception, another source of interindividual differences in generalized responding is the inductive rule that participants derive and use to guide behavior. Lee, Hayes, & Lovibond (2018, see also Wong & Lovibond, 2017; Lovibond, Lee, & Hayes, 2020) reported that after differential training with colored stimuli located on the green-blue color spectrum, participants reported qualitatively different generalization strategies. Some participants reported a *similarity* rule (i.e., the more similar the stimulus is to the CS+, the higher the likelihood of the outcome), while others adopted a *linear* or relational rule based on the relational difference between the CS+ and CS- (i.e. the bluer/greener the stimulus, the higher the likelihood of the outcome). Participants in these rule subgroups showed different generalization patterns, with a more bell-shaped gradient in the similarity subgroup and a more sigmoidal gradient in the linear subgroup.

The factors that predispose participants to different generalization strategies in a differential learning procedure is currently unclear. A recent review proposed that perceptual accuracy might determine the type of rule since rules are extrapolated from experience (Zaman, Chalkia, et al., 2020). One way to describe the difference between linear and similarity rules is that the former only requires coarse stimulus information (about the contrasted dimension), while the latter requires detailed attention to

stimulus features. This leads to the prediction that more accurate stimulus perception should lead to a higher likelihood of deriving a similarity over a linear generalization rule.

Testing the above hypothesis requires the assessment of generalization, perception, and rules in the same participants. A recent study by Lovibond, Lee, and Hayes (2020) provides just this. Lovibond et al. investigated the extent to which gradients of stimulus perception and US expectancy ratings were influenced by the type of rule that participants adopted (linear vs. similar) and the type of conditioning procedure administered in training (simple vs. differential). In an online predictive learning task, participants first learned about the relationship between the occurrence of a hypothetical shock and colored rectangles. Next, stimulus perception (categorization) and response generalization (outcome expectancy) were assessed in separate blocks (counterbalanced for order). They demonstrated distinct US expectancy gradients depending on the type of adopted rule and type of training procedure, but both factors did not produce differences in stimulus perception. However, while the study examined individual differences in induction, it did not consider individual differences in perception. Therefore, whether or not perceptual accuracy was related to inductive rules is currently unknown. Furthermore, in their study, perception and generalized responding were assessed in separate blocks, unlike other studies that have measured both concurrently on the same trial (e.g., Zaman, Ceulemans, et al., 2019; Zaman, Grant, et al., 2020; Zaman, Struyf, et al., 2019). This difference in testing procedure provides a unique opportunity to validate further the close relationship reported between perceptual accuracy and generalization patterns in previous studies with congruent assessment, as any observed effects this time can not be attributed to the close temporal relationship between both measures.

The overall aim of this study was to test whether differences in stimulus perception predict patterns of response generalization, and also, the type of reported rule. We achieved this aim by performing a series of cluster analyses to identify individual differences in stimulus perception using six datasets, four of which were reported in Lovibond et al. (2020), and two of which are unpublished.

2. Design and procedure

In total, six different datasets are reported, four of which have been previously published in Lovibond et al. (2020), and were retrieved from OSF: <https://osf.io/xnva4/>. The full details of the protocol used to generate those datasets (1, 2, 5, and 6) can be found in Lovibond et al. (2020). We only report the main procedural details needed to interpret our analyses below. Table 1 provides an overview of the different datasets and their main methodological differences.

2.1. Participants

We adopted the same exclusion criteria as in the original paper, excluding participants that were colorblind or that failed to meet the learning criteria (see Table 1). All experiments were approved by the University of New South Wales Human Research Ethics Advisory Panel C (HREAP 2615). All participants provided their written informed consent.

Table 1. Overview of the different datasets

Dataset	Exp. in Lovibond et al. (2020)	Conditioning	Nr. of test stimuli	Perceptual Task	Initial N	Color blind	Non-learners*	Final sample
1	1	Simple	11	2AFC (CS+)	179	6	7	166
2		Differential	11	2AFC (CS+)	221	10	30	181
3	-	Simple	11	VAS (CS+)	188	7	12	169
4		Differential	11	VAS (CS+)	211	6	35	170
5	2	Differential	17	2AFC (CS+)	212	15	7	190
6		Differential	17	2AFC (CS-)	185	10	59	116

**To meet the learning criterion, participants had to predict the shock outcome correctly on at least 2 of the last 4 CS+ trials during acquisition. In addition, participants in the differential conditioning group had to predict no shock correctly on at least 3 of the last 4 CS- trials. 2AFC = two-alternative forced choice task where they had to categorize the stimulus as the CS+ (CS-) or as different, VAS = visual analog scale.*

2.2. Stimuli

The stimuli consisted of colored rectangles measuring 236 x 118 pixels within a 472x472 pixel white canvas. In total, 17 different stimuli were created (see Fig. 1A) across the green-blue spectrum. The color changes were equally spaced (in steps of 0.02 between stimuli with an integer number) between hue values of 0.4 and 0.6, keeping saturation and brightness constant. In experiments 1-4, only the stimuli with an integer number were used (in total 11), whereas, in the remaining experiments, all 17 stimuli were used. Color allocation to the stimulus dimension was counterbalanced (with S1 having either a hue value of 0.4 or 0.6). S6 always served as the CS+. In the differential group, S4 was the CS-. Thus depending on color counterbalancing, CS- was either slightly greener or bluer than the CS+ (Fig. 1A).

2.3. Protocol

The experiments listed in Table 1 were run online using Amazon's Mechanical Turk platform. They all consisted of an acquisition phase and a test phase that included a separate generalization test and perceptual assessment, with their order counterbalanced across participants. The entire protocol

lasted 12-min for which participants received US\$ 2. The experiments were identical except for the minor procedural details noted below.

2.3.1. Acquisition

The acquisition phase consisted of a predictive learning task. Participants learned to predict the occurrence of a hypothetical shock (unconditioned stimulus, US) based on visual cues (conditioned stimuli, CS). On each trial, a CS was presented along with a prompt asking participants to indicate whether they expected the US by pressing one of two keys ("L" for shock or "A" for no shock). Participants had unlimited time to respond during which the CS remained on screen. Upon responding, feedback about the correct outcome (either "SHOCK!" or "no shock occurred") was displayed along with the CS for 5s, after which a blank 2s inter-trial interval (ITI) began. During simple conditioning, the CS+ was presented on each of the 12 trials, with the CS+ reinforced at 75%. In addition to these 12 CS+ trials, in the differential conditioning group, 12 CS- trials were included that were never reinforced. Thus, this phase consisted of either 12 trials or 24 trials for the simple and differential conditioning groups, respectively.

2.3.2. Test

Participants were randomly allocated to start either with the perceptual assessment or the generalization test, after which they completed the remaining task.

2.3.2.1. Generalization test

In the generalization test, participants were instructed that they should continue to make predictions, but that they would no longer receive feedback. On each trial, participants were presented with the test stimulus along with a text prompt ("What is the likelihood of this symbol leading to a shock?"), and a visual analog scale (VAS) underneath. Instead of making a keypress choice as in training, participants indicated their US expectancy on the VAS using the mouse. The scale was labeled "Certain NO SHOCK (0% of shock)" on the left and "Certain SHOCK (100% of shock)" on the right. All 11 (Datasets 1-4) or 17 (Datasets 5-6) stimuli, including the CSs, were presented once in randomized order. After participants confirmed each VAS rating by clicking the "Continue" button, a 2s blank ITI started before the next trial. All expectancy ratings ranged between 0-100.

2.3.2.2. Perceptual assessment

Here, participants were asked to categorize each test stimulus as either being the same or different to the exact CS+ (datasets 1-5) or CS- presented in training (dataset 6). In datasets 1, 2, 5, and 6, this was done through a two-alternative forced-choice (2AFC) task where participants were asked to press the

"S" key for same or the "D" key for different. The only difference between datasets 5 and 6 was whether participants were asked to identify the test stimuli as the same as the CS+ (dataset 5) or CS- (dataset 6). In datasets 3 and 4, participants were asked to respond using a VAS scale ranging from "Definitely different" to "Definitely the same" (see Fig. 1). Similar to the generalization test, all 11 (datasets 1-4) or 17 (dataset 5-6) stimuli, including the CSs, were presented once in randomized order. Once participants responded either to the 2AFC task via a button press or clicked on a button to confirm their VAS rating, the blank 2s ITI started after which the next trial began. Like the generalization test, participants were given unlimited time to respond, and there was no feedback. The order of stimulus presentation was randomized independently from the order in the generalization test. Identification responses were either coded as 0 (different) or 1 (same) for the 2AFC task (datasets 1, 2, 5, and 6) or ranged between 0-100 if the VAS was used (datasets 3 and 4).

Fig. 1. Overview of the protocol. (A) In experiments 1-4, only the stimuli with an integer number were used (in total 11), whereas, in the remaining experiments, all stimuli (in total 17) were used. **(B)** The protocol and trial structure. Note that task order during the test phase and color allocation to the stimuli (the CS- was either slightly bluer or slightly greener than the CS+) was counterbalanced. CS = conditioned stimuli, S = stimulus, diff = different

2.3.3. Rule Assessment

Rule assessment was conducted at the end of the experiment for datasets 3 and 4. In the other datasets, rule assessment was conducted immediately after the generalization task (and potentially

prior to the perceptual task, depending on the allocated test order). Participants were asked a 3AFC question *"During Phase ONE (where there was feedback), which of the following options did you think was most true concerning the likelihood of a shock occurring?"*. They selected from a similarity option (*"The more similar the symbol to an aqua (green-blue) color, the higher the likelihood of shock"*), and two linear options (*"The greener/bluer the symbol, the higher the likelihood of shock"*). In the simple conditioning group, both linear options were equally consistent with the training contingencies since participants only ever saw one CS+ and therefore selection of either option was classified as simply "linear". However, for those who received differential conditioning, one of the linear options was consistent (classified as "linear") with their experienced training contingencies while the other was inconsistent (classified as "linear incon").

Fig. 2. Acquisition data for all six datasets. Error bars indicate standard error of the mean.

2.4. Analyses

All datasets and the full reproducible code can be found at the Open Science Framework (OSF) repository (<https://osf.io/g8bjy/>).

Figure 2 shows the acquisition data for all datasets (see Lovibond et al., 2020 for analyses of the training data for datasets 1,2 and 5,6). The data of the generalization phases consisted of perceptual judgements (binary categorizations or continuous ratings) and US expectancy ratings, assessed in separate blocks. We investigated the relationship between perception and US expectancy gradients using a similar technique to Zaman et al. (2019). This technique involves using cluster analysis to identify distinct subgroups based on the performance during the perceptual task, and then examining how the subgroups differ in their generalization gradients.

All analyses were performed in R 4.0.1 (2020-06-06) (R Core Team, 2013). To perform the cluster analysis, we used the k-means algorithm (from the stats package v3.6.2 (R Core Team, 2013)) which repeatedly divides the whole group of participants based on their perceptual data into an incrementing number of clusters until the maximum number of clusters is reached. We set the maximum number of clusters to 10. Each participant is allocated to the cluster for which its squared Euclidean distance to the cluster centroid is minimal. The centroids are computed as the mean score profile per cluster. We ran the analysis 10000 times, each time using a different random initialization of the centroid matrix (Hofmans, Ceulemans, Steinley, & Van Mechelen, 2015), to prevent ending in a local optimum. The optimal number of clusters was based on the largest relative decrease in the sum of squared Euclidean distances and the automated scree test (Ceulemans & Kiers, 2006; Wilderjans, Ceulemans, & Meers, 2013). In all datasets, two subgroups emerged apart from dataset 5 where a three-cluster solution was preferred.

Next, we compared US expectancy gradients between the identified subgroups using mixed models (using the lme4 package v1.1.23, (Bates, Mächler, Bolker, & Walker, 2015)), which provide powerful and flexible approaches to analyze repeated measures data (Blackwell, de Leon, & Miller, 2006). The following model was tested in all datasets, with cluster (2 or 3 levels), stimulus (11 or 17 levels), and their interaction as categorical predictors¹. US expectancy was the dependent variable. In all models, the random effect consisted of a subject-dependent intercept to account for repeated measures. Omnibus tests of main effects were obtained using anova (from the car package v3.0.8. (Fox &

¹ The use of stimulus as categorical rather than scale predictor was preferred based on the Bayesian Information Criterion (BIC) with lower values indicating better model fit

Weisberg, 2019)). Post hoc contrasts using the models' estimated means were conducted with the emmeans package (v.1.5.0. (Lenth, 2020)).

Finally, we tested whether the probability of reporting each rule was dependent upon perceptual accuracy. Depending on the number of levels of the dependent variable rule (two or three), we used either bi- or multinomial logistic regression using multinom (from the nnet package v.7.3.14, (Venables & Ripley, 2002)). Logistic regression expresses the logarithm of the odds (or logit-transformed probability) of a categorical dependent variable as a linear relationship with the predictor variables. In these models, the estimated beta for the intercept (β_i) is the log odds that express the probability that a certain outcome occurred relative to the probability of another outcome (that serves as the reference category). In all models, the linear rule served as the reference category. Beta estimates for the predictor cluster are log odds-ratios that convey to what extent the odds change if one belongs to cluster 2 compared to cluster 1. In other words, the estimated parameter for cluster (β_c) indicates whether the odds for *bad perceivers* are higher ($\beta > 0$), similar ($\beta = 0$) or lower ($\beta < 0$) than the odds for the *good perceivers* (or vice versa depending on which group is used as baseline).

Foreshadowing our results, since many of the effects are found consistently in all datasets, we grouped the datasets according to conditioning protocol and report the analyses in this order rather than by dataset. First, we report those datasets that included a simple conditioning protocol. Next, we report the results from the differential conditioning datasets. The dataset comprising the CS- categorization responses is reported separately at the end since in all other datasets the perceptual assessment was about the perception of the CS+.

Fig. 3. Experiments using a simple conditioning protocol, dataset 1 (top) and 3 (bottom): (A) Probability of being categorized as CS+ across the stimulus range per cluster. Black and red markers indicate correct and incorrect responses, respectively. (B & E) Averaged US expectancy ratings per stimulus per cluster (thick lines) and individual gradients (thin lines). (C & F) The number of participants reporting each generalization rule per cluster. (D) Averaged similarity ratings per cluster (thick lines) and individuals (thin lines). The dashed line indicates which stimulus was used as CS+. Error bars represent standard errors of the mean. US = unconditioned stimulus, S = stimulus, sim. rat. = similarity rating to the CS+.

3. Results

3.1. Simple conditioning - datasets 1 and 3

3.1.1. Perceptual assessment

Despite differences between datasets in whether binary or continuous measures were used, it is clear that the perceptual gradients for both good and bad clusters are very similar (compare Figure 3A and 3D). In dataset 1, 47 participants (28.31%) were allocated to a cluster representing poor performance on the perceptual categorization task (i.e., *bad perceivers*). Their response probability varied around the chance level of 0.5 for the entire range of test stimuli. They made many mistakes for both the CS+ and all the test stimuli (Fig. 3A). For the remaining 119 participants (71.69%), the allocated cluster indicated good categorization accuracy (i.e., *good perceivers*). Here, response probabilities peaked at

the location of the CS+ (i.e., correct responses) and immediately started to decline (symmetrical) as test stimuli differed from the CS+ (Fig. 3A). Response probabilities were around 0.5 for the adjacent test stimuli (S5 and S7), and much lower (i.e., few mistakes) for more distant test stimuli (S1:S4 and S8:S11).

In dataset 3, almost identical subgroups were identified based on similarity ratings instead of categorization responses. Now, 54 participants (31.95%) were allocated to the *bad perceivers* cluster with overall lower, flatter, and slightly asymmetrical similarity ratings across the stimulus range (Fig. 3D). The remaining 115 participants (69.05%, *good perceivers*) that were clustered together had average similarity ratings that peaked at the CS+ and steeply decreased as test stimuli became less similar (Fig. 3D).

3.1.2. US expectancy

As visible in Figures 3B and 3E, there were large differences in US expectancy across the stimulus spectrum in both datasets [dataset 1: stimulus effect: $F(10,1640) = 22.292, p < .001$, dataset 3: stimulus effect: $F(10,1670) = 37.256, p < .001$] with distinct gradients between subgroups of participants that were identified as *good* compared to *bad perceivers*. *Bad perceivers* had a flatter US expectancy gradient that was overall located higher compared to the gradient of the *good perceivers*, which was a symmetrical gradient with a narrow peak at the CS+ location [dataset 1: cluster effect: $F(1,164) = 50.559, p < .001$; cluster $\times$ stimulus effect: $F(10,1640) = 9.154, p < .001$; dataset 3: cluster effect: $F(1,167) = 32.077, p < .001$; cluster $\times$ stimulus effect: $F(10,1670) = 13.984, p < .001$].

Based on the points (S5 and S6) where categorization probabilities crossed between both clusters, we created four custom contrasts to test for difference in US expectancy rating between clusters. We found in dataset 1 that (1) responding to the CS+ was higher in the *good* compared to *bad perceivers* group [S6: $t(1712) = 2.157, p = .031$]; (2) did not differ between groups for the test stimuli most proximal to the CS+ [(S5+S7)/2: $t(1205) = 1.286, p = .199$]; and (3) was lower for more distant test stimuli on either side in the *good* compared to *bad perceivers* group [(S1+...+S4)/4: $t(565) = 6.965, p < .001$; (S8+...+S11)/4: $t(565) = 8.920, p < .001$]. In dataset 3, we found exactly the same pattern of differences between subgroups in US expectancy ratings [S6: $t(1454) = 3.279, p = .001$; (S5+S7)/2: $t(822) = 1.811, p = .071$; (S1+...+S4)/4: $t(394) = -5.235, p < .001$; (S8+...+S11)/4: $t(394) = -9.716, p < .001$].

3.1.3. Rules

Cluster allocation significantly predicted which response rule people selected for their US expectancy ratings at the end of the task (dataset 1, Fig. 3C). *Bad perceivers* were not more likely to select either

of the rules [$\beta_i = -0.48$, $SE = 0.30$, $z = -1.59$, $p = .112$] while *good perceivers* compared to the *bad perceivers* were more likely to select the similarity rule than the linear rule [$\beta_c = 1.85$, $SE = 0.38$, $z = 4.91$, $p < .001$]. We replicated these findings in dataset 3 [$\beta_i = -0.07$, $SE = 0.27$, $z = -0.27$, $p = .786$; $\beta_i = 1.57$, $SE = 0.36$, $z = 4.322$, $p < .001$].

Fig. 4. Experiments with a differential conditioning protocol, datasets 2 (top), 4 (middle) and 5 (bottom): (A & G) Probability of being categorized as CS+ across the stimulus range per cluster. Black and red markers indicate correct and incorrect responses, respectively. (B, E, H) Averaged US expectancy ratings per stimulus per cluster (thick lines) and individual gradients (thin lines). (C, F, I) Rule frequency per cluster. (D) Averaged similarity ratings per cluster (thick lines) and individuals (thin lines). The dotted and dashed lines indicate which stimulus was used as CS+ and as CS-. Error bars represent

standard errors of the mean. US = unconditioned stimulus, S = stimulus, sim. rat. = similarity rating to the CS+.

3.2. Differential conditioning - datasets 2, 4 and 5

3.2.1. Perceptual categorization

As visible in Fig. 4, a two cluster solution was preferred in datasets 2 and 5 (Fig. 4A and 4G)². The pattern for the *good perceivers* was very similar in both datasets. Response probabilities peaked at the CS+ (i.e., correct responses) and started to decline immediately as the physical proximity of the test stimulus to the CS+ decreased. 70.71% (N = 128) of the sample in dataset 2 and 64.21% (N = 122) of the participants in dataset 5 were allocated to this cluster. The remaining 29.29% (N = 53, dataset 2) and 35.79% (N = 68, dataset 5) were allocated to the other cluster. The gradient in this cluster suggests poorer performance although it seems to differ slightly from the poor perceivers' gradients observed in the datasets that used a simple conditioning protocol (see Figure 3A and 3D). Response probabilities reached a plateau from the CS- till the CS+ after which response probabilities started to decrease albeit to a lesser extent compared to the *good perceivers*.

In dataset 4 where we clustered participants based on similarity ratings instead of the categorization responses, a 3-cluster solution was preferred (Fig. 4D). In a first group of *bad perceivers* (N = 47, 27.65%), similarity ratings reached a plateau from CS- till CS+, and decreased as test stimuli approached the end of the spectrum. In the second group of *bad perceivers* (N = 43, 25.29%), similarity rating did not peak at the CS+ but at S7 and remain relatively high for the test stimuli on this side of the CS+ compared to test stimuli on the CS- side of the CS+. The final cluster was the *good perceivers* (N = 80, 47.06%) with similarity ratings peaking at the CS+ (i.e., correct responses) and declining immediately after that in both directions away of the CS+.

3.2.2. US expectancy

We found differences in US expectancy ratings across that the stimulus range with distinct gradients between clusters in all three datasets (see table 2 for omnibus effect and figures 4B, E and H). Like we did for the simple conditioning datasets, we created four custom contrasts to test for points (or areas) in the stimulus space where US expectancy ratings were expected to differ (or not) based on the ratios of the categorization accuracies (or similarity ratings) between the subgroups.

In dataset 2, paralleling the groups' differences in perceptual categorizations, good compared to *bad perceivers*, had higher US expectancy at the CS+ [S6: $t(1814) = 2.802$, $p = .005$], no difference in US

² Both the two and three cluster solution yielded very similar scree-ratios in dataset 5. For parsimony, we used the two-cluster solution for the analyses.

expectancy ratings for the adjacent test stimuli $[(S5+S7)/2: t(1204) = 1.127, p = .260]$, and lower ratings for most of the remaining stimuli on either side of the CS+ $[(S1+...+S4)/4: t(559) = -6.154, p < .001; (S8+...+S11)/4: t(559) = -3.177, p = .002]$.

In dataset 4, the first group of *bad perceivers* compared to the *good perceivers* had significantly higher US expectancy rating from S1-S4 $[(S1+...+S4)/4: t(492) = 7.05, p < .001]$, while they did not differ at the CS+ $[S6: t(1656) = 0.10, p = .92]$, or the stimuli located further on the right side of the CS+ $[(S8+...+S11)/4: t(492) = 1.19, p = .234]$. For the second group of *bad perceivers*, compared to the *good perceivers*, US expectancy ratings were lower at the CS+ $[S6: t(1656) = -3.007, p = .0027]$ and higher at both outsides of the stimulus spectrum $[(S1+...+S4)/4: t(492) = 5.08, p < .001; [(S8+...+S11)/4: t(492) = 6.36, p < .001]$. Comparison between both groups of *bad perceivers*, revealed higher US expectancy ratings at the CS+ $[S6: t(1656) = -2.604, p = .0093]$ and lower US expectancy ratings for the stimuli located on the most right side of the CS+ $[(S8+...+S11)/4: t(492) = 4.663, p < .001]$ for the first compared to the second group. US expectancy ratings were not different across S1 to S4 $[(S1+...+S4)/4: t(492) = 1.59, p = .113]$.

In dataset 5, analogous to the findings in dataset 2, US expectancy ratings were higher for the *good* compared to *bad perceivers* at the CS+ $[S6: t(\text{inf}) = 4.079, p < .001]$, did not differ between group for S5 and S7 $[(S5+S7)/2: t(\text{inf}) = -0.281, p = .778]$, and were lower in the *good perceivers* group for more distant test stimuli in either direction of the CS+ $[(S1+...+S4.5)/6: t(\text{inf}) = -10.630, p < .001; (S7.5+...+S11)/6: t(\text{inf}) = -2.641, p = .008]$ (Fig. 4H). In summary, for stimuli where one of the groups had a higher probability of (mis)identifying the stimulus as the CS+ (or similarity rating), higher generalized responding was observed.

Table 2. F-tests of omnibus effects

	Dataset 2			Dataset 4			Dataset 5		
	df	F	p	df	F	p	df	F	p
Stimulus	10, 1790	81.41	<.001***	10, 1670	85.70	<.001***	16, 3008	61.46	<.001***
Cluster	1, 179	18.82	<.001***	2, 167	14.49	<.001***	1, 188	21.77	<.001***
Cluster × Stimulus	10, 1790	4.43	<.001***	10, 1670	5.74	<.001***	16, 308	12.85	<.001***

df = degrees of freedom, *** = $p \leq .001$

3.2.3. Rules

In dataset 2, *bad perceivers* had similar probabilities of reporting the inconsistent linear rule compared to the linear rule $[\beta_i = -0.51, SE = 0.42, z = -1.21, p = .226]$, while *good perceivers* were significantly less likely to report the inconsistent linear rule $[\beta_c = -2.46, SE = 0.84, z = -2.93, p = .003]$. The bad perceivers

were more likely to report the similarity rule than the linear rule [$\beta_i = 0.66$, $SE = 0.32$, $z = 2.07$, $p = .038$]. This was also the case for the *good perceivers* group as their odds did not differ from the odds in the *bad perceivers* group [$\beta_c = 0.14$, $SE = 0.37$, $z = 0.38$, $p = .70$] (Fig. 4C).

In dataset 4, *good perceivers* were significantly less likely to report the inconsistent linear rule compared to the linear rule [$\beta_i = -2.88$, $SE = 1.03$, $z = -2.813$, $p = .005$] while both groups of *bad perceivers* had similar odds [$\beta_{c_2} = 1.50$, $SE = 1.14$, $z = 1.32$, $p = .189$; $\beta_{c_3} = 1.84$, $SE = 1.13$, $z = 1.63$, $p = .102$]. Furthermore, *good perceivers* were more likely to report the similarity rule than the linear rule [$\beta_i = 1.22$, $SE = 0.27$, $z = 4.55$, $p < .001$], while both *bad perceivers* groups had lower odds of reporting the similarity rule [$\beta_{c_2} = -1.33$, $SE = 0.42$, $z = -3.15$, $p = .002$; $\beta_{c_3} = -0.88$, $SE = 0.42$, $z = -2.11$, $p = .035$].

Final, a similar set of results was found in dataset 5. *Bad perceivers* had similar probabilities of reporting the inconsistent linear rule compared to the linear rule [$\beta_i = -0.18$, $SE = 0.30$, $z = -0.60$, $p = .55$], while *good perceivers* were significantly less likely to report the inconsistent linear rule [$\beta_c = -1.54$, $SE = 0.57$, $z = -2.69$, $p = .007$] (Fig. 4I). Furthermore, *bad perceivers* were equally likely to report the similarity rule compared to the linear rule [$\beta_i = 0.00$, $SE = 0.29$, $z = 0.01$, $p = .99$], while *good perceivers* were much more likely to report the similarity rule of the linear rule compared to *bad perceivers* [$\beta_c = 1.16$, $SE = 0.367$, $z = 3.21$, $p = .001$].

3.3. Differential conditioning - dataset 6

3.3.1. Perceptual categorization (CS- categorization)

Using cluster analysis, we identified two subgroups of participants with homogenous performance on the CS- categorization task (Fig. 5A). 68.97% ($N = 80$) were classified into the *good perceivers'* group, as their accuracy peaked at the CS- location and decreased across the spectrum. For the remaining 31.03% of the participants ($N = 36$), called the *bad perceivers*, categorization responses peaked around the CS+ instead of the CS and were much lower for the actual CS-.

3.3.2. US expectancy

US expectancy ratings dependent upon stimulus [stimulus effect: $F(16,1824) = 48.09$, $p < .001$] and cluster [cluster effect: $F(1,114) = 4.8134$, $p = .030$; cluster $\times$ stimulus effect: $F(16,1824) = 6.76$, $p < .001$]. Based on the point (S4.5) where categorization probabilities crossed between both clusters, we created three custom contrasts to test for differences in US expectancy rating between clusters. As visible in Figure 5B, US expectancy ratings were higher in the *bad* compared to the *good perceivers* for stimuli left of S4.5 [(S1 + ... + S4)/5: $t(395) = 4.766$, $p < .001$], significantly lower for stimuli right of S4.5 [(S5 + ... + S11)/12: $t(153) = -5.439$, $p < .001$] and did not differ at S4.5 [S4.5: $t(1725) = 1.215$, $p = 0.23$].

3.3.3. Rules

Consistent with the other datasets, we found that *good perceivers* were more likely to report the similarity rule [$\beta_i = 0.55$, $SE = 0.24$, $z = 2.362$, $p = .018$], and less likely to select the inconsistent linear rule [$\beta_i = -2.23$, $SE = 0.61$, $z = -3.67$, $p < .001$], compared to the linear rule. The odds for the *bad perceivers* relative to the odds for the *good perceivers* did not differ for the similarity rule [$\beta_c = -0.01$, $SE = 0.44$, $z = -0.03$, $p = .98$] while they were significantly higher for the inconsistent linear rule [$\beta_c = 1.63$, $SE = 0.79$, $z = 2.06$, $p = .040$].

Fig. 5. Experiment with CS- categorization as perceptual task. Dataset 6: (A) Probability of being categorized as CS- across the stimulus range per cluster. Black and red markers indicate correct incorrect responses, respectively. **(B)** Averaged US expectancy ratings per stimulus per cluster (thick lines) and individual gradients (thin lines). **(C)** Rule frequency per cluster. Error bars represent standard errors of the mean. US = unconditioned stimulus, S = stimulus, sim. rat. = similarity rating to the CS+.

4. Discussion

In this study, we tested how individual differences in perceptual ability related to differences in inductive rules and patterns of response generalization. We performed cluster analysis on a large sample and found highly consistent results amongst the six datasets we examined. In all datasets, we identified distinct subgroups concerning perceptual accuracy across the stimulus spectrum representing participants who were good vs. poor at perceptually discriminating the CS+ (or CS-) from the other stimuli on the dimension. More importantly, we found a close relationship between perceptual accuracy and US expectancy gradients. Differences between the *good* and *poor perceivers*

in US expectancy ratings closely paralleled the subgroups' differences in perceptual accuracy across the stimulus range. Furthermore, more accurate stimulus perception was associated with a higher likelihood to adopt the similarity rule compared to a linear rule.

We observed distinct patterns of generalized responding between *good* and *bad perceivers* in all data sets, that closely paralleled their perceptual accuracies across the stimulus range. More precisely, based on the ratio of the categorization accuracies between the groups on the stimulus spectrum, points where US expectancy ratings should differ could be predicted as well as the direction of the difference. For stimuli where one of the groups had a higher probability of (mis)identifying the stimulus as the CS+, higher generalized responding was observed. This led to higher US expectancy ratings for the *bad perceivers* at many of the test stimuli, while this led to higher US expectancy ratings at the CS+ for the *good perceivers*. Similarly, for stimuli where the groups had similar categorization accuracies, they also did not differ on US expectancy. We found the same close relationship when perceptual accuracy was about identifying the CS-, with this time lower responding between groups when stimuli were more likely to be categorized as CS-. Furthermore, this relationship was found when perceptual accuracy was assessed via a 2AFC task as well as via continuous ratings, suggesting that this finding is robust to procedural differences between datasets.

These findings, together with our previous work (Struyf et al., 2017; Zaman, Ceulemans, et al., 2019; Zaman, Struyf, et al., 2019; Zaman et al., 2020), provide compelling evidence that differences in stimulus perception is an important source for variation in generalized responding. It is noteworthy that we found the relationship in all six datasets, despite having only a rough estimate of participants' perceptual accuracy (only one trial per stimulus). Whether the relationship between perception and conditioned responding stems because the former is a determinant for the latter, or alternatively, because both are measures of the same latent process is currently unclear. We argue, however, that evidence is accumulating in favor of the former, with perception as a determinant of generalized responding. For instance, the observation that the type of learning (simple vs. differential) influenced patterns of generalized responses and but not perception (Lovibond et al., 2020), suggests that both are not driven by the same process. Furthermore, a close relationship between perception and generalized responding has now been found in studies that assessed perception and generalized responding on the same trial (Struyf et al., 2017; Zaman, Ceulemans, et al., 2019; Zaman et al., 2020) or as in the current study in separate blocks (Holt et al., 2014; Lovibond et al., 2020; Tuominen et al., 2019), suggesting that their relationship also does not merely emerges due to the temporal relationship in their assessment.

Cluster analyses is an explorative, data-driven approach to search for patterns, with the scientific relevance of the identified subgroups depending on many aspects, including their cross-validation. We found not only similar subgroups but also similar proportion of people allocated to each cluster in all datasets, suggesting these are robust findings that reflect stable interindividual differences. For instance, the prevalence of the *good perceivers* was around 70% in all studies, apart from the dataset with the 3-cluster solution. Compared to a previous study that identified different groups of *bad perceivers* (Zaman, Ceulemans, et al., 2019), here we found only one for the majority of the datasets. It is possible that procedural differences between previous studies and the datasets here are responsible for this discrepancy. For example, the fact that each stimulus was only presented once may have prevented a more detailed differentiation between subgroups of bad perceivers, or it could be due to differences in the stimulus sets or the type of perceptual task used.

Although cluster analysis enabled us to identify these consistent subgroups of participants, it does not explain what caused these groups to differ. As the data were collected online, it is possible that groups merely differed due to differences in the quality of the participants' display, affecting the ability to discriminate between the different hues. It is worth noting, however, that other studies have identified similar (and other) subgroups in controlled experimental settings (Zaman, Ceulemans, et al., 2019; Zaman, Grant, et al., 2020; Zaman, Struyf, et al., 2019). Alternatively, subgroups may have emerged due to individual variation in cognitive factors such as attention, memory, or perceptual decision-making. For instance, the adopted perceptual task has a strong memory component as presented stimuli are compared to a retrieved memory of the CS. Therefore, poor performance on the current task may actually reflect poor CS memory. In a recent study, we demonstrated that increased perceptual memory uncertainty about the CS was associated with broader generalization gradients, much like those we found for the *bad perceivers* (Zenses et al., 2020). Other studies have demonstrated that associative learning is capable of altering the ability to discriminate between stimuli (Laufer & Paz, 2012; Resnik, Sobel, & Paz, 2011; Schechtman, Laufer, & Paz, 2010) with inter-individual differences in the extent of learning-induced changes (Laufer, Israeli, & Paz, 2016). These changes in perceptual discrimination have been linked to altered neural processing in sensory regions (Aizenberg & Geffen, 2013; Li, Howard, Parrish, & Gottfried, 2008). Future studies could further investigate the role of memory and perceptual processes in a context of response generalization.

Finally, we found a relationship between perceptual accuracy and the type of rule people reported. That is, in the majority of datasets the similarity rule was not only reported by those with a more accurate perception, but it was more likely than the linear rule in those participants. This study is novel in reporting a relationship between perception and induction using cluster analysis, and suggests one avenue by which individual differences in inductive rules reported in recent work with these stimuli

(Lee et al., 2018; Lovibond et al., 2020) could arise. It may be that participants who are better able (or happen) to attend to the features of the CS+ during training go on to generalize using these features and therefore report a similarity rule. An intriguing possibility is that participants who derive a linear rule during training do not encode the physical stimulus features to the same extent as other participants, since they become focused on the relational features of the stimuli (i.e., which one is greener or bluer) during training. This would imply that attention to these two different stimulus features compete, and that participants selectively attend to the feature that they believe helps them solve the task most effectively. Future work could explicitly manipulate perception prior to training in order to establish the direction of causation between these factors.

In sum, our study highlights the importance of perception in understanding how generalization behavior differs between individuals, as we showed that better perceptual accuracy was associated with a greater likelihood of adopting a similarity rule over a relational one, as well as sharper generalization gradients. These relationships were robust; we found highly consistent results amongst the six datasets comprising in total a large sample size. These findings suggest that greater consideration should be given to inter-individual variability in perception and induction (and their relationship) in explaining response generalization.

5. Acknowledgments

JZ is a Postdoctoral Research Fellow of the Research Foundation Flanders (FWO, 12P8619N) and received funding from the Efic-Grunthal grant (EGG ID 358254826) and a Special Research Funds (FWO, 1500620N). The research was sponsored by the "Asthenes" long-term structural funding (METH/15/011) – Methusalem grant by the Flemish Government. JL received salary support from the Australian Research Council (DP190103738)

6. Contributions

JL: Conceptualization, Methodology, Software. **JZ:** Formal analysis. **JZ:** Visualization. **JL and JZ:** Investigation. **JL and JZ:** Supervision. **JZ:** Writing - Original Draft. **JL:** Writing- Reviewing and Editing, **JL:** project administration

7. References

- Aizenberg, M., & Geffen, M. N. (2013). Bidirectional effects of aversive learning on perceptual acuity are mediated by the sensory cortex. *Nature Neuroscience*, 16(8), 994–996.
<https://doi.org/10.1038/nn.3443>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1). <https://doi.org/10.18637/jss.v067.i01>
- Blackwell, E., de Leon, C. F. M., & Miller, G. E. (2006). Applying mixed regression models to the analysis of repeated-measures data in psychosomatic medicine. *Psychosomatic Medicine*, 68(6), 870–878. <https://doi.org/10.1097/01.psy.0000239144.91689.ca>
- Ceulemans, E., & Kiers, H. A. L. (2006). Selecting among three-mode principal component models of different types and complexities: A numerical convex hull based method. *British Journal of Mathematical and Statistical Psychology*, 59(1), 133–150.
<https://doi.org/10.1348/000711005X64817>
- Enquist, M., & Ghirlanda, S. (2005). *Neural Networks and Animal Behavior* (Vol. 33). Princeton: Princeton University Press. <https://doi.org/10.1515/9781400850785>
- Fox, J., & Weisberg, S. (2019). *An R Companion to Applied Regression* (3rd ed.). Thousand Oaks, California: Sage. Retrieved from <https://socialsciences.mcmaster.ca/jfox/Books/Companion/>
- Ghirlanda, S., & Enquist, M. (2003). A century of generalization. *Animal Behaviour*, 66(1), 15–36.
<https://doi.org/10.1006/anbe.2003.2174>
- Guttman, N., & Kalish, H. I. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology*, 51(1), 79–88. <https://doi.org/10.1037/h0046219>
- Hofmans, J., Ceulemans, E., Steinley, D., & Van Mechelen, I. (2015). On the Added Value of Bootstrap Analysis for K-Means Clustering. *Journal of Classification*, 32(2), 268–284.
<https://doi.org/10.1007/s00357-015-9178-y>
- Holt, D. J., Boeke, E. A., Wolthusen, R. P. F., Nasr, S., Milad, M. R., & Tootell, R. B. H. H. (2014). A parametric study of fear generalization to faces and Non-Face objects: Relationship to discrimination thresholds. *Frontiers in Human Neuroscience*, 8(SEP), 1–12.
<https://doi.org/10.3389/fnhum.2014.00624>
- 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(3),

- 405–445. <https://doi.org/10.1901/jeab.1981.36-405>
- Laufer, O., Israeli, D., & Paz, R. (2016). Behavioral and neural mechanisms of overgeneralization in anxiety. *Current Biology*, 26(6), 713–722. <https://doi.org/10.1016/j.cub.2016.01.023>
- Laufer, O., & Paz, R. (2012). Monetary loss alters perceptual thresholds and compromises future decisions via amygdala and prefrontal networks. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 32(18), 6304–6311. <https://doi.org/10.1523/JNEUROSCI.6281-11.2012>
- Lee, J. C., Hayes, B. K., & Lovibond, P. F. (2018). Peak Shift and Rules in Human Generalization. *Journal of Experimental Psychology: Learning Memory and Cognition*, (January). <https://doi.org/10.1037/xlm0000558>
- Lee, J. C., Mills, L., Hayes, B. K., & Livesey, E. J. (2020). Modelling generalisation gradients as augmented Gaussian functions. *Quarterly Journal of Experimental Psychology*, 174702182094947. <https://doi.org/10.1177/1747021820949470>
- Li, W., Howard, J. D., Parrish, T. B., & Gottfried, J. A. (2008). Aversive learning enhances perceptual and cortical discrimination of indiscriminable odor cues. *Science*, 319(5871), 1842–1845. <https://doi.org/10.1126/science.1152837>
- Lovibond, P. F., Lee, J. C., & Hayes, B. K. (2020). Stimulus discriminability and induction as independent components of generalization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 46(6), 1106–1120. <https://doi.org/10.1037/xlm0000779>
- Mednick, S. A., & Freedman, J. L. (1960). Stimulus generalization. *Psychological Bulletin*, 57(3), 169–200. <https://doi.org/10.1037/h0041650>
- Purtle, R. B. (1973). Peak shift: A review. *Psychological Bulletin*, 80(5), 408–421. <https://doi.org/10.1037/h0035233>
- R Core Team. (2013). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <http://www.r-project.org/>
- Resnik, J., Sobel, N., & Paz, R. (2011). Auditory aversive learning increases discrimination thresholds. *Nature Neuroscience*, 14(6), 791–796. <https://doi.org/10.1038/nn.2802>
- Russell Lenth. (2020). emmeans: Estimated Marginal Means, aka Least-Squares Means. Retrieved from <https://cran.r-project.org/package=emmeans>
- Schechtman, E., Laufer, O., & Paz, R. (2010). Negative valence widens generalization of learning. *The*

- Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 30(31), 10460–10464. <https://doi.org/10.1523/JNEUROSCI.2377-10.2010>
- Struyf, D., Zaman, J., Hermans, D., & Vervliet, B. (2017). Gradients of fear: How perception influences fear generalization. *Behaviour Research and Therapy*, 93, 116–122. <https://doi.org/10.1016/j.brat.2017.04.001>
- 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. <https://doi.org/10.1016/j.neubiorev.2015.11.004>
- Tuominen, L., Boeke, E., DeCross, S., Wolthusen, R. P., Nasr, S., Milad, M., ... Holt, D. (2019). The relationship of perceptual discrimination to neural mechanisms of fear generalization. *NeuroImage*, 188(November 2018), 445–455. <https://doi.org/10.1016/j.neuroimage.2018.12.034>
- Venables, W. N., & Ripley, B. D. (2002). *Modern Applied Statistics with S* (4th ed.). New York: Springer. Retrieved from <http://www.stats.ox.ac.uk/pub/MASS4>
- Wilderjans, T. F., Ceulemans, E., & Meers, K. (2013). CHull: A generic convex-hull-based model selection method. *Behavior Research Methods*, 45(1), 1–15. <https://doi.org/10.3758/s13428-012-0238-5>
- Wong, A. H. K., & Lovibond, P. F. (2017). Rule-based generalisation in single-cue and differential fear conditioning in humans. *Biological Psychology*, 129(September), 111–120. <https://doi.org/10.1016/j.biopsycho.2017.08.056>
- Zaman, J., Ceulemans, E., Hermans, D., & Beckers, T. (2019). Direct and indirect effects of perception on generalization gradients. *Behaviour Research and Therapy*, 114(January), 44–50. <https://doi.org/10.1016/j.brat.2019.01.006>
- Zaman, J., Chalkia, A., Zenses, A.-K., Bilgin, A. S., Beckers, T., Vervliet, B., & Boddez, Y. (2020). Perceptual variability: Implications for learning and generalization. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-020-01780-1>
- Zaman, J., Grant, E., Zaman, J., Zaman, J., Struyf, D., Ceulemans, E., ... Beckers, T. (2020). Perceptual errors are related to shifts in generalization of conditioned responding. *Psychological Research*, 1–26. <https://doi.org/10.1007/s00426-020-01345-w>
- Zaman, J., Struyf, D., Ceulemans, E., Beckers, T., & Vervliet, B. (2019). Probing the role of perception in fear generalization. *Scientific Reports*, 9(1), 10026. <https://doi.org/10.1038/s41598-019->

46176-x

Zenses, A., Lee, J. C., & Zaman, J. (2020). Differences in perceptual memory influence generalization patterns. *Psyarxiv*, 1–36. <https://doi.org/10.31234/osf.io/tc8f5>