Does error-driven learning occur in the absence of cues? Examination of the effects of updating connection weights to absent cues

Jessie S. Nixon (jessie.nixon@uni-tuebingen.de)
Quantitative Linguistics, Eberhard Karls University of Tübingen, Germany

Sanne Poelstra (sannepoelstra@home.nl)
University of Groningen, Netherlands

Jacolien van Rij (j.c.van.rij@rug.nl)
University of Groningen, Netherlands

Abstract

The Rescorla-Wagner model has seen widespread success in modelling not only its original target of animal learning, but also several areas of human learning. However, despite its success, a number of studies with humans have found effects that are not predicted by the model, thus inspiring proposals for modifications to the model. One such proposal, by Van Hamme and Wasserman (1994, VHW), is that humans not only learn from present cues to all (present and absent) outcomes, as in the original model, but also learn from the absence of cues. They found evidence for this hypothesis in a causal rating experiment. However, behaviour in learning studies may depend on the task. We propose that error-driven learning should be considered to be a form of *implicit* learning and that the results of VHW's contingency judgement task might stem from explicit strategies involving logic and reasoning. The present study investigates this question by a) running simulations with both the original and modified versions of the model; b) replicating the VHW experiment (Experiment 1); and c) extending the experiment with new stimuli and a post-test of learned and unseen stimuli (Experiment 2). Simulations show that the VHW modified model predicts that cues learnt at the beginning will be unlearnt when absent over the following blocks, so that they become negative predictors over time. In contrast, the original RW predicts that the absent cues remain steady (positive) predictors over the blocks. Results showed no significant difference in cue assignment between training and test, in line with the original RW model. Moreover, predictive cues in the training phase showed significantly higher ratings than a new cue introduced in the test phase, at least in some cases, also partially supporting the original RW. The results suggest that when an overt response is required to absent cues, participants adjust ratings. But in later blocks when no response was required, there did not appear to be learning from the absent cues. We propose that in the development of human learning theory, attention should be paid to whether the behaviour (or other learning data) to be modelled results from implicit learning or involves higher level cognitive processes. We suggest that the RW may best capture implicit error-driven learning.

Keywords: error-driven learning; discriminative learning; associative learning; Rescorla-Wagner model; delta rule; absent cues; cue competition

Introduction

The Rescorla-Wagner learning equations (Rescorla & Wagner, 1972, also independently developed by Widrow & Hoff, 1960) were initially developed to explain findings from several decades of animal learning research. However, in the half century since its publication, the model has had a vast influence, not only in animal learning, but also in several other areas of psychology and human learning (see e.g. Siegel & Allan, 1996; Miller, Barnet, & Grahame, 1995, for reviews). While earlier models had assumed that learning ('conditioning' in animal learning parlance, also 'associative learning') resulted from contiguity or co-occurrence of stimuli, a number of findings demonstrated that contiguity was neither necessary nor sufficient for learning (Kamin, 1968, 1969b; Rescorla, 1988). The Rescorla-Wagner model was developed to capture the observation that rather than contiguity, learning was instead driven by surprise / prediction error and uncertainty (Kamin, 1969a; Rescorla, 1988), hence the term 'error-driven learning'. The Rescorla-Wagner model has proven remarkably successful in predicting human category learning (Gluck & Bower, 1988) and has recently been proposed as an account of human language acquisition (Ramscar & Yarlett, 2007; Ramscar, Yarlett, Dye, Denny, & Thorpe, 2010; Ramscar, Dye, & McCauley, 2013), predicting many linguistic phenomena (Baayen, Shaoul, Willits, & Ramscar, 2016; Baayen, Milin, Đurđević, Hendrix, & Marelli, 2011; Ellis, 2006; Lentz, Nixon, & van Rij, 2022; Nixon, 2020; Nixon & Tomaschek, 2020, 2021). In a study investigating the learning mechanisms underlying second language speech sound acquisition, Nixon (2020) demonstrated that a number of key principles of error-driven learning also apply to human learning of speech, including Kamin's 'blocking effect' (Kamin, 1968, 1969b), cue competition, prediction and unlearning (Ramscar et al., 2010).

But are responses in causal judgement tasks captured by error-driven learning models? Van Hamme and Wasserman (1994) noted that, in causal judgement tasks, previous work had suggested that human participants are able to take into account both occurrence and non-occurrence of potential causal factors (Arkes & Harkness, 1983; Levin, Wasserman, & Kao, 1993; Wasserman, Dorner, & Kao, 1990). They proposed that human learning in causal judgement tasks not only involves changes in judgement of the causal relation between cues and outcomes for the particular cues that occur on a given trial, but also for cues that do not occur on that trial. They set out to test this hypothesis with a causal rating task. On each trial, participants were presented with two out of three food types. One food type occurred on every trial: that is, trials were either AX or BX. Participants were told whether or not there was an allergic reaction on that trial and were asked to give a rating (on a scale of 0-8) how likely they thought each of the three food types was the cause of the allergy. Importantly, although only two foods occurred on each trial, participants had to give a rating for all three foods – so in this sense, the

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.

third food type was not entirely absent from the trial. As described in more detail below, Van Hamme and Wasserman concluded from their results that, at least in causal judgement tasks, humans learn from absent cues.

In language, on the other hand, error-driven learning has been shown to be a predictive process, in which temporally earlier *cues* predict temporally later *outcomes* (Hoppe, van Rij, Hendriks, & Ramscar, 2020; Nixon, 2018, 2020; Ramscar et al., 2010). That is, earlier events may be used to predict later events, but not the other way around. This would suggest that in this case, predictions would be unlikely to be generated from absent cues.

Explaining what accounts for the differences in these observed results is important for understanding human learning mechanisms. The question of whether we learn only from present cues or also from absent cues is of fundamental importance to learning theory. It has been proposed that factors other than error-driven learning may also play a role. For example, Anderson (2009) proposes an effect of forgetting over time. It may be that, rather than a predictive relationship between absent cues and occurring vs. non-occurring outcomes, the predictive value of cues instead simply gradually decreases during periods where the cues are not encountered.

A point we believe important in this respect is that different tasks recruit different cognitive processes. Perhaps the different results between experiments stem from different strategies participants use in these different tasks. In this paper, we propose that *implicit* learning does not involve learning from absent cues, due to its predictive, discriminative nature. On the other hand, when instructed to respond to items labelled as 'not present', this requires generating a mental representation of the absent cue and reasoning about its relation to the outcome. For example, on a trial in which an allergy occurs, if a participant believes Cue A (present) to be the culprit, it is logical to suppose that Cue B (absent) was not cause. In contrast, error-driven learning is driven not by logic but by information (Ramscar, Dye, & Klein, 2013). We therefore propose that the Van Hamme and Wasserman task invokes higher-level cognitive processes, which participants can also learn from, but which represent a different type of task to implicit learning and may be outside the scope of phenomena predicted by error-driven learning models.

In order to test this, below we first introduce the computational simulations with the Rescorla-Wagner (RW) model and the adjusted version proposed by Van Hamme and Wasserman (VHW). We then present two experiments based on the VHW study. Finally, we discuss implications of the results for learning theory.

The specification of present cues and outcomes is straightforward; however, less obvious is what counts as an absent cue. Taken literally, a potentially infinite number of cues are absent in every learning event. VHW stated that relevant cues are limited to those that have some existing association with the outcome in question, due to previous learning. "A cue would become relevant after acquiring some level of positive or negative association. For animal subjecs, it would be necessary to present the stimulus at least once in the experimental context followed by reinforcement or presented together with a cue that had previously been reinforced. For human subjects, the relevance of a cue as a potential cause could be established with verbal instructions" (footnote, page 132). Therefore, in our simulations, adjustments are only made to cues that have occurred at least once; cues from later blocks are not included as absent in earlier blocks.

Computational modelling

Original Rescorla-Wagner learning equations

We will first describe the learning equations as proposed independently by both Widrow and Hoff (1960) and Rescorla and Wagner (1972) and then the adaptation proposed by Van Hamme and Wasserman (1994). The Rescorla-Wagner equations estimate the connection *weights* $\mathcal{W}$, between the model inputs, or *cues* $\mathcal{C}$ and a set of *outcomes* $\mathcal{O}$. The training produces a network that consists of a matrix of connection weights; k cues and n outcomes produces a $k \times n$ matrix. During each trial of the training, weights are adjusted between all cues present on that trial and all outcomes present or encountered previously. No adjustment is made to cues not present on a given trial. Adjustments made to the connection weight between cues c_i and outcome o_j on a given trial t, is given by the Rescorla-Wagner equations:

$$w_{ij}^{(t)} = w_{ij}^{(t-1)} + \Delta w_{ij}^{t}. \tag{1}$$

The connection weight at the end of trial t is equal to the weight at the end of the previous trial, $t-1$, plus any change during the current trial.

The change in weights during the current trial, Δw_{ij}^{t}, is given by the Rescorla-Wagner equations:

$$\Delta w_{ij}^{t} = \begin{cases} \text{a) } \alpha_i \beta_j \left(\lambda - \sum_{[Present(c_k,t)]} w_{kj}\right) & \text{if } Present(c_i,t) \\ & \text{and } Present(o_i,t), \\ \text{b) } \alpha_i \beta_j \left(0 - \sum_{[Present(c_k,t)]} w_{kj}\right) & \text{if } Present(c_i,t) \\ & \text{and } Absent(o_i,t), \\ \text{c) } 0 & \text{if } Absent(c_i,t). \end{cases} \tag{2}$$

in which λ is the maximum learnability of the outcome; and α_i and β_j refer to cue and outcome salience, respectively.

The above equation can be summarised as a) if a cue and an outcome are both present, the cue-outcome weight increases; b) if a cue is present and an outcome is not present, cue-outcome weight decreases; and most crucially for the present study, c) for any cues not present on that trial, there is no change. This is the aspect challenged by VHW.

Due to the theoretical importance of learning from negative evidence (Nixon, 2020; Ramscar et al., 2010), β_j is the same in a) and b), meaning that learning is potentially equivalent on positive and negative trials. However, the amount of adjustment (in a and b) depends on connection strengths developed during previous learning. The non-occurrence of a fully expected outcome is as surprising as the occurrence of a fully unexpected outcome.

Adaptation proposed by Van Hamme & Wasserman

$$\Delta w_{ij}^{t} = \begin{cases} \text{a) } \alpha_i \beta_j \left(\lambda - \sum_{[Present(c_k,t)]} w_{kj}\right) & \text{if } Present(c_i,t) \\ & \text{and } Present(o_i,t), \\ \text{b) } \alpha_i \beta_j \left(0 - \sum_{[Present(c_k,t)]} w_{kj}\right) & \text{if } Present(c_i,t) \\ & \text{and } Absent(o_i,t), \\ \text{c) } \alpha_i \beta_j \left(0 - \sum_{[Present(c_k,t)]} w_{kj}\right) & \text{if } Absent(c_i,t), \\ & \text{and } Present(o_i,t), \\ \text{d) } \alpha_i \beta_j \left(\lambda - \sum_{[Present(c_k,t)]} w_{kj}\right) & \text{if } Absent(c_i,t), \\ & \text{and } Absent(o_i,t). \end{cases} \tag{3}$$

The first two rows of Eq. 3 (greyed out) are the same as the original Rescorla-Wagner equations (Eq. 2). The third and fourth rows differ. In the original Rescorla-Wagner equations (Eq. 2), when a cue is absent on a trial, no weight adjustments are made to that cue (expression c). Van Hamme and Wasserman (1994) suggest that this part of the equation should be changed as shown in Eq. 3: when a cue is absent and the outcome occurs, weights are reduced (expression c); when a cue is absent and the outcome does not occur, weights are increased (expression d).

Simulations

Simulations and visualisations were run using the `edl` package (van Rij & Hoppe, 2020) in R (R Core Team, 2020). The combined alpha and beta parameter (the learning rate) was set to 0.001 and lambda (the maximum connection strength to the outcome) was set to 1 (both the default parameters). Note that in the simulation descriptions and figures presented here, we use the stimuli from the original VHW experiment and our Experiment 1, namely foods and allergic reaction; however, since Experiment 2 used an identical experiment design, the simulation also applies to Experiment 2 (replacing cues such as ‘bran’ with ‘footprint’ and the outcome ‘allergy’ (or not) with ‘diamond’ (or not)). Cues were created for each of the nine foods that occurred during the experiment - three food types for each of three blocks (see Table 1). Note that VHW had six food conditions per block, rather than three: conditions 4-6 were repetitions of 1-3, but with different foods. This was to account for any pre-existing biases with respect to the likelihood of allergies to various foods. As this is not applicable in the simulations, the simulations had only three conditions. On each trial of the simulation, two cues (foods) were presented, as well as a background cue that was present on all trials to model the experiment environment, as specified in the Rescorla-Wagner model (Rescorla & Wagner, 1972). There were two outcomes, allergic reaction and no allergic reaction, one of which occurred on every trial. The order of cues and outcomes followed the order set out in Van Hamme and Wasserman (1994): the order of cues was the same for all participants and all blocks, namely AX, BX, BX, AX, AX, BX, AX, BX, BX, AX, BX, AX, AX, BX, AX, BX. The order of outcomes depended on the contingency condition. All three blocks were first run separately to model the results presented in Van Hamme and Wasserman (1994) and then also run as a single simulation to model the learning of one participant over the experiment. This was done for each of the contingency condition orders.

Simulation results

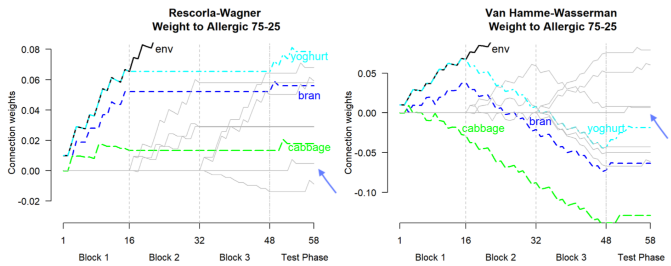

Figure 1: Simulations of the predicted response to ‘allergic’ in Experiment 1 (or ’diamond’ in Experiment 2, if food cues are replaced with forest cues) using the original Rescorla-Wagner equations and the adaptation proposed by Van Hamme & Wasserman. **Left column**: simulations using Rescorla-Wagner equations. **Right column**: simulations using Van Hamme & Wasserman’s adaptation. The RW predicts that cue-outcome connections that developed in the first block (e.g. ‘bran’) will be retained over the course of the later experiment blocks. The VHW adaptation predicts that these cues will reduce substantially and develop negative connection weights to the allergic reaction. In order to test these predictions, we can introduce a new cue after training (blue arrow). The RW model predicts that ‘bran’ will have a higher rating than this new cue; the VHW model predicts that ‘bran’ will have a lower rating than this new cue. (Note that in order to maximise space in the figure, the scales are shifted relative to the vertical axis in these plots – zero is higher in VHW).

Figure 1 shows the simulation results for the original Rescorla-Wagner model (left) and the adjusted Van Hamme and Wasserman version (right). The simulations indicate the relative strength and the pattern of change of the various cues, as an estimate of the rating responses.

The Rescorla-Wagner model (left panel) predicts that cue-outcome connections that developed in the first block (e.g. ‘bran’) will be retained over the course of the later experiment blocks. The Van Hamme-Wasserman adaptation (right panel), on the other hand, predicts that these cues will reduce substantially and develop negative connection weights to the allergic reaction. (Note that in order to maximise space in the figure, the scales are shifted relative to the vertical axis in these plots – zero is higher in VHW).

In order to test these predictions, we compare the ratings during the training block to the responses in the test phase at the end of the experiment. In addition, we also introduce a new cue after training (indicated in the figure with an arrow). The RW model predicts that 'bran' will have a higher rating than this new cue; the VHW model predicts that 'bran' will have a lower rating than this new cue. We will conduct these comparisons in Experiment 2. However, we will first run a direct replication of the VHW experiment to ensure that our online testing environment is able to replicate the conditions of the original in-lab, pen-and-paper experiment. We do this in Experiment 1.

Experiment 1

Experiment 1 was a direct replication of the VHW study, except that it was run online. The purpose of this experiment was to replicate the VHW results and test our online experimental paradigm.

Participants

Twenty participants aged 19-25 were paid for their participation. Participants were recruited on Prolific.

Stimuli

The cues consisted of food items and the outcomes were whether the allergic reaction occurred that day (i.e. experimental trial) or not. A complete list of the cues and contingency conditions is presented in Table 1.

Table 1: *Experimental stimuli. The second column shows the probability of the allergy occurring on AX and BX trials, respectively. For example, in the 75-25 condition the allergy occurs on 75% of the AX trials and 25% of the BX trials. Note that food conditions 1-3 are identical to 4-6; only the specific items differ. Therefore, in the simulations, only three conditions apply.*

food cond	prob. of allergy AX-BX	cues (food types)		
		X	A	B
1	50-50	shrimp	strawberries	peanuts
2	75-25	yoghurt	bran	cabbage
3	100-0	bananas	chicken	mustard
4	50-50	wheat	walnuts	peaches
5	75-25	corn	horseradish	lobster
6	100-0	blueberries	cheese	pork

Experiment design and procedure

Each participant was presented with three different foods in each block (see Table 1). Two foods and one outcome (allergic reaction or no allergic reaction) were presented on a computer screen for up to 15 seconds (as in VHW) or until participant response. Participants were asked to give a rating (0-8) how likely each of all *three* of the foods caused the allergy, not just the two foods presented on screen. VHW did not provide anchors for the response scale, but pre-test ratings suggest the mid-point (rather than 0) was taken as the neutral point (i.e. 'maybe'; maximum uncertainty).

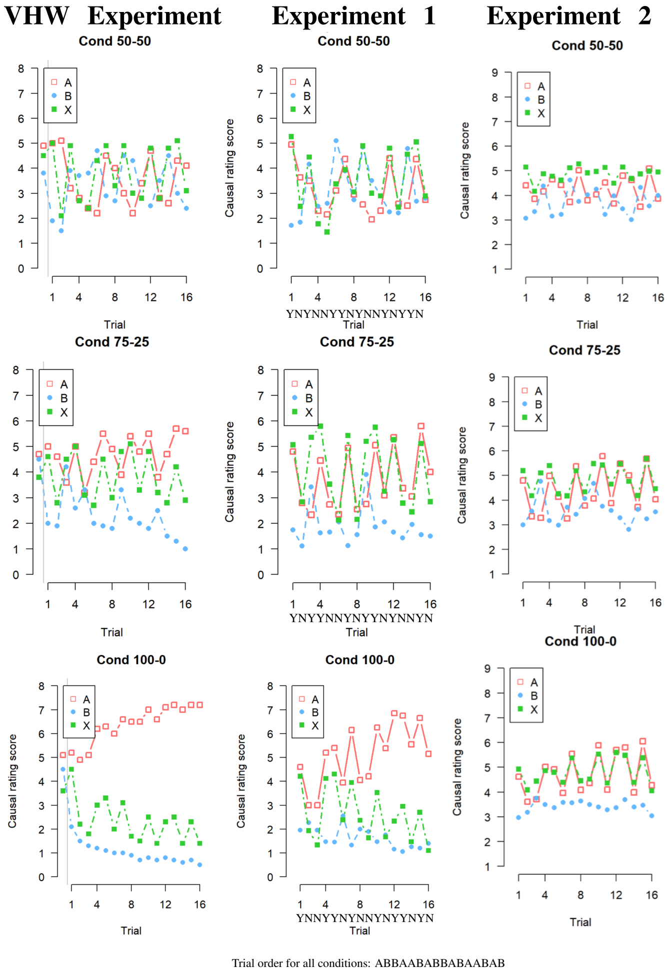

Figure 2: Results of VHW experiment (left), Experiment 1 (centre) and Experiment 2 (right) for the three different contingency conditions: AX and BX both leading to allergy on 50% of trials (top); AX 75% allergy, BX 25% (middle row); and AX 100% allergy, BX 0% (bottom). Trial is on the x-axis; rating score (how likely [CUE] leads to allergy/diamond) is on the y-axis. The Xs and Ys show the order of outcomes, whether the allergy occurred/diamond was found (Y) or not (N) for all experiments.

One of the foods occurred on every trial in the block, so that each trial was either AX (e.g. strawberries and shrimp) or BX (e.g. peanuts and shrimp). All foods differed between blocks. The order of the cues was always AX, BX, BX, AX, AX, BX, AX, BX, BX, AX, BX, AX, AX, BX, AX, BX. The order of the outcomes depended on the food-allergy probability condition (second column, Table 1). Three within-participant probability conditions were used, one per block. The order of conditions varied between participants. The three conditions were: AX and BX both predicted the allergy on 50% of the trials (50-50 condition); AX predicted the allergy on 75% of trials and BX on 25% of trials (75-25 condition) and AX predicted the allergy on 100% of trials and BX on 0% of trials (100-0 condition).

Results

Results are shown in Figure 2 (centre column). The results of VHW are also given for comparison (left column). The pattern of responses is similar between experiments. The average responses are a bit more variable in Experiment 1, probably due to the smaller number of participants. However, overall, the results show that are online experimental set up is sound.

Discussion

The main purpose of Experiment 1 was to replicate the findings of VHW in an online setting. Results replicated the original study, thus verifying our procedure. We now turn to a different set of stimuli and include a test phase to investigate long term changes over time to absent cues when no response to is required from the participants to the absent cues.

Experiment 2

Experiment 2 investigated the effect of task by using a different set of stimuli. Additional post-tests were also included to investigate the conditions required for learning from absent cues. Above, we noted that, during a block, participants were required to respond to the 'absent cue'. Requiring participants to make a judgment on the cue might lead to changes in their representation of the cue. Including a post-test of the same cues allows us to test for learning of the absent cue over time during a period when *no overt response* to the given cue was required. We do this by a) comparing the cue during Block 1 to the same cue in the post-test and b) comparing the learned stimuli to new unseen stimuli in the post-test.

Participants

Sixty participants aged 19-25 were paid for participation.

Stimuli

The cues consisted of objects found in a forest, such as acorn, footprint, flower. The participants were asked to use the items as clues to finding treasure in an alien world. On each trial, they dug for treasure. The outcomes were whether they found a diamond or not. The probability conditions were the same as Experiment 1 (now probability of finding treasure).

Experiment design and procedure

The experiment design and procedure were the same as Experiment 1, except that the participants were given a cover story that they were in an alien world, searching for treasure. In addition, there was also a test phase at the end of the experiment that tested participant responses to cues encountered in the first block. Because these cues had not occurred in blocks 2 and 3, the VHW model predicts that the cue weight becomes negative over time (see Figure 1). However, the RW model predicts the cue weight remains positive. In addition, a new unseen cue was tested. The VHW model predicts these cue weights would be higher than the cues from the first block; the RW model predicts the weights to unseen cues would be lower than previously learned cues.

Results

Training phase Figure 2 (right column) shows the results of the Training Phase in Experiment 2. A roughly similar pattern emerges in the two experiments, although there is less separation of the cues in Experiment 2. One question is whether participants changed their responses on every trial, as found by Van Hamme and Wasserman (1994). Visual inspection of Figure 2 shows that this does occur on some trials. For example in the 75-25 condition, trial 4 is an AX trial with a diamond found; the rating of A rises and the rating of B falls. On the other hand, this was not always the case. Trial 5 is an AX trial with no diamond found; the rating of A falls, but the rating of B does not rise. Nevertheless, we can say that responses to the absent cues change in at least some cases.

We compared the scores on the no diamond trials to those on the diamond trials with paired t-tests. As we did multiple t-tests over the same data-set, we adjusted the p-values according to the Bonferroni-Holm method (Holm, 1979). In the 50-50 Condition, there was no significant difference between 'Diamond' and 'No diamond' trials: A ($t(59) = 2.49$, $p = .08$), B ($t(59) = 1.73$, $p = .27$, X ($t(59) = 2.28$, $p = .10$). In contrast, there was a significant difference between 'Diamond' and 'No diamond' trials for the 75-25 condition (A: $t(59) = 5.16$, $p < .001$; X: $t(59) = 4.31$, $p < .001$) and the 100-0 condition (A: $t(59) = 4.87$, $p < .001$; X: $t(59) = 3.05$, $p = .02$). Cue B did not differ (cue B ($t(59) = 0.81$, $p = .84$; $t(59) = 0.25$, $p = .84$). These results suggest that when there appears to be random variation, participants do not change their responses on every trial. However, when there is a pattern of some cues providing information about the (treasure) outcome, participants adjust cue weights.

Post-test Of most interest in the present study is whether participants also adjust their representations of the cues when they are not required to make an overt response. Here we compared the ratings of cues in Block 1 to the same cue in the post-test. Predictions are shown in Figure 1 as the predicted rating score of 'bran' to 'allergic reaction'. (As explained above, the figues show the simulations for food cues and allergy outcomes as in VHW and Experiment 1; however, since the experiment design is identical, the simulation predictions apply to both experiments.) The RW model predicts no difference between Block 1 and post-test. The VHW model predicts a lower rating in the post-test.

We tested the rating scores between Block 1 and the post-test using Bonferroni-Holm corrected Wilcoxon tests. There was no significant difference between Block 1 and the post-test (50-50 Condition $v=79$, $p=1$; 75-25 Condition, $v=31$, $p=.28$, 100-0 Condition, $v=43$, $p=.32$). Neither was there any significant difference between Block 3 and the Test phase (50-50 Condition v=90, p=1; 75-25 Condition, v=145, p=1, 100-0 Condition, v=82, p=1).

We also introduced a new cue after the training phase (see Figure 1, blue arrow). We used Wilcoxon tests to test differences in scores between Cue A from the Training phase and this new cue introduced in the post-test. Results showed no differences between Block 1 and the new cue (50-50 Condition v=914, p=.37; 75-25 Condition, v=891, p=.73, 100-0 Condition, v=939, p=.32). However, the difference between Block 3 and the new cue was significant (50-50 Condition v=320, p=.001; 75-25 Condition, v=608, p=.002, 100-0 Condition, v=772, p=.005).

Discussion

Experiment 2 investigated whether participants learn from absent cues when learning what clues in the environment could help them find treasure. We introduced a test phase to compare the long term changes to cues that were absent for two blocks of the experiment. If participants learn from absent cues, the weights to these cues should drop off substantially, as shown in the simulation for the VHW model. In contrast, the RW model predicts that participants make predictions based on cues available in the environment, so no changes in cue weights are predicted over the blocks. Results supported the RW model, as no significant difference was found in the ratings between Block 1 and test.

Experiment 2 also tested whether a new cue would receive a higher rating than previously learned cues, as predicted by VHW, or a lower rating, as predicted by RW. Results here were less conclusive. Comparison of Block 1 cue to the new cue was not significant in either direction. This supported neither the RW model, nor the VHW adaptation. However, comparison of Block 3 cue to the new cue showed a lower rating for the new cue, as predicted by RW.

General Discussion

This study tested the predictions of the Rescorla-Wagner equations and the adaptation of the equations proposed by Van Hamme and Wasserman (1994) for predicting participant responses in two experiments based on the design of Van Hamme and Wasserman (1994). Experiment 1 aimed to replicate Van Hamme and Wasserman (1994) in an online environment. Results showed that the main results were replicated and our testing paradigm was sound.

Experiment 2 extended the study to a new task, using clues in the environment to help 'find buried treasure in an alien world'. Van Hamme and Wasserman observed that on specific trials in which a particular cue did not occur and the reaction occurred, the rating went down on average, and when the reaction occurred, the rating went up on average. From this they concluded that humans learn from absent cues in causal attribution. Experiment 2 also added a test phase to test whether there were long term changes to the absent cues over two blocks of the experiment. Figure 1 shows the predictions of the two models. The test phase showed that there was no difference in response to cues between Block 1 and the test phase, as predicted by the RW model. This is counter to the predictions of the VHW adaptation, which suggests the absent cues should develop negative weights. The second comparison was between learned cues and new cues (purple arrow in Figure 1). These results were not entirely consistent, as there was no significant effect in either direction for Block 1, i.e. neither model was supported. However, the Block 3 cues did show support for the RW model, as the already learned cue had a higher rating than the new cue.

The present results raise the question of how it is that although participants at least sometimes change their responses on individual trials during the training phase, as reported by Van Hamme and Wasserman (1994), we did not find long-term effects of these responses. We suggest that the within-trial adjustments reported by Van Hamme and Wasserman stem from the process of consciously processing the absent cue in order to give a rating for all three outcomes. Recall that during each block, participants were required to give a rating for all cues, even the 'absent' cue. In contrast, in Blocks 2 and 3, no reponse was required to the Block 1 cues. In this case, the participants did not learn from these absent cues.

We suggest that error-driven learning is best understood as an implicit learning process and that this involves predictions based on present cues. On the other hand, when a response is required to any 'absent' information, this affects learning through an additional reasoning process. Weights are adjusted up or down for the absent cue during the rating process when required to make responses to the absent cue. In object label learning, children learn through error-driven informativity of cues; however, when a small number of items is presented, adults tend to adopt a logical exclusion strategy (Ramscar, Dye, & Klein, 2013). Given that the rating task included the absent cue as well as the present cues, perhaps an exclusion strategy may have affected participants' responses in this task.

In summary, the present results shed new light on the question of learning from absent cues. When participants are asked to give ratings to both present and absent cues, as in the training phase, then they make adjustments to the absent cues. However, results from the test phase suggest that when participants are no longer asked to explicitly give ratings to the absent cues, it appears that they no longer learn from these absent cues. This suggests the possibility that people use different learning mechanisms – explicit logical reasoning versus implicit learning – depending on the task.

Acknowledgments

This paper is based on the Masters dissertation of the second author which she wrote while at the University of Groningen. She has since gone on to pursue a career outside academia. We would like to thank four anonymous reviewers for helpful feedback on an earlier version of this paper. This research was supported by an iCog Research Grant to the first author, a Netherlands Organisation for Scientific Research NWO Veni grant (275-70-044) to the third author and an ERC Advanced Grant (Grant number 742545).

References

Anderson, J. R. (2009). *How can the human mind occur in the physical universe?* Oxford University Press.

Arkes, H. R., & Harkness, A. R. (1983). Estimates of contingency between two dichotomous variables. *Journal of Experimental Psychology: General*, *112*(1), 117.

Baayen, R. H., Milin, P., Đurđević, D. F., Hendrix, P., & Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological review*, *118*(3), 438.

Baayen, R. H., Shaoul, C., Willits, J., & Ramscar, M. (2016). Comprehension without segmentation: A proof of concept with naive discriminative learning. *Language, Cognition and Neuroscience*, *31*(1), 106–128.

Ellis, N. C. (2006). Selective attention and transfer phenomena in L2 acquisition: Contingency, cue competition, salience, interference, overshadowing, blocking, and perceptual learning. *Applied Linguistics*, *27*(2), 164–194.

Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*(3), 227.

Hoppe, D. B., van Rij, J., Hendriks, P., & Ramscar, M. (2020). Order matters! influences of linear order on linguistic category learning. *Cognitive Science*, *44*(11), e12910.

Kamin, L. J. (1968). "Attention-like" processes in classical conditioning. In *Miami symposium on the prediction of behavior: Aversive stimulation* (pp. 9–31).

Kamin, L. J. (1969a). Predictability, surprise, attention and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279–296). New York: Appleton-Century-Crofts.

Kamin, L. J. (1969b). Selective association and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 42–64). Halifax, Canada: Dalhousie University Press.

Lentz, T. O., Nixon, J. S., & van Rij, J. (2022, April). Temporal response modelling uncovers electrophysiological correlates of trial-by-trial error-driven learning. *Language, Cognition & Neuroscience*. Retrieved from `https://www.psyarxiv.com/dg5mw` doi: 10.31234/osf.io/dg5mw

Levin, I. P., Wasserman, E. A., & Kao, S.-f. (1993). Multiple methods for examining biased information use in contingency judgments. *Organizational Behavior and Human Decision Processes*, *55*(2), 228–250.

Miller, R. R., Barnet, R. C., & Grahame, N. J. (1995). Assessment of the rescorla-wagner model. *Psychological bulletin*, *117*(3), 363.

Nixon, J. S. (2018, September). Effective acoustic cue learning is not just statistical, it is discriminative. In *Interspeech 2018 – 19th Annual Conference of the International Speech Communication Association* (pp. 1447 – 1451). Hyderabad, India.

Nixon, J. S. (2020). Of mice and men: speech acquisition as discriminative learning from prediction error, not just statistical tracking. *Cognition*, *197*, 104081.

Nixon, J. S., & Tomaschek, F. (2020). Learning from the acoustic signal:Error-driven learning of low-level acoustics discriminates vowel and consonant pairs. In S. Denison, M. Mack, Y. Xu, & B. C. Armstrong (Eds.), *Proceedings of CogSci2020, the 42nd Annual Meeting of the Cognitive Science Society* (pp. 585–591). Toronto: Cognitive Science Society.

Nixon, J. S., & Tomaschek, F. (2021). Prediction and error in early infant speech perception: A speech acquisition model. *Cognition*, *212*, 104697.

R Core Team. (2020). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from `https://www.R-project.org/`

Ramscar, M., Dye, M., & Klein, J. (2013). Children value informativity over logic in word learning. *Psychological science*, *24*(6), 1017–1023.

Ramscar, M., Dye, M., & McCauley, S. M. (2013). Error and expectation in language learning: The curious absence of mouses in adult speech. *Language*, *89*(4), 760–793.

Ramscar, M., & Yarlett, D. (2007). Linguistic self-correction in the absence of feedback: A new approach to the logical problem of language acquisition. *Cognitive Science*, *31*(6), 927–960.

Ramscar, M., Yarlett, D., Dye, M., Denny, K., & Thorpe, K. (2010). The effects of feature-label-order and their implications for symbolic learning. *Cognitive Science*, *34*(6), 909–957.

Rescorla, R. A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist*, *43*(3), 151.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning ii: Current research and theory* (Vol. 2, pp. 64–99). New-York: Appleton-Century-Crofts.

Siegel, S., & Allan, L. G. (1996). The widespread influence of the rescorla-wagner model. *Psychonomic Bulletin & Review*, *3*(3), 314–321.

van Rij, J., & Hoppe, D. (2020). *edl: Toolbox for error-driven learning simulations with two-layer networks.* (R package version 0.3.)

Van Hamme, L. J., & Wasserman, E. A. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning and motivation*, *25*(2), 127–151.

Wasserman, E. A., Dorner, W., & Kao, S. (1990). Contributions of specific cell information to judgments of interevent contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*(3), 509.

Widrow, B., & Hoff, M. E. (1960). Adaptive switching circuits. *1960 WESCON Convention Record Part IV*, 96–104.