

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18

Supplementary Materials Supporting:

G. R. Esber

&

M. Haselgrove

Reconciling the influence of predictiveness and uncertainty on stimulus salience:

A model of attention in associative learning

19
20

S1. The Mackintosh (1975) model

21 Mackintosh [2] suggested that the change in the association (ΔV) between a cue (i) and
22 the outcome proceeds according to the equation:

23

$$24 \quad \Delta V_i = \alpha_i \theta (\lambda - V_i) \quad \text{Equation 1.1}$$

25

26 In which θ is a parameter that determines the rate of learning, $(\lambda - V_i)$ is an error term that
27 represents the discrepancy between the asymptote of learning that can be supported by
28 the outcome (λ) and the current strength of the association between cue i and outcome
29 (V_i). Most importantly, α_i is the attention paid to cue i , which varies according to the
30 following rules:

31

$$32 \quad \Delta \alpha_i > 0 \text{ if } |\lambda - V_i| < |\lambda - V_p| \quad \text{Equation 1.2}$$

33 and

$$34 \quad \Delta \alpha_i < 0 \text{ if } |\lambda - V_i| \geq |\lambda - V_p| \quad \text{Equation 1.3}$$

35

36 in which V_p is the associative strength of all cues other than i present on that trial.
37 Equation 1.2 ensures that attention to cue i (α_i) will increase if its error term is smaller
38 than the error term generated by all other cues (i.e. if cue i is the best available predictor
39 of the outcome). Equation 1.3 ensures that attention to cue i (α_i) will decrease if its error
40 term is, at best, no better than the error term generated by all other cues (i.e. if cue i is not
41 the best available predictor of the outcome). In this way, Mackintosh's theory predicts
42 that attention will be higher for predictive than non-predictive cues.

43 S2. The Pearce-Hall (1980) model

44

45 According to Pearce and Hall [5], the change in the association (ΔV) between a cue (i)
46 and the outcome proceeds according to the equation:

47

48 $\Delta V_i = S_i \alpha_i \lambda$ Equation 2.1

49

50 The terms within this equation that are common to Mackintosh's theory refer to the same
51 quantities. The novel S_i is a parameter similar to Mackintosh's θ and determines the rate
52 of learning, and depends upon the intensity of cue i . Pearce & Hall suggested that a cue
53 can also be associated with a representation of the outcome's absence, a so-called no
54 outcome (\bar{V}), and the change in the association ($\Delta \bar{V}$) between a cue (i) and the no
55 outcome proceeds according to the equation:

56

57 $\Delta \bar{V}_i = S_i \alpha_i \bar{\lambda}$ Equation 2.2

58

59 The magnitude of the absence of the outcome ($\bar{\lambda}$) is determined by the difference between
60 the expectation of the outcome and λ itself (which on trials with the absence of the
61 outcome will equal zero). Thus, this magnitude is equal to:

62

63 $\bar{\lambda} = (\Sigma V_i - \Sigma \bar{V}_i) - \lambda$ Equation 2.3

64

65 Finally, the attention that is paid to cue i on the subsequent conditioning trial (α_i^{n+1}) is
66 equal to the current absolute difference between the magnitude of the US, and the total
67 associative strength of the cue:

68

69 $\alpha_i^{n+1} = |\lambda - (\Sigma V_i - \Sigma \bar{V}_i)|$ Equation 2.4

70

71 Consequently, a cue that is paired, on some trials, with an outcome but on other trials
72 with the absence of the outcome will enjoy a maintenance of its attention as the value of
73 α will remain positive. However, a cue that is consistently followed by an outcome will
74 suffer a loss of attention as the value of α will tend towards zero.

75

76

77

S3. The importance of distinguishing between salience and associability

78

79

80

81

82

83

84

85

The model we advocate envisages that reinforcers of opposite affective polarity should be capable of enhancing the salience of their predictors independently, each in proportion to its emotional significance. This implies that the acquired salience of a cue depends directly upon, but is detached from, its acquired affective value. While a common assumption in attentional theories of learning [2,4,5,7], this notion has been challenged by demonstrations that the associability of a cue—a widely used index of its salience—is partly determined by its affective value.

86

87

88

89

90

91

92

93

94

95

96

97

98

99

Using a human causal learning paradigm, Le Pelley, Oakshott, Wills, and McLaren [S1] required participants to play the part of a dietician and learn the effects that ingesting certain types of food (e.g. carrots, eggs) would have on fictitious patients. Participants were divided into two groups, one in which foods were followed by appetitive post-ingestive consequences (enjoyment reactions), and another in which they were followed by aversive post-ingestive consequences (allergic reactions). During Stage 1, some of the foods were highly predictive of their effects, whereas others had no predictive value. In Stage 2, half of the participants in each group were faced with a new problem. Novel combinations of previously predictive and nonpredictive foods now signalled, with identical correlation, a set of novel consequences drawn from the same affective category (e.g. novel types of allergy if allergies had been used in Stage 1). During a final test, participants were asked to rate the likelihood that each of these foods would by themselves cause the effects they were paired with in Stage 2. Replicating previous studies [16], Le Pelley et al. [S1] found that foods that had been highly

100 predictive in Stage 1 scored significantly higher than foods that had been nonpredictive,
101 despite their being equally predictive of Stage-2 outcomes. These results thus confirm
102 that a history of predictiveness can influence the associability of a cue, presumably
103 through an enhancement of its salience.

104 The critical results came from the other half of the participants. They received
105 similar training in Stage 2, except for the fact that post-ingestive consequences belonged
106 to the opposite affective class (e.g. enjoyment reactions if allergies had been used in
107 Stage 1). In the test, these participants did not rate foods that had been predictive during
108 Stage 1 any different from foods that had been nonpredictive. Thus, it appears that for the
109 acquired salience of a cue to translate into greater associability, the new learning
110 experience must involve a reinforcer of the same affective class as the initial learning
111 experience. Le Pelley et al. [S1] took their results to show that salience is not a general,
112 value-detached property of the cue, but rather one that is modified separately within (and
113 therefore specific to each motivational system. This interpretation is incompatible with
114 the processes of salience modification we propose here.

115 How might the model be salvaged in the face of Le Pelley et al. [S1]’s results? In
116 our view, the force of this challenge is undermined when we abandon the tacit
117 assumption that the associability of a cue is entirely determined by its salience. In the
118 associative learning literature, these terms are often used interchangeably, presumably
119 because differences in associability remain one of the few tools at the investigator’s
120 disposal for inferring corresponding differences in salience. It is evident, however, that
121 the influence that the salience of a cue exerts on its associability may be obscured by
122 other factors, such as the cue’s similarity to the outcome [S2]; its biological relevance or

123 *preparedness* [S3,S4], the degree of generalization from other cues [S5] and the
124 inhibitory nature of appetitive-aversive interactions [33,S6].

125 The antagonistic relationship between appetitive and aversive motivational
126 systems has been well documented [for a review: 36]. Most relevant here are classical
127 studies showing that it is rather difficult to establish a predictor of shock as a signal for
128 food (and vice versa), by comparison with a previously nonreinforced cue [S7,S8]. These
129 results join Le Pelley et al.'s [S1] in demonstrating that the associability of a cue with a
130 history of predictiveness may under some circumstances appear no different or even
131 inferior to that of a relatively neutral stimulus. Learning models that assume hardwired
132 inhibitory associations between appetitive and aversive affective representations, such as
133 that proposed here, can readily accommodate this pattern of results. Briefly, these models
134 predict that a cue previously associated with shock will have no difficulty in entering into
135 an association with food. For a good part of this training, however, this new association
136 will not be readily expressed in performance, because the cue will simultaneously
137 activate the representation of shock, which will in turn suppress activation of the
138 representation of food. As a consequence, acquisition of the cue-food association will
139 appear retarded. Note that the influence of this process on associability is orthogonal to
140 that of stimulus salience, which leaves the possibility intact that the acquired salience of a
141 cue might be detached from its value¹.

142 It should now be possible to recognize that the experiments of Le Pelley et al.
143 [S1] likely involved such appetitive-aversive interactions. In their procedure, an accurate
144 predictor of, say, a certain allergy may not only have acquired substantial salience by the
145 end of Stage 1, but should also have developed a strong association with that outcome.

146 When this cue is next paired with an enjoyment reaction during Stage 2, it is plausible
147 that appetitive-aversive interactions might dampen the expression of this association,
148 offsetting the advantage conferred by the cue's high salience. By contrast, appetitive-
149 aversive interactions should have less of an influence on performance to a nonpredictive
150 cue, for its association with allergy should be rather weak to begin with. This would
151 explain why predictive and nonpredictive cues seemed to enter into the second
152 association with equal readiness.

153 Thus, the distinction between salience and associability is crucial in the context of
154 the current model, as it allows us to account for these and other findings that, on the
155 surface, appear problematic (e.g. [S9]). Indeed, the main implication from the foregoing
156 analysis is that some caution is advisable when making inferences about the salience of a
157 cue on the basis of its associability.

158

159 S4. Parameters in the model

160 V and \bar{V} : In simulations of the model, an asymptote at 1 was imposed on the
161 cue→reinforcer (V) and cue→no-reinforcer (\bar{V}) associations. This was done in order to
162 prevent the run-away growth they would otherwise undergo in partial reinforcement
163 schedules. Unbound increments in the strength of these associations is a consequence of
164 the fact that they both contribute to the prediction errors used to adjust their values.

165 ϵ : Although we must remain agnostic as to the nature of the function f in Equations 1.2
166 and 1.3, for the purpose of simulating the model we took it to be the identity function.

167 This means that, in our simulations, ε is simply the sum of associative strengths of the
168 cue. In partial reinforcement, therefore, ε equals the strength of the cue→reinforcer and
169 cue→no-reinforcer associations ($V + \bar{V}$). Because we capped the asymptotes of V and \bar{V}
170 at 1, the maximum value that ε can take on in our simulations of partial reinforcement is
171 2.

172 α : allowing ε to take on values greater than 1 implies that α can in principle also be
173 greater than 1 (see Equation 5). In order to keep the product of our learning rate
174 parameters ($\alpha \times \beta$) between 0 and 1 in Equations 3.1 and 3.2, we stipulated for β to be in
175 the range $0 < \beta < 0.1$. Alternatively, we could have normalized α by dividing its actual
176 value by the maximum possible value that α could take in that specific training situation.
177 Although this solution may appear more elegant, notice that it involves adjusting the
178 normalization coefficient for each specific training situation, since the maximum possible
179 value for ε in the denominator will vary depending on the number of reinforcer (and no-
180 reinforcer) representations involved.

181 β : Because in partial reinforcement each of the cue→reinforcer and cue→no-reinforcer
182 associations can either strengthen or weaken on a given trial, 4 types of changes in
183 associative strength may take place as a result of this training: 1) an increment or 2)
184 decrement in the cue→reinforcer association, and 3) an increment or 4) decrement in the
185 cue→no-reinforcer association. As noted above, we acknowledge the possibility that
186 these changes might occur at different rates, which prompts the question of how to go
187 about selecting the 4 corresponding β values (see Equations 3.1 and 3.2). To begin to
188 address this question, we set out to find the combinations of β values that satisfied what

189 we regard to be the central theoretical problem faced by the model. This problem,
190 succinctly, is the fact that partial reinforcement, relative to continuous reinforcement,
191 endows the cue with higher salience *if no better predictors of the trial outcomes are*
192 *present*, but lower salience *if better predictors of the trial outcomes are present* [for a
193 direct comparison, see 14]. For example, if a cue X is partially reinforced with food
194 (X→food, X→nothing), then evidence suggests that as a result of this training the
195 salience of X should be greater than that of a cue A that is continuously reinforced
196 (A→food). In the simulations discussed below, we refer to this pattern of results as
197 Condition 1. By contrast, if partial reinforcement with X is embedded in a true
198 discrimination of the form AX→food, BX→nothing, where the outcome of each trial is
199 unambiguously predicted by A and B, then evidence suggests that the salience of X
200 should now be less than that of the continuously reinforced A. We shall refer to this
201 pattern of results as Condition 2.

202 In order to explore the β -parameter space and identify the β combinations that
203 satisfy Conditions 1 and 2, we simulated the above scenarios across multiple iterations
204 that systematically varied the values of the 4 β s (β values were restricted to 0.01, 0.03,
205 0.05, and 0.07). Out of the 256 possible β combinations, 114 ($\approx 45\%$) were found to
206 satisfy Condition 1; i.e., these combinations allowed the model to predict that the salience
207 of X following X→food, X→nothing training is greater than that of A following
208 A→food training. A lawful relation across these 114 combinations could be identified:
209 the product of the 2 β s for increments in the cue→food and cue→no-food associations
210 was always greater than the product of the 2 β s for decrements in these associations².
211 This relation was further investigated with a wide range of β combinations ($0 < \beta < 0.1$),

212 which confirmed that only when this product rule is violated is the salience of A higher
213 than that of X.

214 It is straightforward to see why this should be the case. According to the model,
215 the acquired salience of X (i.e. ϵ) results from the sum of the X→food and X→no-food
216 associations. If these associations are eroded during training through substantial
217 extinction, then their combined values will not overcome that of the A→food association,
218 and consequently the salience of X will be lower than that of A. In the opposite extreme,
219 if little extinction of these associations is allowed to occur, then the model erroneously
220 predicts that, following training of the kind AX→food, BX→nothing, the salience of X
221 will be higher than that of A. Such was the case with 33 β combinations that violated
222 Condition 2 ($\approx 13\%$ of the initial 256). In all of them, the sum of the X→food and
223 X→no-food associations overcame the value of the A→food association. We proceeded
224 therefore to eliminate these 33 cases in order to obtain the list of 81 β combinations that
225 satisfied Conditions 1 and 2 ($\approx 32\%$ of the initial 256).

226 A final, refining step was taken, which consisted in testing the remaining 81
227 combinations across a number of simulations of salience-related problems investigated in
228 the Pearce and Haselgrove labs [14, S12]. Fourteen combinations were then selected for
229 providing the best fit for our data, and one of these was used in the simulations presented
230 in Figure 1. The novel predictions generated from the model were also based on this
231 particular set of parameters.

232

233

234

Supplementary notes

235

1. That the influence of salience on associability is orthogonal to that of other factors, such as appetitive-aversive interactions, is suggested by more recent experiments from the Le Pelley lab [S10]. They showed that cues that are highly predictive of an affectively neutral outcome subsequently enter into associations with either *positive or negative* outcomes more readily than previously nonpredictive cues.

236

237

238

239

240

241

2. It is notable that existing theories of associative learning (e.g. [32]) adopt the assumption that the acquisition of associative strength proceeds more rapidly than the loss of associative strength, and that Rescorla [S11] has provided evidence that the acquisition of conditioned responding proceeds more readily than the extinction of conditioned responding.

242

243

244

245

246

247
248

Supplementary references

- 249 S1 Le Pelley, M. E., Oakeshott, S. M., Wills, A. J. & McLaren, I. P. L. 2005 The
250 outcome specificity of learned predictiveness effects: parallels between human casual
251 learning and animal conditioning. *J. Exp. Psychol.: Anim. Beh. Proc.* **31**, 226–236.
- 252 S2 Rescorla, R. A. & Furrow, D. R. 1977 Stimulus similarity as a determinant of
253 Pavlovian conditioning. *J. Exp. Psychol.: Anim. Beh. Proc.* **3**, 203–215.
- 254 S3 Garcia, J. & Koelling, R. A. 1966 Relation of cue to consequence in avoidance
255 learning. *Psychon. Sci.* **4**, 123–124.
- 256 S4 Rozin, P. & Kalat, J. W. 1971 Specific hungers and poison avoidance as adaptive
257 specializations of learning. *Psychol. Rev.* **78**, 459–486.
- 258 S5 Pearce, J. M. 1987 A model of stimulus generalization for Pavlovian
259 conditioning. *Psychol. Rev.* **94**, 61–73.
- 260 S6 Dickinson, A. & Pearce, J.M. 1977 Inhibitory interactions between appetitive and
261 aversive stimuli. *Psychol. Bull.* **84**, 690–711.
- 262 S7 Konorski, J. & Szwejkowska. G. 1956 Reciprocal transformations of
263 heterogeneous conditioned reflexes. *Ada Biologiae Experimentalis* **17**, 141–165.
- 264 S8 Scavio, M. J. 1974 Classical-classical transfer: Effects of prior aversive
265 conditioning upon appetitive conditioning in rabbits. *J. Comp. Physiol. Psychol.* **86**, 107–
266 115.
- 267 S9 Dickinson, A. & Mackintosh, N. J. 1979 Reinforcer specificity in the
268 enhancement of conditioning by posttrial surprise. *J. Exp. Psychol.: Anim. Beh. Proc.* **5**,
269 162–177.

270 S10 Le Pelley, M. E., Reimers, S. J., Beesley, T., Spears, R., Murphy, R. A. &
271 Calvini, G. 2010 Stereotype formation: Biased by association. *J. Exp. Psychol.: Gen.* **139**,
272 138–161.

273 S11 Rescorla, R. A. 2002 Comparison of the rates of associative change during
274 acquisition and extinction. *J. Exp. Psychol.: Anim. Beh. Proc.* **28**, 406–415.

275 S12 Pearce, J. M., Esber, G. R., George, D. N., & Haselgrove, M. 2008 The nature of
276 discrimination learning in pigeons. *Learn. Behav.* **36**, 188–199.

277

278