

Categorization and the Ratio Rule

A. J. Wills (ajw43@cam.ac.uk)

Mark Suret (mbs22@cam.ac.uk)

I. P. L. McLaren (iplm2@cus.cam.ac.uk)

Department of Experimental Psychology, University of Cambridge,
Downing St., Cambridge. CB2 3EB. U.K.

Abstract

Many formal models of categorization adhere to two basic principles. First, the extent to which a stimulus is subjectively characteristic of a particular category can be represented by a single number. Second, the probability with which people choose a particular category label for a stimulus can be derived from these numbers via the Ratio Rule a.k.a. the Luce choice axiom (Luce, 1959). A categorization experiment employing artificial visual stimuli is presented and is shown to be problematic for these two principles. We demonstrate that, for the data presented here, the first principle can be retained if one replaces the Ratio Rule with a simple connectionist model.

Introduction

Category learning is the task of acquiring the correct category label for each of a set of presented stimuli. The ability to categorize is central to cognition, and it has been the subject of a large number of studies. Over the last thirty years, these studies have typically involved abstract stimuli grouped into categories not necessarily definable in terms of a simple rule (e.g. Homa, Sterling, & Trepel, 1981; Medin & Schaffer, 1978; Posner & Keele, 1968). Psychologists have proposed a variety of formal models of our ability to learn and make decisions about such categories. The models differ in many respects - for example, the Generalized Context Model (Nosofsky, 1986) proposes the memorization of presented examples, whilst a number of other theories propose the formation of feature-category associations (e.g. Gluck & Bower, 1988; Kruschke, 1996; McClelland & Rumelhart, 1985). Despite such diversity, a great many theorists seem to agree on two fundamental principles. First, the extent to which a stimulus is subjectively characteristic of a particular category can be represented by a single number. We will refer to such numbers as *category magnitude terms*. Second, the probability with which a participant decides that a stimulus belongs to a particular category is determined by the Ratio Rule, a.k.a. Luce's Choice Axiom (Luce, 1959). In the current context, the Ratio Rule can be stated

$$P(i) = \frac{v_i}{\sum_{j=1}^n v_j}$$

where $P(i)$ is the probability of choosing category i from n alternative categories and v_j is the category magnitude term for the j th alternative.

Theorists seldom justify their adoption these principles. Of greater concern is the fact that, as far as we are aware, there have been no direct tests of the Ratio Rule in the context of categorization. The evidence for the Ratio Rule, such as it is, comes from pair-comparison experiments and identification experiments (Bradley, 1954; Clarke, 1957; Hopkins, 1954). The evidence provided by such studies is equivocal at best, and some studies provide direct evidence against the Ratio Rule (e.g. Burke and Zinnes, 1965; Lamming, 1977).

Previously, we had made an unsuccessful attempt to disprove the Ratio Rule in the context of two-choice categorization decisions (Jones, Wills, & McLaren, 1998). The reason for this might have been that the Ratio Rule was basically correct for categorization decisions, but we suspected that it was because the predictions made by the Ratio Rule in a two-choice situation tend to be numerically close to the predictions of a number alternative accounts. Therefore in the experiment described here we tested a property more characteristic of the Ratio Rule - its predictions about probability ratios.

A long appreciated feature of the Ratio Rule is that it predicts that the ratio in which two alternatives are chosen is unaffected by the addition of a third alternative. For example, in a taste preference test between Coke and Pepsi, participants might choose Coke with a probability of 0.8. The Ratio Rule predicts that whilst the addition of lemonade might change the probability with which either Coke or Pepsi is chosen, it does not change the 4:1 ratio of probabilities.

The ratio we concentrate on in the current study is directly related to this property. It is the ratio between 1) the probability with which a particular response is made to a stimulus when three category labels are available and 2) the probability with which the same response is made to an equivalent stimulus when only two of the labels are available. For example, let's call the three labels A, B, and C, and say that A is the option which is disallowed in the two-choice example. Under the assumption that category magnitude terms for allowed alternatives are not affected by the number of alternatives available, it can be shown that the Ratio Rule predicts

$$\frac{P(B : B, C)}{P(B : A, B, C)} = \frac{v_A}{v_B + v_C} + 1 \quad \text{Measure 1}$$

This is not much use as it stands because we do not have any direct way of measuring the category magnitude terms, and different theories of categorization do not generally agree on how one might estimate the terms from observable data. The utility of Measure 1 lies in the fact that the Ratio Rule's predictions for the probability with which category A is chosen (when it is allowed) are similar in form. Specifically

$$P(A : A, B, C) = \frac{v_A}{v_A + v_B + v_C} \quad \text{Measure 2}$$

This correspondence means that, in situations where v_a is constant, any given change in $(v_B + v_C)$ will produce the same direction of change in these two measures. One no longer needs to know what values the magnitude terms take. Instead one just needs to set up a situation where it is reasonable to assume v_A is constant across a set of stimuli. Then, to the extent different stimuli result in different values of $P(A:A,B,C)$ and the 2 choice to 3 choice ratio (Measure 1), these differences must be in the same direction for both measures if the Ratio Rule is correct.

A number of similar correspondences can be set up, but we employ just one further here. Consider a second set of stimuli which are comparable to the first, except in their relative similarity to one of the three categories. In this situation it may be reasonable to assume that the magnitude terms for these two sets of stimuli differ only in respect to that category. Taking the category on which they differ as A, and the two magnitude terms as v_A and $v_{A'}$, the ratio of probabilities with which category B (or C) is chosen in response to these otherwise comparable stimuli is

$$\frac{P(B : A', B, C)}{P(B : A, B, C)} = \frac{v_A + v_B + v_C}{v_{A'} + v_B + v_C} \quad \text{Measure 3}$$

Note that in a situation where the two magnitude terms for category A can be assumed to be constant, and $v_A > v_{A'}$ this third measure must exhibit the same direction of change as the other two. We investigated whether all three measures do indeed show the same direction of change in the context of a simple categorization task.

Experiment

The experiment had two phases. In the training phase, all participants learned about the category membership of a set of novel, artificial stimuli. Each training stimulus belonged to one of three categories - A, B or C. In the transfer phase which followed participants were asked to determine the category membership of a set of test stimuli. Some participants were allowed to respond A, B or C, whilst for others the option A was disallowed. The stimuli presented in the transfer phase were designed to vary smoothly from being characteristic of category B and uncharacteristic of category C through to being characteristic of C and uncharacteristic of B. They were designed in this way so that (hopefully) the

three measures we were interested in comparing would be relatively smooth functions of the number of category B (or category C) elements. If reliable functions were found for our measures and these functions all exhibited the same direction of change then we would have evidence in support of the Ratio Rule. If the functions found exhibited different directions of change then this would be strong evidence against the Ratio Rule.

Participants with three response alternatives were presented with one of two sets of test stimuli. The members within a set were designed to be equally characteristic of category A. For one set they were somewhat characteristic of category A, whilst for the other they were uncharacteristic. All participants with two response alternatives received the test stimuli somewhat characteristic of A.

Method

Participants and Apparatus. 36 Cambridge University students participated. They were tested individually in a quiet cubicle on a Acorn Risc PC microcomputer with a 14" color monitor. The computer's screen was at eye level, approximately 90 cm directly in front of where the participant sat. Responses were recorded via the "X", "B" and "M" keys of a standard PC keyboard. For this experiment the keys were re-labeled "A", "B" and "C" using bold red letters against a white background.



Figure 1: An example stimulus.

Stimuli. Each stimulus was a collection of twelve different small pictures (hereafter elements), arranged on an invisible four-by-three grid inside a 4.5 cm by 3.5 cm rectangle outline (see Figure 1 for an example). Every stimulus contained twelve elements drawn from a pool of 40 that we have used in a number of previous experiments (see Jones et al., 1998). No stimulus contained more than one copy of any particular element. At the beginning of the experiment, and separately for each subject, 12 elements from the pool were randomly designated as category A elements, a different 12 as category B elements, and a different 12 again as category C elements. The remaining four elements were designated as novel elements and were not employed in the training phase. Each training stimulus for each category was constructed by starting with all 12 elements characteristic of that category (e.g. category A elements for a category A training stimulus). Each element in the training stimulus then underwent a 10% chance of being replaced by a randomly chosen element from one of the other two sets (e.g. replaced by a B or C element in the case of a category A training stimulus). It was these modified stimuli that were

presented to subjects as training stimuli. This procedure produces training examples which are composed predominantly of elements characteristic of a particular category but also exhibit considerable variability in terms of the specific elements they contain. Ninety training examples were created for each subject, thirty from each of the three categories.

Participants received one of two sets of test stimuli - a *familiar-elements* set or a *novel-elements* set. Each stimulus in a familiar-elements set contained four A elements, x B elements and $(8-x)$ C elements where x could be 0, 1, 2, 3, 4, 5, 6, 7 or 8. Ten examples of each of these nine types of test stimulus were created for each participant receiving a familiar-elements test set. The specific elements used to create each test stimulus were chosen randomly within the constraints provided by the number of A, B and C elements the stimulus was to contain. Ten examples of each of four dummy stimuli were also created, these stimuli being (8 A, 0 B, 4 C), (8 A, 4 B, 0 C), (0 A, 4 B, 8 C) and (0 A, 8 B, 4 C). The purpose of the dummy stimuli was to obscure from the participants that all test stimuli of interest (from the perspective of the experimenter) were constant in terms of the number of elements from category A they contained. Stimuli in the novel-element test sets were constructed in the same manner as familiar-element stimuli, except that the four novel elements (see above) were used instead of four randomly selected A elements.

The position of elements within a stimulus was randomly determined for each stimulus presented, with the constraint that exactly one element occurred at each location in the four-by-three grid. Where stimuli were accompanied by a category label, this was presented as a large sans-serif capital A, B or C in an outline rectangle (4.5 by 3.5 cm) immediately to the right of the stimulus itself.

Procedure. Participants were allocated to one of three groups such that an equal number (12) participated in each. The three groups, referred to hereafter as the *two-choice*, *three-choice* and *novel-elements* groups, differed in the number of response alternatives available in the test phase and the stimuli presented during the test phase.

The training phase was the same for all participants. After some general instructions the ninety training stimuli were presented sequentially and in a random order. Each training stimulus was presented for five seconds in the center of the monitor, accompanied by the appropriate category label. Two seconds of plain mid-gray mask in the stimulus and label rectangles preceded the next example. Participants were not required to respond during the training phase. They were simply asked to concentrate on the examples shown as they would later be asked to classify new, unlabelled examples. This training procedure had proved effective in a number of previous experiments (Jones et al., 1998; Wills & McLaren, 1997).

The training phase was followed by a test phase. There were 130 stimuli in the test phase (90 target stimuli and 40 dummy stimuli) which, again, were presented sequentially and in a random order. Test stimuli were not accompanied by a category label. Participants in the two-choice and three-choice conditions received a familiar-elements test set whilst participants in the novel-elements condition received a

novel-elements test set (see *Stimuli*). On the presentation of each test stimulus, participants were asked a question. Participants in the two-choice condition were asked "Is this a B or a C?". Participants in the three-choice and novel-elements conditions were asked "Is this an A, a B or a C?". In all conditions they responded by pressing the appropriate key on the computer keyboard. They then pressed the "Y" key, whereupon the next stimulus was immediately presented. There was no time limit for these decisions, and participants were put under no pressure to respond quickly.

The allocation of the category labels "A", "B", and "C" to the logical categories A, B and C was counter-balanced.

Results

Figure 2a shows the probability with which participants responded with the category A label to test stimuli (Measure 2) as a function of the number of category B elements they contained (the conclusions of this study are unaffected if one plots against category C elements instead). The functions for the three-choice and novel elements conditions both appeared to show an inverted-U trend. The significant fit of a second-order polynomial to the nine mean data points confirmed this appearance for the three-choice condition, $F(2, 6) = 5.6$, $p < 0.05$, but not for the novel-elements condition, $F(2, 6) = 3.2$, $p > 0.1$. The quadratic co-efficient for the three-choice condition was significantly different from zero, $b^2 = -0.006$, $t(7) = 2.4$, $p < 0.05$.

The data points in Figure 2b are the average of the probability with which participants responded with their category B label to stimuli with x category B elements and the probability with which they responded with their category C label to test stimuli with x category C elements. In other words, it shows response probability as a function of the number of *category-appropriate* elements. Averaging these two probabilities is appropriate because, across subjects, there is no factor that determines which of the two categories providing variable numbers of elements to test stimuli should be described as category B and which as category C. A replication of this experiment with non-counterbalanced category labels failed to reveal any significant response bias. Number of category-appropriate elements in Figure 2 reduces from left to right in order to follow the convention that generalization functions (such as those shown in Figure 2b) are plotted as slopes with negative gradients.

For our current purposes it is not the data presented in Figure 2b which are of central interest, but the ratios calculated from the mean points it displays (Measures 1 and 3). These ratios are presented as a function of category-appropriate elements in Figure 2c. Inspection of this figure shows that the 2 choice to 3 choice ratios (Measure 1) appear to exhibit an increasing, accelerating trend whilst the 3 choice novel-elements to 3 choice ratios (Measure 3) exhibit a decreasing, accelerating trend. The significant fit of a second-order polynomial to the nine points of Measure 1, $F(2, 6) = 803$, $p < .0005$, with a best-fit line for which all three co-efficients were significantly different from zero, $b^2 = 0.049$, $t(7) = 14$, $p < .005$; $b = -0.674$, $t(7) = 24$, $p < .0005$; $a = 3.48$,




Figure 2: **a)** Probability of producing a category A response. **b)** Mean response probability (see text). **c)** Two ratios calculated from the data in Figure 2b. Plot symbols = empirical data. Lines = predictions of the winner-take-all model presented in the Modeling section.

$t(7) = 2.5, p < .05$, supports this conclusion. The nine points of Measure 3 were also a significant fit to a second-order polynomial, $F(2, 6) = 17, p < .005$, and all three coefficients differed significantly from zero, $b^2 = -0.021, t(7) = 3.0, p < .05$; $b = 0.244, t(7) = 4.3, p < .005$; $a = 0.632, t(7) = 3.8, p < 0.01$.

Discussion

Our results pose two central problems for the Ratio Rule as it is currently employed in many formal models of categorization.

First, the Ratio Rule predicts that the two ratios represented in Measures 1 and 3 should show the same direction of change over any interval of category-appropriate elements. However, the best-fitting quadratics for the corresponding functions show opposite directions of change (Figure 2c). Second, the Ratio Rule predicts that the probability of choosing category A in the three-choice condition (Measure 2) should show the same direction of change over any interval of category-appropriate elements as the other two measures. However, the best-fitting quadratics for measures 1 and 2 are of opposite shape (compare Figure 2a with Figure 2c). One might argue that these findings are of relatively little consequence because the discrepancies are in derived measures with no straightforward psychological interpretation, rather than in the response probabilities themselves. Such a position is disingenuous. The predictions under test arise naturally and unavoidably from a central (some would say defining) feature of the Ratio Rule - the fact that the ratio in which two alternatives are chosen is unaffected by the addition of a third alternative. These data provide evidence against that central tenet and hence bring the formulation into question.

If any one step in a chain of inferences is incorrect then the conclusions drawn from that process must be brought into question. Consequently, theoretical conclusions about the nature of categorization must be re-examined if our conclusion is found to be generally valid. Conversely, if the assumptions we have made in coming to our conclusions can be shown to be invalid then the Ratio Rule is not necessarily incorrect. Below we briefly consider some possible criticisms of our conclusion.

First, one could argue that we have disproved the Ratio Rule for means across participants, but this does not disprove the formulation for individual participants. This is a valid point, but as most formal theories of categorization have been applied to group means our conclusion still stands for these theories. Second, it is true that our stimuli are rather more complex than those typically used in category learning experiments. It may be the case that our results do not generalize to simpler stimuli, or that our stimuli are unusual in some other way. This seems to be an empirical matter, and one which is worthy of investigation. A third, substantial criticism is that we have assumed that category magnitude terms are, for our stimuli, univariate functions of the number of category-appropriate elements the stimulus contains (i.e. the magnitude term is determinable solely from this property). There are at least two distinct ways in which this assumption could be incorrect.

First, for specific models of categorization it may be possible to show that the category magnitude term for category A is not invariant under changes in the magnitude terms for categories B and C. For example, one might be able to demonstrate for the GCM model (Nosofsky, 1986) that the test stimuli were not at a fixed distance (in psychological similarity space) from category A examples. The difficulty here is that the procedure which Nosofsky uses to derive the psychological similarity space assumes that the Ratio Rule is correct. Some way around this circularity would have to be devised.

Second, one could quite reasonably argue that category magnitude terms are importantly affected by what response alternatives are available (as a number of theorists outside of the categorization literature have argued e.g. Restle, 1961; Tversky, 1972). If this were the case in our experiment then the derivation of Measure 1 would be invalid because it is directly based on this assumption.

Therefore one response to our results might be to retain the Ratio Rule but introduce a mechanism by which category magnitude terms can be affected by the alternatives available for decision. However, for most formal models of categorization this would require considerable revision of the basic principles upon which they were based. We wondered whether there was a direct replacement for the Ratio Rule that could accommodate our results without having to modify the rest of the theory.

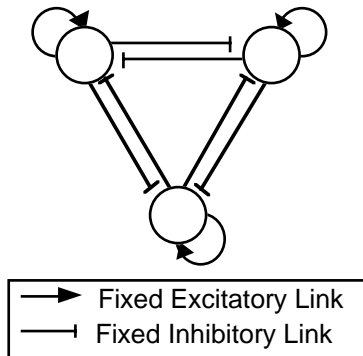


Figure 3: The winner-take-all model.

Modeling

Previously we have proposed that response probabilities in categorization might be modeled by a simple winner-take-all connectionist system employing category magnitude terms as input activations (Wills & McLaren, 1997). Such a system is illustrated in Figure 3. In addition to the magnitude-term inputs, each unit has a fixed excitatory connection to itself and fixed inhibitory connections to the other units. These connections can cause the units to “compete” with one another until only one has a non-zero activation. In our system a decision is deemed to have been made when the highest activation exceeds its nearest competitor by some threshold value, S . This general architecture has been proposed previously by Grossberg (1976) amongst others, and has been employed in the modeling of a number of other

psychological phenomena (e.g. Houghton, 1990; Usher & McClelland, 1995).

For the purposes of this simulation we assume that category magnitude terms are defined by the function $v = 0.047c + 0.012$, where c is the number of category-appropriate elements the stimulus contains. The exact form of this equation is not critical. It was chosen because it describes the behavior of a simple localist delta-rule network with a learning rate of 0.0025. This learning rate was previously found to be successful in modeling the rate of learning in similar experiments (Wills & McLaren, 1997). The important thing to note is that we are preserving the assumption that category magnitude terms are independent of the response alternatives available.

The magnitude term input activations (r) are assumed to be noisy and, for simplicity, this noise is assumed to be rectangular, have a mean of zero, and a range from $-N$ to $+N$. Magnitude input activations are also constrained to lie between 0 and 1. The specific shape of the noise distribution is not critical and similar mean behavior could be produced with a Gaussian distribution. The output activations of the units are governed by the equations

$$o = \frac{o + En}{1 + En + D}, \text{ if } n > 0 \text{ and } o = \frac{o + En}{1 - En + D} \text{ otherwise,}$$

where n is the total input the unit, and E and D are excitation-rate and decay-rate constants respectively. These are standard activation equations with properties similar to those used by, for example, McClelland & Rumelhart (1985). Output activations in our model are constrained to be non-negative. Total input (n) for a given unit is the sum of r and o for that unit, minus the sum of the outputs (o) of the other units. For the current simulation $E = 0.2$, $D = 0.1$ and $N = 1.1$. The threshold parameter S was set to 0.18 for the two-choice condition, 0.65 for the three-choice condition and 0.72 for the novel-elements condition.

In the two-choice condition of our experiment, participants were not allowed to make category A responses. In our WTA model this was simulated by fixing the output activation of the category A unit at zero.

The results of our simulation are shown as lines in Figure 2. Note that the model respects all the major trends in the experiment and is numerically close to the observed data. A detailed discussion of the principles underlying the success of this model is not possible here, but it is important to note that the exact details of the implementation are not critical. Indeed, not even the expression in connectionist terms is essential. The model simply provides a mechanism by which a decision similar in principle to Thurstonian choice (Thurstone, 1927) can be made. We have demonstrated in other analyses that simply choosing the noisy alternative which is instantaneously the biggest does reasonably well in predicting the trends in Measures 1 and 2 (whether one employs Gaussian or rectangular noise).

However, only the connectionist system correctly predicts the trend in Measure 3. This is because it employs different decision thresholds in the three-choice and novel-elements conditions, which allows it to predict that Measure 3 falls below unity without having to make the counter-intuitive

assumption that the category A magnitude term for a stimulus containing no category A elements is greater than for a stimulus containing four category A elements (the only way a simple Thurstonian choice process could predict ratios smaller than one).

Conclusion

The Ratio Rule as generally applied in formal models of categorization was shown to be incorrect for the experiment presented. Whilst further investigation is necessary, we suggest that our results may indicate a need to replace the Ratio Rule as currently employed with an alternative system (perhaps still based around the Ratio Rule). One possibility would be to substantially revise existing models so that they provided a mechanism by which category magnitude terms could be affected by the alternatives available for decision. We have shown that our results do not require that this modification be made. Rather, one simply needs to directly substitute the Ratio Rule with a decision mechanism based on the principles of Thurstonian choice. The noise employed in this mechanism may have one of a number of distributions. As a caveat, one distribution it is unlikely to have is a double exponential distribution because this would make it indistinguishable from the Ratio Rule (Yellott, 1977).

Acknowledgments

The authors would like to thank Fergal Jones, Koen Lamberts, Donald Laming and Thomas Palmeri for their helpful comments. Thanks are also due to Stian Reimers and Neil Stewart who helped out with similar studies for which there was insufficient space in this article. This research was supported by a grant from the ESRC to I.P.L. McLaren.

References

- Bradley, R. A. (1954). Incomplete block rank analysis: On the appropriateness of the model for a method of paired comparison. *Biometrics*, *10*, 375-390.
- Burke, C. J. & Zinnes, J. L. (1965). A paired comparison of pair comparisons. *Journal of Mathematical Psychology*, *2*, 53-76.
- Clarke, F. R. (1957). Constant-ratio rule for confusion matrices in speech communication. *Journal of the Acoustical Society of America*, *29*(6), 715-720.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal Of Experimental Psychology: General*, *117*(3), 227-247.
- Grossberg, S. (1976). Adaptive pattern classification and universal recoding: Part I. Parallel development and coding of neural feature detectors. *Biological Cybernetics*, *23*, 121-134.
- Homa, D., Sterling, S., & Trepel, L. (1981). Limitations of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 418-439.
- Hopkins, J. W. (1954). Incomplete block rank analysis: Some taste test results. *Biometrics*, *10*, 391-399.
- Houghton, G. (1990). The problem of serial order: A neural network model of sequence learning and recall. In R. Dale, C. Mellish, & M. Zock (Eds.), *Current Research in Natural Language Generation*. London: Academic Press.
- Jones, F. W., Wills, A. J., & McLaren, I. P. L. (1998). Perceptual categorization: Connectionist modelling and decision rules. *The Quarterly Journal of Experimental Psychology*, *51B*(3), 33-58.
- Kruschke, J. K. (1996). Base rates in category learning. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *22*(1), 3-26.
- Laming, D. (1977). Luce's choice axiom compared with choice-reaction data. *British Journal of Mathematical and Statistical Psychology*, *30*, 141-153.
- Luce, R. D. (1959). *Individual Choice Behavior*. New York: John Wiley & Sons.
- McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal Of Experimental Psychology: General*, *114*(2), 159-188.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*(3), 207-238.
- Nosofsky, R. M. (1986). Attention, similarity and the identification-categorisation relationship. *Journal Of Experimental Psychology: General*, *115*(1), 39-57.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*(3), 353-363.
- Restle, F. (1961). *Psychology of judgement and choice*. New York: Wiley.
- Thurstone, L. L. (1927). A law of comparative judgement. *Psychological Review*, *34*, 273-286.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*(4), 281-299.
- Usher, M., & McClelland, J. L. (1995). *On the time course of perceptual choice: A model based on principles of neural computation*. (Technical Report PDP.CNS.95.5): Carnegie Mellon University.
- Wills, A. J., & McLaren, I. P. L. (1997). Generalization in human category learning: A connectionist explanation of differences in gradient after discriminative and non-discriminative training. *The Quarterly Journal of Experimental Psychology*, *50A*(3), 607-630.
- Yellott, J. I., Jr. (1977). The relationship between Luce's choice axiom, Thurstone's theory of comparative judgement, and the double exponential distribution. *Journal of Mathematical Psychology*, *15*, 109-144.

From 1st September 2000, A.J.Wills should be contacted at: University of Exeter, School of Psychology, Washington Singer Laboratories, Perry Rd., Exeter. EX4 4QG. United Kingdom. <http://www.ex.ac.uk/Psychology/>