

Is Competitive Learning an Adequate Account of Free Classification?

Jan Zwickel (jzwickel@ix.urz.uni-hd.de)
Department of Psychology, Hauptstraße 47-51
69117 Heidelberg, Germany

A.J. Wills (a.j.wills@ex.ac.uk)
School of Psychology, University of Exeter
Perry Road, Exeter. EX4 4QG. England.

Abstract

Rumelhart & Zipser's (1986) competitive learning algorithm is an account of unsupervised learning and, as such, might be considered a potential model of free classification behavior in humans. However, selective learning effects (e.g. Dickinson, Shanks & Evenden, 1984) suggest that human learning, at least under conditions of feedback, may be better characterized by an error-correcting system. An experiment is reported that provides preliminary evidence for the existence of a selective learning effect in free classification. Simulations indicate that Rumelhart & Zipser's algorithm does not provide an adequate account of the behavior observed, whilst an error-correcting variant of competitive learning does.

Introduction

Free classification, or free sorting as it is also called, is a procedure in which human participants are presented with a set of stimuli and are asked to group them in any way that seems sensible or reasonable to them (e.g. Bersted, Brown & Evans, 1969; Regehr & Brooks, 1995; Wills & McLaren, 1998). It may be contrasted with the more standard experimental task of category learning via trial-specific feedback that has been the dominant mode of enquiry into humans' categorization abilities for the last fifty years (e.g. Bruner, Goodnow & Austin, 1956; Medin & Schaffer, 1978; Wills, Reimers, Stewart, Suret & McLaren, 2000).

The study of categorization under conditions where each decision receives immediate feedback from a totally reliable source has allowed psychologists great control over the structure of the categories participants acquire. As a methodology, it has been successful in broadening our understanding of the category learning process. However, the level of feedback available in such tasks seems higher than that available in many real-world situations, begging the question of whether what we have learned about the categorization process will generalize to situations where the feedback is absent or scarce.

An interesting parallel may be drawn with the sort of connectionist systems that have been proposed for learning in the presence or absence of feedback. For

example, Rumelhart & Zipser's (1986) competitive learning model is an unsupervised system. It extracts statistical regularities in the input to form categorical representations, and does so in the absence of feedback. In contrast, McClelland & Rumelhart's (1985) model is a supervised system. It can be taught multiple categories (cat vs. dog vs. bagel in their example) but learns to categorize because each stimulus is accompanied by an externally-provided category label.

One of the differences between these two models is the nature of the weight-update algorithms they employ. McClelland & Rumelhart (1985) employ an error-correcting algorithm, where the size of the weight change is proportional to the mismatch between an external teaching signal and internal inputs. In other words, learning only occurs when the system fails to fully predict the teaching signal. Specifically,

$$\Delta w_{ij} = \eta(e_i - i_i)a_j \quad 1$$

where Δw_{ij} is the change in the strength of the connection from unit j to unit i , e_i is the external teaching signal to unit i , a_j is the activity of unit j , and i_i is the total internal input to unit i , this being calculated as

$$i_i = \sum_j a_j w_{ij} \quad 2$$

In contrast, Rumelhart & Zipser's algorithm does not employ error-correction in this sense. Rumelhart & Zipser use the internal input to determine which unit is the "winner" and then change weights to the winning unit by an amount proportional to the difference between the current weight of that connection and an asymptote¹. Specifically, the change in weight from unit j to the winning unit is

¹ Rumelhart & Zipser (1986) also discuss a variant where connections to the losing unit are also changed via Equation 3, but with a much lower learning rate. The current article concentrates on the "winner-only" version, although the conclusions drawn are valid for both variants.

$$\Delta w_j = \eta \left(\frac{a_j}{n} - w_j \right) \quad 3$$

where n is the number of active input units, and the winning unit is the one with the highest internal input. It is assumed in the current paper that active input units have an activity of 1 and inactive input units have an activity of zero.

Error-correction is assumed by some investigators to be a fundamental aspect of human learning in the presence of feedback, as evidenced by the phenomenon of selective learning (see below). If human learning in a free classification task fails to show evidence of selective learning, concerns would arise as to the generality of an empirical research program heavily based on learning with feedback. On the other hand, if selective learning is found to occur in free classification, the sort of unsupervised system proposed by Rumelhart & Zipser may not be an appropriate model for free classification behavior.

Selective learning

Probably the best-known example of selective learning is Kamin's (1969) "blocking" effect. Kamin's study involved rats but, as will be discussed later, there is now abundant evidence that corresponding effects can also be found in human learning (with feedback).

Kamin taught hungry rats that pressing a lever would result in food. Following this, pressing the same lever whilst a noise was present resulted in a mild electric shock. Unsurprisingly, rats learned to not press the lever whilst the noise was present.

Next, the auditory tone was accompanied by a light and pressing the lever whilst this tone-light compound was present also resulted in mild shock. The rats learned not to press the lever whilst the tone-light compound was present.

Group	Stage One	Stage Two	Test
Expt.	N→Shock	LN→Shock	L
Ctrl.		LN→Shock	L

Figure 1: Kamin's (1969) blocking experiment. "N" is an auditory stimulus and "L" is a visual stimulus.

In the test phase, just the light was presented, and the rats' behavior observed. The rats in the experimental condition pressed the lever quite a lot, whilst control rats (which had participated in stage two but not stage one) pressed the lever very little indeed. The design of this experiment is summarized in Figure 1.

The rats in the experimental group appear not to have learned the relationship between light and shock even though the control rats, which received an equal amount

of training with the light-noise compound, have learned the relationship. Why might this be?

A number of animal learning theorists (e.g. Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972) have essentially argued that it happens because learning is driven by surprise. For the rats in the experimental group, the shock is not a particularly surprising event because it is predicted by the noise. The rats therefore don't bother to learn about the light in stage two. Similar effects have been demonstrated with undergraduates using computer-based tasks, including simple computer games (e.g. Dickinson, Shanks & Evenden, 1984), stock market simulations (e.g. Chapman & Robbins, 1990) and simulated medical diagnosis tasks (e.g. Gluck & Bower, 1988).

As Rescorla and Wagner (1972) noted, the notion of surprise-driven learning is well-captured by an error-correcting learning rule such as the one given in Equations 1 and 2. In fact, the learning theory they proposed is basically a variant of this learning rule.

An error-correcting system reproduces the blocking effect in the following way. For simplicity, consider that there are three units - the noise unit, the light unit and the shock unit. The external input produced by the presence of a stimulus is 1, and the external input produced by its absence is zero.

Initially, shock is not expected, so the link between noise and shock, and light and shock, are small. During stage one of the experiment, the strength of the link between noise and shock increases and eventually reaches 1. In the second stage, noise, light and shock are all present. However, the internal input to the shock unit is already 1 because of the strength of the link between the noise and the shock. Therefore, no weight change can occur via Equation 1. The light does not become associated with shock, even though it clearly would in the control condition. Under the non-error-correcting algorithm given in Equation 3, the light→shock association would reach an equivalent level in both conditions.

Experiment

It is reasonably clear from previous research that humans and other animals engage in selective learning under conditions of trial-specific, informative feedback. Do humans also display selective learning in a task without such feedback? The experiment reported in this paper represents a first attempt to address this question.

In our experiment, participants had to make up their own categories, although they were constrained by the fact they were only allowed two groups. Previous research demonstrates that category learning can proceed successfully in the absence of feedback (e.g. Homa & Cultice, 1984; Wills & McLaren, 1998).

Our participants received intermittent, non-trial-specific, feedback about their overall level of performance following every 24 stimuli, in order to maintain motivation and encourage adoption of the experimenter-defined categories. We believe that such a procedure is still properly described as "free classification" as no single response can be considered correct or incorrect. Situations where all forms of feedback are entirely absent are probably almost as rare outside the laboratory as situations where feedback is always immediate and trial-specific.

Abstract, novel stimuli were employed in this experiment as we wished to study category learning with adult participants - with such participants the category learning process is probably complete or far-advanced for most realistic stimuli.

Stimulus presentation was brief and followed by a mid-gray mask. The time available for a decision was also very limited. Both of these procedures were employed to encourage participants to rely on relatively non-analytic, similarity-based categorization processes, rather than analytic, rule-based processes.

The basic design of the experiment is shown in Figure 2. The letters A to J each represent sets of features present in the stimuli shown to participants.

In the first phase of the experiment, participants were presented with examples from two different categories. Examples from category 1 were created from a base pattern that contained feature sets G and H. Examples from category 2 were created from a base pattern that contained feature sets I and J. Note that the labels "category 1" and "category 2" are essentially arbitrary in a free classification task - they could be reversed without changing anything in the design or execution of the experiment.

As Figure 2 illustrates, once the participant had mastered the GH vs. IJ categorization they were transferred to a second categorization. Participants proceeded through all five categorizations in this way, at which point the experiment was over.

The datum of central importance in this design is the category to which the first stimulus presented in phase five is allocated. The first stimulus is chosen because subsequent decisions in phase five may be contaminated by learning on previous phase five trials.

Phase	1	2	3	4	5
Cat. 1	GH	GE	AB	AE	CE
Cat. 2	IJ	IF	CD	CF	AF

Figure 2: Design of the experiment. Letters represent sets of features, hence category 2 in phase 3 contains feature sets C and D.

The Rumelhart & Zipser system provides the null hypothesis for this experiment because it predicts that

either key is equally likely to be used. It is perhaps not immediately apparent why this should be. To elucidate, one first needs to note that in each of the first four phases, all features are equally predictive of category membership. This means that, for a system such as Rumelhart & Zipser, in each phase learning should end with two features (A and E in phase four) being equally associated to one category representation, and two features (C and F in phase four) being equally associated to the other category representation. Hence, the first stimulus presented in phase five will activate both category representations equally and so the choice of which category to place it into must be arbitrary. This conclusion is confirmed by simulation in a later section.

Why might one expect anything other than a null result with this design? One possible reason would be if people exhibited selective learning in free classification. Note that, across phases 1 to 4, E and F only occur in situations where the information they provide is partially redundant. In phase 2 the stimuli can be categorized on the basis of whether they contain G or I features, a categorization already learned in phase one. In phase four, the stimuli can be categorized on the basis of whether they contain A or C features, a categorization already learned in phase three. Hence, through analogy to selective learning effects in tasks with feedback, one might consider that E and F develop little control over responding.

Method

Due to space limitations, we are unable to report the pilot studies performed. Reports may be found in McCooe(2000) and Zwicker (2001).

Participants and apparatus

Sixteen first-year Psychology students from the University of Exeter participated to fulfil a course requirement. Participants were tested in groups in a quiet computer room. Stimulus presentation was on 17" color monitors connected to Tiny Pentium III PCs running the DMDX software package (Forster & Forster, 2000, version 2). Responses were collected via the left and right CTRL keys on standard PC keyboards. Participants sat approximately 50cm from the screen.

Stimuli

Each stimulus was made up of 12 small pictures (hereafter "elements") taken from a set of 72 that have been used in a number of previous experiments (see Jones, Wills & McLaren, 1998 for the full set). For any given stimulus, the 12 elements were randomly arranged in a square of 3 rows with 4 icons in each row, and were surrounded by a gray rectangle outline 5.5cm in height and 4cm in width. Figure 3 shows an example

stimulus. Throughout all five phases, no stimulus contained more than one copy of any given element.




Figure 3: An example stimulus

Each of the letters A to J in Figure 2 represent a set of six elements. The assignment of elements to letters was randomly determined for each of 8 pairs of participants, with the remaining 12 elements (72 elements - 10 letter sets x 6 elements per set) being used for practice trials.

In order to control for possible effects of differential salience of the elements, one participant in each pair received the stimuli described in Figure 2 whilst the other received a design where E was transposed with A, and F was transposed with C. Hence, the putatively redundant elements were E and F for one member of the participant pair, whilst they were A and C for the other member. This means that any preference revealed in phase five cannot be due to A and C elements being more salient overall than E and F elements. To aid clarity, all participant data is reported as if E and F were the putatively redundant elements.

The stimuli actually presented to participants were generated by random distortion of the base patterns described in Figure 2. Each element in a base pattern was given a 10% chance of being replaced by a randomly selected element from the other base pattern (no element occurred more than once in any given stimulus).

An example may be helpful. To create an AF stimulus in phase five, the six A elements and the six F elements were randomly arranged in the four-by-three grid of the stimulus. Each element was then given a 10% chance of being replaced by a randomly selected element from set C or E. This method of stimulus construction produces training examples which are composed predominately of elements characteristic of a particular category but which also exhibit considerable variability.

Procedure

The five phases described in Figure 2 were preceded by some general written instructions and a brief practice phase to familiarize participants with the procedure. The experiment then proceeded in blocks of 24 trials.

On each trial, a stimulus was presented for 800ms and followed by a mid-gray mask that was presented for 1200ms. If a response was not detected within 2000ms of stimulus onset, the trial terminated with the message "You responded too slowly, please speed up!" and the participant was moved on to the next trial.

Each block comprised the sequential presentation of 24 stimuli, 12 from each of the two categories. At the end of each block a short message appeared stating the percentage of correct responses made by the participant in that block, and that they needed to score more than 80% to proceed to the next part of the experiment.

Clearly, percent correct has a slightly different interpretation in a free classification task to a task with trial-specific feedback as the relationship between categories 1 and 2 and the two response keys is arbitrary. Hence, percent correct was computed under the assumption that category 1 would receive a particular response, and the resulting number was subtracted from 100 if it was less than 50.

When a participant's score exceeded 80% they were moved on to the next phase of the experiment, after having been presented with the message "You did very well! You are now entering the next phase". If participants completed 10 blocks without ever reaching the 80% criterion they were moved on to the next phase with the message "You are entering the next phase as you have been in the last block of this phase".

Results

Consider Figure 2 again. The central null hypothesis we are attempting to reject is that, in the first trial of phase five, a participant will be no more likely to categorize AF using the response typically made to AE in phase four than the response typically made to CF in phase four. Similarly, they will be no more likely to categorize CE using their typical AE response than their typical CF response.

Of the 16 participants tested, 12 used the same response key for CE that they had typically used for CF (or the AE response key for an AF stimulus). Three participants showed the opposite response, using the CF key for an AF stimulus or the AE key for a CE stimulus. The remaining participant could not be described as having a preference for any in key in response to AE or CF as they scored exactly 50% across phase four. Treating this participant in the manner that makes it hardest to reject the null hypothesis, we can state that at least 12 participants emitted the CF→CE (or AE→AF) response, whilst no more than 4 participants emitted the opposite response. Given the null hypothesis would predict 8 responses of each type, the probability of the null hypothesis being correct is smaller than 0.05, $\chi^2(1) = 4.0$. The effect is also significant with an exact binomial test.

Participants completed a mean of 5.88 blocks in phase one, 4.13 blocks in phase two, 6.12 blocks in phase three, 5.75 blocks in phase four, and 5.56 blocks in phase five. The number of participants failing to achieve more than 80% correct in the five phases were 6, 3, 7, 7 and 7 respectively.

Discussion

The results of the current experiment appear to be problematic for those that would attempt to explain free classification behavior in terms of the competitive learning algorithm of Rumelhart & Zipser (1986). The model predicts no preference for which of the two categories developed in phase four are used to categorize the first stimulus in phase five, yet a clear preference was observed. The direction of the preference is that predicted if one assumes the presence of selective learning in free classification

One possible defense of the Rumelhart & Zipser algorithm is that its predictions were derived for a situation where learning in each phase is essentially complete before the next phase begins. Given the relatively high numbers of participants failing to reach criterion, it might reasonably be argued that asymptotic predictions are not appropriate. Does this make a difference? This is one of the questions addressed in the following section.

Modeling

We employed simulation techniques to more thoroughly investigate whether Rumelhart & Zipser's (1986) competitive learning algorithm could accurately reproduce the categorization preference observed in our experiment. To this end we set up a network with 72 input units (one for each element) and 2 output units (one for each category). Each input unit had a forward connection to each output unit, and the connection weights were initialized to small, random values.

One network simulation was performed for each participant in the experiment, with weights being initialized for each participant. The nature of the stimuli presented to a simulated participant, and the order in which they were presented, were determined by the specific stimuli presented to a corresponding human participant. After the presentation of each stimulus, the winning category node was determined in the same manner as Rumelhart & Zipser (1986). In other words, it was determined by calculating the total internal input to each unit, and selecting the unit with the larger total. The weights from each of the input units to the winning category unit were then updated in accordance with Equation 3. The weights of the losing unit remained unchanged.

The value of η (the learning rate) employed by Rumelhart & Zipser was 0.05. At this value, no preference in the categorization response to the first stimulus in phase five was found. Six simulated participants made CF→CE or AF→AE responses whilst six made the opposite response. The nature of the response made by four simulated participants could not be determined because in phase four they employed both category nodes equally for both stimulus types.

Hence, unlike the human participants, the networks did not display a categorization preference in phase five.

The Rumelhart & Zipser (1986) algorithm was applied to our data with a wide range of learning rates (0.001 to 0.009 in steps of 0.001, 0.01 to 0.09 in steps of 0.01, and 0.1 to 0.9 in steps of 0.1). In no case did the algorithm display a categorization preference in phase 5.

An error-correcting competitive algorithm

We also attempted to simulate our result using an algorithm that combined the error-correcting principle of Equations 1 and 2 with the basic properties of the competitive learning algorithm of Equation 3. On any trial, the winning unit was determined in the same manner as the Rumelhart & Zipser model. The weight-update algorithm employed on each connection from an input unit j to the winning unit was

$$\Delta w_j = \eta \frac{(1-i)}{n} \quad 4$$

for connections from active input units and

$$\Delta w_j = -\eta(1-i) \times \frac{n}{m} \quad 5$$

for connections from inactive input units. In these equations, η is the learning rate, i is the total internal input to the winning unit, n is the number of active input units and m is the number of inactive input units. The weights from input units to the losing unit are not changed. This chimera of an algorithm is not equivalent in behavior to either of its components but does preserve some of the properties of each.

Removing the weight update algorithm of Equation 3 from our previous simulation, and replacing it with the algorithm described in Equations 4 and 5, we find a dramatic change in behavior. Now, at a learning rate of 0.05, all 16 simulated participants make CF→CE or AF→AE responses. In other words, the simulation now reproduces the behavior observed in our human participants, although the overall level of learning is slightly higher in our simulation. A reliable preference is found for a wide range of learning rates - from 0.01 to about 0.4.

Conclusion

The experiment reported in this paper provides preliminary evidence that the ubiquity of selective learning effects in tasks with immediate, trial-specific feedback extends to some categorization tasks where feedback is scarce and not trial-specific. To the extent this phenomenon is found to be general to free

classification tasks, it casts some doubt on the adequacy of certain types of competitive learning algorithms as accounts of free classification behavior. In particular, an algorithm suggested by Rumelhart & Zipser (1986) was found to have difficulty in reproducing the results found. We suggest that a competitive algorithm which includes some aspect of error-correction may be a more appropriate account. One simple algorithm of this type was described, tested, and found to be able to reproduce our results.

The two main avenues of future research suggested by the results and simulations in this paper are a) investigation of the generality of selective learning effects in free classification, b) consideration of whether other unsupervised systems (e.g. Adaptive Resonance Theory, Grossberg, 1976) are capable of accounting for the results so far found.

References

- Bersted, C. T., Brown, B. R., & Evans, S. H. (1969). Free sorting with stimuli in a multidimensional attribute space. *Perception & Psychophysics*, 6B, 409-413.
- Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A Study of Thinking*. New York: Wiley.
- Chapman, G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition*, 18, 537-545.
- Dickinson, A., Shanks, D., & Evenden, J. (1984). Judgement of act-outcome contingency: The role of selective attribution. *The Quarterly Journal of Experimental Psychology*, 36A, 29-50.
- Forster, K. I., & Forster, J. C. (2000). DMDX version 2. Retrieved from the internet on 12/2000 from <http://www.u.arizona.edu/~jforster/dmdx.htm>.
- 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. & Cultice, J. (1984). Role of Feedback, Category Size, and Stimulus Distortion on the Acquisition and Utilization of Ill-Defined Categories. *Journal of Experimental Psychology: Learning Memory, and Cognition*, 10, 83-94.
- 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.
- Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. Campbell & R. Church (Eds.) *Punishment and aversive behavior*. New York: Appleton.
- 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.
- McCoee, M. (2000). *Learning without being taught* (unpublished project report). Cambridge: Dept. of Experimental Psychology, University of Cambridge.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276-298.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85(3), 207-238.
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87, 532-552.
- Regehr, G., & Brooks, L. R. (1995). Category organisation in free classification: The organising effect of an array of stimuli. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 21(2), 347-363.
- 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* (pp. 64 - 99). New York: Appleton-Century-Crofts.
- Rumelhart, D. E. & Zipser, D. (1986). Feature Discovery by Competitive Learning. In D. E. Rumelhart, J. L. McClelland and the PDP Research Group (Eds.) *Parallel Distributed Processing. Explorations in the Microstructure of Cognition. Volume 1: Foundations*. Cambridge, London: MIT Press.
- Wills, A. J., & McLaren, I. P. L. (1998). Perceptual learning and free classification. *The Quarterly Journal of Experimental Psychology*, 51B(3), 235-270.
- Wills, A. J., Reimers, S., Stewart, N., Suret, M., & McLaren, I. P. L. (2000). Tests of the ratio rule in categorization. *The Quarterly Journal of Experimental Psychology*, 53A(4), 983-1011.
- Zwickel, J. (2001). *Categorisation without feedback* (unpublished report). Exeter: School of Psychology, University of Exeter.

Acknowledgments

This research was supported by a research fund provided by the University of Exeter. The authors would like to thank Manfred Amelang, Meg Davies, Mike McCoee and Joachim Funke for their practical assistance, and Fergal Jones, Timm Lochmann, Ian McLaren and Katrin Scharpf for their insightful comments.