Primacy Effects and Selective Attention in Incremental Clustering

David M. Thau
University of Michigan, Ann Arbor
330 Packard Rd.
Ann Arbor, Michigan 48104
Phone: (313) 764-0318
E-mail: thau@engin.umich.edu

Abstract

Incremental clustering is a type of categorization in which learning is unsupervised and changes to category structure occur gradually. While there has been little psychological research on this subject, several computational models for incremental clustering have been constructed. Although these models provide a good fit to data provided by some psychological studies, they overlook the importance of selective attention in incremental clustering. This paper compares the performance of two models, Anderson's (1990) rational model of categorization, and Fisher's (1987) COBWEB, to that of human subjects in a task which stresses the importance of selective attention. In the study, subjects were shown a series of pictorial stimuli in one of two orders. The results showed that subjects focussed their attention on the first extreme feature they saw, and later used this feature to classify ambiguous stimuli. Both models fail to predict human performance. These results indicate the need for a selective attention mechanism in incremental clustering as well as provide one constraint on how such a mechanism might work.

Introduction

Imagine trying to acclimate yourself to a city you have never visited before. As you wander through the streets, you may begin to notice similarities and differences between the styles of some of the houses. Each new house may remind you of a few others, leading you to group them together. Eventually, you may form fairly well defined categories. The process through which these categories are devised is called incremental clustering. Incremental clustering may be characterized by two qualities. First, learning is unsupervised. In the example above, the houses were divided into categories without feedback from a teacher. Second, changes to the category representation are made incrementally Each new exemplar is incorporated into an already existing category structure. This is in contrast to non-incremental categorization, in which the entire category structure is reconsidered whenever a new exemplar is encountered.

Surprisingly, incremental clustering has received little attention from the cognitive psychology community. While both supervised, incremental category learning (e.g. Posner & Keele, 1968, Smith & Medin, 1981), and unsupervised, non-incremental category learning (e.g. Ahn & Medin, in press, Bersted, Brown & Evans, 1969) have been studied in detail, there have been few experiments on unsupervised, incremental category learning (Fried & Holyoak, 1984, Homa & Cultice, 1984).

In the machine learning literature, on the other hand, incremental clustering has received a good deal of attention. The combinatoric explosions that result in computer systems that try to organize categories in a non-incremental fashion have lead machine learning researchers to study incremental learning. This, coupled with the need for systems that learn without constant and consistent feedback, has lead to several models of incremental clustering. I will briefly describe two of the more recent computational models of incremental clustering, Anderson's (1990) rational model of categorization, and Fisher's (1987) COBWEB model. These descriptions will be followed by a study that demonstrates a flaw shared by these models.

The first model I will describe is Anderson's rational model of categorization. Anderson provides a Bayesian analysis of category structure goodness. When presented with a new stimulus, the model calculates the goodness of the whole category scheme for each possible categorization of the new stimulus. For example, if the model has already constructed three categories, it determines the goodness of four different category structures: one structure for when the new item is placed in each of the already existing categories, and one structure for when a new category containing only the new item is created.

Although the actual formula Anderson uses to determine category structure goodness is not important for the purposes of this paper, it is important to note that goodness is determined by feature counts within each category. Information

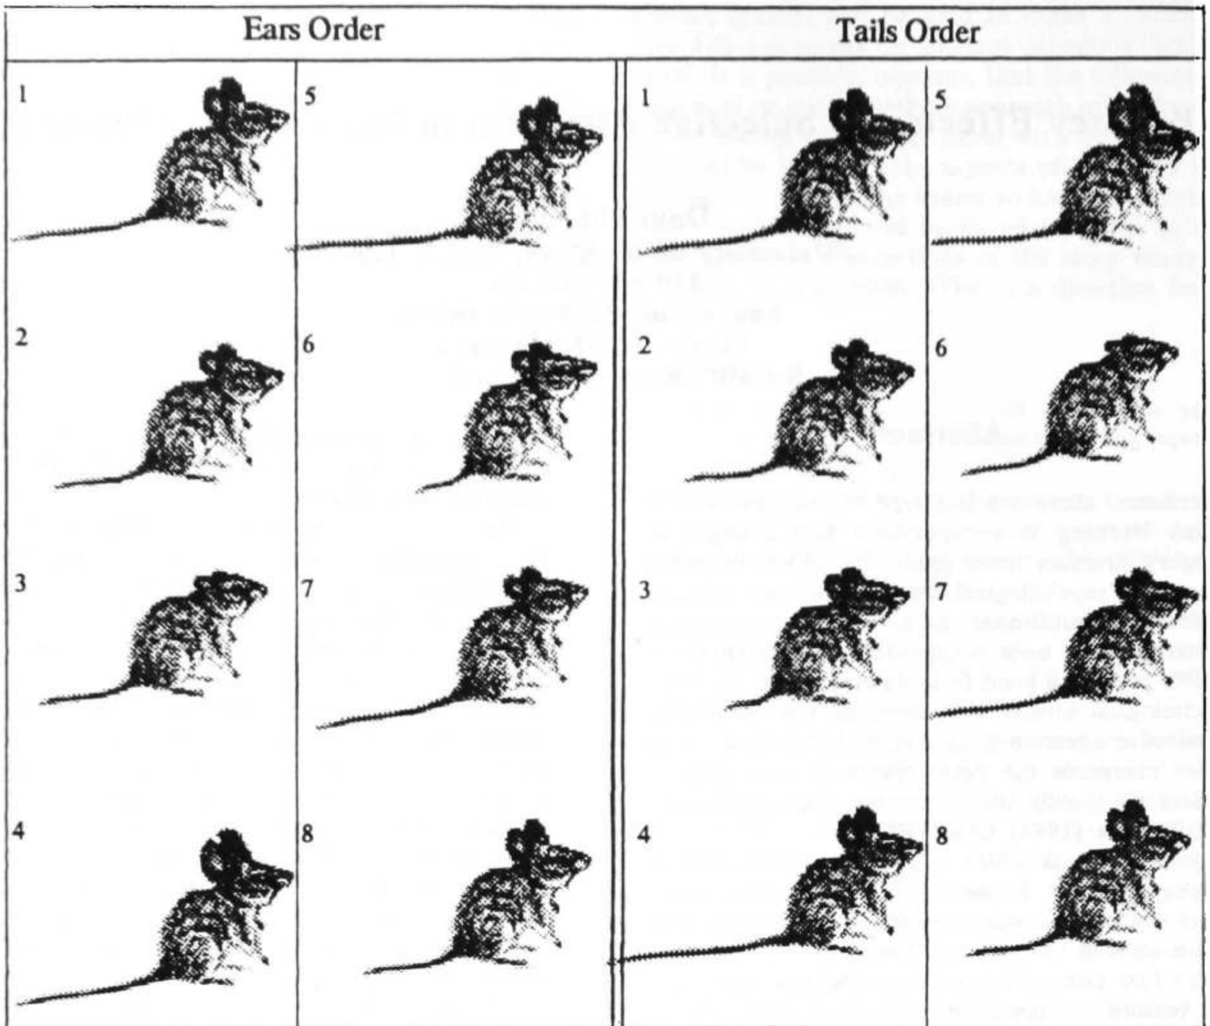

Figure 1: The two orders of the stimuli. The first order is meant to stress ears, the second is meant to stress tails. The stimuli have been reduced to approximately 30% of their actual size.

about the order in which the items were classified is lost. Information about why an item was classified the way it was is also lost. Another interesting limitation of Anderson's model is that it cannot change its partitionings. Once an item is classified, it cannot be reclassified unless it is seen again

While Fisher's COBWEB does allow for reclassification of stimuli, and uses a different category structure goodness function, it is in other ways very similar to Anderson's model. When COBWEB is shown a new stimulus, it, like Anderson's model, considers the goodness of placing the item in each of its categories, or a new category. In addition, COBWEB considers either merging the two best categories, and placing the item in the merged category, or splitting the single best category into two, and placing the item into one of them. Although the model can reclassify items, it shares the information reduction limitations of Anderson's model. All information other than the present membership of each category is lost.

People, however, may represent information beyond category membership. They may, for example, remember why they categorized an exemplar the way they did. Consider two people who have sorted a set of items into the same categories, but have done so for different reasons. Although these different reasons may not have manifested themselves yet, there may be a point in the future where these two people will react differently to a new exemplar. Consider, on the other hand, two runs of COBWEB or Anderson's model. Because these models only represent the current membership of each category, two runs which have formed the same categories will categorize a new exemplar the same way.

The simulations and experiment that follow will demonstrate that this limitation of the models is indeed a problem. Eight pictures of mice with different sized tails and ears were used as stimuli. The mice were actually eight pictures of the same mouse, scanned into a Macintosh computer, and modified using a graphics program. In six out of the eight mice, the ear size and tail size were positively correlated; mice with big ears also had big tails, mice with small ears had small tails. Of these six, four had one relatively extreme feature, either ears that were very big or very small, or a tail that was very big or very small. In the remaining two mice, ear size and tail length were negatively correlated; one mouse had big ears and a small tail, the other had small ears and a big tail. See Figure 1 for a graphical representation of the stimuli.

The mice were put into two orders, each of which is represented in the Figure. The first two mice (mice 1 and 2) were the same in both orders. These were the two mice that maintained a positive correlation between ear size and tail length, but had moderately sized features. The next four mice were either ordered such that the two mice with the extreme ear sizes came next, or the two mice with the extreme tail sizes came next. The next two mice in each order were the other two mice with extreme features. The final mice (mice 7 and 8) were those in which the features violated the positive correlation. One of these mice had a long tail but short ears, the other had a short tail and long ears.

The stimuli were ordered to induce human subjects to pay attention to ears in one condition and tails in the other. The hypothesis was that subjects would attend to the first extreme feature they saw and then focus most of their attention on that feature throughout the sorting. During the sorting of the first six mice, however, this attention weighting has no effect on performance. Because the features of these mice were positively correlated, the sorting will be the same regardless of which feature was more important; mice with big ears and big tails will be sorted into one category, mice with small ears and small tails will be sorted into the other. In terms of the models described above, both orders will produce the same category structure.

The test of the models comes during the sorting of the final two mice. Because the features of these mice violate the positive correlation, the way in which they are sorted provides important information about the subject's sorting strategy. If the subject thinks ear size is more important, he or she will put the mouse with small ears, but a large tail, into the category of mice with small ears and small tails. If tail size is more important, the subject will put the same mouse into the other category. However, the two models can not account for this result. Because the models base their sortings entirely on the current category structure, the order by which that structure was created has no effect. Therefore, because the two orderings produce the same categories, the different ordering of the stimuli should have no effect on the models' sortings.

The preceding observation was supported by simulations using Fisher's and Anderson's models[1]. Both models took two features as input, the size of each mouse's ears (in cm^2) and tail (in cm). The stimuli were presented to the models in the two different orders described above. As can be seen in Table 1, neither Fisher's COBWEB, nor Anderson's rational model showed any effect of order. COBWEB[2] sorted the last two mice according to tail length regardless of the order. Anderson's model[3], which outputs the probability of sorting a stimulus into each possible category, provided the same probabilities regardless of the order.

In summary, neither model was affected by the different orderings of the stimuli. The remainder of the paper compares these results with those of human subjects.

Method

Thirty-five University of Michigan undergraduates participated in the study as part of an introductory psychology course requirement. The subjects, who were tested individually, were told that they would be shown pictures of different mice. They were told that their task was to sort the mice into two different kinds, but that it was up to them to decide how to divide them. They were also told that at any point they could reclassify any of the mice. This reclassification was permitted for two reasons. First, in early stages of categorization, reclassification should be expected (Fried and Holyoak, 1984). Second, the models being tested both take into account the need for reclassification. Fisher's COBWEB model explicitly allows reclassification through its merging and dividing operations. Anderson's rational model requires that the mean and variance of each feature of the stimuli be predetermined, reducing the amount of reclassification necessary.

The experiment proceeded with the experimenter presenting the mice to the subjects one by one. After each mouse was shown, subjects classified it by verbally responding either A or B. The mouse was then placed in front of the subject in a way that allowed subjects to see how each mouse had been classified. Subjects were permitted to see all their

[1]Versions of both models were kindly provided by their authors.

[2]Instead of COBWEB, Fisher's CLASSIT program was used. CLASSIT is a version of COBWEB which allows for features with real values. COBWEB only allows nominal values. Acuity in this simulation was set to 0.5.

[3]The coupling parameter in these runs was set to 0.3.

Table 1
Percentage of Simulated Subjects Who Sorted by Ear Size, Tail Length, or Something Else

Fisher's COBWEB

Order	Sorted by: Ears	Tail	Other
Ears	0	100	0
Tail	0	100	0

Anderson's Rational Model

Order	Sorted by: Ears	Tail	Other
Ears	27	23	50
Tail	27	23	50

classifications in order to diminish reliance on memory. The models of incremental clustering being tested both assume that the system has full memory of the stimuli it has seen. If subjects were not allowed to see the mice they had classified, they would have been working with less information than the models being tested. After all the mice from one order had been presented and classified, the experimenter asked the subject if he or she was satisfied with the sorting, and then asked the subject to describe the categories that were formed.

Results

Table 2 shows the percentage of subjects in each order condition who sorted the mice by ear size, tail length, or something else. As described in the introduction, the stimuli were constructed such that one sorting clearly indicated that the subject was sorting by tail length, and another sorting clearly indicated that the subject was sorting by ear size. Only six of the 35 subjects provided sortings different from the two that were anticipated. These subjects are represented in the 'other' column of Table 2. Most of the subjects who fell into this condition put a mouse that had an extreme feature (such as the mouse with the smallest ears) into one category and the rest of the mice in the other category.

The manipulation clearly had the expected effect, $X^2(2, N=35) = 10.6, p < 0.005$. Subjects who were presented with the stimuli that were ordered to emphasize ear size did in fact sort by ear size. Those who were presented with the stimuli that were ordered to emphasize tail length were more likely to sort by tail length.

Discussion and Conclusion

The results summarized above indicate that models of incremental clustering need to take the role of selective attention into account. Neither Fisher's COBWEB, nor Anderson's rational model provide mechanisms by which different features can become

Table 2
Percentage of Subjects in Each Order Who Sorted by Ear Size, Tail Length, or Something Else

Order	Sorted by: Ears	Tail	Other
Ears	72	0	28
Tail	41	53	6

more or less important in the midst of a categorization task apart from the straight accumulation of features. Although both models can preset the salience of the different features, neither can change feature weights on-line.

The results further provide one constraint on how a selective attention mechanism should work. Subjects in this task focussed on the first feature that clearly differentiated between two categories and later used this feature when classifying ambiguous examples.[4]

How this feature weighting occurs in people has yet to be determined. One simple explanation for the results of this study is that once the subjects found a salient feature along which to classify, they ignored all other features. If, for example, ear size seemed like a diagnostic feature, a subject could simply look at the ears of each stimulus and ignore the other features. This mechanism might be implemented as a rule-like system (if big ears, then category A), which would obviate the need for the Anderson and Fisher models once a satisfactory feature was discovered.

Although this approach to selective attention would fit the results of this study, it is unlikely that subjects completely ignore all but the most diagnostic features. If this were true, people could not adapt to changing circumstances. In addition, there is evidence (Medin, Wattenmaker and Michalski, 1987) that people include redundant information when devising classification rules. An alternative approach would involve learning feature weights across dimensions. While there has been a great deal of work on feature weighting models when feedback is immediate, few models (cf. Gennari, 1991, Kohonen, 1982, Rumelhart and Zipser, 1985, Grossberg, 1987) have been developed which apply to unsupervised learning. Future research will involve applying these models to the present task, and extending them or positing new models where necessary.

In conclusion, current models of incremental clustering must be extended to take into account on-line learning of feature weighting. Although the mechanisms involved in this weighting are still uncertain, this study has provided one constraint; salient differences between features are weighted more heavily when they occur in early examples.

---

[4]Recently, Gennari (1991) described CLASSWEB, an extension of COBWEB that includes a selective attention mechanism. CLASSWEB's selective attention mechanism, however, mainly acts to focus attention away from irrelevant features. In the study described here, both features are relevant. Consequently, CLASSWEB does not predict the results provided by human subjects. Instead, in this situation, CLASSWEB behaves exactly like COBWEB.

Acknowledgements

I would like to thank Douglas Medin for his help throughout all phases of this study. I would also like to thank Woo-Kyoung Ahn, Evan Heit and Joushua Rubinstein for reading an earlier draft of this paper.

References

Ahn, W-K., & Medin, D. L. (In Press). A two stage model of free sorting. *Cognitive Science.*

Anderson, J. R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Erlbaum.

Bersted, C. T., Brown, B. R., & Evans, S. H. (1969). Free sorting with stimuli clustered in a multidimensional space. *Perception & Psychophysics, 6,* 409-414.

Fisher, D. H. (1987). Knowledge acquisition via incremental conceptual clustering. *Machine Learning, 2,* 139-172.

Fried, L. S. & Holyoak, K. J. (1984). Induction of category distributions: A framework for classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 234-257.

Gennari, J. H. (1991). Concept formation and attention. Proceedings of the Thirteenth Annual Congference of the Cognitive Science Society (pp. 724-728). Chicago, Illinois: Earlbaum.

Grossberg, S. (1987). Competitive learning: from interactive activation to adaptative resonance. *Cognitive Science, 11,* 23-63.

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.

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics, 43,* 59-60.

Medin, D. L., Wattenmaker, W. D., & Michalski, R. S. (1987). Constraints and preferences in inductive learning: An experimental study of human and machine performance. *Cognitive Science, 11,* 299-339.

Posner, M. L., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 83,* 304-308.

Rumelhart, D. E. & Zipser, D. (1985). Feature discovery by competitive learning. *Cognitive Science, 9,* 75-112.

Smith, E., & Medin, D. L. (1981). *Categories and concepts.* Cambridge: Harvard University Press.