The effect of pre-exposure on family resemblance categorization for stimuli of varying levels of perceptual difficulty.

Fraser Milton (f.n.milton@exeter.ac.uk) [1], Edward Copestake[1], David Satherley[1], Tobias Stevens[1], & Andy J. Wills[2]

[1]School of Psychology, College of Life and Environmental Sciences, University of Exeter, United Kingdom
[2]School of Psychology, University of Plymouth, United Kingdom

Abstract

This study investigated the effect that pre-exposure to a set of stimuli has on the prevalence of family resemblance categorization. 64 participants were tested to examine the effect that pre-exposure type (same-stimuli vs unrelated-stimuli) and the perceptual difficulty of the stimuli (perceptually similar vs perceptually different) has on categorization strategy. There was a significant effect of perceptual difficulty, indicating that perceptually different stimuli evoked a higher level of family resemblance sorting than perceptually similar stimuli. There was no significant main effect of pre-exposure type; however, there was a significant interaction between pre-exposure type and level of perceptual difficulty. Post-hoc tests revealed that this interaction was the result of an increase in family resemblance sorting for the perceptually different stimuli under relevant pre-exposure but no such effect for perceptually similar stimuli. The theoretical implications of these findings are discussed.

Keywords: free classification; family resemblance; unidimensional; perceptual learning, match-to-standards.

Introduction

Categorization is a fundamental cognitive mechanism that enables us to function effectively in our everyday environment. In particular, it allows us to make inferences about objects that we have never seen before and to treat different objects in the same way, greatly simplifying the environment that we live in. However, in view of the immense number of objects we encounter, this process must necessarily be highly constrained. This is illustrated by the fact that just ten items can be partitioned in more than 100,000 different ways. Thus, a greater knowledge of how we acquire the categories we have is an important pre-requisite for our understanding of human cognition.

One reasonable assumption is that the categories we prefer to create would reflect the underlying structure of objects we encounter outside the laboratory. One early prominent theory - the “classical” view - posited that categories are organised around necessary and jointly sufficient features. According to this theory, if an item has the necessary feature (or features) it can be considered a member of that category regardless of the characteristics of its other features. However, many natural categories do not appear to have such a structure and the seminal work of Rosch and colleagues highlighted the idea that instead they are often organized around “family resemblance” relations (e.g., Rosch & Mervis, 1975), in which categories possess a number of characteristic but not defining features. Under a family resemblance structure, an object does not have to possess any particular feature (or features) but can be considered a member of that category if it possesses enough characteristic features. In other words, a family resemblance structure is organized around overall similarity relations.

It is therefore surprising that despite the plausibility of a family resemblance theory of natural categories, previous work has shown that when people are asked to free classify stimuli (i.e., are given no feedback on their responses) they find it far from natural to sort by family resemblance. In fact, people have a strong preference to free classify unidimensionally (i.e., on the basis of a single feature, e.g. Ahn & Medin, 1992; Ashby, Queller, & Berretty, 1999; Medin, Wattenmaker, & Hampson, 1987) – an approach that seems more consistent with the classical view. Whilst manipulations of the method of stimulus presentation (Regehr & Brooks, 1995), the level of spatial separability of the stimulus dimensions (Milton & Wills, 2004; Milton, Viika, Henderson, & Wills, 2011), the perceptual difficulty of the stimuli (Milton & Wills, 2008), the structure of the categories (Pothos & Close, 2008), time pressure (Milton, Longmore, & Wills, 2008), instructions (Wills, Milton, & Longmore, 2013), and background knowledge (Spalding & Murphy, 1996) have all been shown to influence the extent of family resemblance categorization, such sorting is still far from common. One important question, therefore, is to understand why the categories we prefer to create do not reflect the commonly assumed underlying structure of natural world categories.

One explanation for this anomaly is that participants generally have not seen any of the stimuli prior to classification. This appears atypical of categorization outside the lab where we usually have had a great deal of exposure to the objects we categorize. One consequence of asking participants to sort a set of stimuli they have never seen before may be to encourage them to fall back on a simplistic, unidimensional, strategy as they have had little experience of the stimuli. One possibility, then, is that if participants receive substantial pre-exposure to the stimuli prior to classification they may find a family resemblance response a more intuitive and easier strategy to perform.

Previous work has, in general, used response accuracy to measure the effects of pre-exposure on behavior (e.g., McLaren, 1997). Typically such studies have shown that pre-exposure has a beneficial effect on response accuracy (e.g., McLaren, Leevers, & Mackintosh, 1994). Whilst such work is undoubtedly important, a related question that is arguably just as interesting is the extent to which pre-exposure can actually change the *nature* of the categories that we create. It is, therefore, surprising that there is currently a paucity of research that has addressed this issue. One exception to this is the work of Wills and McLaren (1998) which used a free classification procedure to show that pre-exposure can influence the number of categories people use. The current work extends this finding by investigating the hypothesis that pre-exposure may facilitate family resemblance sorting.

A large body of work has shown that pre-exposure improves stimulus differentiation (e.g., McLaren et al., 1994) and this may be one mechanism by which pre-exposure could facilitate family resemblance sorting. This increased differentiation (which may facilitate discrimination of the values of a particular dimension and also allow easier extraction of the relevant dimensions themselves) should make family resemblance sorting an easier and more viable option as the differences between the various items will be more apparent as, perhaps, will be the dimensional inter-correlation. This may, for example, facilitate the use of a multi-dimensional rule (c.f., Wills et al., 2013). This assumption receives support from recent work which indicates that perceptually different stimuli lead to greater levels of family resemblance sorting than more perceptually similar stimuli (Milton & Wills, 2008).

Previous work indicates that perceptual learning is most pronounced for perceptually similar stimuli (e.g., Oswalt, 1972) and that pre-exposure can even inhibit learning if the stimuli are sufficiently different (e.g., Chamizo & Mackintosh, 1989). One model that can explain this pattern of findings is the MKM model (McLaren, Kaye, & Mackintosh, 1989). The MKM account assumes that stimuli are represented by a number of elements. Items that share many common elements will be more difficult to discriminate than items that share few elements. One of the key assumptions of this model is that when elements co-occur, there will be a reduction in the salience of these elements (due to what is known as latent inhibition). As a consequence, one of the principal effects of pre-exposure is that elements which frequently co-occur reduce in salience more quickly than elements that rarely co-occur. This means that the *unique* elements that discriminate one stimulus from another will tend to be higher in salience than the *common* elements that both stimuli share (because the common elements will have been presented more often). This effect is likely to be greater for items that are perceptually similar because they share many common elements and hence latent inhibition will be more pronounced than for items that are very different (i.e., that have few common elements).

According to this line of reasoning, if perceptual learning is more marked for perceptually similar items than perceptually different items as the MKM (McLaren et al., 1989) model predicts and as previous work indicates (Oswalt, 1972), then one prediction that follows on from our account is that pre-exposure would lead to a greater elevation of family resemblance sorting for perceptually similar stimuli compared to perceptually different stimuli. Correspondingly, if our account is correct and increased differentiation helps encourage family resemblance categorization then perceptually different stimuli should, more generally, result in greater levels of family resemblance sorting than perceptually similar stimuli, as found by Milton and Wills (2008). These predictions are the focus of the current study.

Method

Participants and apparatus

Students from the University of Exeter participated either for course credits or for a payment of £5. There were sixty-four participants (16 in each of the four between-subject conditions) who were tested individually in a quiet testing cubicle. We tested participants using E-prime, on a Dell PC with a 17-in. monitor and a standard computer keyboard.

Stimuli

The four stimulus sets employed in this study had the same abstract structure as previously used by Medin et al. (1987). This structure is shown in Table 1. The structure consisted of five binary-valued dimensions (D1-D5) and the stimuli were organized around two prototypes, each representative of one of the categories. These prototypes were constructed by taking all the positive values on the dimensions for one of the stimuli (1, 1, 1, 1, 1) and all of the zero values on the dimensions (0, 0, 0, 0, 0) for the other category. The rest of the stimuli were mild distortions of the two prototypes in that they had four features characteristic of their category and one atypical feature more characteristic of the other category. In total, there were 12 stimuli in each set. Sorting the stimuli by family resemblance, as shown in Table 1, maximizes within-group similarities and minimizes between-group similarities.

Table 1: The abstract stimulus structure

	Category A						Category B				
	D1	D2	D3	D4	D5		D1	D2	D3	D4	D5
E1	1	1	1	1	1	E7	0	0	0	0	0
E2	1	1	1	1	0	E8	0	0	0	0	1
E3	1	1	1	0	1	E9	0	0	0	1	0
E4	1	1	0	1	1	E10	0	0	1	0	0
E5	1	0	1	1	1	E11	0	1	0	0	0
E6	0	1	1	1	1	E12	1	0	0	0	0

Note. Each row (within each category) describes a different stimulus. D = dimension: 1 and 0 represent the values of each dimension.

Two of the stimulus sets were based on ladybirds and the other two stimulus sets were based on houses (see Figure 1). The two pairs of stimulus sets were identical except that for one of the sets the binary values for each dimension were relatively easy to distinguish (e.g., for the “ladybird” stimuli the difference in the length of the antennae were relatively large) and for the other set the differences were relatively hard to distinguish (e.g., the difference in the length of the antennae was relatively small). We term these sets the “perceptually different” and the “perceptually similar” stimuli respectively. The five dimensions for the ladybird stimuli were: the length of the antennae, the size of the head, the number of dots, the length of the green ovals on the body, and the size of the legs. The five dimensions for the house stimuli were the height of the aerial, the length of the chimney, the number of lines on the roof, the size of the windows, and the height of the door.

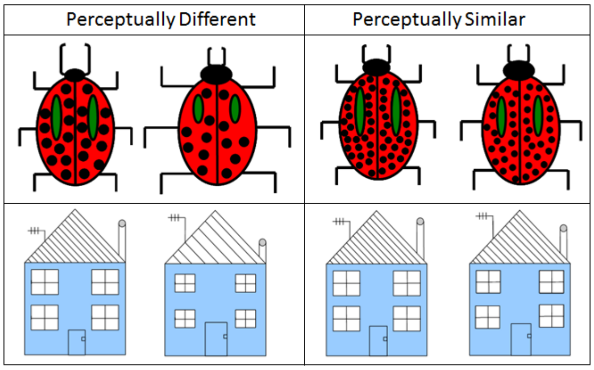

Figure 1. The category prototypes of the four stimulus sets used in this experiment.

Design

The experiment had a 2 x 2 between-subjects factorial design. The first factor was the perceptual difficulty of the stimuli (two levels: *perceptually different*/*perceptually similar*). The second factor was the type of pre-exposure, which also had two levels. In the *same-stimuli* exposure condition, participants were pre-exposed to the same stimuli that they were subsequently asked to free classify (e.g., they were pre-exposed to and then free classified the 'ladybird' stimuli). In the *unrelated-stimuli* exposure conditions, participants were pre-exposed to stimuli different to those that they later free classified (e.g., they were exposed to the 'house' stimuli and then free classified the 'ladybird' stimuli). This led to four conditions: *perceptually different/same-stimuli* exposure; *perceptually similar/same-stimuli* exposure; *perceptually different/unrelated-stimuli* exposure and *perceptually similar/unrelated-stimuli* exposure. In all conditions, the stimulus set (either ladybirds or houses) that participants classified was counterbalanced.

Procedure

The running-recognition phase.

In both the same-stimuli and unrelated-stimuli conditions, participants were pre-exposed to the appropriate set of stimuli via a running-recognition task (e.g., Wills & McLaren, 1998). Each of the twelve stimuli in the set was presented twice in each block in a random order. Each trial began with a black fixation cross presented in the middle of the screen lasting 500ms. This was immediately followed by one of the stimuli from the set appearing in the middle of the screen for 3000ms. Participants were not allowed to respond during this time. Once this time had elapsed participants were asked to say whether they had seen that stimulus before in that particular block (by pressing x) or whether they had not seen it before in that block (by pressing m). This response was self-paced. Following this, the next trial immediately began. At the end of each block, participants were informed of their accuracy in that block. In total, participants completed sixteen blocks of 24 trials.

The free classification phase.

The basic categorization procedure was the same for all four conditions. We used a computer-based variation of Regehr and Brooks's (1995) match-to-standards procedure that was the same as that adopted in Milton, Wills, and Hodgson (2009). Participants were asked to classify a number of stimuli into two categories (they were not told this during the running-recognition task to avoid them categorizing the stimuli during that phase). They were informed that there were many ways in which the stimuli could be split and that there was no one correct answer. Participants were also told that the two groups did not have to be of equal sizes and that they should classify the stimuli in the way that seemed most sensible or natural.

At the beginning of each trial, a black fixation cross was presented for 500ms in the centre of the screen. The two category prototypes were then presented at the top of the screen and below these prototypes one of the twelve stimuli in the set (E1-E12) was displayed. Participants categorized the stimulus into category A by pressing "x" and into category B by pressing "m". This decision was self-paced and no feedback was given on their response. A blank screen was then presented for 1000ms before the next trial began. Each of the stimuli in the set appeared once in each block in a random order. At the end of each block, participants were asked to describe as precisely as possible how they categorized the stimuli in the previous block. In total, there were six blocks of twelve stimuli. The inclusion of multiple blocks provided the opportunity to build up a reliable index of an individual's sorting behaviour rather than relying on a limited number of responses from just one block. Previous work indicates a close correspondence between multiple block procedures and single block procedures (Wills et al., 2013).

Analysis of results

Each participant was classified as having produced one of the sort types described below. These sort types are similar to those employed by Regehr and Brooks (1995) and are identical to those used by Milton and Wills (2004). To be classified as sorting by either family resemblance or unidimensionally, the participant's description of their strategy had to match their behavioural response. As in previous work (e.g., Wills et al., 2013) each block was categorized independently.

A *family resemblance sort*, also commonly known as an "*overall similarity*" sort, has the same structure as shown in Table 1. In this type of strategy, the participant has to place each of the prototypes, along with their derived one-aways, into separate categories without error. Additionally, they have to describe their strategy as being based either on general similarity or on placing each item into the category with which it had more features in common. A *one-away* *family resemblance* sort is similar to the one-away unidimensional sort with the exception that the error occurred in a sort that was otherwise family resemblance.

A *unidimensional* sort is based on a single dimension of the stimulus. It does not matter which of the dimensions is used as the basis of sorting, so long as all of the positive values for the chosen dimension are placed in one category and all of the zero values for that dimension are in the other category. Additionally, to be classified as a unidimensional sort, the participant has to describe their sort as being based on a single dimension. Participants were classified as producing a *one-away unidimensional* sort if they described their sorting as being driven by a single dimension but there was a solitary error in their classification.

Any classifications other than those described above were classified as *other* sorts, even if the description given by the participant fitted one of the sorts described above. The verbal descriptions were clear and easy to categorize into the appropriate group and as in previous work (e.g., Milton & Wills, 2009) there was a very high correspondence between verbal report and behavioral strategy.

Results

Running-recognition phase

The mean accuracy levels for all the conditions across the 16 blocks are displayed in Figure 2. We conducted a mixed-design ANOVA with one within-subjects factor (block, 16 levels) and two between-subjects factors: pre-exposure type (same-stimuli vs unrelated-stimuli) and level of perceptual difficulty (similar vs different).[1] There was a significant effect of perceptual difficulty, $F(1,60) = 6.422$, $p = .014$, $\eta^2_p = .097$, indicating that stimuli in the perceptually different condition were better recognized than those in the perceptually similar condition. There was also a significant effect of block, $F(15,900) = 4.956$, $p < .001$, $\eta^2_p = .076$, indicating that performance on the task improved across the blocks. Unsurprisingly, given that there was no difference between the same-stimuli and unrelated-stimuli exposure conditions in this phase, there was a non-significant effect for this factor, $F(1, 60) = 1.342$, $p = .251$, $\eta^2_p = .022$. None of the interactions approached significance (all Ps>.2).

Free classification phase

For every block, each participant's sorting strategy was classified according to the sort types described above. One-away unidimensional and one-away family resemblance sorts were classified as unidimensional and family resemblance sorts respectively (cf. Milton & Wills, 2004).

The mean proportions of family resemblance and unidimensional categorizations produced in the four conditions are shown in Figure 3. The difference in family resemblance sorting between our conditions was assessed using a 2 x 2 between-subjects ANOVA with the two factors being the level of perceptual difficulty (similar vs different) and pre-exposure type (same-stimuli vs unrelated-stimuli). There was a significant effect of perceptual difficulty, $F(1,60) = 30.601$, $p < .001$, $\eta^2_p = .338$, which indicated that family resemblance categorization was higher for the perceptually different stimuli than the perceptually similar stimuli. There was no significant effect of pre-exposure type, $(F1,60) = 3.211$, $p = .078$, $\eta^2_p = .051$, although there was a trend for the proportion of family resemblance categorization to be greater in the same-stimuli condition than the unrelated-stimuli condition. Finally, there was a significant interaction between level of perceptual difficulty and pre-exposure type, $F(1, 60) = 9.233$, $p = .004$, $\eta^2_p = .133$. Pairwise comparisons revealed that same-stimuli exposure resulted in greater levels of family resemblance categorization than unrelated-stimuli exposure for the perceptually different stimulus set, $t(30) = 2.570$, $p = .015$, but not for the perceptually similar stimulus set, $t(30) = 1.826$, $p = .078$. Indeed, there was a numerical trend for family resemblance sorting to be lower in the same-stimuli condition than in the unrelated-stimuli condition.

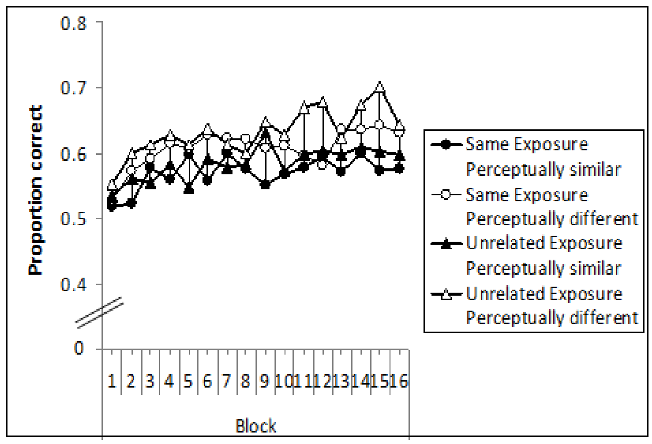

Figure 2. The mean accuracy for the four conditions in the running-recognition pre-exposure task.

For the mean proportion of unidimensional sorts, a similar 2 x 2 between-subjects ANOVA was conducted. There was again a significant effect of perceptual difficulty, $F(1,60) = 6.910$, $p = .011$, $\eta^2_p = .103$, indicating that unidimensional sorting was greater for the perceptually similar stimuli than for the perceptually different stimuli. There was, however, no significant effect of exposure type, $F(1,60), = .925$, $p = .340$, $\eta^2_p = .015$, but there was again a significant interaction between stimulus type and exposure type, $F(1,60) = 6.252$, $p = .015$, $\eta^2_p = .094$. Pairwise comparisons, assessing this interaction, revealed that for the perceptually different stimuli, unidimensional sorting was lower in the same-stimuli exposure condition than in the unrelated-stimuli exposure condition, $t(30) = 2.220$, $p = .034$. In contrast, there was no significant effect of exposure type for the perceptually similar stimuli, $t(30) = 1.229$, $p = .229$.

For "other" classifications, a 2 x 2 between-subjects ANOVA yielded a significant main effect of perceptual

[1] We ran additional ANOVAs that assessed whether there were any differences in the pattern of results between the two types of stimuli (ladybirds/ houses). None of these analyses approach significance and for conciseness we have, therefore, not included this factor in the reported analyses.

difficulty, $F(1,60) = 6.959$, $p = .011$, $\eta^2_p = .104$, with the perceptually similar stimuli resulting in a greater level of "other" sorting. There was, however, no significant main effect of exposure type, $F(1,60) = .278$, $p = .600$, $\eta^2_p =$ = .005, and no interaction between exposure type and perceptual difficulty, $F(1,60) = .031$, $p = .861$, $\eta^2_p = .001$.

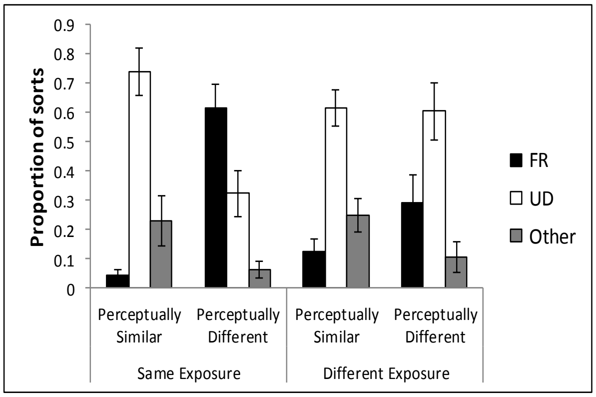

Figure 3. The mean proportion of family resemblance, unidimensional, and other sorts for each condition. FR = family resemblance, UD = unidimensional.

General Discussion

The results of our experiment provide, to our knowledge, the first demonstration that prior exposure to a set of stimuli can elevate the subsequent level of family resemblance categorization. However, the precise pattern of this effect was somewhat different to what we predicted. Specifically, we found that same-stimuli pre-exposure increased family resemblance sorting (and reduced unidimensional sorting) for the perceptually different stimuli but not for the perceptually similar stimuli. Indeed, for the perceptually similar stimuli we found a non-significant trend for an effect in the opposite direction.

The finding that pre-exposure led to a significant increase in family resemblance categorization for the perceptually different stimuli but not for the perceptually similar stimuli is surprising according to the latent inhibition mechanism of the MKM model (McLaren et al., 1989) outlined in the introduction. However, the effect we observed could, potentially, be driven by another mechanism that is a fundamental property of the MKM model – unitization. Unitization is the process by which individual dimensions or units can be bound into a single perceptual configuration (e.g., Schyns, Goldstone, & Thibaut, 1998). If unitization were to occur, then this would be likely to encourage a family resemblance strategy as the dimensions would be bound into a holistic item and it would require effort to analyse the constituent parts separately as would be required by a unidimensional strategy. It could be that the unitization process occurs more rapidly for items that are perceptually more discriminable than those which are more similar and this would explain our pattern of results. This hypothesis would be consistent with the traditional view that family resemblance sorting is the result of a quick, holistic, non-deliberative, process (e.g., Kemler Nelson, 1984). If unitization is driving this effect, then one prediction is that the imposition of a time constraint (or a concurrent load) should lead to increased levels of family resemblance sorting compared to no time constraints (or no concurrent load) for same-stimuli pre-exposure but not for unrelated-stimuli pre-exposure. This latter prediction is based on previous work which indicates that, for novel stimuli, family resemblance sorting is the result of a deliberative, analytic, strategy under the conditions we have used here (cf., Milton et al., 2008; Wills et al., 2013). If this is correct, then it would provide insight into why we appear to apply an effortful and time-consuming family resemblance strategy in the lab whilst in the real world, where one often has substantial prior exposure to the items, family resemblance sorting appears to be relatively automatic. More generally, this issue would be of relevance to the question of whether there are competing implicit and explicit categorization systems as proposed by the COmpetiton between Verbal and Implicit Systems (COVIS) model (Ashby et al., 1998) which has been subject to intense debate in recent years (e.g., Ashby et al., 1998, 1999; Newell et al., 2013).

Contrary to our predictions, same-stimuli pre-exposure did not enhance family resemblance categorization for the perceptually similar stimuli. Our a priori hypothesis was that the pre-exposure would increase the discriminability of the dimension which should, on the basis of previous work (Milton & Wills, 2008), lead to an elevation of family resemblance categorization. One potential explanation for this failure to find the predicted pattern is that pre-exposure may have increased the differentiation of the stimuli as in previous work (e.g., Oswalt, 1972) but not to the extent necessary to encourage family resemblance sorting. Indeed, family resemblance categorization was very low for the perceptually difficult stimuli in both the same-stimuli and unrelated-stimuli pre-exposure conditions (for a related effect, see Milton & Wills, 2004). If this is correct, then a greater level of pre-exposure (perhaps over multiple sessions) should increase family resemblance sorting for the perceptually similar stimuli.

A number of issues are worthy of future research. First, whilst the match-to-standards procedure is often characterized as a form of free classification, it is more restricted than other procedures which do not present the prototypes on each trial and do not specify the number of categories that can be created (e.g., Pothos & Close, 2008). It would be interesting to see the extent to which our findings generalise to other procedures. Second, there were only 12 unique stimuli in the set - it would be useful to examine whether having a greater number of unique stimuli influences the results. Third, whilst we used unrelated stimuli as our baseline exposure condition these stimuli shared the same category structure as the stimuli they later classified. It is, therefore, possible that both groups benefited from pre-exposure to some extent. Future work could use a completely unrelated task as a baseline to assess this possibility. Relatedly, it would be worth examining whether our results generalise to different types of pre-exposure tasks such as, for example, pleasantness ratings.

Nevertheless, whilst the experiment produced a somewhat different pattern of results than was anticipated, we found clear evidence that pre-exposure can elevate family resemblance categorization at least for stimuli that are of relatively low perceptual difficulty. This result appears important as in previous free classification studies participants typically have had little or no pre-exposure to the stimuli before being asked to sort them, an approach that seems atypical of the items that we encounter in the real world. The present study, therefore, provides some explanation for why the prevalence of family resemblance categorization is low in previous free classification experiments. Clearly, much research needs to be conducted to understand the conditions under which this effect will occur and the mechanisms by which it operates. In this regard, the present study should be seen as an important first step that we hope will help motivate future research.

Acknowledgments

We thank Ian McLaren and the anonymous reviewers for their helpful thoughts and suggestions on this work.

References

Ahn, W.-K., & Medin, D. L. (1992). A two-stage model of category construction. *Cogn Sci, 16*, 81-121.

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychol Rev, 105*, 442-481.

Ashby, F.G., Queller, S., & Berretty, P.M. (1999) On the dominance of unidimensional rules in unsupervised categorization. *Percept Psychophys*, *61*, 1178–1199.

Bruner, J.S., Goodnow, J.J., Austin, G.A., 1956. *A Study of Thinking*. Wiley, New York.

Chamizo, V. D., & Mackintosh, N. J. (1989). Latent learning and latent inhibition in maze discriminations. *Q J Exp Psychol B*, *41*, 21–31.

Kemler Nelson, D. G. (1984). The effect of intention on what concepts are acquired. *J Verb Learn Verb Behav*, *23*, 734 – 759.

McLaren, I.P.L. (1997). Categorization and perceptual learning: an analogue of the face inversion effect. *Q J Exp Psychol B*, *50*, 257-273.

McLaren, I.P.L., Kaye, H., & Mackintosh, N.J. (1989). An associative theory of the representation of stimuli: applications to perceptual learning and latent inhibition. In Morris R.G.M. (ed) *Parallel Distributed Processing - Implications for Psychology and Neurobiology*, Oxford: Oxford University Press.

McLaren, I., Leevers, H., & Mackintosh, N. (1994). Recognition, categorization, and perceptual-learning (or, how learning to classify things together helps one to tell them apart). *Attention and Performance XV*, *15*, 889-909.

Medin, D.L., Wattenmaker, W.D. & Hampson, S.E. (1987). Family resemblance, conceptual cohesiveness, and category construction. *Cognit Psychol, 19*, 242-279.

Milton, F. N., Longmore, C. A., & Wills, A. J. (2008). Processes of overall similarity sorting in free classification. *J Exp Psychol Hum Percept Perform*, *34*, 676 – 692.

Milton, F.N., Viika, L., Henderson, H., Wills, A.J. (2011). The effect of time pressure and the spatial integration of the stimulus dimensions on overall similarity categorization. *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*, 795-800.

Milton, F.N., & Wills, A.J. (2004). The influence of stimulus properties on category construction. *J Exp Psychol Learn Mem Cogn*, *30*, 407-415.

Milton, F., & Wills, A.J. (2008). The influence of perceptual difficulty on family resemblance sorting. *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, 2273-2278.

Milton, F., & Wills, A. J. (2009). Long-term persistence of sort strategy in free classification. *Acta Psychol, 130*, 161-167.

Milton, F., Wills, A. J., & Hodgson, T. L. (2009). The neural basis of overall similarity and single-dimension sorting. *NeuroImage, 46*, 319-26.

Newell, B.R., Moore, C.P., Wills, A.J., & Milton, F. (2013). Reinstating the frontal lobes? Having more time to think improves “implicit” perceptual category learning. A comment on Filoteo, Lauritzen & Maddox (2010). *Psych Sci*, *24*, 386-389.

Oswalt, R. M. (1972). Relationship between level of visual pattern difficulty during rearing and subsequent discrimination in rats. *J Comp Physiol Psychol*, *81*, 122–125.

Pothos, E.M., Close, J., (2008). One or two dimensions in spontaneous classification: a simplicity approach. *Cognition* 107, 581–602.

Regehr, G. & Brooks, L.R. (1995). Category organization in free classification: The organizing effect of an array of stimuli. *J Exp Psychol Learn Mem Cogn*, *21*, 347-363.

Rosch, E. & Mervis, C.B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognit Psychol*, *7*, 573-605.

Schyns, P.G., Goldstone, R.L., & Thibaut, J-P. (1998). The development of features in object concepts. *Behav Brain Sci*, *21*, 1-54.

Spalding, T. L., & Murphy, G. L. (1996). Effect of background knowledge on category construction. *J Exp Psychol Learn Mem Cogn, 22*, 525–538.

Wills, A.J., & McLaren, I.P.L. (1998). Perceptual learning and free classification. *Q J Exp Psychol B*, *51*, 235-270.

Wills, A.J., Milton, F., Longmore, C.A., Hester, S., Robinson, J. (2013). Is overall similarity classification less effortful than single-dimension classification? *Q J Exp Psychol*, *66*, 299-318.