Acta Psychologica 130 (2009) 161–167

Contents lists available at ScienceDirect

Acta Psychologica

journal homepage: www.elsevier.com/locate/actpsy

# Long-term persistence of sort strategy in free classification

Fraser Milton *, Andy J. Wills

*University of Exeter, Washington Singer Laboratories, Perry Road, Exeter EX4 4QG, United Kingdom*

ARTICLE INFO

*Article history:*
Received 1 July 2008
Received in revised form 1 December 2008
Accepted 1 December 2008
Available online 12 January 2009

*PsycINFO classification:*
2343

*Keywords:*
Family resemblance
Free classification
Perseveration
Unidimensional

ABSTRACT

Two free classification experiments that investigate the persistence of sort strategy are reported. Participants tend to persist with their initial categorization type (family resemblance or unidimensional) for the remaining sorts, overriding the effects of otherwise influential stimulus properties. Sort type was found to persist even after a one-week delay. Stimulus-driven models of free classification (e.g., the SUSTAIN model, [Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. Psychological Review, 111, 309–332]) cannot predict the sort type persistence effects we observe, but they are naturally accounted for by theories that posit strategic selection of a problem-solving strategy (e.g., Hypothesis theory, [Levine, M. (1971). Hypothesis theory and nonlearning despite ideal S-R-reinforcement contingencies. Psychological Review, 78, 130–140]).



## 1. Introduction

Classification – dividing the world into groups of things – is one of the fundamental building blocks of cognition. However, in view of the immense number of objects we encounter, this process must necessarily be highly constrained; for example, just 10 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.

The majority of research in this area has concentrated on a task where participants get accurate, trial-specific feedback about their categorization decisions, and the thorough study of this task has yielded important information. However, it is increasingly being appreciated that the level of feedback present in this task seldom exists outside the laboratory. As a consequence, there has been growing interest in how people categorize stimuli in the absence of feedback, a task described as *free classification* (Regehr & Brooks, 1995), *spontaneous categorization* (Pothos & Chater, 2002), or *category construction* (Medin, Wattenmaker, & Hampson, 1987).

One of the fundamental phenomena in free classification is that people tend to exhibit a strong preference to sort on the basis of a single stimulus dimension (e.g., Medin et al., 1987), for example, sorting entirely on the basis of color, whilst apparently discounting variability on other dimensions (e.g., size). This result is perhaps surprising given that many natural categories seem to have a "family resemblance" (overall similarity) structure (Rosch & Mervis, 1975), in other words, they possess a number of characteristic but not defining features.

Recent research has demonstrated that people's preference for unidimensional sorting can be substantially reduced, and family resemblance sorting increased, by a number of different factors. For example, the prevalence of family resemblance sorting is increased by sequential stimulus presentation (Regehr & Brooks, 1995), by spatial separation of the stimulus dimensions (Milton & Wills, 2004), and by stimulus dimensions which are easier to perceptually differentiate (Milton & Wills, 2008). In addition, family resemblance sorting can be elevated by increasing the time available to categorize the stimuli (Milton, Longmore, & Wills, 2008), and by relevant background knowledge which allows features to be integrated in a meaningful and coherent way (Spalding & Murphy, 1996; Wattenmaker, 1995). The current article examines another factor that might be expected to influence the prevalence of family resemblance sorting – the sort type the participant has previously produced.

Our hypothesis is that the prevalence of family resemblance sorting will be affected by the sort strategy employed previously in similar situations. This hypothesis seems generally in line with a number of previously reported findings which demonstrate that people commonly persist with an initial strategy, even when a related task can be completed more effectively with a different strategy. The prototypical example of this is perhaps the water jug *Einstellung* problems used to study the perseveration of mental set (e.g., Luchins, 1942; Schultz, Stone, & Christie, 1997). Cognitive rigidity has also been demonstrated in a sequential number reduction task (Woltz, Gardner, & Bell, 2000), the sequential application

* Corresponding author. Tel.: +44 1392 264636.
*E-mail address:* f.n.milton@exeter.ac.uk (F. Milton).

0001-6918/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.actpsy.2008.12.001

of logic gates (Woltz, Gardner, & Gyll, 2000), and in conditional and cumulative risk judgments (McCloy, Beaman, Morgan, & Speed, 2007). A further example of cognitive inflexibility is provided by studies of functional fixedness (e.g., Duncker, 1945). A general overview of cognitive rigidity is presented in Schultz and Searleman (2002).

The most directly related study to our hypothesis, however, is the observation that people tend to sort subsets of stimuli in a self-consistent manner. In two experiments, Murphy (2001) asked participants to sort a set of nine items in which three taxonomic categories (vehicle, profession, and location) were crossed with three thematic categories (air travel, water travel, and auto travel). In Experiment 1, thematic sorting dominated (e.g., *pilot* was grouped with *airport* and *airplane*). However, in Experiment 2, the location subset (airport, dock, and garage) was replaced with an animal subset (skunk, raccoon, and woodchuck), which was difficult to partition thematically but easy to group taxonomically. The other two subsets (vehicle and profession) were identical to those which encouraged thematic sorting in Experiment 1. Participants in Experiment 2 had a strong preference to sort all three subgroups taxonomically. As Murphy (2001) argues, these results suggest that participants have a tendency to sort simultaneously presented subgroups in a consistent manner. Such a conclusion is clearly related to our hypothesis that the sort strategy adopted for one set of stimuli will tend to persist when subsequently sorting other stimuli.

This hypothesis also receives some support from free classification studies in which the requirement to sort more than one stimulus set was an incidental feature of the design. For example, one of the findings of Wattenmaker (1995) is that social stimuli produce more family resemblance sorting than object stimuli. Interestingly, in the two experiments in which family resemblance sorting of social stimuli is most prevalent (Experiments 1a and 2c), object stimuli are more likely to be sorted by family resemblance if the participants sort social stimuli immediately beforehand. Wattenmaker reports this pattern of results, but provides neither statistical analysis nor explanation. One possible explanation is that the family resemblance sort strategy evoked by the social stimuli perseverates to the object stimuli. However, the data presented in Wattenmaker's paper are (understandably) not ideal for assessing this hypothesis – for example, an alternative explanation is that exposure to the sorting task, rather than exposure to social stimuli *per se*, facilitates family resemblance sorting of object stimuli.

To summarize, a number of converging lines of evidence provide general support for the idea that sort type might persist in free classification, but this hypothesis has not been directly investigated. The two studies presented in the current paper begin this investigation. Evidence for perseveration of sort strategy would be of considerable theoretical importance when one considers that current models of free classification have no mechanism to capture such a result. For example, SUSTAIN ( Love, Medin, & Gureckis, 2004), which is arguably the best developed model of free classification, can produce family resemblance or unidimensional sorts depending on the stimulus structure with which it is presented, but currently has no process or parameter that can represent an acquired bias for family resemblance or unidimensional sorting. For this reason, SUSTAIN is, at present, incapable of accounting for the type of sort type persistence we investigate in the current paper. The Simplicity model (Pothos & Chater, 2002), which has had some success in accounting for the relative prevalence of unidimensional and multidimensional sorting in free classification in situations where SUSTAIN makes an incorrect prediction (Pothos & Close, 2008), also cannot account for sort type perseveration. Like SUSTAIN, it has no mechanism that would enable it to represent an acquired bias for a particular sort strategy. In contrast, the type of process suggested by Hypothesis theory (Levine, 1971) provides a better account of sort type persistence in free classification. Hypothesis theory states that when faced with a novel problem, people will tend to select a strategy which has previously been used successfully for similar problems. Hypothesis theory has often been used to explain the results of Einstellung effects (e.g., Levine, 1971; Sweller & Gee, 1978), and has been used to generate novel predictions that have received empirical support (e.g., Levine, 1971; Sweller, 1976; Sweller & Gee, 1978).

Persistence of sort strategy in free classification could be accounted for in at least two ways. First, an individual differences account would assume that some individuals have a pre-experimental bias toward family resemblance sorting and others toward unidimensional sorting. The first sort produced in an experiment would, therefore, predict subsequent sort strategy because that individual had a pre-existing tendency to produce, say, family resemblance sorts. Alternatively, or additionally to an individual differences account, a perseveration account would assume that participants develop a bias for a particular sort type during the experiment itself. If this were the case, the first sort would predict sort strategy in the second sort because of a bias developed during the production of the first sort. Although either reason for sort strategy persistence would be problematic for existing models of free classification such as SUSTAIN, and both would be straightforward to explain in terms of Hypothesis theory, it nevertheless appears important to determine whether sort strategy persistence can result from a single episode of sorting.

In the current studies, we investigate the possibility that persistence of sort strategy can be acquired within the experiment itself by employing a form of overshadowing manipulation. Milton and Wills (2004) demonstrated that stimuli whose features are spatially integrated tend to produce less family resemblance sorting than stimuli whose features are spatially separate. If sort type persistence is at least partially caused by the development of a bias during the experiment, then this acquired bias might be expected to overshadow the spatial integrality effect reported by Milton and Wills. This account assumes a special role for the first sort in the experiment – the first sort affects the sort strategy of subsequent sorts but is not itself influenced in this way. This means that other influences on sort strategy (e.g., spatial integration) should be more pronounced in the first sort (where perseveration does not act) than in subsequent sorts (where it does). In contrast, an individual differences account does not obviously make such a prediction.

## 2. Experiment 1

In our first experiment, participants free sorted four different stimulus sets, one after the other.

### 2.1. Method

#### 2.1.1. Participants

Forty-eight students from the University of Exeter took part either for money or for course credits. Participants were tested individually in a quiet cubicle.

#### 2.1.2. Stimuli

All four sets of stimuli (whose prototypes are shown in Fig 1) took the same abstract structure employed by Medin et al. (1987). This abstract stimulus set is shown in Table 1. Each stimulus set consisted of four binary-valued dimensions (D1–D4), and the stimuli for each category were organized around two prototypical stimuli, each representative of a category. These prototypes were constructed by taking all the positive-valued dimensions for one of the stimuli (1,1,1,1) and all the zero-valued (0,0, 0, 0)

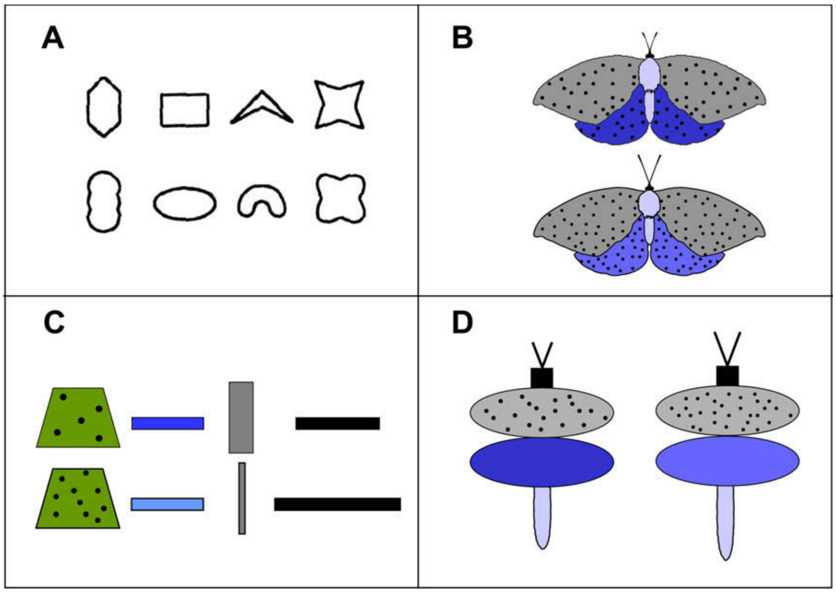


**Fig. 1.** The prototypes for the four stimulus sets used in Experiment 1.

dimensions for the other stimuli. The other stimuli (one-aways) each had three of the four characteristic features of their category and one atypical feature characteristic of the other category. Sorting the stimuli by a family resemblance structure, as shown in Table 1, maximizes within-group similarities and minimizes between-group similarities. In total, there were 10 stimuli in each stimulus set.

The spatially separable stimuli, shown in Fig 1A and C, are known to evoke family resemblance sorting (Milton & Wills, 2004; Regehr & Brooks, 1995), whilst the spatially integrated stimuli, shown in Fig 1B and D, are known to evoke unidimensional sorting (Milton & Wills, 2004). For the current purposes, it is not critical whether spatial separation is the crucial variable producing this difference in sort type prevalence.

*2.1.3. Procedure*

Participants were required to classify each of the four stimulus sets in turn. Before classifying each set, participants were introduced to the stimuli with a pre-sort procedure employed previously by Milton and Wills (2004). Two copies of each of the 10 stimuli in the set were spread out randomly in an array. Participants then had to match these stimuli into identical pairs without feedback. If participants made any mistakes, the pairs had to be matched again. The purpose of this task was to ensure that participants could fully distinguish the four feature-pairs.

After this pre-sort procedure, the method of stimulus presentation for the categorization task was similar to that used by Regehr and Brooks (1995) and identical to that used in Milton and Wills (2004). The two prototypes were placed side-by-side on a table throughout the sort and participants were informed that these were characteristic of category A and category B (which category each prototype represented was randomized). Participants were then given the 10 stimuli, which had been shuffled into a random order, and were asked to look at each stimulus individually and place it into the group to which they felt it was most representative. Once they had made their decisions, participants placed each card face down directly below their chosen category. Participants were told that there were many ways in which the stimuli could be split and that there was no one correct answer. They were also told that the two groups did not have to be of equal size and that the only constraints were that each stimulus had to be placed into one of the two groups and that they were not allowed to look through the cards that remained to be sorted or to return to cards already sorted. There was no time limit for completing the task. Participants then classified the remaining three stimulus sets after being given the same instructions as before. In other words, the pre-sort procedure followed by the classification procedure was repeated four times, once with each of the stimulus sets.

After each sort, participants were asked to describe as precisely as possible the strategy they had used to sort that stimulus set. The order in which the four stimulus sets were presented was counterbalanced, with each of the 24 permutations completed by two participants. Hence, each of the four sets was presented an equal number of times in each of the four possible positions, and each stimulus set followed and preceded the other three sets an equal number of times.

**Table 1**
The abstract stimulus set employed in Experiments 1 and 2.

| Category A | | | | | Category B | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Stimulus | *D1* | *D2* | *D3* | *D4* | Stimulus | *D1* | *D2* | *D3* | *D4* |
| 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 |
| 2 | 1 | 1 | 1 | 0 | 2 | 0 | 0 | 0 | 1 |
| 3 | 1 | 1 | 0 | 1 | 3 | 0 | 0 | 1 | 0 |
| 4 | 1 | 0 | 1 | 1 | 4 | 0 | 1 | 0 | 0 |
| 5 | 0 | 1 | 1 | 1 | 5 | 1 | 0 | 0 | 0 |

*2.1.4. Classification of sort behaviour*

As in previous work (e.g., Milton & Wills, 2004; Milton & Wills, 2008; Milton et al., 2008), two sources of information were used to classify the sorts participants produced: the description the participant gave and the categories they constructed.[1] The categories that sorts were placed into were closely modeled on those used in Regehr and Brooks (1995), and identical to those used in Milton and Wills (2004).

*A unidimensional* sort was defined as a sort based on a single dimension of the stimulus. It did not matter which dimension was used as the basis of sorting, providing all the positive-valued

[1] The overlap between the verbal and behavioral measures across both Experiments 1 and 2 is extremely high (.99).

**Table 2**
Contingency tables showing the relationship between sorts in Experiment 1.

| (A) | | Sort 2 FR | Sort 2 UD |
|---|---|---|---|
| Sort 1 | FR | 16 | 5 |
| | UD | 3 | 23 |

| (B) | | Sort 3 FR | Sort 3 UD |
|---|---|---|---|
| Sort 2 | FR | 15 | 3 |
| | UD | 4 | 23 |

| (C) | | Sort 4 FR | Sort 4 UD |
|---|---|---|---|
| Sort 3 | FR | 17 | 2 |
| | UD | 1 | 26 |

*Note*: FR, family resemblance; UD, unidimensional.

features for the chosen dimension were in one category and all the zero-valued features were in the other category. Participants also had to describe their sort as based on that particular dimension. Sorts were also classified as unidimensional if participants described their classification as based on a single dimension but there was a solitary error in their classification. In other words, nine of the items were classified on the basis of a single dimension but the other item was placed in the wrong category.

A *family resemblance* sort had a structure identical to that shown in Table 1. In this type of sort, each of the prototypes along with their derived one-aways were placed in separate categories without error. Additionally, participants had to describe their sort as being either based on overall similarity or by indicating that they placed each item into the category with which it had more features in common. Sorts described in this way, but which contained a solitary sorting error, were also classified as family resemblance sorts.

Any other sorts were placed into an *other* category, even if the description given by the participant fitted one of the sort types described above.

### 2.2. Results and discussion

The 2.1% of sorts that were classified as *other* were excluded from analysis. Of the remaining sorts, 6.9% each included a solitary error (see Section 2.1). Table 2A shows, in contingency table format, the relationship between the classifications produced by participants on the first sort and classifications they went on to produce in the second sort.[2] For example, Table 2A shows that of the 21 participants who produced a family resemblance classification on sort 1, 16 also produced a family resemblance classification on sort 2. Inspection of Table 2A indicates that there is a tendency to persist in sort 2 with the classification type produced in sort 1. The strength of this relationship can be quantified by the $\Delta P$ statistic (e.g., Allan, 1980):

$$\Delta \mathrm{P} = \mathrm{P}(\mathrm{UD}_2|\mathrm{UD}_1) - \mathrm{P}(\mathrm{UD}_2|\mathrm{FR}_1)$$

where $P(UD_2|UD_1)$ is the probability of the second sort being unidimensional given that the first sort was unidimensional, and $P(UD_2|FR_1)$ is the probability given that the first sort was family resemblance. If $\Delta P$ is 1, the previous sort perfectly predicts the current sort, if it is −1, the previous sort perfectly predicts the opposite to the previous sort, and if it is 0, there is no contingent relationship between the two sorts. Note that $\Delta P$ can also equivalently be described as $P(FR_2|FR_1)-P(FR_2|UD_1)$, and results in an identical value.

The $\Delta P$ for the relationship between classification type in sort 1 and classification type in sort 2 is 0.64. Each participant contributes a maximum of one count to Table 2A, hence the observations that make up Table 2A are independent of each other, and a contingency chi-square calculation can be employed to assess the statistical reliability of the relationship.[3] The relationship is statistically reliable, $\chi^2$ $(1, N = 47) = 20.16$, $p < .001$. Table 2B shows the relationship between the second and the third sorts ($\Delta P = 0.68$), and Table 2C shows the relationship between the third and the fourth sorts ($\Delta P = 0.85$). Both relationships were significant, $\chi^2$ $(1, N = 45) = 20.79$, $p < .001$, and $\chi^2$ $(1, N = 46) = 34.48$, $p < .001$, respectively. Whilst the $\Delta P$ was higher for the relationship between the third and fourth sorts compared to that between the first and second sorts, log-linear analysis[4] revealed that this trend did not reach significance (likelihood ratio = 2.44, df = 1, $p = .12$).

The magnitude of the spatial integrality effect can also be expressed in $\Delta P$ terms

$$\Delta P_{\mathrm{stim}} = P(UD_{\mathrm{sort}}|\mathrm{Sep}) - P(UD_{\mathrm{sort}}|\mathrm{Int})$$

where $P(UD_{sort}|\mathrm{Sep})$ is the probability that a stimulus set is sorted unidimensionally, given it is one of the spatially separate sets, and $P(UD_{sort}|\mathrm{Int})$ is the probability given that it is one of the spatially integrated sets. For the first sort of the current experiment, $\Delta P_{\mathrm{stim}} = 0.54$, $\chi^2$ $(1, N = 48) = 14.31$, $p < .001$, whilst for the second, third, and fourth sorts it is −0.02, 0.01 and 0.10, respectively. Analyses for the effects of stimulus type in sorts 2–4 did not approach significance (Maximum $\chi^2 = 1.68$, $p > .5$).

In summary, Experiment 1 demonstrates reliable sort type persistence. The spatial integration effect, previously reported by Milton and Wills (2004), is found in the first sort, but not in subsequent sorts. Taken together, these results suggest that people develop a bias during the first sort, which then persists for subsequent sorts, and this consequently leads to the reduction in the magnitude of the spatial integration effect. An individual differences account (i.e., that participants have a pre-experimental bias for, say, family resemblance sorting) is unable to provide a full account for this pattern of results, as it cannot explain why the spa-

[2] An alternative approach might have been to group participants according to the stimulus set they were initially assigned to (i.e., family resemblance or unidimensional inducing) rather than according to the strategy they used in the preceding block. One issue with such an approach, however, is that a significant result would be expected when there is no sort type persistence as well as when there is. This is due to the anticipated spatial integration effect which would be present in all the blocks (rather than just the first) in the absence of sort type persistence. The analysis we have used will only be significant in the presence of sort type persistence.

[3] A McNemar test applied to the current data would investigate the hypothesis that the ratio of FR to UD sorts changed from sort 1 to sort 2. This is not the hypothesis under examination, in fact, it is doubly dissociated from the hypothesis currently under test. Yates's correction has not been applied because the assumption of fixed marginals is not appropriate for this data set. No corrections have been applied for the low expected frequencies of some of the cells, as small expected frequencies do not increase the chance of type I errors (Bradley, Bradley, McGrath, & Cutcomb, 1979).

[4] A general discussion of the use of log-linear analysis can be found in Howell (2002, pp. 655-690).

tial integration effect would appear on sort 1, but not on subsequent sorts.

## 3. Experiment 2

Experiment 1 indicated that previous sorts can have a significant influence on future categorization behavior; Experiment 2 investigated the persistence of this effect. In Experiment 1, there was a short delay between sorts (due to the pre-sort matching pairs procedure), but Hypothesis theory (Levine, 1971) would predict that sort type persistence should be observable over longer intervals; in fact, this interval should only be limited by the individual's ability to retrieve which sort type they used previously. In Experiment 2, we extended the delay between sorts to one-week. A condition was also included where the delay was comparable to that of Experiment 1, so that the effect of an increased delay on sort type persistence could be assessed.

### 3.1. Method

#### 3.1.1. Participants

Forty-eight students from the University of Exeter were recruited to take part for course credits. Half the participants completed the two sorts in a single, uninterrupted session (the no-delay condition) and the other half had a one-week delay between sorts.

#### 3.1.2. Stimuli

In Experiment 2, participants were required to sort the two sets of stimuli shown in Fig. 1A and B.

#### 3.1.3. Procedure

Participants were randomly allocated into one of the two between-subject conditions. The no-delay condition had a procedure identical to that used in Experiment 1 with the exception that, in this experiment, participants sorted only two stimulus sets. The delay condition was identical to the no-delay condition except that there was a delay of at least seven days between sorts. In both conditions, the order in which the two stimulus sets were presented was counterbalanced. As in Experiment 1, participants were asked to describe, as precisely as possible, the strategy they had used to sort that particular stimulus set at the end of each sort.

### 3.2. Results and discussion

The 3.1% of sorts that were classified as *other* were excluded from analysis. Of the remaining sorts, 3.2% each included a solitary error. Table 3A shows, for the no-delay condition, the relationship between the classifications produced by participants on the first sort and the classifications they made on the second sort. The relationship is statistically reliable, $\Delta P = 0.47$, $\chi^2$ $(1, N = 23) = 5.05$, $p < .05$. Table 3B shows the corresponding relationship in the one-week delay condition; this relationship is also reliable, $\Delta P = 0.43$, $\chi^2$ $(1, N = 22) = 5.31$, $p < .05$. Log-linear analysis failed to reveal any significant difference in the magnitude of the sort type persistence effect found in the no-delay and one-week delay conditions (likelihood ratio = 0.11, df = 1, $p > .5$). As in Experiment 1, spatial integrality had a significant effect on sort type in the first sort, $\Delta P_{stim} = 0.40$, $\chi^2$ $(1, N = 46) = 7.77$, $p < .01$, but not in the second sort, $\Delta P_{stim} = 0.06$, $\chi^2$ $(1, N = 47) = 0.181$, $p > .5$.

For those participants who perform both sorts unidimensionally, one potential explanation for the perseveration observed is that it is due to pre-existing dimension compatibilities across the stimulus sets. For instance, a participant may sort on dimension 1 of the first stimulus set and then shift attention to dimension 2 of the second stimulus set because, in some way, these two dimensions resemble each other more than the other possible pairings. If so, the probability of selecting a particular dimension in a second unidimensional sort should be affected by the dimension employed in the first unidimensional sort. A contingency chi-square test of this hypothesis is non-significant, $\chi^2$ $(9, N = 23) = 7.51$, $p > .5$. Comparable analyses of the Experiment 1 data were also non-significant. Our results, therefore, provide no evidence that the perseveration of sort strategy was due to pre-existing dimension compatibilities across the stimulus sets.

In summary, Experiment 2, in agreement with the results from Experiment 1, shows that individuals are significantly influenced by previous sorts when they free classify. Experiment 2 extends the findings of Experiment 1 by demonstrating that the sort type persistence effect remains even when there is a one-week delay between the two sorts; indeed, there was no significant reduction in the magnitude of the sort type persistence effect when the delay was increased to one-week.

**Table 3**
The relationship between sorts in the no-delay and one-week delay conditions of Experiment 2.

**(A) No delay**

| | | Sort 2 FR | Sort 2 UD |
|---|---|---|---|
| Sort 1 | FR | 7 | 3 |
| | UD | 3 | 10 |

**(B) 1-week delay**

| | | Sort 2 FR | Sort 2 UD |
|---|---|---|---|
| Sort 1 | FR | 4 | 4 |
| | UD | 1 | 13 |

*Note*: FR, family resemblance; UD, unidimensional.

## 4. General discussion

Experiments 1 and 2 show that under a free classification procedure, people tend to persist with their initial sort type (family resemblance or unidimensional) across subsequent sorts of different stimulus sets. Experiment 1 shows the basic effect and Experiment 2 shows that this effect is robust even when there is a delay of one-week between sorts. Indeed, this substantial delay did not have any reliable effect on sort type persistence. The fact that the previously reported spatial integrality effect, where spatially separable stimuli evoke greater family resemblance sorting than more spatially integrated stimuli (Milton & Wills, 2004), is seen in the first sort, but not in the second or subsequent sorts, is consistent with the idea that participants acquire a bias for a particular sort type as a result of the first sort in the experiment, and this bias overshadows the otherwise influential effects of stimulus properties in subsequent sorts. Our results, therefore, are in line with the perseveration account of sort type persistence. In contrast, an individual differences account of sort type persistence is unable to provide a full explanation of our findings because it cannot explain why the physical stimulus properties reliably affect the prevalence of family resemblance sorting during the first sort, but have no detectable effect on subsequent sorts.

The perseveration effect we report is in line with previous research that shows that people often demonstrate inflexibility in switching strategies across a variety of domains (e.g., Luchins, 1942; Schultz & Searleman, 2002). More specifically, our findings are related to the results of Murphy (2001), who showed that individuals have a strong tendency to sort multiple sets of simultaneously presented items in a self-consistent manner, and to those of Pothos and Chater (2005), who found that unsupervised classifications can be influenced by previous supervised categorization decisions. However, as far as we are aware, ours is the first direct demonstration that sort type persistence occurs in free classification. As such, these results add to the existing work that shows that stimulus presentation technique (Regehr & Brooks, 1995), spatial integration (Milton & Wills, 2004), time pressure (Milton et al., 2008), category structure (Pothos & Close, 2008), and background knowledge (e.g., Spalding & Murphy, 1996) all have a significant impact on sorting behavior.

The current demonstration of sort type persistence has important methodological implications. It is relatively common for participants to complete multiple sorts of different stimulus sets (e.g., Berger & Hatwell, 1995; Handel & Rhodes, 1980; Ward, 1983; Ward, Foley, & Cole, 1986). Our findings suggest that the results of some of these experiments could have been influenced, at least to some extent, by participants completing multiple sorts. On the basis of these findings, it is suggested that it is unwise for participants in future free classification studies to complete sorts from different experimental manipulations, unless it can be demonstrated for the particular procedures employed that no significant perseveration occurs (or that the perseveration cannot account for the results obtained).

The presence of sort type persistence in free classification provides a significant challenge to stimulus-driven models of free classification, such as SUSTAIN (Love et al., 2004) and the Simplicity model (Pothos & Chater, 2002). SUSTAIN has been shown to be capable of producing either unidimensional or family resemblance sorts, dependent on the form of input representations it receives, but there is no process or parameter in SUSTAIN that can represent an acquired bias for unidimensional versus family resemblance sorting strategy. Rather, the sorts produced by such models are the product of the application of a fixed set of principles to the stimuli being presented and to the memory of the stimuli previously classified. Models of this type almost certainly cannot accommodate, for example, the perseveration of a family resemblance sort strategy from the sorting of a set of abstract geometric shapes (Fig. 1A) to the subsequent family resemblance sorting of a set of semi-naturalistic depictions of butterflies (Fig. 1B). Perseveration of this type can only be predicted by such models if the stimulus representations encode the correspondence between the abstract structures of the two stimulus sets. It is not hard to imagine that stimulus representations could achieve this for the sets shown in Fig. 1B and D. If such a correspondence exists, models such as SUSTAIN can produce sort type persistence because the correspondence of the stimulus representations in the second sort with the memory representations of the stimuli in the first sort will mean that the second stimulus set will tend to be classified into the same groups as were formed for the first stimulus set. However, it does not seem possible to produce representations that could credibly be derived from the presented stimuli that would code the required correspondences between the stimulus set depicted in Fig. 1A, and the set depicted in Fig. 1B; for example, representations that could encode a correspondence between the upper stimulus in Fig. 1A and the upper stimulus in Fig. 1B, and the lack of correspondence between the upper stimulus in Fig. 1A and the lower stimulus in Fig. 1B.

An alternative approach, and one that involves greater modifications, would be to provide SUSTAIN with a system that allows it to modulate sort strategy independently of the stimulus structure. Previous work has shown that by modulating its attention parameters, SUSTAIN can account for the finding that presentation order of stimuli within a particular set influences the dimension participants employ when performing a unidimensional sort (Gureckis & Love, 2002). In a similar manner, SUSTAIN may also be able to adjust its general bias toward family resemblance and unidimensional sorts via its attentional parameters. In principle, SUSTAIN could be modified in such a manner that a strategy selection process, outside the current SUSTAIN model, could modulate the distribution of attention, and hence accommodate the sort type persistence effects observed in the current paper. Such an account has something in common with the Hypothesis theory account discussed later in this section.

The Simplicity model would also, in its current form, struggle to explain the sort type persistence effect. In Simplicity, categorization is driven by the abstract stimulus structure and items are assigned to the category which minimizes the information required to describe the object's structure. The Simplicity model can successfully account for a variety of different data sets; however, it has no memory system or other parameters that would enable it to account for the sort type persistence effect. It is possible that with the introduction of a memory system and additional parameters, the Simplicity model may be able to account for sort type persistence, but the specific nature of the required modifications is not immediately apparent.

The sort type persistence we report can be accommodated if one assumes the operation of a process of strategy selection, such as that proposed by Hypothesis theory (e.g., Levine 1971). Hypothesis theory states that when faced with a novel problem, people will tend to select a strategy which has previously been used successfully for similar problems. For our current experiments, this perceived similarity of problem type may be quite general, and could include aspects such as the test context. One avenue for future research would be to investigate the factors that might be effective in modulating the magnitude of the sort type persistence effect. Hypothesis theory suggests that both the similarity of the presented stimuli and the similarity of the testing contexts could be effective. Another possibility would be to introduce some aspect of summary feedback (e.g., "Expert botanists rate your classification of these butterflies as good/poor"). Hypothesis theory indicates that negative summary feedback should reduce sort type persistence.

Whilst Hypothesis theory can explain the sort type persistence effect, one could question whether it would also be able to account for the substantial body of data on the factors that influence the first sort (e.g., Milton & Wills, 2004; Milton et al., 2008; Pothos & Close, 2008; Regehr & Brooks, 1995). A theory that can explain both types of phenomena appears to be an appropriate goal for future research. The suggested modifications of SUSTAIN or Simplicity may be helpful in achieving this goal.

To summarize, stimulus-driven theories such as SUSTAIN propose a single-process account of classification, with differences in sort behavior resulting from differences in the stimuli that are presented. The results of the current experiments cannot easily be predicted by accounts of this type. In contrast, the current results seem straightforwardly accommodated by a multi-process account of classification in which the classifiers have a number of classification processes at their disposal, and in which the process selected in a given situation is determined by its previous success in similar situations.

Future work should assess the generality of the perseveration effect across different methodologies. For instance, one could examine whether manipulations such as the introduction of continuous-valued dimensions or allowing participants to create a greater number of categories have any impact on sort type persistence. In the future, it may also be preferable to randomly select the prototypes from the four pairs of 1-away stimuli and construct the remaining items round these new prototypes. On a somewhat different note, the descriptions that participants produced at the end of each sort are likely to have encouraged some kind of verbal reasoning processes. An interesting question is, therefore, whether these verbal processes might have encouraged sort type persistence. This could be assessed by asking participants to sort the sets without providing a verbal description of their strategy.

Finally, it is worth noting that the presence of acquired long-term sort type persistence in free classification provides one potential process by which background knowledge may operate. It is well known that the availability of relevant background knowledge can increase the prevalence of family resemblance sorting (e.g., Spalding & Murphy, 1996). One reason this might happen is that the similarity of the experimental materials in some background knowledge categorization studies to stimuli participants have classified pre-experimentally might lead to the selection of a family resemblance sort strategy (on the basis that this type of strategy had been effective pre-experimentally). If this analysis is correct, it should be possible to find situations where the availability of relevant background knowledge can decrease the prevalence of family resemblance sorting (because there are real-world categories for which family resemblance sorting is not appropriate). Although, to our knowledge, this experiment has not been done, the idea is supported in general terms by some of the findings of Wattenmaker (1995). Wattenmaker demonstrated that categories for which family resemblance sorting is likely to have been effective pre-experimentally (e.g., classifying people as introverted or extroverted) produce more family resemblance sorting in a free classification experiment than categories for which family resemblance sorting was likely to have been less effective pre-experimentally (e.g., something that is easy to grasp, has a flat surface, is one foot long, and is made of cardboard, has many of the features of a good hammer, but will not be effective as a hammer). In summary, one of the things background knowledge may provide is an idea of the type of classification strategy that is likely to be effective.

## Acknowledgements

This research was supported by ESRC Grants PTA-030-2003-00287, PTA-026-27-1256 and a Great Western Research Fellowship awarded to the first author, and by ESRC Grant RES-000-22-1779, BBSRC Grant 9/S17109 and EC Framework 6 project Grant 516542 (NEST) to the second author. We thank Chris Longmore, Alice Welham, Thomas Beesley, Mark Suret, Gareth Croft, Ian McLaren, Emmanuel Pothos, and two anonymous reviewers for their helpful comments and suggestions on earlier versions of this article.

## References


Allan, L. G. (1980). A note on the measurement of contingency between two binary variables in judgment tasks. *Bulletin of the Psychonomic Society, 15*, 147–149.

Berger, C., & Hatwell, Y. (1995). Development of dimensional vs. global processing in haptics: The perceptual and decisional determinants of classification skills. *British Journal of Developmental Psychology, 13*, 143–162.

Bradley, D. R., Bradley, T. D., McGrath, S. E., & Cutcomb, S. D. (1979). Type I error rate of the Chi-Square test of independence in R × C tables that have small expected frequencies. *Psychological Bulletin, 86*, 1290–1297.

Duncker, K. (1945). On problem-solving. *Psychological Monographs, 58,* 5(Whole No. 270).

Gureckis, T., & Love, B. C. (2002). Who says models can only do what you tell them? Unsupervised category learning data, fits, and predictions. In W. D. Gray, & C. D. Schunn (Eds.), *Proceedings of the 24th Annual Conference of the Cognitive Science Society* (pp. 399–404). Hillsdale, NJ: Erlbaum.

Handel, S., & Rhodes, J. W. (1980). Free classification: Element-level and subgroup-level similarity. *Perception and Psychophysics, 28*, 249–253.

Howell, D. C. (2002). *Statistical methods for psychology* (5th ed.). Boston: Duxbury Press.

Levine, M. (1971). Hypothesis theory and nonlearning despite ideal S-R-reinforcement contingencies. *Psychological Review, 78*, 130–140.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. *Psychological Review, 111*, 309–332.

Luchins, A. S. (1942). Mechanization in problem solving. *Psychological Monographs, 54*, 6(Whole No. 248).

McCloy, R., Beaman, C. P., Morgan, B., & Speed, R. (2007). Training conditional and cumulative risk judgements: The role of frequencies, problem-structure and Einstellung. *Applied Cognitive Psychology, 21*, 325–344.

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

Milton, F. N., & Wills, A. J. (2008). The influence of perceptual difficulty on family resemblance sorting. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 2273–2278). Austin, TX: Cognitive Science Society.

Milton, F. N., Longmore, C. A., & Wills, A. J. (2008). Processes of overall similarity sorting in free classification. *Journal of Experimental Psychology: Human Perception and Performance, 34*, 676–692.

Milton, F. N., & Wills, A. J. (2004). The influence of stimulus properties on category construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30*, 407–415.

Murphy, G. L. (2001). Causes of taxonomic sorting by adults: A test of the thematic-to-taxonomic shift. *Psychonomic Bulletin and Review, 8*, 834–839.

Pothos, E. M., & Chater, N. (2002). A simplicity principle in unsupervised human categorization. *Cognitive Science, 26*, 303–343.

Pothos, E. M., & Chater, N. (2005). Unsupervised categorization and category learning. *Quarterly Journal of Experimental Psychology, 58A*, 733–752.

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. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 347–363.

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

Schultz, P. W., & Searleman, A. (2002). Rigidity of thought and behavior: 100 years of research. *Genetic, Social, and General Psychology Monographs, 128*, 165–207.

Schultz, P. W., Stone, W. F., & Christie, R. (1997). Authoritarianism and mental rigidity: The Einstellung problem revisited. *Personality and Social Psychology Bulletin, 23*, 3–9.

Spalding, T. L., & Murphy, G. L. (1996). Effect of background knowledge on category construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*, 525–538.

Sweller, J. (1976). The effect of task complexity and sequence on rule learning and problem solving. *British Journal of Psychology, 67*, 553–558.

Sweller, J., & Gee, W. (1978). Einstellung, the sequence effect, and hypothesis theory. *Journal of Experimental Psychology: Human Learning and Memory, 4*, 513–526.

Ward, T. B. (1983). Response tempo and separable-integral responding: Evidence for an integral-to-separable processing sequence in visual perception. *Journal of Experimental Psychology: Human Perception and Performance, 9*, 103–112.

Ward, T. B., Foley, C. M., & Cole, J. (1986). Classifying multidimensional stimuli: Stimulus, task, and observer factors. *Journal of Experimental Psychology: Human Perception and Performance, 12*, 211–225.

Wattenmaker, W. D. (1995). Knowledge structures and linear separability: Integrating information in object and social categorization. *Cognitive Psychology, 28*, 274–328.

Woltz, D. J., Gardner, M. K., & Bell, B. G. (2000). Negative transfer errors in sequential cognitive skills: Strong-but-wrong sequence application. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 601–625.

Woltz, D. G., Gardner, M. K., & Gyll, S. P. (2000). The role of attention processes in near transfer of cognitive skills. *Learning and Individual Differences, 12*, 209–251.