Original Article

QJEP

# Individual and age differences in block-by-block dynamics of category learning strategies

Quarterly Journal of Experimental Psychology
2020, Vol. 73(4) 578–593

DOI: 10.1177/1747021819892584
qjep.sagepub.com
SAGE

**Reshma Gouravajhala[1], Christopher N Wahlheim[2] and Mark A McDaniel[1]**

**Abstract**

The present experiment examined individual and age differences in the dynamics of category learning strategies. Participants learned categories determined by a disjunctive rule with relational features through a feedback training procedure. During training, participants responded to strategy probes following each block to provide online assessment of the extent to which rule- and exemplar-based strategies were used throughout the training period. We introduced this measure as an alternative to model-based approaches to assessing individual differences in strategy use during training. Following training, participants classified ambiguous transfer objects that were assumed to distinguish between earlier use of rule- and exemplar-based learning strategies. We included this measure to obtain a relatively objective index of strategy use during training. Next, participants provided global ratings of their use of rule- and exemplar-based strategies during training. Results showed that strategy preferences expressed on the final training block predicted categorisation of ambiguous transfer objects better than global strategy reports. In addition, we utilised the block-by-block strategy reports to investigate the dynamics of learners' strategy preferences over the course of training. The findings revealed greater fluidity in strategy preferences for both younger and older adults than has been previously documented in the category learning literature. The novel block-by-block strategy reports in conjunction with the transfer-based approach allowed for a more nuanced examination of individual and age differences in strategy use and categorisation performance.

**Keywords**

Category learning; aging; individual differences; metacognition; strategies

Received: 22 May 2018; revised: 2 October 2019; accepted: 6 October 2019

Categorisation is a fundamental cognitive ability that allows people to group objects and events based on shared features. This ability allows both younger and older adults to learn and adapt to novel stimuli in the environment. Over the past 60 years, a central focus of category learning research has been to characterise the strategies and representational structures integral to this process. Much of the research on this topic has examined two approaches to such learning: rule- and exemplar-based strategies. Rule-based learning involves developing a rule that can be subsequently applied to categorise novel exemplars, whereas exemplar-based learning involves using memory for prior instances as basis for classification. Research has shown that the use of both strategies varies across categorisation tasks (Anderson & Betz, 2001; Ashby et al., 1998; Kruschke, 1992; D. R. Little et al., 2011), and across people within tasks (e.g., Craig & Lewandowsky, 2012; J. L. Little & McDaniel, 2015; McDaniel et al., 2014; Regehr & Brooks, 1993; Wahlheim et al., 2016).

In the present study, we extend previous work on individual variation in rule- and exemplar-based strategy use within a single task. Our novel contribution here is that we characterise individual and age differences in the dynamics of self-reported strategy selection when younger and older adults learn complex rule-based categories across multiple trials. To examine such strategy dynamics, we used a strategy-probe procedure that enabled us to assess self-reported use of exemplar- and rule-based strategies

[1]Department of Psychological and Brain Sciences, Washington University in St. Louis, St. Louis, MO, USA
[2]The University of North Carolina at Greensboro, Greensboro, NC, USA

**Corresponding author:**
Reshma Gouravajhala, Department of Psychological and Brain Sciences, Washington University in St. Louis, 1 Brookings Drive, Campus Box 1125, St. Louis, MO 63130, USA.
Email: rgouravajhala@wustl.edu

across multiple learning trials. Importantly, this approach allowed us to examine several theoretical possibilities regarding the stability of strategy selection that have as yet not been directly tested. We describe those possibilities below after describing our approach and its advantages for investigating these possibilities.

## Approaches to assessing category learning strategies

A common approach to assessing the use of rule- and exemplar-based categorisation strategies is to fit computational models to category learning performance during a training period (Ashby et al., 1998; Kruschke, 1992; Love et al., 2004; Nosofsky et al., 1994). In this model-based approach, the processes underlying categorisation performance are characterised by comparing model fits, yielding what is presumed to be an objective assessment of strategy use. This approach has been heavily preferred in part because it requires precise specification of hypothesised computations during learning. However, a shortcoming of this approach is that conclusions about individual differences in strategy use might be heavily dependent on the details of the model itself (Donkin et al., 2015; Edmunds et al., 2018). It is therefore important to seek converging evidence for individual variation in strategy selection using multiple methods.

A complementary approach to model-based evaluation of strategy use is to ask participants to report their use of rule- and exemplar-based strategies and compare their reports to performance on a transfer task (e.g., J. L. Little & McDaniel, 2015; McDaniel et al., 2014; Regehr & Brooks, 1993; Wahlheim et al., 2016). In this transfer-based approach, participants complete a training phase and then categorise novel transfer objects for which responses that are based on either the correct rule or on memory for training exemplars are in opposition. These transfer objects are ambiguous in that they include both relationships among features that afford categorisation according to the correct rule, and perceptual characteristics that are highly similar to training exemplars from the category opposite of which should be chosen based on the correct rule. Thus, categorisation on the basis of memory for training exemplars is shown when very few objects are categorised according to the correct rule. Finally, participants provide retrospective reports of the extent to which they used rule- or exemplar-based learning strategies during the training period.

Using this approach, J. L. Little and McDaniel (2015) and Wahlheim et al. (2016) recently found that some learners categorised ambiguous transfer items according to perceptual similarity (perhaps reflecting an exemplar-based representation), whereas others categorised those same items according to an abstracted rule. Important for the present study, Little and McDaniel, and Wahlheim et al. also found variation in the extent to which people reported using one strategy or the other that corresponded with their performance on ambiguous transfer objects. These results raise the possibility that some people varied in their strategies across blocks during training, but this potential heterogeneity could not be detected with a global retrospective report. Consequently, a more fine-grained approach to evaluating strategy use across blocks is needed.

## Strategy-probe methodology (block-by-block reports)

Prior research has used self-reported global strategy orientations (reported after a transfer test phase) to classify people's strategy preferences (J. L. Little & McDaniel, 2015; Wahlheim et al., 2016; see also Bourne et al., 2010). Global strategy questionnaires require participants to reflect upon a categorisation training phase, and report which strategy (rule- or exemplar-based) most closely aligned with their approach to the task. However, this method is ill-suited to address the primary research questions in the present study for two reasons.

First, a global retrospective strategy questionnaire places a sizable memory demand on learners because they are asked to reflect on their strategy preferences for a long series of earlier training blocks. To be accurate on this task, people must be able to remember their strategies across many blocks, which would be especially difficult if such strategy use varied. Moreover, these heavy episodic memory demands might be especially troublesome for older adults, who experience memory deficits when tasks require self-generated reinstatement of contextual features to cue past experiences (Balota et al., 2000; Zacks et al., 2000). These considerations suggest caution when interpreting global self-reports and attendant performance on transfer tasks, especially for older adults.

Second, a global strategy report cannot sufficiently inform the range of theoretical possibilities we will outline below with regard to strategy dynamics during category learning. We believe that a global strategy questionnaire is only useful insofar as learners adopt a particular learning process at the outset of category learning, and consistently use it throughout the task (e.g., initiate hypothesis testing to learn a rule (Levine, 1975; Trabasso & Bower, 1968) or focus on learning exemplar-category label pairings (Kruschke, 1992; Medin & Schaffer, 1978; Nosofsky, 1984). Therefore, a global strategy self-report may not be unambiguously sensitive to any shifting in strategies that might occur. Specifically, for any learners who might use multiple strategies during the task, it would be unclear whether their global ratings reflect a strategy preference at the start of the task, at the end of the task, or some mixture (e.g., when ratings are given that are intermediate to rule or exemplar-based strategies). Further complicating matters, these learners might differ in how they comprehend the judgement (e.g., whether to report initial preferences, final preferences, or a blending of the two).

To improve upon the global-self report method used by Little and colleagues (J. L. Little & McDaniel, 2015; Wahlheim et al., 2016), we developed a procedure that should be more sensitive to strategy use dynamics during training. In the current procedure, strategy probes appeared after each training block that asked participants to report the extent of their rule- and exemplar-based strategy use on only the immediately preceding block. By asking people to report their strategy use after each block, we minimised episodic memory demands. We believe this procedure is better suited to evaluate strategy dynamics and to inform the central theoretical issues of this study, which we will detail next.

## Strategy preference dynamics

One primary goal of the present study was to characterise individual and age differences in strategy dynamics during training. As previously mentioned, some prior work has shown that preferences for rule- and exemplar-based category learning strategies vary across people (J. L. Little & McDaniel, 2015; McDaniel et al., 2014; Regehr & Brooks, 1993). Moreover, such individual variation in preferences has been shown to correspond reasonably well with transfer performance for both younger and older adults (Wahlheim et al., 2016). One possible interpretation of such results is that people have stable strategy preferences, and those preferences vary across people.

Another possibility is that global self-reports and transfer performance reflect an amalgamation of the fluctuation in approaches to learning during training. Consistent with this, there are several lines of evidence suggesting that people can exhibit representational shifts during training (Anderson & Betz, 2001; Hoffmann et al., 2013; Johansen & Palmeri, 2002; Kalish et al., 2005; Nosofsky et al., 1994; Raijmakers et al., 2014). Relevant to the present study, Bourne et al. (1999) used a trial-by-trial self-report strategy procedure to find that, over the course of training, participants transitioned from a rule-based strategy to an exemplar-based strategy (see also Dixon & Bangert, 2002, for similar findings using a problem-solving task). Finally, Kalish et al. (2005) more recently extended this line of research to investigate potential factors underlying strategy shifting. They found evidence of strategy shifting between two rules (a single-dimensional rule vs. a complex bi-dimensional rule) and theorised that participants shifted in the face of high performance error during training. Here, we leveraged the strategy-probe procedure to directly test Kalish et al.'s (2005) claim that high performance error precedes strategy shifting.

It is also possible that within a single category-learning task, there exist individual differences in strategy dynamics such that some people persist with a preferred strategy (even when faced with high-performance error), whereas others shift among multiple strategies during the task. To date, these potential individual differences in strategy shifting have not been examined using a transfer-based approach (see Figure 1).

A final more exploratory possibility that we consider is that nearly all participants are fluid in their strategy preferences, which leads to at least one shift during training. This perspective draws inspiration from Levine's (1975) theory that during training, people select a focal hypothesis to guide each decision, but they also monitor all available hypotheses not currently in use. The idea is that people can evaluate multiple hypotheses during a category learning task, but for any given decision, they have a focal hypothesis that guides responding (also see Bourne, 1974). Here, we test the possibility that a similar process might also operate at a broader level. In particular, when an accurate solution can be reached using rule-based or exemplar-based strategies, learners might adopt a focal strategy for a block of items, but not in an exclusive manner. Learners might also notice information relevant to an alternative strategy, and for other blocks, adopt that alternative as a focal strategy to find a solution. (For a formal model that assumes that both rule-based and exemplar-based solutions are considered, but one system guides responding on any particular trial, see Erickson & Kruschke, 1998; ATRIUM model.) This view diverges from previous ideas that participants tend to favour one or the other strategy and persist with that strategy throughout the task. Accordingly, unlike the "persistent" strategy view, this new hypothesis suggests that the majority of learners will exhibit repeated strategy switching as a function of their focus on any given set of trials.

## Age-related differences in strategy preferences and dynamics

A second objective of the present study was to explore whether healthy older adults would exhibit one of the four patterns of strategy preferences outlined above and to determine how such preferences compare to those of younger adults. In the past, comparisons of younger and older adults' strategy preferences showed that a similar number of participants from both groups used rule-based strategies when learning complex bi-dimensional categories defined by rules that were difficult to verbalise (e.g., see Filoteo & Maddox, 2004; Maddox et al., 2010). Furthermore, more older than younger adults used rule-based learning in a multiple-cue categorisation task that included categories that were presumably of lower complexity (Mata et al., 2012). Most relevant to the present study, however, Wahlheim et al. (2016) used a transfer-based approach combined with global strategy reports to find comparable frequencies of younger and older adult rule-abstractors and memorisers. Given this, it is possible that younger and older adults in the present study do not differ in their strategy preferences throughout the category learning task.

However, there are other age-related factors that we believe might influence the strategy dynamics of older, relative to younger, adults during the category learning task. For example, older adults experience declines in executive function, resulting in poorer performance on tasks require fast processing, cognitive flexibility, and mental transformation (Hasher & Zacks, 1988; Lezak et al., 2012; Murman, 2015). Thus, it is possible that older adults exhibit more strategy persistence (even in the face of high performance error), and subsequently report fewer strategy switches than younger adults.

Alternatively, we believe it is possible that older adults will switch strategies more often than younger adults. According to the Competition Between Verbal and Implicit Systems model (COVIS; Ashby et al., 1998), older adults experience a verbal system deficit which results in an impaired rule-learning ability. By this account, older adults might experience more task success if they choose to endorse an exemplar-based strategy. But, older adults also have a well-documented memory decline that could hinder their ability to form and remember associations between objects and their category labels during training, thus impacting exemplar-based learning (Balota et al., 2000; Zacks et al., 2000). Owing to these cognitive deficits that impact their ability to effectively use either rule- or exemplar-based strategies, older adults might face higher performance error than younger adults (see e.g., Wahlheim et al., 2016). If high performance error precedes strategy shifting (as theorised by Kalish et al., 2005), then we would expect that older adults might be more likely, relative to younger adults, to exhibit strategy shifting throughout the task. Importantly, this issue of age-related differences in strategy preference dynamics has been largely unaddressed in the literature.

Before presenting the study, one additional feature warrants mention. We thought it possible that aforementioned age-related deficits found in Wahlheim et al. (2016) may have occurred because older adults require more training trials than provided to develop a rule or memorise all exemplars. To this end, we increased the amount of training in the present study by 33%, resulting in a total of 16 training blocks (relative to Wahlheim et al.'s 12 blocks). Note that because younger adults' training performance in Wahlheim et al. was at ceiling by the final training block, we expected that additional training would not markedly improve their performance.

## The present experiment

As a brief overview of the present paradigm, participants were trained on the category membership of eight objects (4 per category) using a feedback training procedure over 16 blocks. Shape and colour features of the objects were varied such that they could be categorised based on a disjunctive rule applied to relational features (as in J. L. Little & McDaniel, 2015; Wahlheim et al., 2016). We retained the previously used global strategy rating scale, wherein participants were asked to think back to the training phase and to judge whether they primarily engaged in a rule-abstraction strategy or a memorisation strategy. However, the primary focus was on the block-by-block ratings (similar to those successfully used in Bourne et al. [1999] for another type of conceptual-learning task) that participants provided at the end of every training block.

Here, we first demonstrate that the strategy-probe procedure has little if any reactive effect on strategy choices across category learning trials. We then leverage a combination of strategy probe reports and transfer performance to inform the following issues. First, we examined whether learners prefer a particular strategy (rule- or exemplar-based) and use that strategy consistently throughout training or whether they shift strategies, perhaps in response to performance error. To anticipate, many people shifted their strategies during training. Consequently, we examined whether people generally shift back and forth between strategies as they try to arrive at a solution (cf. Erickson & Kruschke, 1998), or if some participants choose to shift strategies while others persist. Finally, we examine whether these patterns differ between younger and older adults, as some theories would predict (e.g., Ashby et al., 1998).

## Method

Below we report how we determined our sample size, all data exclusions, all manipulations, and all measures in this experiment (Simmons et al., 2011). The stimulus materials and raw anonymised data can be found at the following URL: https://osf.io/ps87g/. The experiments were administered using E-Prime 2 software (Psychology Software Tools, Pittsburgh, PA, USA). The research reported here was approved by the Institutional Review Board of Washington University in St. Louis.

### *Participants*

The participants were 60 younger adults (20 men; years of age: mean *[M]* = 19.75; standard deviation [*SD*] = 1.32) and 60 older adults (20 men; years of age: *M* = 73.02; *SD* = 5.92). The younger adults were undergraduates at Washington University in St. Louis who received partial course credit for their participation. The older adults were members of the St. Louis community who received $10/hr for their participation. Though we recognise that our sample size yielded a slightly underpowered experiment, we matched the present sample sizes with those used by Wahlheim et al. (2016) to compare key findings across studies. All participants were tested individually. Participants completed the vocabulary subtest of the Shipley Institute of Living Scale (Shipley, 1940).

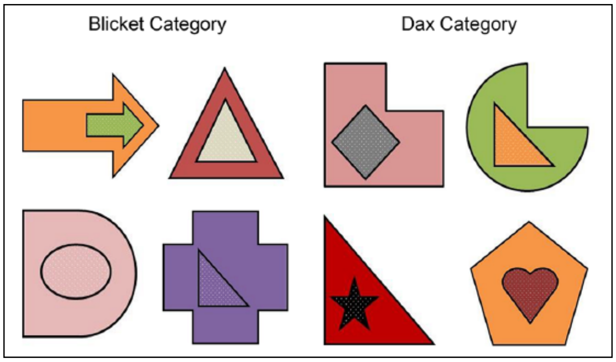


**Figure 1.** The eight objects above were presented in the category learning phase of the experiment. The four objects on the left represent the Blicket category and include inner and outer shapes that are the same in either colour or form. The four objects on the right represent the Dax category and include inner and outer shapes that differ in both colour and form.

Vocabulary scores were significantly higher for older adults ($M = 36.67$, $SD = 3.32$) than younger adults ($M = 34.43$, $SD = 2.81$), $t(118) = 3.99$, $p < .001$, $d = 0.73$. Older adults also reported significantly more years of education ($M = 16.53$, $SD = 2.83$) than younger adults ($M = 13.53$, $SD = 1.29$), $t(118) = 7.49$, $p < .001$, $d = 1.36$.

## *Procedure and materials*

The experiment lasted approximately 1 hr. Participants completed several tasks in the following order: (a) training, (b) ambiguous object categorisation (transfer), (c) global strategy questionnaire, (d) rule object categorisation (transfer), (e) memorisation object categorisation (transfer), and (f) verbal cued recall.[1] We describe these tasks in more detail below.

*Training.* We modelled the training phase procedure on that used by Wahlheim et al. (2016). Participants completed a feedback learning procedure during which they learned to categorise objects into one of two categories (see Figure 1). Objects were made up of two coloured shapes with one shape inside the other. The categories were determined by a disjunctive rule that related the inner shape to the outside shape. If objects' inner and outside shapes matched on either colour or form, that object belonged to the "Blicket" category. If the inner and outer shapes differed in both colour and form, the object belonged to the "Dax" category.

Participants received 16 blocks of training that each comprised eight objects (four from each category) presented individually in the centre of a computer screen against a white background. For two of the objects from the Blicket category, the form of the inner and outer shapes matched, and for the other two, the colour of the inner and outer shapes matched (Figure 1, left panel). For all four objects from the Dax category, the inner and outer shapes differed in both form and colour (Figure 1, right panel). Participants were instructed to assign each object to a category by pressing the "S" key for the Blicket category and the "L" key for the Dax category. The labels "Blicket" and "Dax" appeared on the computer screen below each object on the left and right side of the screen, respectively. The eight objects appeared in a different predetermined random order in each of the 16 training blocks, with the constraint that no more than two objects from the same category could appear sequentially within a block. To examine individual differences, each participant received the same predetermined order. The objects remained on the screen until the participant pressed either the "S" or "L" key. After each response, the correct category label ("Blicket" or "Dax") appeared on the screen for 2 s as feedback.

After each training block, participants were probed about the extent to which they used rule-based, exemplar-based, and other undefined strategies to categorise the objects in the immediately previous block. The following three questions appeared in random order after each block: "How often did you apply a rule?"; "How often did you memorise objects and their category names?"; "How often did you use a strategy other than rule-learning or object-memorizing?" Participants rated their strategy use on a Likert-type scale ranging from 1 (*never*) to 5 (*always*). A prompt reminding participants of this scale appeared below each of the questions. Participants were instructed to give an extreme rating only when they used a strategy exclusively during a block, and intermediate ratings to indicate the extent to which they used each strategy when they did not exclusively use one strategy.

*Ambiguous object categorisation (transfer).* Participants categorised a set of eight new ambiguous objects into the Blicket and Dax categories that were similar in form to training objects, but, according to the rule, belonged to the opposite category (see Figure 2, left panel, for examples). Objects and category label prompts appeared in the same manner as in the training phase. Participants made their categorisation decisions using the same keys as during training. No feedback was provided during this phase to minimise any additional learning that could affect categorisation decisions in the subsequent transfer tasks.

*Global Strategy Questionnaire.* After categorising ambiguous objects, we informed participants that we were interested in how their learning processes worked. We told them that we would assess this by asking them a series of questions (see Appendix A, for the exact wording of each question). We first instructed participants to think back to the training phase and remember whether they were mostly trying to memorise the objects or mostly trying to establish a rule. Participants reported their strategy using a Likert-type scale ranging from 1 (*memorisation strategy*) to 7

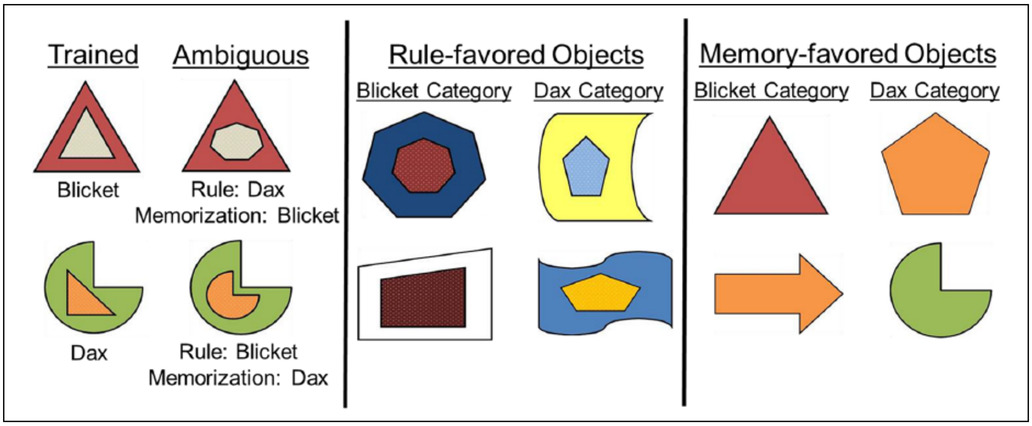


**Figure 2.** Panel 1: Examples of objects from the training phase (left column) and their respective ambiguous objects from the transfer phase (right column). The ambiguous objects place rule-abstraction and memorisation strategies in opposition to one another. Categorisation of ambiguous objects on the basis of rules results in assignment to the category opposite that to which the perceptually similar training items belonged. Panel 2: examples of novel objects used to assess rule-abstraction for categories of training objects. Panel 3: examples of novel objects used to assess memory for categories of training objects. The interested reader can find the full set of transfer objects on OSF.

(*rule-establishing strategy*). As with the block-by-block ratings during the training phase, participants were instructed to give an extreme rating only if they used that strategy exclusively during training. Participants were instructed to give an intermediate rating to indicate the extent to which they had used each strategy when they did not rely exclusively on one strategy. Participants were also told to give an exact intermediate score (4) if they had used both strategies equally or if they were unsure about which strategy they had used.

Participants were then instructed to think back to the most recent phase, where they had categorised ambiguous objects, and to make ratings similar to those they had made for the training phase. In this case, however, the memorisation strategy was framed in terms of categorising an object based on its similarity to one of the objects in the training phase. More specifically, participants were instructed to indicate the extent to which they had focused on the similarity between the "new" shapes in the transfer phase and the "old" shapes in the training phase. Participants made their ratings on a Likert-type scale ranging from 1 (*similarity*) and 7 (*rule*), and again, participants were instructed to only give extreme ratings if they exclusively relied on that particular strategy during the ambiguous categorisation transfer phase.

After providing these ratings, participants were probed about their specific understanding of the rules that distinguished the Blicket and Dax categories. First, they were asked to explain the rule that they used to classify objects, if they developed one. Next, if they attempted to develop a rule but could not do so perfectly, they were instructed to write down the best rule they came up with. Finally, if they stated a rule, they were asked to indicate when (either during the training phase or the transfer phase) they had developed the rule.

To reiterate, we retained this global strategy measure to compare the findings of the present study to those in Wahlheim et al. (2016) when necessary. However, we were primarily interested in the aforementioned block-by-block strategy probes.

*Rule-favoured object categorisation (transfer).* In this phase, participants were presented with eight new objects whose inner and outer shapes were perceptually dissimilar to any objects they had seen before (see Figure 2, middle panel), and told to classify each object as a member of either the Blicket or Dax category. Four of these objects followed the rule from the Blicket category, and the other four followed the rule from the Dax category. Object and category labels appeared as in earlier phases, the same keys were used for categorisation decisions, and no feedback was provided during this phase.

*Memory-favoured object categorisation (transfer).* Participants were presented with a final set of eight new objects that included only the outer shape of the objects shown in the initial training phase with their original colour (see Figure 2, right panel), and again instructed to categorise them. Object and category labels appeared as in earlier phases, the same keys were used for categorisation decisions, and no feedback was provided during this phase.

## Results

### *Overview*

We first present Bayesian analyses testing whether introducing block-by-block probes produced a reactive influence on training accuracy, relative to Wahlheim et al. (2016). We then classified individuals according to both

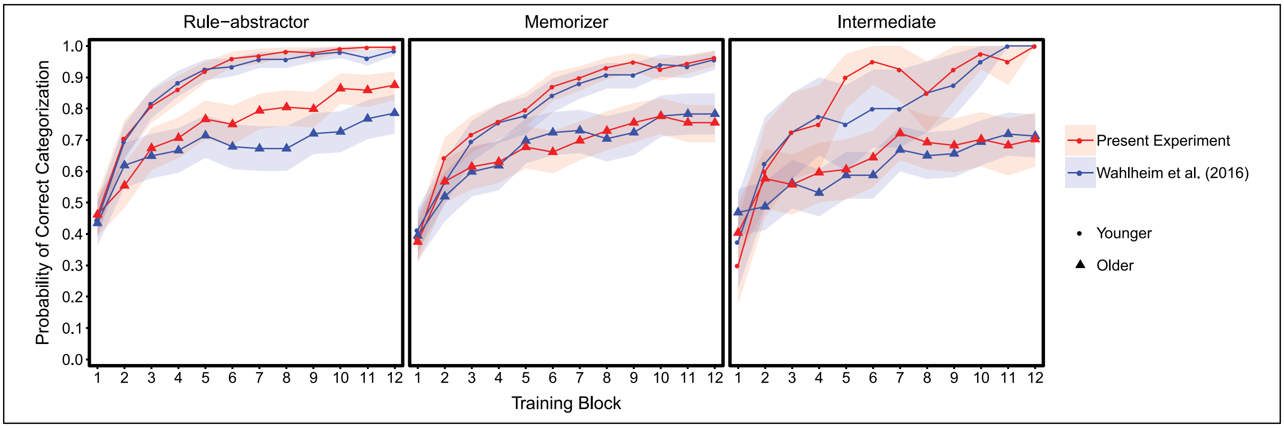


**Figure 3.** Probability of correct categorisation across training blocks as a function of age and global strategy preference in the present experiment (in red) and in Wahlheim et al. (2016; in blue). Shaded regions are 95% confidence intervals.

their block-by-block probes and their global strategy preferences to determine how well their transfer performance was captured by each classification method (block probes vs. global probe). Next, we address the questions raised in the introduction regarding age-related differences in strategy preference and training performance, and then finally, explore the possibilities surrounding strategy dynamics.

Bayesian analyses were conducted using the *anovaBF* function in the *R* package *BayesFactor* (Morey et al., 2015). Logistic mixed-effect regression models were created using the *glmer* function in the *R* package *lme4*, and subsequent linear combination tests were conducted using the *glht* function in the *R* package *multcomp* (Bates et al., 2015). Regression tables were created using the *sjPlot* package in *R* (Lüdecke, 2017; please find full regression tables on OSF). The level of significance was set at $\alpha = .05$.

## *Probe reactivity check*

To determine if the introduction of the block-by-block strategy probes led to reactive effects on training performance, we compared the present learning curves with those reported by Wahlheim et al. (2016) (see Figure 3) using a Bayesian approach. Specifically, we computed Bayes factors (BF) for all possible combinations of fixed factors and interactions, with an intercepts-only model as our null and experiment as a predictor (described in further detail below). We set *whichModels* in the *anovaBF* function to "bottom," which created models by adding single factors and/or interactions to the null model, allowing us to determine each factors' effect.

In our analysis of variance (ANOVA) model, we predicted participants' training accuracy (averaged across each block) by the following factors: four main effects (Experiment [coded 0 for the present study, 1 for Wahlheim et al., 2016], Age Group [coded 0 for younger adults, 1 for older adults], Strategy[2] [coded 0 for rule-abstractor, 1 for memoriser, 2 for intermediate], and Block [1 – 12]). We also included all possible interactions in the model, resulting in a total of 15 predictors. To address the issue of reactivity of strategy probes across experiment, we calculated the BF values for the main effect of Experiment, as well as all interactions including the Experiment factor. The addition of Experiment to the null model yielded a BF value of 1.09, indicating a lack of strong evidence in favour of rejecting an a priori assumption that block-by-block probes should not alter participants' training performance. The BF values of the interaction terms were more conclusive, however. The three two-way interactions of Block x Experiment (BF = .0000016), Strategy x Experiment (BF = .04), and Age Group x Experiment (BF = .13) all corresponded to low BF values. For the three-way interaction terms, Strategy x Age Group x Interaction yielded the highest BF value (BF = 3.3), while the BF values for other two interaction terms including Experiment as a factor (BF[Block x Strategy x Experiment] = .0000055 and BF[Block x Age Group x Experiment] = .00013) were both less than 1. Critically, the ratio of the BF for the four-way interaction (Block x Strategy x Age Group x Experiment) to the BF of the three-way interaction without the Experiment term (Block x Strategy x Age Group) was considerably less than 1 (BF ratio = .0000002). Taken together, these findings indicate that the block-by-block strategy probes did not produce a reactive influence on participants' training performance in the present study, relative to Wahlheim et al. (2016).

## *Alignment of participants' block-end responses with transfer performance*

After establishing that the block-by-block probes did not have a reactive influence on individuals' training performance, we then examined whether participants' block-level preferences would be accurately reflected in their subsequent transfer performance. We believed that strategies used during transfer would correspond to participants' final representations, and so we developed a simple heuristic to classify individuals according to their strategy preference on the 16th training block. Specifically, we compared participants' response on the rule and memory probes[3] that followed the final block of training. If participants exhibited a difference between these strategies, they were classified according to the strategy they had assigned a numerically higher value. Because participants were presented with the strategy probes sequentially, we believed any numerical difference between them to be meaningful. For example, if any individual reported a 4 (on a 1–5 scale) on the rule probe, and a 3 on memory, we classified them as having a "rule-end" preference. Individuals who endorsed rule and memory strategies equally on the final block were determined to belong to the "no preference end" category. The frequencies of participants classified according to this block-end method are displayed in Table 1.

To determine whether participants' self-reported block-end strategies corresponded to their categorisation of ambiguous, rule, and memory objects (Figure 4), we conducted a no-intercept mixed-effects logistic regression model containing 18 fixed effects with accuracy coded using 0s and 1s. We chose to fit a no-intercept model because we were primarily interested in the effects of age and self-reported learning strategy (during the last training block) on transfer performance.[4] Each of the 18 fixed effects represented a group (e.g., older adults with a rule-end strategy categorising ambiguous objects) that was dummy coded (using 0s and 1s) for group membership. Overall, the model accounted for 53.35% of the total variance (conditional $R^2 = .53$).

**Table 1.** Frequencies of self-reported block-end training strategies.

| Age group | Block-end strategy | | |
|---|---|---|---|
| | Rule | Memory | No preference |
| Younger | 19 | 30 | 11 |
| Older | 10 | 27 | 23 |

We conducted a series of linear combination tests to determine significance of any effects (see Table 2). Younger adults outperformed older adults in categorisation of rule ($z = -5.23$, $p < .001$) and memory ($z = -6.59$, $p < .001$) objects, and categorised ambiguous objects significantly more according to the rule ($z = -3.93$, $p < .001$), collapsed across all block-end strategy preferences. Relative to those with a memory-end strategy preference, participants with a rule-end strategy preference categorised ambiguous objects significantly more according to the rule ($z = 9.33$, $p < .001$), and exhibited better categorisation performance on rule ($z = 5.66$, $p < .001$) but not memory objects ($z = -2.45$, $p = .21$). Similarly, relative to participants in the no-preference-end group, those with a rule-end preference categorised ambiguous objects more according to the rule ($z = 5.13$, $p < .001$), and had better categorisation performance on rule ($z = 4.05$, $p = .001$), but not memory objects ($z = -.88$, $p = .99$).

Compared to those in the no preference-end group, those with a memory-end strategy preference categorised ambiguous objects significantly less according to the rule ($z = -4.55$, $p < .001$), and showed equal performance on rule ($z = -2.17$, $p = .38$) and memory objects ($z = 1.41$, $p = .91$). The age x rule-end vs. memory-end strategy preference was significant for rule ($z = -3.40$, $p = .01$) and ambiguous ($z = -5.00$, $p < .001$) objects, but not for memory objects ($z = 2.16$, $p = .39$). Specifically, younger adults' higher accuracy relative to their older adult counterparts was especially true for those with a rule-end preference. The age x rule-end vs. no-preference-end strategy interaction was significant only for ambiguous objects ($z = -3.48$, $p = .01$), such that, again, younger adults especially outperformed older adults in categorisation of ambiguous objects according to the rule when they exhibited a rule-end strategy preference. The age x memory-end vs. no-preference-end was not significant for any of the transfer object types, all $p$s $> .05$.

**Table 2.** Categorisation performance on transfer tasks as a function of block-end training strategy.

| Object type | Age | Block-end strategy | | |
|---|---|---|---|---|
| | | Rule-end | Memory-end | No preference-end |
| Ambiguous | Younger | 0.95 [0.87, 1.00] | 0.25 [0.13, 0.37] | 0.52 [0.26, 0.78] |
| | Older | 0.75 [0.51, 1.00] | 0.24 [0.14, 0.34] | 0.42 [0.33, 0.52] |
| Rule-favoured | Younger | 0.99 [0.97, 1.00] | 0.65 [0.58, 0.73] | 0.77 [0.63, 0.92] |
| | Older | 0.80 [0.60, 1.00] | 0.55 [0.46, 0.65] | 0.60 [0.52, 0.67] |
| Memory-favoured | Younger | 0.89 [0.83, .95] | 0.98 [0.95, 1.00] | 0.93 [0.85, 1.00] |
| | Older | 0.80 [0.57, 1.00] | 0.76 [0.68, 0.84] | 0.71 [0.64, 0.78] |

*Note.* Performance on ambiguous and rule-favoured objects reflect probabilities of categorisation according to the correct rule. For ambiguous objects, performance of 1.00 would indicate perfect categorisation according to the correct rule, and performance of 0.00 would indicate perfect categorisation according to memory for perceptually similar features of training objects. Performance on memory-favoured objects reflects probabilities of correct categorisation based on memory for training objects with perceptually similar features. Margins of error for 95% confidence intervals are displayed in parentheses.

Taken together, these findings illustrate that individuals' self-reported block-end strategy preferences generally captured their category-learning approach and representation by the end of training. We also addressed the issue of whether participants' block-end strategies more accurately captured their transfer performance than their global training strategy preferences (see Table 3 for frequencies). We re-classified individuals according to their global strategy preferences, and conducted the same analyses as above (Figure 5, top panel). Because these analyses are secondary to the overarching goals of the present experiment, we report these findings in Appendix B. However, to briefly compare the two classification methods, it appears that transfer performance was better predicted when people were classified by their block-end than global strategy reports. Particularly, there are clear differences between the two classification methods' alignment with participants' performance on both ambiguous and rule objects. Finally, we assessed whether memory load was an underlying factor for these global strategy reports (which required participants to think back to the training phase before selecting a strategy preference) to be worse predictors of transfer performance, relative to block-end strategy reports. To do so, we re-conditionalised participants according to their global strategy reports focused on the *transfer* phase (using the same heuristic as was used for global training strategy reports), and then re-analysed transfer performance. As shown in Figure 5 (bottom panel), this classification method better aligned with learners' transfer performance, relative to global strategy reports focused on the training phase. While global transfer strategy reports corresponded to transfer performance in a similar manner to the block-end strategy reports, only the latter allowed us to examine both training and transfer strategy preferences and performance. Given that the primary focus of this study is on establishing the efficacy of a block-by-block methodology, we do not mention participants' global strategy reports hereafter.

**Table 3.** Frequencies of self-reported global training strategies.

| Age group | Global training strategy | | |
|---|---|---|---|
| | Rule-abstraction | Exemplar memorisation | Intermediate |
| Younger | 28 | 24 | 5 |
| Older | 23 | 27 | 13 |

### *Age-related differences in block-end strategy preferences*

Having established that block-end strategy preferences generally reflected the approach that participants arrived at to learn and represent the trained category, we are now in a position to address whether there were age differences in learning strategies at the end of training (see Table 1). We initially computed a 2 (age: younger vs. older) x 3 (block-end strategy: rule-end vs. memory-end vs. no preference-end) chi-square test of independence,

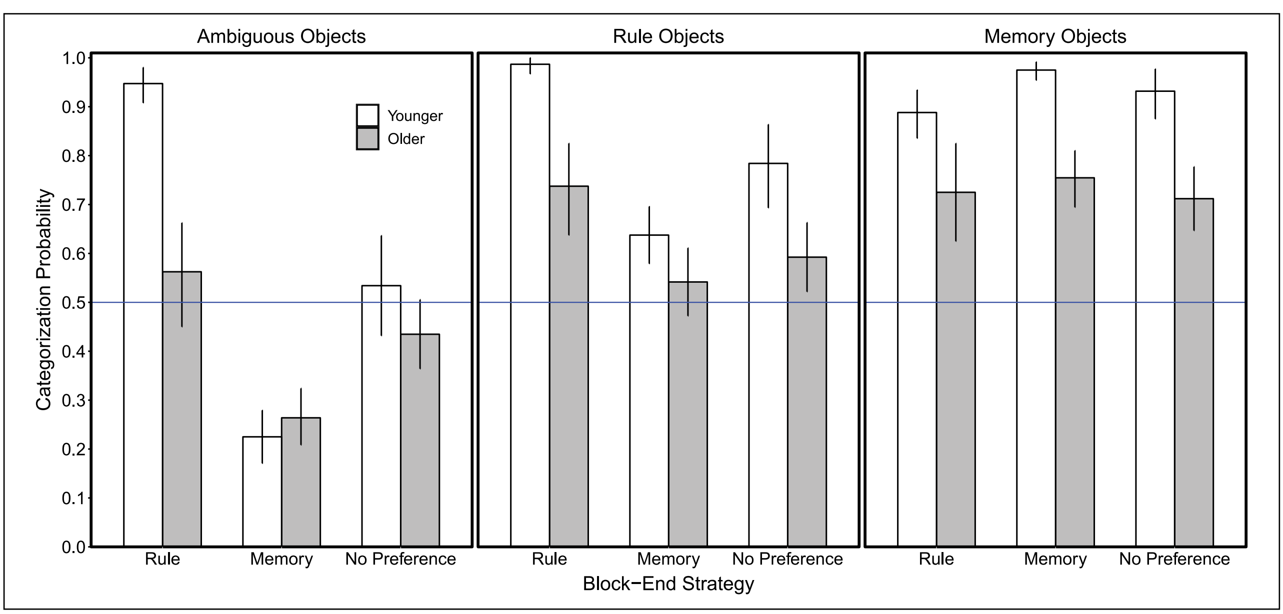


**Figure 4.** Categorisation performance on ambiguous, rule, and memory transfer objects as a function of age and block-end strategy preference. For ambiguous and rule objects, categorisation probability refers to the probability of classification according to the correct rule. For memory objects, categorisation probability refers to recognition accuracy for the outer shape of a trained object.

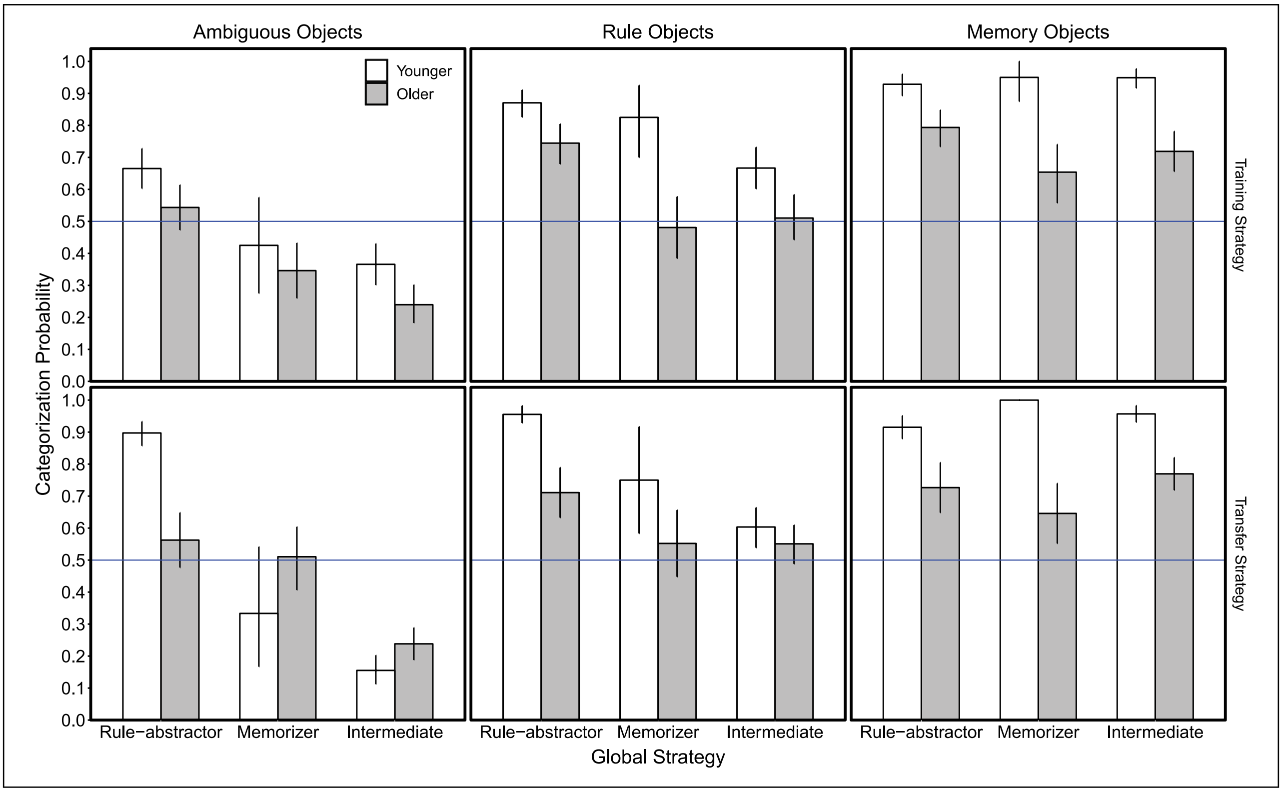


**Figure 5.** Categorisation performance on ambiguous, rule, and memory transfer objects as a function of age and global training strategy preference (top panels) and global transfer strategy preference (bottom panels). For ambiguous and rule objects, categorisation probability refers to the probability of classification according to the correct rule. For memory objects, categorisation probability refers to recognition accuracy for the outer shape of a trained object.

which revealed a significant age x block-end strategy interaction, $\chi^2(2) = 7.19$, $p = .03$.

However, given that the chi-square test could not reveal the underlying reasons for this interaction (given the 2 x 3 design), we then conducted a generalised linear model (using a Poisson error distribution) to predict frequency by age (dummy coded using 0s and 1s) and block-end strategy (with group membership dummy coded using 0s and 1s). Along with these dummy-coded main effects, we also included three age x block-end strategy interaction terms to create a fully saturated model that accounted for all variability in frequencies. Next, we conducted a series of theoretically motivated linear combination tests of significance to determine if there were age-related differences in each of the three strategy groups. As predicted, we found an age group difference in the no-preference-end strategy, such that older adults were more likely to belong to that group than younger adults, $z = -2.01$, $p = .04$. However, despite the sizable numerical difference between younger and older adults in the rule-end strategy group, we found no statistically significant age-related difference in participants' endorsement of either rule-end or memory-end strategies, both $ps > .05$.

## *Age-related differences in training performance*

As described in the Introduction, one of the changes in the present study relative to Wahlheim et al. (2016) was the addition of four training blocks (i.e., a 33% increase). The purpose of this procedural change was to investigate whether additional experience would diminish previously reported age differences in training performance. We first examined whether older adults were able to achieve the same level of training performance as younger adults by the final block (see Figure 6). With participants classified according to their block-end strategies, we conducted a no-intercept logistic mixed-effects regression model to predict training accuracy (coded using 0s and 1s) by age and strategy. Six groups representing each combination of age and block-end strategy groups were dummy coded for group membership using 0s and 1s. Note that the no-intercept model allowed for the linear and non-linear effects of block (1–16) to affect each of these groups differently. Thus, the final model contained 18 total fixed effects (six group terms, six group x block interaction terms, and six group x block-squared interaction terms). Overall, this model accounted for 44.78% of the variance (conditional $R^2 = .45$).

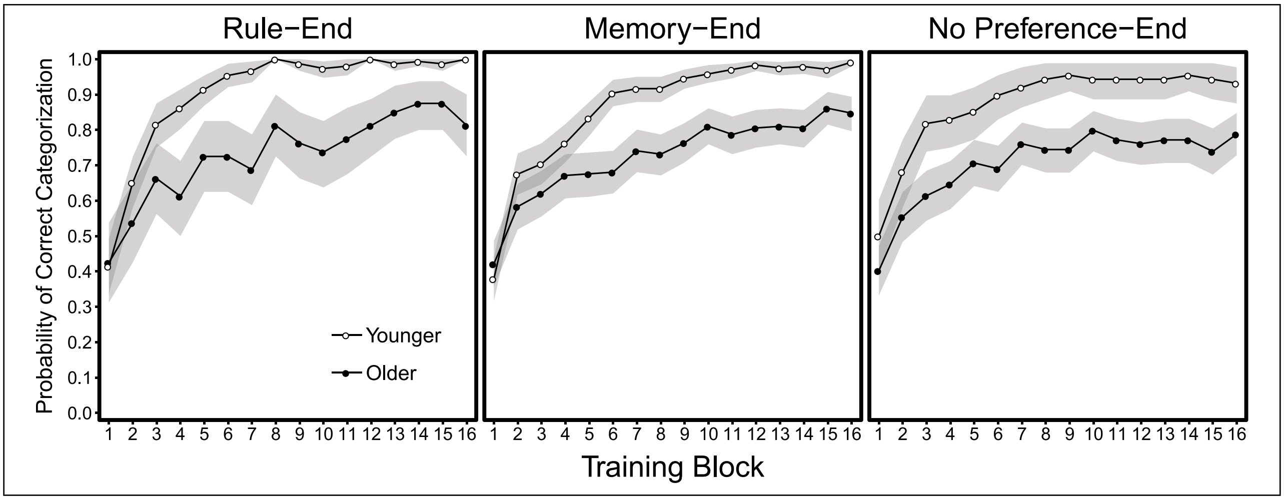


**Figure 6.** Probability of correct categorisation across training blocks as a function of age and block-end strategy preference. Shaded regions are 95% confidence intervals.

We then conducted a series of linear combinations to examine the significance of each effect. The linear combination test of age was significant ($z$=-11.154, $p<.001$), indicating that older adults performed worse than younger adults in the training phase. None of the linear combination tests comparing the three block-end strategies (in a pairwise fashion) was significant, all $p$s $>.05$. The linear combination test examining an age difference in slopes was significant ($z$=-9.54, $p<.001$), indicating that, across all strategy groups, younger adults exhibited faster learning than did older adults over the course of training. The linear combination test of an age difference in the extremity of non-linearity was also significant ($z=5.89$, $p<.001$), such that, collapsing across strategy groups, younger adults reached asymptote in learning after fewer trials than older adults. Moreover, this pattern remained true within each strategy group (lowest $z=3.04$, highest $p=.02$).

We were also interested in determining if older adults in the present study showed higher learning at the end of 16 blocks of training compared to older adults at the end of 12 blocks of training in Wahlheim et al. (2016). We compared training accuracy on the final block across experiments using an independent samples t-test. There was no significant difference between older adults from the present study ($M=.82$, $SD=.20$) and from Wahlheim et al. ($M=.79$, $SD=.13$), $t(118)=.95$, $p>.05$.

## Evidence and patterns of strategy shifting

To address the question of whether strategy shifting existed during our category learning task, we first identified each participant's strategy preference on each block using the same point-differential heuristic used for participants' block-end strategy preferences (see Figure 7). Next, we determined whether individuals switched strategy preferences (rule to memory, rule to no preference, memory to no preference, or vice versa) in consecutive blocks. We considered all six of these patterns to be strategy switches, as they indicate a change in preference (whether to the opposite strategy, or to equal use of strategies) from the previous block. Finally, we calculated the total number of switches for each participant. As shown in Figure 8, there was clear evidence of strategy shifting, as most participants (93.3%) switched at least once during training.

Having established the existence of strategy shifting in the present study, we were next interested in whether there was individual variation in strategy persistence. In other words, for people who reported the same strategy (rule, memory, no preference) in both Block 1 and Block 16, would they exhibit differences in strategy preference dynamics in the intervening blocks? There were 16 younger adults and 21 older adults who reported the same preference at the start and end of training. Of these participants, only five younger adults and three older adults reported zero switches throughout training. Thus, it appears that the majority of people who started and ended the task with the same preference also exhibited some degree of strategy shifting in the task.

Finally, we were interested in potential age differences in strategy shifting. As discussed in the Introduction, we believed it possible that older adults, who experience both verbal system and episodic memory deficits, would shift more often than younger adults. Consequently, we conducted an independent samples $t$-test comparing shift frequency across all younger and older adult participants. As predicted, older adults shifted strategies more frequently than younger adults (older: $M=4.20$, $SD=2.74$ vs. younger: $M=2.68$, $SD=2.28$), $t(114)$=-3.30, $p=.001$.

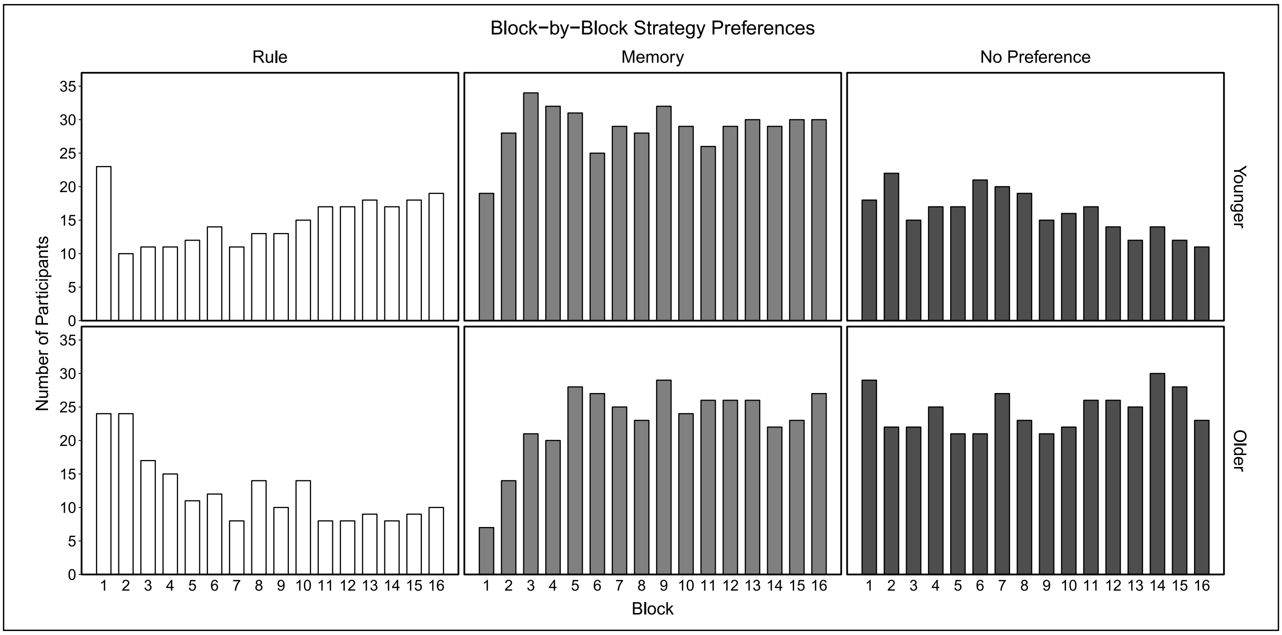


**Figure 7.** Histogram of younger and older adults' block-by-block strategy preferences over the course of training.

### Performance error and strategy shifting

Finally, we tested Kalish et al.'s (2005) claim that strategy switching occurs when people experience high performance error and are aware that there is an alternative strategy. To do this, we fitted hierarchical logistic regression models to participants' block-by-block strategy reports to predict strategy shifts (coded using 0s and 1s) using standardised error rates. In other words, participants' switching behaviour on blocks 2–16 was predicted by their standardised error rates on blocks 1–15. We allowed both intercepts (participants' baseline likelihood to switch) and slopes (variability in participants' response to error rates) to vary randomly. We conducted a model comparison between a model where slopes were not allowed to vary (which would indicate that all participants responded to changes in error rate in the same way) and one in which they were, and found that the latter model better fit the data, $F(3, 5)=91.31$, $p<.001$.

Because we did not have strong theoretical predictions about whether younger or older adults would be more susceptible to high performance error, we collapsed across age in the model. As shown in Table 4, participants were 1.33 times more likely ($OR=1.33$, $p=.001$) to switch strategies with one SD increase in error on the previous block. This represented a 57.0% probability to switch with an increase in error, relative to a 24.0% probability of switching at average levels of error. Thus, it appears as though strategy switching was precipitated by higher-than-average levels of performance error.

## Discussion

In the present study, we examined individual and age differences in category learning strategies and subsequent representations using a complex rule-based category learning task. We assessed strategy use dynamics using a self-report method in which people rated their strategy use after each training block. This approach extends previous work by Bourne et al. (1999), and departs from the global retrospective report method that has typically accompanied a transfer-based approach (cf. J. L. Little & McDaniel, 2015; Wahlheim et al., 2016). We had three main goals. First, we sought to determine the effectiveness of the strategy-probe procedure in assessing individual and age differences in strategy dynamics during training. Second, we used the strategy-probe procedure to test theoretically motivated hypotheses regarding strategy shifting during training (Kalish et al., 2005). Third, we were interested in how strategy dynamics characterised using the strategy-probe procedure may differ between younger and older adults. We discuss the present findings that inform each of these goals in turn below.

### The utility of the strategy-probe method

Global strategy reports used in previous transfer-based studies have provided useful information about individual variation in strategy preferences and their consequences for category representations (Bourne et al., 2010; J. L. Little & McDaniel, 2015; Wahlheim et al., 2016). However, global reports have two key limitations that we sought to overcome using the strategy-probe method in the present study. One key limitation is that global reports place a substantial burden on episodic memory retrieval, and this may be exacerbated for older adults, given their well-established episodic memory deficit (e.g., (Balota et al., 2000; Zacks et al., 2000). Another key limitation is that global

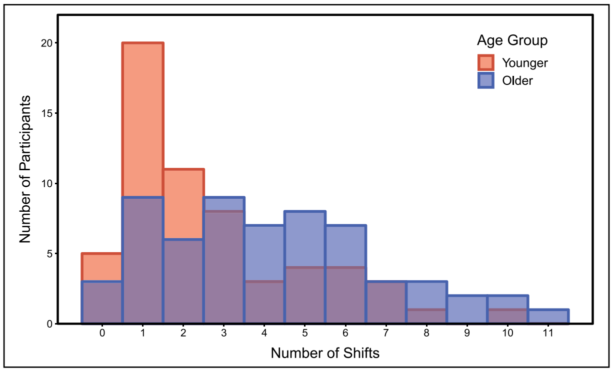


**Figure 8.** Histogram of younger and older adults' strategy shifts over the course of training.

**Table 4.** Standardised error predicting shifts.

| Predictors | Shift | | |
|---|---|---|---|
| | Odds ratios | CI | *p* |
| (Intercept) | 0.32 | 0.20–0.49 | <0.001 |
| Z Error | 1.33 | 1.12–1.57 | 0.001 |
| Random effects | | | |
| $\sigma^2$ | 3.29 | | |
| $\tau_{00\ \mathrm{Subject}}$ | 2.80 | | |
| $\tau_{11\ \mathrm{Subject.\ BlockLag}}$ | 0.10 | | |
| $\rho_{01\ \mathrm{Subject}}$ | –0.86 | | |
| Interclass correlation | 0.46 | | |
| $N_{\mathrm{Subject}}$ | 120 | | |
| Observations | 1,800 | | |

strategy reports cannot directly inform whether strategy use during training varies within and across people. We addressed these limitations by implementing a strategy-probe method that required people to report their strategy use following each training block. We believed that collecting more reports each after fewer trials would both alleviate episodic memory demands and enable a more nuanced examination of strategy dynamics.

A comparison of training performance between the present study and Wahlheim et al. (2016) revealed comparable category learning dynamics, despite the fact that the present study included strategy-probes during training, whereas Wahlheim et al. did not. The similarity in training performance between studies strongly suggests that the strategy-report method did not have reactive effects on training performance. Having verified that strategy-probes are a somewhat non-intrusive method for assessing strategy dynamics, we then leveraged reports on the final block to more accurately characterise the category representations that participants developed during training and subsequently used during transfer. Our assumption here was that the lower memory demand for the final block as compared to global reports would enable more precision in participants' assessment of their most recent representational state. Consistent with this assumption, we found that block-end preferences corresponded better with categorisation of ambiguous and rule-based objects than did global reports. These results showed that block-end reports are more sensitive than global reports to individual and age differences in category representations formed during training. More generally, these results indicate that the strategy-probe method is a viable way to identify the factors that lead to variations in category learning strategies. We believe that future experiments comparing the outcomes of transfer- and model-based approaches could benefit from the inclusion of frequent strategy reports during training.

## *Dynamics of category learning strategies*

We addressed the central issue of individual variation in strategy dynamics and tested the four theoretical possibilities outlined in the Introduction by assessing learning preferences for each training block. To briefly review, the first possibility was that individual differences in rule- and exemplar-based learning would emerge at the outset of training and persist across blocks (J. L. Little & McDaniel, 2015; McDaniel et al., 2014). The second possibility was that learners would switch between strategies, perhaps in adaptation to performance error (Kalish et al., 2005). The third possibility was that aggregate differences in strategy dynamics would reflect a balanced mix between learners who persisted with a preferred strategy, and others who shifted among strategies. The fourth possibility was that most participants would exhibit fluidity in their strategy preferences, switching repeatedly between the two strategies as a function of their focus on any given block (cf. Erickson & Kruschke, 1998; Levine, 1975).

The patterns of strategy dynamics observed in the present study strongly aligned with the fourth possibility: Most learners showed fluid strategy switching throughout training. Indeed, 93.3% of participants shifted at least once and 32.5% of learners shifted at least five times over the course of the 16 training blocks. These findings lend support to the idea that learners might adopt a focal strategy on any given block, but switch if they notice information pertinent to the alternative strategy (e.g., Levine, 1975). Strategy preferences are not decided in an exclusive manner and can be adapted to fit the learners' point of focus on any particular block. In contrast with previous work that suggests that learners experience single strategy shifts (e.g., from a rule-based strategy to an exemplar-based strategy, or a complex rule to a simple rule) over the course of training (e.g., Hoffmann et al., 2013; Johansen & Palmeri, 2002; Kalish et al., 2005), our findings strongly suggest that strategy shifting is a much more fluid and common occurrence in category learning tasks than has been previously considered or appreciated. Future studies examining participants' preferences (whether using computational modelling or self-report strategy reports) should be sensitive to this possibility.

In addition to revealing the dynamics of strategy use, the strategy-probe method also allowed us to directly test the idea that high performance error precedes strategy shifting (Kalish et al., 2005). Kalish et al. found evidence of strategy shifting in their task where the two categorisation strategies (a single-dimensional rule of varying validity and a bi-dimensional rule) differed in accuracy. Participants who experienced high performance error with the single-dimensional rule, the authors reasoned, were more motivated to shift to the more accurate bi-dimensional rule. The present study extended the Kalish et al. finding to demonstrating categorisation-strategy shifts when two strategies, specifically rule- and exemplar-based strategies, were equally viable—as indicated by similar error rates associated with each strategy during training (see Figure 3, left and middle panels). Of issue is whether strategy shifts between equally viable strategies would be driven by performance error, at least in part. Consistent with the Kalish et al. finding, strategy shifting in the present study was largely precipitated by performance error. Learners who experienced higher than average error rates on any given block were more likely to switch strategies on the subsequent block. It is important to note that there were some individual differences in this pattern (i.e., not all participants who were faced with high performance error on a given block shifted strategies on the subsequent block), and so more work must be conducted to help identify the root of these individual differences. Nevertheless, the present results reinforce that strategy shifting within a category learning task can be stimulated by performance error (Kalish et al.) and suggests that this is the case even when shifting between two equally effective strategies (on average).

### *Age-related differences in strategy preferences and training performance*

We found no significant differences in the endorsement of rule-end and memory-end strategies between younger and older adults (see Table 1 for frequencies). However, older adults were significantly more likely, relative to younger adults, to exhibit no strong strategy preference (i.e., belong in the no preference end group) at the end of training (consistent with Wahlheim et al., 2016, global strategy reports). We believe that because older adults experience deficits in both their episodic memory and verbal system (Ashby et al., 1998; Balota et al., 2000), they showed equal reliance on both strategies to glean any useful information in learning to categorise the stimuli. Relatedly, older adults exhibited a higher frequency of strategy shifts relative to younger adults, perhaps reflecting the need to rely on both strategies to improve categorisation performance (see Figure 8).

While older adults' strategy preferences were largely similar to those of younger adults, their performance on the category learning task was considerably worse. The present study's 33% increase in training blocks relative to Wahlheim et al. (2016) was not effective in achieving the goal of improving older adults' performance. Older adults in the present study did not show significant improvement in training performance relative to older adults in the previous study, suggesting that simply increasing the number of training blocks was not sufficient to mitigate previously identified age-related performance differences. We found that less than half of the older adults who ended with a strong strategy preference showed perfect accuracy at the end of training, whereas almost all younger adults did. Furthermore, only 5 older adults (compared to 25 younger adults) were able to fully develop the rule by the end of training.[5] Overall, younger adults showed greater and faster learning during the training phase, and outperformed older adults on categorisation of all three types of transfer objects. Given these findings, older adults may require even more training (or coaching of relevant features; cf. Miyatsu et al., 2019) to show improvement, perhaps because their impairments in episodic memory and a verbal rule-learning system are too great to overcome.

## Conclusion

In the present study, we focused on individual and age differences in strategy dynamics and performance on a rule-based category learning task and developed a novel strategy-probe methodology that was highly sensitive to largely unaddressed nuances in learners' strategy preferences. Classification of participants according to their strategy probes yielded greater alignment with objective measures of transfer (relative to the more commonly used global strategy reports), leading us to encourage their use in future studies. Having established their utility, we then used learners' strategy reports to examine strategy dynamics over the course of training and found that participants were more fluid with their strategy shifts than has been previously theorised or revealed, to our knowledge. Furthermore, building on and extending prior work (e.g., Kalish et al., 2005), we found that strategy shifting was, in part, a response to performance error. Finally, we confirmed previously documented age differences in strategy preferences and categorisation performance (see Wahlheim et al., 2016), and extended these findings to report how older adults differ from younger adults in their strategy dynamics. Importantly, older adults appear to rely on strategy shifting more so than younger adults in attempting to learn the category. Doing so, however, did not in the present category-learning task enable older adults to achieve the learning levels displayed by younger adults.

### Author note

The data reported here were presented in part at the 2016 Cognitive Aging Conference in Atlanta, Georgia. The data and

stimulus materials are available through the follow URL: https://osf.io/ps87g/.

### Acknowledgements

We express appreciation to Jack Basse and Melanie Marcille for their assistance with data collection.

### Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

### Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported in part by a training grant from the National Institute on Aging: T32-AG000030.

### ORCID iD

Reshma Gouravajhala https://orcid.org/0000-0002-9370-0246

### Supplementary material

The Supplementary Material is available at: qjep.sagepub.com

### Notes

1. We initially included this measure to address hypotheses about age-related individual differences that were not central to the main focus of the paper. Consequently, we do not report a description of the task or participants' performances; the interested reader can find these on OSF.
2. Following earlier studies, we segmented participants within each age group according to the training strategy that they reported on the global strategy questionnaire that they completed immediately following the ambiguous object categorisation task. Participants who gave ratings of 5, 6, or 7 were classified as rule-abstractors; participants who gave ratings of 1, 2, or 3 were classified as memorisers; and participants who gave a rating of 4 were classified as intermediate learners. See Table 3 for frequencies.
3. Because it is unclear what strategies participants may have been using when answering the "other" probe question and because participants consistently gave low ratings for "other" strategy use, we did not include those strategy reports in determining a rule or memorisation preference.
4. In our no-intercept models, there is no single global intercept for everyone (as is traditional). Instead, each of the six groups get their own intercept (fixed effect), and each individual subject has their own random intercept which varies around their group's fixed effect. In other words, random intercepts are included in the models.
5. To identify participants' ability to learn the correct rule, the first and second authors independently coded the rules that participants reported using. The raters showed good initial agreement (88% of responses) resulting in Cohen's Kappa=.79, and resolved disagreements through discussion. Reported rules were categorised as correct, partial, and incorrect. Correct responses were those in which participants articulated the exact rules for both Blickets and Daxes, or for one of the categories and noted that all other objects belonged to the other category. Partial responses were those in which participants stated a rule that would correctly classify some, but not all, of the objects. Finally, responses were categorised as being incorrect when they consisted of rules that did not work, ambiguous utterances, or when no response was offered. The interested reader can find a table of frequencies on OSF.

### References

Anderson, J. R., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin & Review*, *8*, 629–647.

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

Balota, D. A., Dolan, P. O., & Duchek, J. M. (2000). Memory changes in healthy older adults. In E. Tulving & F. I. M. Craik (Eds.), *The Oxford handbook of memory* (pp. 395–409). Oxford University Press.

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). lme4: Linear mixed-effects models using Eigen and S4 (R Package 1.1-10). *Journal of Statistical Software*, *67*, 1–48.

Bourne, L. E. Jr. (1974). Stimulus selection and memory strategies in concept problems. *Memory & Cognition*, *2*, 270–274.

Bourne, L. E. Jr., Healy, A. F., Parker, J. T., & Rickard, T. C. (1999). The strategic basis of performance in binary classification tasks: Strategy choices and strategy transitions. *Journal of Memory and Language*, *41*, 223–252.

Bourne, L. E. Jr., Raymond, W. D., & Healy, A. F. (2010). Strategy selection and use during classification skill acquisition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*, 500–514.

Craig, S., & Lewandowsky, S. (2012). Whichever way you choose to categorize, working memory helps you learn. *Quarterly Journal of Experimental Psychology*, *65*, 439–464.

Dixon, J. A., & Bangert, A. S. (2002). The prehistory of discovery: Precursors of representational change in solving gear system problems. *Developmental Psychology*, *38*, 918–933.

Donkin, C., Newell, B. R., Kalish, M., Dunn, J. C., & Nosofsky, R. M. (2015). Identify strategy use in category learning task: A case for more diagnostic data and models. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *41*, 933–948.

Edmunds, C. E., Milton, F., & Wills, A. J. (2018). Due process in dual process: Model-recovery simulations of decision-bound strategy analysis in category learning. *Cognitive Science*, *42*, 833–860.

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, *127*, 107–140.

Filoteo, J. V., & Maddox, W. T. (2004). A quantitative model-based approach to examining aging effects on information-integration category learning. *Psychology and Aging*, *19*, 171–182.

Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension, and aging: A review and a new view. In K. D.

Federmier (Ed.), *Psychology of Learning and Motivation* (vol. 22, pp. 193–225). Academic Press.

Hoffman, J. A., von Helversen, B., & Rieskamp, J. (2013). Deliberation's blindsight: How cognitive load can improve judgments. *Psychological Science*, *24*, 869–879.

Johansen, M. K., & Palmeri, T. J. (2002). Are there representational shifts during category learning? *Cognitive Psychology*, *45*, 347–354.

Kalish, M. L., Lewandowsky, S., & Davies, M. (2005). Error-driven knowledge restructuring in categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 846–861.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychology Review*, *99*, 22–44.

Levine, M. (1975). *A cognitive theory of learning*. Lawrence Erlbaum.

Lezak, M. D., Howieston, D. B., Bigler, E. D., & Tranel, D. (2012). *Neuropsychological assessment* (5th ed.). Oxford University Press.

Little, D. R., Nosofsky, R. M., & Denton, S. E. (2011). Response-time tests of logical-rule models of categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*, 1–27.

Little, J. L., & McDaniel, M. A. (2015). Individual differences in category learning: Memorization versus rule abstraction. *Memory & Cognition*, *43*, 283–297.

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

Lüdecke, D. (2017). *sjPlot: Data visualization for statistics in social science* (R Package Version 2.3.3). Retrieved from https://rdrr.io/cran/sjPlot/

Maddox, W. T., Pacheco, J., Reeves, M., Zhu, B., & Schnyer, D. M. (2010). Rule-based and information-integration category learning in normal aging. *Neuropsychologia*, *48*, 2998–3008.

Mata, R., von Helversen, B., Karlsson, L., & Cüpper, L. (2012). Adult age differences in categorization and multiple-cue judgment. *Developmental Psychology*, *48*, 1188–1201.

McDaniel, M. A., Cahill, M. J., Robbins, M., & Wiener, C. (2014). Individual differences in learning and transfer: Stable tendencies for learning exemplars versus abstracting rules. *Journal of Experimental Psychology: General*, *143*, 668–693.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.

Miyatsu, T., Gouravajhala, R., Nosofsky, R. M., & McDaniel, M. A. (2019). Feature highlighting enhances learning of a complex natural-science category. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *45*, 1–16.

Morey, R. D., Rouder, J. N., & Jamil, T. (2015). *BayesFactor: Computation of Bayes factors for common designs* (R Package 0.9.12-2). Retrieved from https://rdrr.io/cran/BayesFactor/

Murman, D. L. (2015). The impact of age on cognition. Seminars in Hearing, *36*, 111–121.

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 104–114.

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, *104*, 266–300.

Raijmakers, M. E., Schmittmann, V. D., & Visser, I. (2014). Costs and benefits of automatization in category learning of ill-defined rules. *Cognitive Psychology*, *69*, 1–24.

Regehr, G., & Brooks, L. R. (1993). Perceptual manifestations of an analytic structure: The priority of holistic individuation. *Journal of Experimental Psychology: General*, *122*, 92–114.

Shipley, W. C. (1940). A self-administering scale for measuring intellectual impairment and deterioration. *The Journal of Psychology: Interdisciplinary and Applied*, *9*, 371–377.

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, *22*, 1359–1366.

Trabasso, T., & Bower, G. H. (1968). *Attention in learning: Theory and research*. Wiley.

Wahlheim, C. N., McDaniel, M. A., & Little, J. L. (2016). Category learning strategies in younger and older adults: Rule abstraction and memorization. *Psychology and Aging*, *31*, 346–357.

Zacks, R. T., Hasher, L., & Li, K. Z. H. (2000). Human memory. In T. A. Salthouse & F. I. M. Craik (Eds.), *Handbook of aging and cognition* (2nd ed., pp. 293–357). Lawrence Erlbaum.