Title

A Trade-Off in Learning Across Levels of Abstraction in Adults and Children

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

Authors
Wharton-Shukster, Erika
Finn, Amy Sue

Publication Date
2019

Peer reviewed

A Trade-Off in Learning Across Levels of Abstraction in Adults and Children

Erika Wharton-Shukster (e.wharton.shukster@mail.utoronto.ca)
Department of Psychology, 100 St George St
Toronto, ON M5S 3G3 Canada

Amy Sue Finn (amy.finn@utoronto.ca)
Department of Psychology, 100 St George St
Toronto, ON M5S 3G3 Canada

Abstract

Learning about novel objects not only involves noticing information that makes the object unique, but also what makes objects the same. Yet, these two levels of learning involve different pieces of information, meaning that learning one well could come at the cost of the other. Moreover, children may categorize in a fundamentally different way, resulting in these levels of learning interacting differently. To investigate this, we had adults and children perform a categorization task followed by an item recognition test. We found that adults showed a trade-off, such that the ability to categorize items came at the cost of memory for those items. Using a subset of more unique lures, children's memory trended towards a trade-off with category learning. However, this was only observed among the older children. This suggests that adults' efficient learning comes at a cost, and this trade-off may start to appear in the elementary school years.

Keywords: cognitive development; category learning; abstraction; generalization; memory; selective attention

Background

Learning about objects goes beyond simply identifying unique features. We also learn abstract information, picking up on consistencies across objects. These different levels of abstraction are each informative in different ways, as the former provides details specific to an individual object while the latter facilitates categorization and generalization across objects. Given the discrepancy in the information learned, it seems likely that learning information really well at one level of abstraction might impede learning at the other, resulting in a trade-off in learning.

In fact, we know that categorical knowledge can impact memory in important ways. For instance, the Deese-Roediger-McDermott (DRM) paradigm has been found to show memory distortions in adults for words on a list when the presented words are categorically related. In this case, the category-level information that connects the words on the list causes memory distortions, a loss of memory of the specific words themselves (item-level information) and false memory for words that did not occur. This memory distortion is not seen when the word list is not categorical, suggesting that it is the abstraction of the category that is obscuring the details and producing distortions (see Brainerd, Reyna, & Ceci, 2008).

In addition, Sloutsky and Fisher (2004) found a drop in adults' memory for specific animals after participants sorted them categorically. When left uncategorized, memory for the individual animals was good, however following an induction task that required sorting the animals into species, memory for the individuals dropped to chance levels.

One possible explanation for this pattern of behaviour is the longstanding fuzzy-trace theory, in which abstraction necessarily involves a distillation of information to a vague, detail-free, gist representation (Brainerd & Reyna, 1990). In other words, abstraction is facilitated by a lack of detailed, item-level information. Similarly, the schema literature would suggest a comparable process in the organization of our knowledge, as schemas are abstract representations of something (be it a place, animal, social interaction) accumulated through experience that create expectations for the future (Mandler, 1984). Details from specific experiences are filtered out and consistencies are used to create a generic representation. Again, it is the loss of detailed information that makes a schema so generalizable.

Given this previous work and theory, there could be a trade-off in learning item-level and category-level information. Yet, work exploring this relationship has to date only included pre-existing categories and hasn't tackled whether a trade-off might occur *during* or in the service of category *learning*. It remains unclear how novel category learning would affect the interaction of item- and category-level information.

An eye-tracking study that assessed attention during a novel category learning task may provide some insight into this question (Rehder & Hoffman, 2005). In this study, as participants were learning to categorize the stimuli, they were found to fixate on all features of a stimulus. However, once they had successfully learned to categorize, they were found to fixate only on the diagnostic feature. This narrowing of focus would likely result in better categorization behavior, likely at the expense of learning about non-category relevant features of the objects, thereby producing a trade-off in object and category learning.

In contrast, when it comes to category learning, an abstract representation could still be formed not by ignoring irrelevant features of objects, but by learning *all* of the features of objects—both relevant and irrelevant—well. After all, abstracting to learn a category only necessitates learning the relevant information for that category and not necessarily ignoring what is irrelevant. This approach to category learning would of course result in a different relationship with item memory: there would not be a trade-off.

There is some research suggesting that children may not demonstrate a trade-off in item-level and category-level information. For instance, in the DRM paradigm discussed above, adults consistently fall prey to false memories when the word lists are thematic. Interestingly, children do not succumb to the same memory distortions. In fact, this paradigm finds that young children have few false memories, with that number steadily increasing across the elementary school years and peaking in adulthood (Brainerd et al., 2008). In this case, children seem to have no thematic intrusions and remember the item-level information despite their category membership, resulting in no trade-off.

A similar result was found by Sloutsky and Fisher (2004), discussed above. Whereas the adults' memory for individual animals dropped after categorization, children's memory remained consistent. Here again, children maintained memory for the item information despite categorization. Given these findings, it is likely that children will similarly not show a trade-off during category *learning*, although the answer to this question is as yet unknown.

Along these lines, it has been suggested that instead of utilizing only the diagnostic dimensions, children categorize by including all item-level information (Sloutsky, 2010). Indeed, support for a holistic approach to categorization was found in children but not adults; adults were found to use only diagnostic features to categorize, while children were found to use the entirety of the item, basing their categorization on overall similarity (Smith & Kemler, 1977). This difference may reflect children's developing ability to selectively attend, as they are less successful at suppressing irrelevant information (Rueda, Posner, & Rothbart, 2005).

Furthermore, these differences in attention are likely to impact memory. For instance, during change detection and search tasks, children have also been found to have superior memory for task-irrelevant information compared to adults, suggesting that children's distributed attention facilitates memory for task-irrelevant information (Plebanek & Sloutsky, 2017). Together, these findings raise a final question. If children are attending to all available information when learning to categorize, will they retain item-level information despite successfully learning to categorize? In other words, might children be immune to the trade-off in item-level and category-level learning that we expect to see in adults?

To answer these questions, two experiments were performed, one with adults and one with children. In each, participants performed an A/B categorization task to measure category learning and a recognition memory test to measure item memory. To assess how specifically the items were remembered, half of the recognition foils were similar to the categorization stimuli along an orthogonal (not categorically-diagnostic) dimension, and half were dissimilar.

Experiment One

Methods

Participants Participants included 60 undergraduate students from the University of Toronto participating for course credit (*M* = 19.73 years, 76% female).

Materials and Procedure The category learning task consisted of 60 trial-unique trials of a feedback driven A/B sort task. Participants were instructed to sort "amoebas" into one of two categories based on the feedback given. They were not told what features defined category membership. Each stimulus was presented for 1.5 seconds or until a response was given, and stimulus presentation order was randomized between participants. The task was conducted on an Apple desktop computer using PsychoPy (Peirce, 2008).

The stimuli were designed to vary categorically along one dimension and orthogonally along two dimensions. Category membership was defined by distortions of two prototypical dot patterns shown below (Fried & Holyoak, 1984; Seger et al., 2000; Figure 1a). The 84 exemplars were generated by allowing dots a 7% chance of differing from the original. No exemplars were repeated across stimuli.

The orthogonal dimensions included colour and shape. Unique colours were randomly assigned, and shapes were created by making two extremely different shapes—generated from images of paint splatter—and morphing them together to varying degrees to create a series of related shapes (Figure 1b). All three dimensions were combined by placing the dot pattern in black on the coloured shape to create a total of eighty-four unique items (Figure 1d).

The item memory task consisted of a surprise item recognition test that always took place after the categorization task. In this recognition task, participants were asked if stimuli were present in the categorization task (old) or were new. Of the 48 stimuli presented at test, 24 were old, 12 were novel-shape lures, and 12 were same-shape lures. *Same-shape lures* were generated from the morphing procedure that was used to generate the categorization stimuli, but all twelve were unique and had not occurred during the categorization phase. Novel shape lures were created outside of the shape space used to generate the categorization stimuli, but were likewise generated from paint splatter images (see Figure 1c for examples). Order of presentation was randomized, and there was no time limit for response.

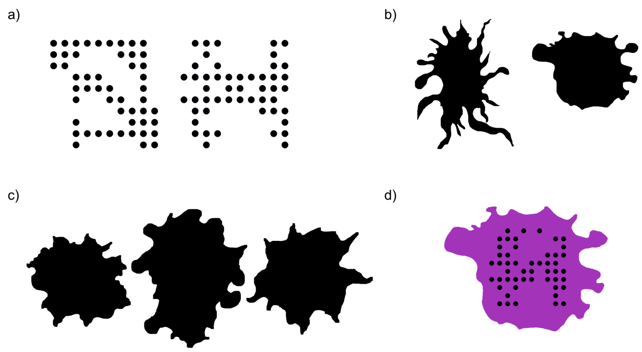

Figure 1: a) Diagnostic features defining category membership, b) Two shapes were morphed together to create a shape space, c) Novel-shape lures: Shapes were created outside of the shape space, d) An example of a complete stimulus: unique colours, shapes, and dot patterns were combined to create a set of completely unique stimuli.

Statistical Analysis We conducted all statistical analyses in R (R Core Team, 2017). Categorization accuracy was operationalized by calculating percentage correct in the category learning task. Item memory was calculated using d' (Z(hit rate) – Z(false alarm rate)), and scores were compared with chance using an independent-samples t-test. The general linear model was applied to all basic correlations, and general linear mixed-effects models were applied for analyses involving trial number using the lme4 package in R (Bates, Mächler, Bolker, & Walker, 2015). The fixed effects were categorization trial number and accuracy. Models contained random intercepts and slopes grouped by stimulus.

Results

Participants demonstrated learning in the categorization task, with accuracy increasing across trial number ($\beta = 0.025$, $z = 9.793$, $p < .001$) and an average overall accuracy of 76% ($SD = 16.827$). Similarly, participants demonstrated memory for the items at test, successfully distinguishing old items from new at a rate significantly different from 0 ($M = 0.268$, $SD = 0.374$; $t(59) = 5.558$, $p < .001$). As predicted, a tradeoff was also observed such that participants' categorization scores were negatively related to their recognition scores ($F(1,58) = 14.31$, $p < .001$, Figure 2). Thus, individuals who performed the categorization task better, had worse memory for exemplars.

To determine if memory was different for items that were categorized accurately from those that were not, we performed a t-test comparing the memory (d') for correctly and incorrectly categorized items. Memory was equivalent across correctly and incorrectly categorized items ($t(495.18) = 0.463$, $p = 0.644$). However, this relationship shifted over time, such that there was an interaction between memory for correctly and incorrectly categorized items and trial number ($\beta = -0.016$, $z = -2.062$, $p = 0.039$) such that memory for incorrectly categorized items moderately increased with an increasing number of trials ($\beta = 0.001$, $z = 1.873$, $p = 0.061$), and memory for correct trials moderately decreased with an increasing number of trials ($\beta = -0.006$, $z = -1.648$, $p = 0.099$).

To determine how specifically items were remembered, item memory was analyzed separately for novel-shape and same-shape lures by calculating d' using each as a unique false alarm score. A paired sample t-test determined that the two sets of scores were significantly different ($t(59) = -7.643$, $p < .001$). Using novel-shape lures, memory was significantly different from 0 ($M = 0.7272$, $SD = 0.627$; $t(59) = 8.986$, $p < .001$). Using same-shape lures, however, memory did not differ from 0 ($M = -0.103$, $SD = 0.538$; $t(59) = -1.488$, $p = 0.142$). Relating this to categorization performance across individuals, d' calculated using novel-shape lures was negatively correlated with categorization scores ($F(1,58) = 25.7$, $p < .001$), while d' calculated using same-shape lures was not ($F(1,58) = 0.57$, $p = 0.453$). Thus, the trade-off is only observed when using the novel-shape lures, for which there is evidence of memory.

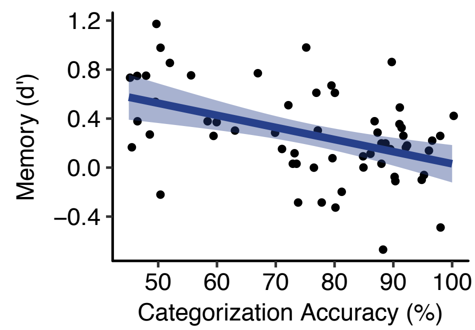

Figure 2: Item memory (d') by category accuracy (%). Each dot is an individual, the line signifies the slope and shading indicates standard error of the mean.

Discussion

As a group, participants successfully learned to categorize and remembered the items at post-test. Individual difference scores showed a trade-off between levels of learning such that those who performed well at the categorization task, performed more poorly at the item recognition task.

Participants failed to exhibit memory in comparison to the same-shape lures, but they did demonstrate memory when the lures were more distinct. Taken together, these data show that for adults, learning to categorize well impedes memory for items.

Experiment Two

Methods

Participants Participants included 61 children between the ages of 5- and 8- years old ($M = 6.42$ years, 48% female) recruited at a science museum. Exclusion criteria included lack of English skills to understand instructions, with one child meeting exclusion criteria.

Materials and Procedure The same two tasks were completed as in Experiment One, with a different, age-appropriate cover task. For the category learning task, participants were told to sort two alien families onto their correct spaceship. Stimuli were presented for 3 seconds or until the participant responded. For the item memory test, 32 stimuli were presented randomly. Again, half were previously seen and half were new. Of the new, half were novel-shape lures and half were same-shape lures. Children were asked to verbally confirm their response after each button press. Tasks were completed on an Apple laptop using PsychoPy (Peirce, 2008).

Statistical Analysis We conducted the same analyses as Experiment One, as well as independent samples t-tests comparing adult and child scores.

Results

Children demonstrated learning in the categorization task, with accuracy increasing across trial number (β = 0.015, z = 6.619, p <.001). However, accuracy was significantly lower than adults (M = 65.889, SD = 17.664; β = -0.073, t = -3.399, p <.001, Figure 3). Children's item memory was poor: d' did not differ from 0 (M = 0.065, SD = 0.515; t(59) = 0.984, p = .329). Their item memory (d') was also significantly worse than adults (t(110) = 2.743, p = .007, Figure 4).

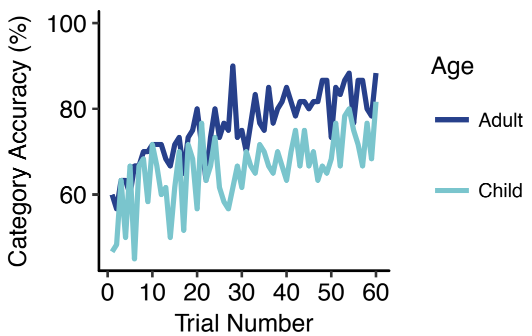

Figure 3: Category accuracy (%) by trial number and age group. The lines signify average group accuracy by trial. The dark blue line signifies the adult group, and the light blue line signifies the child group.

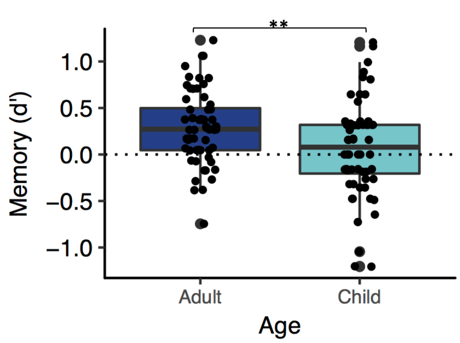

Figure 4: Item memory (d') by age group. d' is plotted separately for adults (dark blue) and children (light blue). The boxes signify the interquartile range and the whiskers signify the first quartile and below and the third quartile and above, respectively. The dotted line signifies chance, or no evidence of memory, and individual dots represent participants.

No trade-off was found between children's categorization accuracy and item memory (d') (F(1,58) = 1.766, p = .1891, Figure 5). Moreover, children's age was not found to interact with this relationship (F(3,56) = 0.647, p = 0.5885).

To determine if memory was different for items that were categorized accurately from those that were not, we performed a t-test comparing memory (d') for correctly and incorrectly categorized items. While not significant, we found a marginal difference in memory for correctly and incorrectly categorized items (t(690.2) = 1.950, p = .052), such that incorrectly categorized items were remembered moderately better. In addition, this relationship showed a trend toward shifting over time, with a trending interaction between time and accuracy (β = 0.013, z = 1.660, p = .097). Incorrectly categorized items were moderately better remembered at the beginning of the task, and correctly categorized items at the end of the task. While not significant, it is important to note that this relationship is the opposite pattern of that observed in adults.

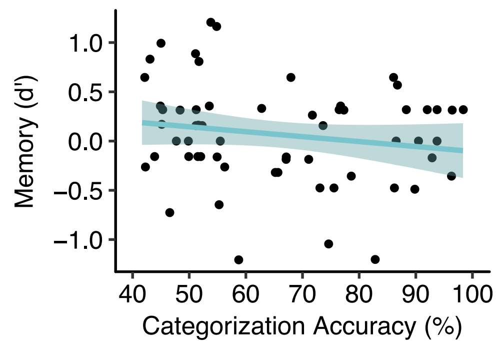

Figure 5: Child item memory (d') by category accuracy (%). Each dot is an individual, the line signifies the slope and shading indicates standard error of the mean.

To determine how specifically items were remembered, d' was calculated twice, once with each type of lure. Using a paired sample t-test, the two sets of scores were found to be significantly different (t(59) = -3.302, p = 0.002). When calculated with only novel-shape lures, memory was significantly different from 0 (M = 0.2701, SD = 0.822; t(59) = 2.547, p = 0.013), but memory did not differ from 0 when calculated with only the same-shape lures (M = -0.092, SD = 0.526; t(59) = -1.361, p = .179). While this is a similar pattern to that found in the adults, when compared to adults, the former was significantly lower (t(110.22) = 3.43, p <.001, Figure 6). Thus, children did show memory, but needed more distinct lures to demonstrate it, and it was poorer than that observed in the adults.

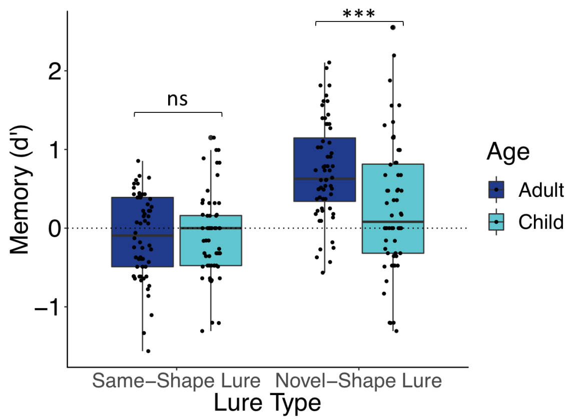

Figure 6: Item memory (d') calculated with novel- and same-shape lures across age group. The boxes signify the interquartile range and the whiskers signify the first quartile and below and the third quartile and above, respectively. The dotted line signifies chance, or no evidence of memory, and individual dots represent participants. The two left-hand boxes signify scores calculated with same-shape lures and those on the right signify scores calculated with novel-shape lures.

Looking at individual differences, d' calculated using novel-shape lures was not significant but had a moderate effect trending towards a trade-off with categorization accuracy, as more successful category learners had worse memory ($F(1,58) = 3.504$, $p = .0663$). No relationship was found between d' calculated with same-shape lures and categorization accuracy ($F(1,58) = 0.07289$, $p = .788$). The trade-off did not interact with children's age for d' calculated with the novel-shape lures ($F(3,56) = 1.618$, $p = 0.196$) or same-shape lures ($F(3,56) = 0.591$, $p = 0.623$). Nonetheless, to further assess the impact of children's age on the trade-off, we broke the children into two age groups: 35 5- and 6-year old children (young) and 25 7- and 8-year old children (old). Upon analyzing the trade-off in each group, the moderate effect found in the novel-shape lures continued in the old children ($F(1,23) = 3.303$, $p = 0.082$), but disappeared in the younger children ($F(1,33) = 0.790$, $p = 0.380$, Figure 7). No relationship between categorization accuracy and d' calculated with the same-shape lures was found in the old ($F(1,23) = 0.094$, $p = 0.762$) or young children ($F(1,33) = 0.182$, $p = 0.672$).

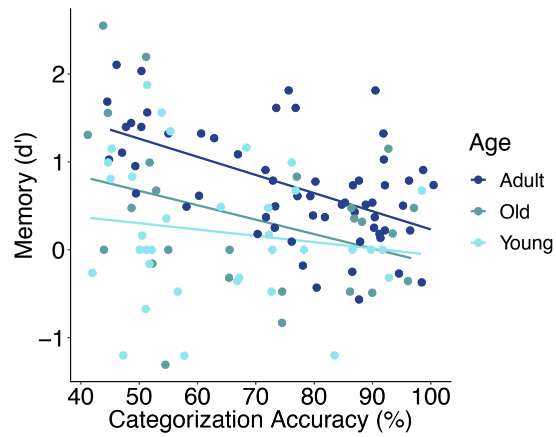

Figure 7: Item memory (d') by category accuracy (%) across adults, 7-8 year olds ("Old"), and 5-6 year olds ("Young"). Each dot is an individual, and the lines signify the slope. Adults are represented in dark blue, older children in teal, and young children in light blue.

Discussion

Although children learned to categorize, their memory for items was very poor overall. Like adults, children had no memory when compared to the same-shape lures but demonstrated memory when compared to the novel-shape lures, suggesting that they could only distinguish new and old items when the lures were distinctive.

Item memory did not significantly predict category learning, which is not surprising given how low children's memory was. However, a moderate effect was observed when d' was calculated with only the novel-shape lures, such that category learning scores were negatively correlated with memory (d'). Finally, this trend was only present in the older, 7- and 8-year old children but was not found in the younger, 5- and 6-year old children.

General Discussion

Building off of prior research showing a trade-off in pre-existing category knowledge and item memory, we found that a trade-off also occurs during the process of learning new categorical structures. For the adults, there was a cost to category learning, as those who performed well at the category learning task demonstrated worse memory for the items at post-test. This effect was driven by the more distinct lures, as memory for the similar lures was overall quite poor. In comparison, the children did not show a trade-off in category learning and item learning. However, when d' was calculated using only the distinct lures, a moderate effect was observed. This effect was found to be driven by the older, 7- and 8-year old children. When divided into two age groups, the moderate trade-off was observed in the older children, but there was no trade-off found in the younger, 5- and 6-year olds. These data may point to adults and young children approaching categorization in different ways.

While the mechanism at play remains unclear, a selective attention account provides an explanation for the pattern of data observed in adults. Selective attention is a top-down process that not only directs attention to *relevant* information, but also suppresses *irrelevant* information (Pashler, Johnston, & Ruthruff, 2001). In a categorization task, the relevant information is the diagnostic feature, while the remaining information is irrelevant. Indeed, as discussed above, upon learning what defines a category, learners have been shown to fixate their gaze on the diagnostic feature (Rehder & Hoffman, 2005). In the current task, successful category learners likely fixated on the dot pattern while suppressing the irrelevant, item-level information. Given that unattended information is not remembered well (Simons, 2000), this could explain the successful learners' poor memory performance. Conversely, the poor category learners may have failed to learn which feature was diagnostic, thereby never selectively attending to it and, thereby continuing to attend to item-level information.

In comparison, it is less clear what mechanism best accounts for the children's pattern of behaviour, as they displayed only a moderate trade-off and, more specifically, only in the older children when calculating memory using the most distinctive lures. One possibility is that the younger children are utilizing a holistic, similarity-based categorization style, leading to a lack of trade-off. In comparison, the older children may be beginning to shift from this categorization style towards a more adult-like style focused on a diagnostic feature. A developmental shift in this age group would account for the moderate trade-off observed.

Prior research suggests that young children may categorize based on overall item similarity (Smith & Kemler, 1977), and if this were the case, children would attend to all features equally instead of selectively attending to a single feature. As such, their item memory would not drop upon categorization. Interestingly, Smith and Kemler (1977) found that this approach to categorization was consistently used among 5-year olds, but results were more ambiguous among 8-year olds. Perhaps the ambiguity reflects the beginning of adult-like categorization, and explains the moderate trade-off that we see here.

Indeed, a shift away from holistic processing would reflect the developmental course of selective attention, as the ability to filter irrelevant information has been found to improve across the elementary school years (Enns & Akhtar, 1989). An increase in selective attention would facilitate a more adult-like approach and result in a trade-off between category learning and item memory. Future research assessing the role of selective attention and its developmental course on the trade-off would help clarify the mechanism behind the patterns of behaviour observed in each age group.

Interestingly, children's memory was quite poor overall, which is not aligned with a holistic processing approach. Prior studies found children to have superior memory to adults for item-level information since they processed more information overall (e.g., Sloutsky & Fisher, 2004; Plebanek & Sloutsky, 2017). However, it is also well established that children have poor memory compared to adults (Ghetti, Angelini, & Annunzio, 2008; Rubin, 2000), and the adult group's item memory was not particularly strong either. It may be the case that children learned to categorize in a different way than the adult group but did not have the memory capacity to demonstrate it.

Alternatively, children's poor memory may not reflect poor memory overall, but may be symptomatic of poor pattern separation. Work by Ngo, Newcombe, & Olson (2018) found that 4-year olds were significantly worse than 6-year olds and adults at distinguishing old items from very similar items, irrespective of overall memory scores. Due to the similarity across items in the current study, it is possible that the younger children were disproportionately unable to discriminate the items. While unclear at this time, boosting children's memory in the future by increasing discriminability between items would help us to better understand how category learning and item memory interact across development.

The different patterns of learning across categorization trials in the adult and child groups suggest that the groups could be using different learning strategies. First, adults remembered incorrectly categorized items better towards the end of the task, while children remembered them better towards the beginning and moderately better overall than correctly categorized items. It may be the case that with increased learning across trials, errors became rare and surprising to adults and were, thus, remembered better. In comparison, children's heightened memory for errors throughout may reflect their tendency to respond more reactively than adults (Chatham, Frank, & Munakata, 2009), the surprise of which could have a memory boosting effect throughout.

Second, adults remembered correctly categorized items better towards the beginning of the task while children remembered them better towards the end. Given our assertion of increased selective attention with category learning in the adult group, it follows that correctly categorized items would be remembered more poorly towards the end of the task after learning had occurred and irrelevant information became unattended. Since children showed the reverse pattern in memory, this may provide further support for a more holistic approach to categorization than one of selective attention. This pattern would suggest that children maintain distributed attention throughout the task, as their memory for task-irrelevant information does not drop. The boost in memory observed towards the end may be a simple recency effect. Whatever the explanation, these divergent patterns of learning show that adults' and children's online categorization performance impacts memory.

Across the lifespan, our approach to learning changes as our needs change. Children are still figuring out what information is important, so it makes sense that they would attend to much of the information available. On the other hand, adults have a good sense of what information to prioritize and so attend to only what they deem informative. Inevitably, this means that a certain amount of information is always going to be missed. These findings make clear that we are always only seeing a piece of the picture, but perhaps we did not all start out that way.

Acknowledgments

We would like to thank Meg Schlichting, Katherine Duncan, Tess Forest, Michael Dubois, Alexandra Decker, and the rest of the Finn lab for all your help and support. We would also like to thank the Ontario Science Centre, Child Study Centre, NSERC and SSHRC, as well as the reviewers for their helpful commentary and feedback. Finally, a special thank you to all the parents and children who made this research possible.

References

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using **lme4**. *Journal of Statistical Software*, *67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01

Brainerd, C. J., & Reyna, V. F. (1990). *Gist Is the Grist: Fuzzy-Trace Theory and the New lntuitionism. DEVELOPMENTAL REVIEW* (Vol. 10). Retrieved from https://ac-els-cdn-com.myaccess.library.utoronto.ca/027322979090003M/1-s2.0-027322979090003M-main.pdf?_tid=5ee0e9d6-0629-4cf8-845e-a1078571a776&acdnat=1535355885_93d0c33e7a0ae9b5767f328477ae24f1

Brainerd, C. J., Reyna, V. F., & Ceci, S. J. (2008). Developmental Reversals in False Memory: A Review of Data and Theory. https://doi.org/10.1037/0033-2909.134.3.343

Chatham, C. H., Frank, M. J., & Munakata, Y. (2009). Pupillometric and behavioral markers of a developmental shift in the temporal dynamics of cognitive control. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(14), 5529–33. https://doi.org/10.1073/pnas.0810002106

Enns, J. T., & Akhtar, N. (1989). A Developmental Study of Filtering in Visual Attention. *Child Development*, *60*(5), 1188–1199. Retrieved from https://www.jstor.org/stable/pdf/1130792.pdf?casa_token=Ocv6L18S3SAAAAAA:xIZmcqlfE9qULHcqa8ACTrYA34m3RD_JdC53X9APaBKZoVZou0T2AIre-14C5ygR4zo_lIZ9FWAG8KbIEIrei5a4sjvR0c6LxYYa_XwEa3EXeT7aTu8

Fried, L. S., & Holyoak, K. J. (1984). Induction of Category Distributions: A Framework for Classification Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*(2), 234–57. https://doi.org/10.1037/0278-7393.10.2.234

Ghetti, S., Angelini, L., & Annunzio, G. D. (2008). The Development of Recollection and Familiarity in Childhood and Adolescence : Evidence From the Dual-Process Signal Detection Model, *79*(2), 339–358.

Mandler, J. M. (1984). Stories, scripts, and scenes: Aspects of schema theory. *Hillsdale: Erlbaum*.

Ngo, C. T., Newcombe, N. S., & Olson, I. R. (2018). The Ontogeny of Relational Memory and Pattern Separation. *Developmental Science*, *21*(2). https://doi.org/10.1111/desc.12556

Pashler, H., Johnston, J. C., & Ruthruff, E. (2001). Attention and Performance. *Annual Review of Psychology*, *52*(1), 629–651. https://doi.org/10.1146/annurev.psych.52.1.629

Peirce, J. W. (2008). Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, *2*, 10. https://doi.org/10.3389/neuro.11.010.2008

Plebanek, D. J., & Sloutsky, V. M. (2017). Costs of Selective Attention: When Children Notice What Adults Miss. *Psychological Science*, *28*(6), 723–732. https://doi.org/10.1177/0956797617693005

Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, *51*(1), 1–41. https://doi.org/10.1016/J.COGPSYCH.2004.11.001

Rubin, D. C. (2000). The distribution of early childhood memories. *Memory*, *8*(4), 265–269. https://doi.org/10.1080/096582100406810

Rueda, M. R., Posner, M. I., & Rothbart, M. K. (2005). The Development of Executive Attention: Contributions to the Emergence of Self-Regulation. Retrieved from https://www.researchgate.net/profile/Maria_Rueda/publication/7617847_The_Development_of_Executive_Attention_Contributions_to_the_Emergence_of_Self-Regulation/links/09e41505b54a74a51f000000.pdf

Seger, C. a, Poldrack, R. a, Prabhakaran, V., Zhao, M., Glover, G. H., & Gabrieli, J. D. (2000). Hemispheric asymmetries and individual differences in visual concept learning as measured by functional MRI. *Neuropsychologia*, *38*(9), 1316–1324. https://doi.org/10.1016/S0028-3932(00)00014-2

Simons, D. J. (2000). Attentional capture and inattentional blindness. *Trends in Cognitive Sciences*, *4*(4), 147–155. https://doi.org/10.1016/S1364-6613(00)01455-8

Sloutsky, V. M. (2010). From Perceptual Categories to Concepts: What Develops? *Cognitive Science*, *34*(7), 1244–1286. https://doi.org/10.1111/j.1551-6709.2010.01129.x

Sloutsky, V. M., & Fisher, A. V. (2004). When Development and Learning Decrease Memory. *Psychological Science*, *15*(8), 553–558. https://doi.org/10.1111/j.0956-7976.2004.00718.x

Smith, L. B., & Kemler, D. G. (1977). Developmental trends in free classification: Evidence for a new conceptualization of perceptual development. *Journal of Experimental Child Psychology*, *24*(2), 279–298. https://doi.org/10.1016/0022-0965(77)90007-8