Greater discrimination difficulty during perceptual learning leads to stronger and more distinct representations

Authors: Vencislav Popov^{1,2} & Lynne M. Reder^{1,2}

¹ Department of Psychology, Carnegie Mellon University, Pittsburgh, PA

² Center for the Neural Basis of Cognition, Pittsburgh, PA

Correspondence: vencislav.popov@gmail.com or redler@cmu.edu

Cite as: Popov, V. & Reder, L. (2020). Greater discrimination difficulty during perceptual learning leads to stronger and more distinct representations. *Psychonomic Bulletin & Review*. Advance Online Publication.

Abstract

Despite the conventional wisdom that it is more difficult to find a target among similar distractors, this study demonstrates that this disadvantage is short-lived, and that high target-to-distractor (TD) similarity during visual search training can have beneficial effects for learning. Participants with no prior knowledge of Chinese performed 12 hour-long sessions over 4 weeks, where they had to find a briefly presented target character among a set of distractors. At the beginning of the experiment, high TD similarity hurt performance, but the effect reversed during the first session and remained positive throughout the remaining sessions. This effect was due primarily to reducing false alarms on trials in which the target was absent from the search display. In addition, making an error on a trial with a specific character was associated with slower visual search RTs on the subsequent repetition of the character, suggesting that participants paid more attention in encoding the characters after false alarms. Finally, the benefit of high TD similarity during visual search training transferred to a subsequent n-back working memory task. These results suggest that greater discrimination difficulty likely induces stronger and more distinct representations of each character.

Keywords: perceptual learning, memory, similarity, representation, visual search

Introduction

Intuitively, the greater the similarity between two different objects, the more difficult it should be to distinguish them. Research using the visual search task has repeatedly confirmed this intuition - when a target is more similar to distractors in the search array, accuracy decreases, response times increase, and more errant saccades are made to the highly similar distractors (Bichot & Schall, 1999; Duncan & Humphreys, 1989; Treisman & Gormican, 1988). This negative target-to-distractor (TD) similarity effect appears to be robust - it holds regardless of whether participants can anticipate the similarity on each trial (Pashler, 1987),

regardless of whether they perform feature or conjunction-based search (Phillips, Takeda, & Kumada, 2006), and it holds with both simple laboratory constructed-stimuli (Duncan & Humphreys, 1989) as well as with more visually complex real-world objects (Alexander & Zelinsky, 2012). The apparent robustness of this negative visual similarity relationship has played a pivotal role in the development of many visual search theories (Alexander & Zelinsky, 2012; Duncan & Humphreys, 1989; Treisman & Gelade, 1980).

The ubiquity of the negative TD similarity effect, however, is based predominately on studies that use familiar stimuli with pre-existing representations in long-term memory (letters, lines, familiar objects). In this study, we explored how the TD similarity effect changes as people gain visual expertise with novel stimuli over an extended period of time. US undergraduates with no previous knowledge of Chinese performed a visual search task with 64 novel Chinese characters for 12 hour-long visual search sessions over 4 weeks. We found a striking pattern - the search time and accuracy advantages for a target among dissimilar distractors were short-lived and *reversed* after only a single session of training, such that *greater* TD similarity led to better performance over time. After documenting this unexpected finding, we will propose an explanation and test some of its predictions.

Method

The results presented below come from a novel reanalysis of data from Reder, Liu, Keinath and Popov (2016). That experiment was concerned with the effects of frequency of exposure on learning and memory; character similarity was not analyzed in that report and all findings presented here are novel. We describe the full method for completeness. The data and analysis code are available at <https://github.com/venpopov/similarity-discrimination-learning>

Participants

Nineteen U.S. college students with no prior experience learning Chinese participated in this experiment. One participant's data was excluded from the analyses because closer inspection revealed that the subject number was miscoded and that the resulting files contained a mixture of two different participants' partial data.

Materials

The stimuli for the visual search task were 64 Chinese characters. The characters were grouped based on their visual similarity into 16 sets of four characters, such that characters

within a set had a higher similarity to each other compared to characters from other sets¹. We subsequently confirmed this by analyzing orthographic vector representations of the characters (Xing & Li, 2004). Highly similar distractors were used in order to force participants to encode the entire character rather than a subset of diagnostic features. For each participant, half of the sets were randomly assigned to be presented 20 times more often during the visual search task. This frequency manipulation was the main focus in the Reder et al (2016), but those results will not be reported here.

Procedure

Participants performed 4 different tasks with the same characters over the course of six to eight weeks. In this report, we will focus on two of those tasks - a visual search training task, and a working memory N-back task². Performance on the visual search task was used as a measure of training and to explore how the TD similarity effect changes over time. With the N-back task we tested whether similarity during visual search training had transferable effects to novel memory tasks.

Visual search task.

Participants performed a visual search task for three hour-long sessions per week for a total of four weeks and each session consisted of 672 trials. Participants had to search for a different target on each trial. Half of the trials were “absent” trials. Figure 1 illustrates a single visual search training trial. Each high-frequency character was presented as a target on 20 trials in each session while each low frequency character was presented once. Each trial showed a single target character followed by a display of 3-5 characters and participants had to indicate whether the target character was present in the display. The visual search display always contained exactly three of the four characters from a target’s similarity set along with between zero and two additional characters from different character sets from the same frequency class.

N-back task

Two to four weeks after the final visual search session, participants performed an N-back working memory task using the same Chinese characters. Participants were shown a series of individual Chinese characters one at a time for 2.5 seconds each and they had to indicate whether the current stimulus matched the stimulus that appeared N presentations prior,

¹ These were selected by Xionan Liu, a graduate student in the lab at the time and a native Chinese speaker. We are also grateful to him for pointing us to the work describing the characters’ orthographic vector representations (Xing & Li, 2004), which made the analyses in this paper possible.

² Participants also performed two cued-recall tasks with these characters. We could not analyze performance in those tasks as a function of visual search similarity, because the identity of the characters was not recorded.

where N varied from one to three in different blocks of trials. The 3-back task is particularly demanding since it involves holding three stimuli in working memory so that the identities of the stimuli that are “3-back,” “2-back,” and “1-back” can be updated with each presentation, as well as simultaneously determining the correct response and pushing the button. There were a total of 24 blocks of 17 trials each – 8 blocks for each level of the N-back task. Half of the blocks contained only HF characters, and the other half contained only LF characters. The order of blocks was randomly determined for each participant.

Figure 1. Illustration of the visual search procedure

Data Analysis

In order to perform the analyses presented below, we calculated two continuous similarity measures for each trial - how similar was the target to the distractors (TD similarity) and how similar were the distractors to each other (DD similarity). We calculated the similarity based on vector representations obtained from Yang et al. (2009). Each character was represented as a vector of 270 binary features for five dimensions - simple features, shapes, structure, position and strokes. These vector representations are based on an orthographic analysis of the characters (Xing & Li, 2004) and have been used to model print-to-sound mappings in Chinese (Yang et al., 2009). For each target character, we calculated the mean Euclidean distance between the feature vectors for the target and each distractor in the search arrays for all trials in which the target had to be found. We also calculated, for each target character, on average how dissimilar were the distractors to one another, by computing the mean of all pairwise distractor-to-distractor Euclidean distance scores. The resulting distance measures were reversed and scaled to form similarity metrics with means of 0 and standard deviations of 1. All analyses were performed with these continuous similarity measures with all of the data. Even though the analyses were performed with all of the data, for the plots we defined high and low similarity groups as being below or above one SD around the mean similarity.

All the analyses focused only on the high frequency characters because each low frequency character appeared only once per session as a target. We excluded from the analyses trials

where the RT was more than 3 median absolute deviations above or below the median RT, calculated separately for each participant, session and condition (2.26%; total number of remaining observations = 130,538). For the RT analyses, we considered only trials with correct responses (8.19 % error; total number of remaining observations = 119,488). Since similarity varied randomly from trial to trial, it was important to ensure that any results were not due to a confound with different trial types (target-present or target-absent), the number of distractors, the session number, etc. To account for the effect of all variables, we fit a sequence of linear-mixed effects regression models with a random intercept for each subject in the experiment. Before fitting any of the similarity measures, we identified the maximal model from all the control variables and their interactions, using likelihood ratio tests. Next, we fit the model by including the TD and DD similarity measures. Finally, we tested for an interaction between these similarity measures and session number. To determine the significance of each factor, we compared a model with the effect in question and a reduced model without it using likelihood ratio tests.

Results and discussion

The TD similarity effect reverses over the course of training

Consistent with the existing literature, during the first training session, greater TD similarity initially hurt performance, leading to slower search RTs ($\Delta \text{AIC} = -6$, $\chi^2(1) = 8.70$, $p = .003$). As can be seen from Figure 2 (upper panels), the RT benefit for low TD similarity trials was short-lived - when we looked at the effect of TD similarity as a function of how many times each target character had been presented thus far during Session 1, we found that after the first five repetitions of a target, the effect disappeared, and it reversed after 15 repetitions such that low TD similarity started to hurt RTs. The interaction between character repetition number and TD similarity was significant ($\Delta \text{AIC} = -9$, $\chi^2(1) = 10.28$, $p = .001$). Even though numerically the effects of TD similarity on accuracy were similar, neither the main effect of TD similarity, nor its interaction with character repetition *during session 1* were significant (both $p > .11$).

By the beginning of Session 2, higher TD similarity actually helped visual search performance. Importantly, this benefit of higher TD similarity, continued over subsequent sessions, leading to overall faster RTs ($\Delta \text{AIC} = -16$, $\chi^2(1) = 17.67$, $p < .001$), and better accuracy ($\Delta \text{AIC} = -32$, $\chi^2(1) = 34.10$, $p = .001$). These effects increased over sessions, leading to a significant interaction between session number and TD similarity for RTs ($\Delta \text{AIC} = -45$, $\chi^2(1) = 46.98$, $p < .001$) and accuracy ($\Delta \text{AIC} = -6$, $\chi^2(1) = 7.76$, $p = .005$). In summary, we found that initially high TD similarity hurt performance, but this effect was short-lived - it reversed during session 1, and the learning benefit for high TD similarity increased over subsequent sessions.

Before explaining the reversal of the TD similarity effect, we should discount the possibility that it is due to a confound with DD similarity, as the two similarity measures correlated

moderately with each other, $r(118) = 0.52$, $p < .001$. Research has shown that when there is high similarity among the distractors, the target is easier to detect (Duncan & Humphreys, 1989). To test this possibility, we repeated all mixed-effects regression models after including DD similarity as an additional predictor of performance. Consistent with prior research, higher DD similarity lead to faster RTs ($\Delta AIC = -943$, $\chi^2(1) = 944.24$, $p < .001$) and better accuracy ($\Delta AIC = -97$, $\chi^2(1) = 96.62$, $p < .001$) from the start. Importantly, the effects of TD similarity described above remained even after accounting for DD similarity in the model, excluding the possibility that the learning benefit of high TD similarity is due to a confound.

Figure 2. Response times (left) and accuracy (right) in the visual search task as a function of the mean group similarity for each character group. Top panels present performance in the first session depending on which repetition of a character the participant was seeing, and bottom panels present performance across sessions.

The learning benefit of high TD similarity is due to decreasing partial matching

What causes the reversal of the TD similarity effect as people gain visual expertise? We used unfamiliar Chinese characters, for which our participants had no preexisting representations. Initially, when a target character appears, participants have no choice but to encode it as a configuration of simpler visual features. They have to store this composite representation in short-term memory long enough so that they can compare it to each character in the search array. During these initial trials, TD similarity would be a function of the amount of feature

overlap between characters (Tversky, 1977). As a result, the greater the feature overlap is between a target and distractors on target-absent trials, the more likely people would be to make false alarms by incorrectly recognizing one of the distractors as the target. This mechanism, also known as *partial matching* (e.g., Reder & Kusbit, 1991), would predict that the initial benefit for low TD trials should be observed primarily in lower false alarm rates on target-absent trials, which is consistent with existing data (Alexander & Zelinsky, 2012).

Partial matching explains the initial benefit for low TD similarity trials, but why does the effect reverse over time? As people develop perceptual expertise, they tend to chunk features into unitary representations of each stimulus (Gobet et al., 2001; Palmer, 1977; Simon, 1974). Importantly, Feigenbaum and Simon (1984) argued that rather than incorporating vivid detail in each representation, people tend to only learn those features that are relevant for distinguishing one category from another. Based on this idea, we propose that *greater discrimination difficulty early in the visual search task lead to the development of richer and more distinctive chunked representations of each character*. This would occur because of the demands of the task: the memory system aims to make highly similar patterns more distinct from one another, so as to be better suited to support future performance.

The hypothesis that greater discrimination difficulty leads to more distinct representations allows us to make several testable predictions. Developing a richer representation over time should reverse the amount of partial matching errors on target-absent trials, because people would no longer depend on basic feature overlap to guide performance. In contrast, this would not affect target-present trials, because the target always has the highest match with itself, regardless of the nature of the representation. Indeed, this is exactly what we found. As can be seen from **Figure 3**, the negative effect of high TD similarity in the first session is entirely due to increased false alarms on absent trials. More importantly, the subsequent benefit for high TD similarity occurred primarily on target-absent trials, consistent with the idea that richer representations prevented the partial matching of shared features between the target and distractors. These findings were supported by a significant interaction between TD similarity and whether the target was present/absent ($\Delta AIC = -5$, $\chi^2(1) = 7.78$, $p = .005$).

Figure 3. Accuracy in the visual search task for each session and TD similarity groups, depending on whether the target was absent (left panel) or present (right panel) in the search array.

Initial discrimination difficulty causes participants to pay more attention after false alarms, leading to stronger encoding

What is the mechanism through which discrimination difficulty leads to improved representations? People learn primarily from making errors, especially when feedback is provided as it was in this study. One possibility is that the error feedback after each false-alarmed trial is associated with the target character and, when the character is presented again later, participants pay more attention when encoding it and in comparing it to the distractors. The increased attention during encoding and/or the subsequent search process should lead to the stronger and more distinct representations. If this conjecture is accurate, we should see that, after participants make an error with a given target, the next time the same target is presented, they should be slower to perform the task.

To test this idea, for each correct trial we calculated response times as a function of how accurate participants were on average for the previous three presentations of the target. As can be seen from Figure 4 (left panel), participants get progressively slower the more times they were previously incorrect with a given character ($\Delta AIC = -259$, $\chi^2(1) = 260.51$, $p < .001$). Note that this slow down occurs on the next presentation of the *same target*, not the next trial in the sequence of trials. Trials with the same target were separated by multiple other trials ($M = 32$, $SD = 30$). Thus, this slow-down is character-specific and is the result of learning, and not merely due to slowing down immediately after making an error.

Figure 4. Left - RT on the current presentation of a target, depending on the average accuracy on the last 3 repetitions of the same target. Right - post-error slow down on the current occurrence of a target depending on whether the target is present in the search array on the current trial and whether it was present on the previous repetition. Slowdown is the difference between cases when the participant made an error on the previous trial vs when they were correct on the previous trial.

Consistent with the proposal that learning should be greatest from making partial matching errors, we found that the slow-down was greatest when the target was absent on the current trial and its previous occurrence (see Figure 4, right panel; the effect was supported by

significant three-way interaction, $\Delta AIC = -14$, $\chi^2(1) = 17.28$, $p < .001$). This interaction suggests that making a false alarm forces participants to pay more attention in encoding the character the next time they see it in order to avoid subsequent false alarms. Additionally, participants might be more cautious during the subsequent search process itself, which could also benefit learning.

The learning benefits due to higher discrimination difficulty transfer to subsequent memory tasks

Finally, if the representations of targets that were paired with highly similar distractors were indeed strengthened, then we should expect to see transfer to other tasks performed with the same stimuli. Specifically, having a stronger representation should make it easier to manipulate and maintain the character representations in working memory. We tested performance in the final N-back task, which was performed at the end of training, as a function of whether each character appeared with highly similar or weakly similar distractors *during the visual search training*. As can be seen from Figure 5, it was easier to maintain and update the characters in working memory in the N-back task when they had highly similar distractors during the visual search training ($\Delta AIC = -6$, $\chi^2(2) = 9.87$, $p = .007$). Thus, greater discrimination difficulty during training lead to representations which were more resistant to interference in working memory from other concurrently held representations.

Figure 5. Performance on the N-Back task for each level depending on whether the two characters were seen with high or low similarity distractors during the visual search task.

General discussion

Our findings undermine the apparent robustness of the negative effect of TD similarity on visual search performance and demonstrate that initial high discrimination difficulty could have beneficial effects for learning. Prior research had shown negative effects of high TD similarity on visual search (Duncan & Humphreys, 1989; Treisman & Gormican, 1988), but those studies used familiar stimuli and did not provide extended training with them. We

found that this high TD similarity disadvantage was short-lived when the stimuli are novel and relatively complex - visual search for previously unfamiliar Chinese characters suffered when the target was highly similar to distractors only during the beginning of the training. After 1 session of training the effect reversed and high TD similarity facilitated visual search performance for the remaining sessions.

We suggest that the reason for this reversal is that when it is more difficult to discriminate a target from distractors during learning, participants are forced to develop richer and more detailed representations of the novel characters to be able perform the task better in the future. An alternative possibility is that TD similarity does not by itself lead to better learning; rather, the increased search duration times improved learning indirectly. We did not find evidence for this alternative explanation (see Appendix A). We also showed that the positive effect of TD similarity was mostly due to reducing false alarms on absent trials, suggesting that partial matching based on feature overlap was eliminated. This explanation is consistent with learning models like Feigenbaum and Simon's (1984) discrimination net, in which people learn only those key features of stimuli that would allow them to distinguish them from other categories. This is a rational strategy because it helps preserve limited cognitive resources (Popov & Reder, 2019), but it also means that performance would suffer if suddenly discrimination demands are increased, as demonstrated by the fact that, for example, people who used pennies every day of their lives for decades failed to discriminate the correct representation of a penny from plausible foils (Nickerson & Adams, 1979). Ellis and Turk-Brown (2019) recently made similar arguments concerning visual complexity, namely, that even though complex stimuli demand more cognitive resources for processing, complexity can eventually lead to the development of richer representations, thus facilitating perceptual sensitivity. Nevertheless, visual complexity and similarity are orthogonal concepts, and it is theoretically significant to know that both lead to similar learning benefits. In summary, both complexity and early discrimination difficulty during training could be considered examples of what Bjork calls "desirable difficulty" (Bjork, 1994).

Other results from this study indicate that this desirable difficulty improves future learning by increasing attention. We showed that if participants made a false alarm on one trial with a certain target character, they slowed down on the next occurrence of the target. This slow-down suggests that error feedback was associated with the target's representation. This caused people to be more cautious in representing the character the next time it appeared and/or to be more cautious in comparing it to distractors during the subsequent search phase. The current study does not allow us to determine whether the increased attention operates during the target encoding stage, during the subsequent search phase, or both. It is possible that both stages are affected by attention and determining which one contributes more to learning is an excellent venue for future research.

Finally, the fact that the TD similarity effect transferred from the visual search to the N-back task is crucial - it suggests that rather than improving task-specific performance, the effects of discrimination difficulty indeed cause representations of the characters to be stronger and more distinct.

This research is related to the concept of dimensional modulation in category learning, the idea that learning causes “stretching” of representations along discrimination-relevant dimensions (Folstein, Palmeri, Van Gulick, & Gauthier, 2015; Goldstone & Steyvers, 2001). Learning category-relevant dimensions makes stimuli more distinguishable along those simple dimensions, a concept also known as differentiation in cognitive models (Criss & Koop, 2015) or as pattern separation in neuroimaging studies (Yassa & Stark, 2011). Our findings extend these results by showing that in addition to occurring at the level of individual dimensions, differentiation can also occur with more complex representations.

Conclusion

We suggest that increasing the discrimination difficulty during learning improves future performance through a six stage process:

1. High discrimination difficulty initially leads to more partial matching errors
2. The post-error feedback becomes associated with the target’s representation
3. When the target is encountered again, the associated error feedback causes participants to pay more attention
4. As a result of increased attentiveness, participants slow down and engage in two processes that could improve learning. First, they might refine the representation of the target, encoding it more precisely; second, they might be more cautious when comparing the character representation in visual WM to the subsequently presented distractors, also leading to refinement of the target’s representation.
5. Over time the more cautious encoding and comparison to distractors leads to the development of stronger and more distinct long-term representations
6. Stronger representations support better future performance regardless of the task

Appendix A: Additional analyses

One could ask whether the superior learning associated with high TD similarity is merely a byproduct of longer search durations for trials with high TD similarity. Prior work has shown that searching for an object results in better memory for that object compared to simply attempting to remember it (Draschkow, Wolfe, & Vö, 2014), and it is possible that longer search would benefit memory even more. This explanation is not inconsistent with the claim that TD similarity improves learning, but suggests that whatever mechanism is responsible for the learning benefit, it occurs during the search process for the current trial. This view contrasts with our proposal that initial high TD similarity leads to errors, and the resulting

error feedback causes participants to slow down and pay more attention during the *subsequent* presentation of that character.

If the learning benefit associated with high TD similarity is a byproduct of longer search times, then the following two patterns should be observed in the data. First, longer search for a particular trial with a target character X should be associated with better search performance (lower RTs and higher accuracy) on the next trial that the same character appears (could be far in the future). Second, the effect of TD similarity should disappear if the RT for the previous occurrence of the target is included as a covariate in the regression models.

We tested both predictions in the following way. For each trial in which a certain character appears as a target, we calculated the search RTs for the same target in the last trial that included it. For example, if a target character X appears on trials 1, 7, 11, 23, etc, then for trials 7, 11 and 23 we included as a predictor the RTs for trials 1, 7 and 11 respectively.

First, we looked at performance during Session 1, because that is the session in which the TD effect begins as negative and then reverses. The previous occurrence RTs significantly predicted both RTs ($\Delta AIC = -162$, $\chi^2(1) = 163$, $p < .001$) and accuracy ($\Delta AIC = -3$, $\chi^2(1) = 4.96$, $p = .0259$) for the current trial; however, the effect in the opposite of the predicted direction – slower search RTs on the preceding occurrence of the target are associated with slower and less accurate search performance for the current trial with that target. These are not subject-based artifacts (i.e. a slower subjects tends to remain slower), because the mixed-effects regression model already accounts for different by-subject intercepts in RTs and accuracy. Furthermore, including this predictor in the model did not change the effect of TD similarity, nor its interaction with character repetition during Session 1.

We also looked at whether the average RTs for each character during the previous session N-1 would predict performance in Session N – if the alternative explanation is correct, then slower RTs during Session 1 should lead to better learning and as a result to better performance during Session 2, etc. Like in the previous analysis, we found exactly the opposite – slower RTs for a particular character during Session N-1 were associated with slower ($\Delta AIC = -2335$, $\chi^2(1) = 2336$, $p < .001$) and less accurate ($\Delta AIC = -304$, $\chi^2(1) = 307$, $p < .001$) search performance for that character during Session N. The effect was the same even when limiting the analysis only to Sessions 1 and 2, which is when the TD similarity effect reverses.

In summary, we found no evidence for the possibility that longer search times, *per se*, lead to better learning. In fact, longer search times were associated with worse future performance. Thus, this cannot explain the beneficial effect of TD similarity. Our proposal that high TD similarity leads to initial errors, which then cause participants to slow down during

subsequent trials with the same character and pay more attention, which in turn improves learnings, is more consistent with the data.

Open Practices Statement

The data and analysis code are available at <https://github.com/venpopov/similarity-discrimination-learning>

References

- Alexander, R., & Zelinsky, G. (2012). Effects of Part-based Similarity on Visual Search: The Frankenbear Experiment. *Vision Research*, 54, 20-30.
- Bichot, N., & Schall, J. (1999). Effects of similarity and history on neural mechanisms of visual selection. *Nature Neuroscience*, 2(6), 549-554.
- Bjork, R. A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe and A. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 185-205). Cambridge, MA: MIT Press.
- Criss, A., & Koop, G. (2015). Differentiation in episodic memory. *Psychology Press Festschrifts. Cognitive Modeling in Perception and Memory: A Festschrift for Richard M. Shiffrin*, 112-125.
- Draschkow, D., Wolfe, J. M., & Vo, M. L. H. (2014). Seek and you shall remember: Scene semantics interact with visual search to build better memories. *Journal of Vision*, 14(8), 10-10.
- Duncan, J., & Humphreys, G. (1989). Visual search and stimulus similarity. *Psychological Review*, 96(3), 433-458.
- Ellis, C. T., & Turk-Browne, N. B. (2019). Complexity can facilitate visual and auditory perception. *Journal of Experimental Psychology: Human Perception and Performance*, 45(9), 1271-1284.
- Feigenbaum, E. A., & Simon, H. A. (1984). EPAM-like Models of Recognition and Learning. *Cognitive Science*, 8(4), 305-336
- Folstein, J., Palmeri, T., Van Gulick, A., & Gauthier, I. (2015). Category learning stretches neural representations in visual cortex. *Current Directions in Psychological Science*, 24(1), 17-23.
- Gobet, F., Lane, P., Croker, S., Cheng, P., Jones, G., Oliver, I., & Pine, J. (2001). Chunking mechanisms in human learning. *Trends in Cognitive Sciences*, 5(6), 236-243.
- Goldstone, R., & Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General*, 130(1), 116.
- Nickerson, R., & Adams, M. (1979). Long-term memory for a common object. *Cognitive Psychology*, 11(3), 287-307.
- Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, 9(4), 441-474.

- Pashler, H. (1987). Target-distractor discriminability in visual search. *Perception & Psychophysics*, 41(4), 285-292.
- Phillips, S., Takeda, Y., & Kumada, T. (2006). An inter-item similarity model unifying feature and conjunction search. *Vision Research*, 46(22), 3867-3880.
- Popov, V., & Reder, L. M. (2020). Frequency effects on memory: A resource-limited theory. *Psychological Review*, 127(1), 1-46
- Reder, L., & Kusbit, G. (1991). Locus of the Moses illusion: Imperfect encoding, retrieval, or match? *Journal of Memory and Language*, 30(4), 385-406.
- Reder, L., Liu, X., Keinath, A., & Popov, V. (2016). Building knowledge requires bricks, not sand: The critical role of familiar constituents in learning. *Psychonomic Bulletin & Review*, 23(1), 271-277.
- Simon, H. A. (1974). How big is a chunk? *Science*, 183(4124), 482-488.
- Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97-136.
- Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, 95(1), 15-48.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4), 327.
- Xing, H., & Li, P. (2004). The acquisition of Chinese characters: Corpus analyses and connectionist simulations. *Journal of Cognitive Science*, 5(1), 1-49.
- Yang, J., McCandliss, B. D., Shu, H., & Zevin, J. D. (2009). Simulating Language-specific and Language-general Effects in a Statistical Learning Model of Chinese Reading. *Journal of Memory and Language*, 61(2), 238-257.
- Yassa, M., & Stark, C. (2011). Pattern separation in the hippocampus. *Trends in Neurosciences*, 34(10), 515-525.