Brief Communication

A multistage retrieval account of associative recognition ROC curves

Olya Hakobyan and Sen Cheng

Institute for Neural Computation, Ruhr University Bochum, 44081 Bochum, Germany

Despite its name, associative recognition is a paradigm thought to rely on memory recall. However, it remains unclear how associative information may be represented and retrieved from memory and what its relationship to other information, such as item memory, is. Here, we propose a computational model of associative recognition, where relational information is accessed in a generic, multistage retrieval process. The model explains the relative difficulty of associative recognition compared with item recognition, the difference in experimental outcomes when different types of lures are used, as well as the conditions leading to the emergence of associative ROC curves with different shapes.

Associative recognition is the process of distinguishing novel configurations of items from those encountered before. It combines elements of both recognition and recall. Consider the following case: After studying pairs of words, you have to distinguish the studied pairs from recombined pairs; that is, novel pairings of the studied words. Knowing whether the individual words are old or new (that is, item recognition) is not sufficient to solve this task, since all words have been studied. Rather, you need to recall whether the two words had been presented in that combination or not. However, unlike free recall or cued recall, associative recognition tasks provide a larger number of retrieval cues (Squire et al. 2007). Due to its unique place in memory paradigms, associative recognition has served as a tool to test statistical models of memory processes, albeit not without causing some controversies (Yonelinas et al. 2010).

One such controversy focuses on the involvement of two putative, distinct memory processes: familiarity and recollection (Mandler 1980; Yonelinas et al. 2010). The former is fast and effortless, and represents a vague feeling of having encountered the information before. The latter is a slow and conscious process and retrieves specific details. The relative contributions of these processes to recognition performance can be quantified using the receiver operating characteristics (ROC) analysis. In this approach, participants indicate whether they have encountered a test stimulus before, and their confidence in their memory judgement on a (typically six-point) scale. The recognition ROC curve is a plot of the cumulative hit rate (correct acceptance of studied stimulus) against the cumulative false alarm rate (incorrect acceptance of unstudied stimulus) for different confidence ratings.

ROC curves for item recognition are typically curvilinear with a pronounced *y* offset (Yonelinas and Parks 2007). There has been much debate about what these features represent with dual-process models mapping them to familiarity and recollection (Yonelinas 1997; Eichenbaum et al. 2007), respectively, and single-process models viewing them as an indication of differences in the distribution of the memory strength (Squire et al. 2007; Wixted 2007). We recently suggested an alternative account that goes beyond this dichotomy (Hakobyan and Cheng 2019, 2021). For associative recognition, the dual-process model predicts linear ROC curves, since familiarity is not useful in this task, and indeed associative ROC curves show little curvilinearity (Yonelinas et al. 1999; Healy et al. 2005). Nevertheless, they are significantly better fit by nonlinear than linear models and become more curvilinear when the memory strength is increased by repeated presentations (Kelley and Wixted 2001).

Several explanations have been suggested to account for curvilinear ROC curves in associative recognition. For instance, Mickes et al. (2010) proposed that recollection is a continuous variable, unlike the all-or-none threshold process suggested by the original dual-process models. Another suggestion is that associative ROC curves reflect a mixture of different signals, which result in a non-Gaussian distribution (DeCarlo 2003). Finally, a so-called unitization process might introduce a familiarity process into associative recognition by encoding the items in a pair as a single entity, wherever possible (Quamme et al. 2007; Haskins et al. 2008). For example, the two words "snow" and "man" can be encoded as a single word: "snowman."

Another interesting question concerns the relationship between item and associative information. Some theories suggest that item and associative information are combined into a single memory trace, while others propose separate storage and retrieval (Clark and Gronlund 1996). Experimental findings indicate that item and relational information are correlated but can also be dissociated. For example, selective attention to item encoding leads to poor associative recognition, while focusing on associations does not have an effect on item recognition (Hockley and Cristi 1996). Conversely, there is ample evidence that strong item memory improves relational memory, such as source retrieval (Glanzer et al. 2004; Slotnick and Dodson 2005). Another example, where item information interferes with associative judgments, is the increased false alarm rate when recombined pairs are used as lures as compared with when novel pairs are used (Gold et al. 2006).

The goal of this study is to account for several phenomena in associative recognition, such as the relative difficulty of associative recognition, the relationship between item and associative information, the difference between associative tests with recombined and novel items as well as the conditions that lead to the emergence of linear and curvilinear ROC curves. To this end, we adapted a computational model that we previously developed for item

Corresponding author: sen.cheng@rub.de

Article is online at http://www.learnmem.org/cgi/doi/10.1101/lm.053432.121.

recognition (Hakobyan and Cheng 2021). The model consists of input, memory and decision-making modules. In brief, we used images of faces as stimuli (Fig. 1A). The images were processed in the HMAX hierarchical neural network (Serre and Riesenhuber 2004), which reflects the processing along the ventral visual stream. The network's output then undergoes principle component analysis (PCA) for dimensionality reduction, mapping each image to a 6D vector. These vectors represent the face stimuli in our model.

When an image pair is "studied," the memory module stores representations of both individual items $x_{i1}$and x_{i2} in the item memory. Their association in a pair is represented by concatenating the representations of the two items into a single 12D pattern (x_{i1}, x_{i2}), which is independently stored in the association memory. Note that this is different from unitization, which would encode a pair as an individual (compound) item—represented by a 6D vector. The memory module uses a memory robustness parameter p to scale distances between incoming patterns $y_i = px_i$ (for more details, see Hakobyan and Cheng 2021).

To mimic the imperfect nature of memory, each pattern is modified by noise according to Equation 1:

$$\tilde{y}_i = y_i + \epsilon_i, \tag{1}$$

where y_i is the original pattern and ε is the noise vector, whose components are independently and normally distributed with zero mean and variance ω.

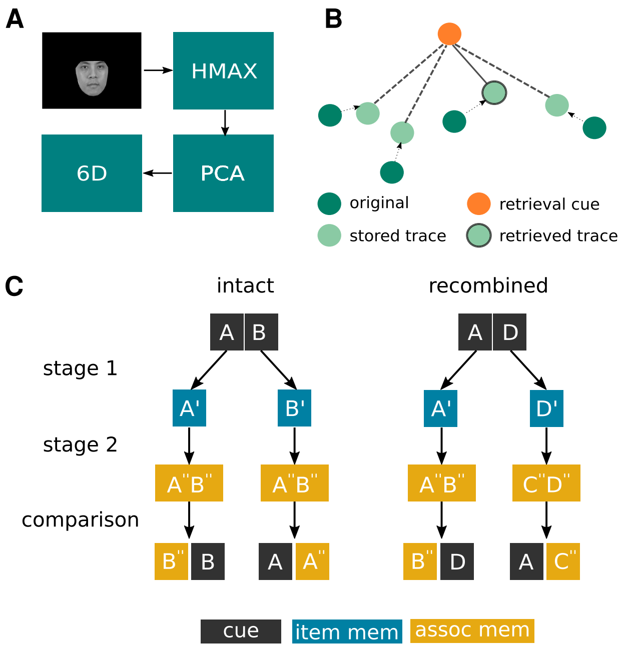

Figure 1. Input and memory retrieval process. (*A*) Naturalistic face images are turned into 6D representations after being processed in the HMAX model and undergoing PCA for dimensionality reduction. (*B*) Memory traces are noisy versions of the original patterns. For memory retrieval, the cue is compared with the stored items and the one with the lowest distance is retrieved. (*C*) Multistage memory retrieval for associative task illustrated for an intact test stimuli and for a recombined pair. Each item in a pair is used as a cue to retrieve the corresponding memory trace at item level (stage 1), which is then used for further retrieval at associative level (stage 2). Finally, the item from the test cue that was not used as the current retrieval is compared with the corresponding part of the retrieved trace to arrive at a memory judgement.

During the memory test, either the studied pairs or one of two lure types (pairs recombining studied items or pairs consisting of novel items) are presented as cues. The test initiates parallel retrieval strands with two stages (Fig. 1C). In the first stage, each individual item is used as a cue y' to retrieve a memory trace of that item. For that purpose, the cue is compared with the memory traces stored at the item level and the trace with the lowest correlation distance (Eqs. 2, 3) is retrieved (Fig. 1B):

$$\hat{\imath} = \underset{j}{argmin}\, D\left(\tilde{y}_j, y'\right) \tag{2}$$

and

$$\hat{y} = \tilde{y}_{\hat{\imath}}. \tag{3}$$

This triggers the second stage of the process, in which the item retrieved in the first stage ($\hat{y}_i$ in Eq. 3) is used as a cue to retrieve an association (Fig. 1C, stage 2). Since items and pairs have different dimensionality, the missing features are padded with zeros to enable a distance comparison. In this stage, too, the pattern with the lowest distance to the cue is retrieved. If that distance is below a threshold γ, associative retrieval is deemed successful.

If associative retrieval is successful for at least one strand (two parallel strands in Fig. 1), the model makes the final associative recognition judgment as follows. It uses the distance d_i between the item in the test probe that was not used as a cue in the given retrieval instance and the corresponding item in the retrieved association. If both strands were successful, the recognition judgment is based on the strand that results in the smaller d_i. If the distance is smaller than a given threshold (θ_j), a positive response is returned:

$$r_i = \begin{cases} 1\ if\ d_i \le \theta_j + \lambda \\ 0\quad otherwise \end{cases}. \tag{4}$$

If associative retrieval fails for both strands of retrieval, the recognition judgment is made based on item information alone: If the distance between retrieved items and their respective cues are both smaller than the item threshold (similar to $\theta_j + \lambda$ in Eq. 4), a positive response is given. This fallback option can be viewed as informed guessing based on item information when associative information is unavailable. It will lead to the wrong result in some cases but might return the correct result in a sufficient number of cases to be useful.

The decision thresholds are adjusted to the associative or item memory strengths (similar to the range model in Hirshman 1995) and the term λ is a small bias term (for details, see Hakobyan and Cheng 2021).

To illustrate how the elements of the model work, we first analyzed retrieval accuracy (that is, whether the correct memory trace is retrieved) separately at each retrieval stage. In the first retrieval stage, accuracy is relatively high; that is, the cue leads to the correct retrieval of the corresponding memory trace in most cases (Fig. 2A). This accuracy decreases in the next stage of the retrieval since the memory trace from the first stage used as retrieval cue is noisy. However, if parallel retrieval processes are initiated with different test probes (two-stranded retrieval in Fig. 1C), the accuracy in the second stage increases since the model has more opportunities to retrieve the correct memory trace. This finding is consistent with experimental results showing that memory performance benefits from the availability of multiple cues (Tulving and Thomson 1971; Humphreys 1976). Both stages of the retrieval benefit when memory robustness (p) is high. Correct retrieval results in a

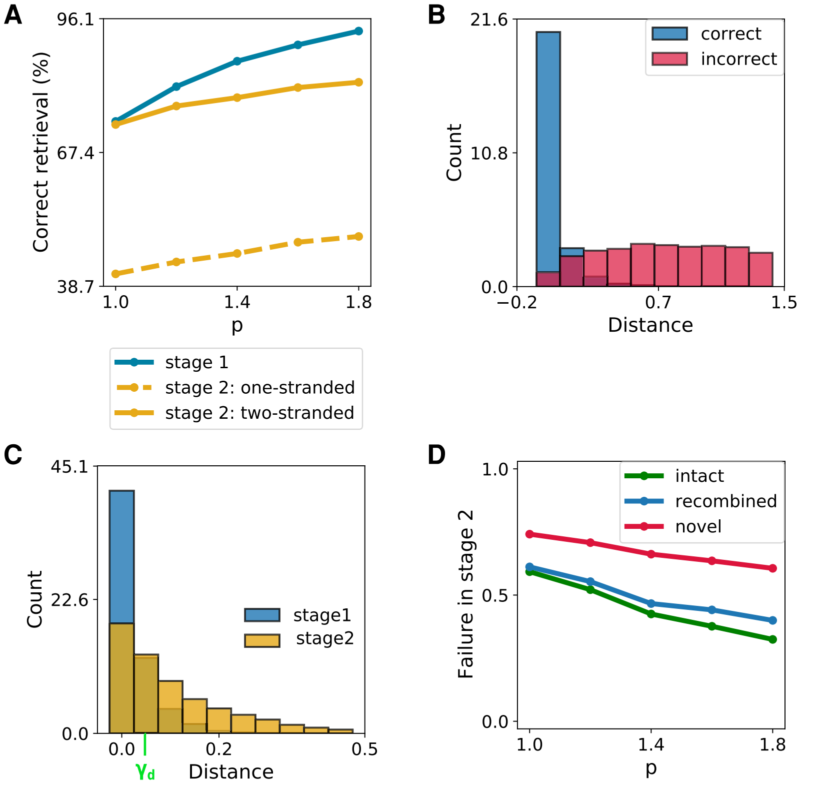

Figure 2. Assessment of retrieval accuracy and precision. (*A*) Retrieval accuracy is high if the memory system is noise-resistant. Nevertheless, the accuracy suffers with each retrieval stage (stage 1: item retrieval; stage 2: associative retrieval) unless the number of available cues increases. (*B*) Correctly retrieved targets are more similar to the test cue, resulting in a higher chance of subsequent recognition. Here and in the next subplot, transparency is used to make both histograms visible even where they partially overlap. (C) Retrieval in stage 2 is less precise than in stage 1 because it is cued with the noisy memory traces retrieved in stage 1. γ_d denotes the associative failure threshold for this particular simulation instance. (*D*) Associative failure (exceeding γ_d) is more likely for novel pairs than for both recombined lures and targets, for which associative failure is about equally likely.

smaller decision distance; that is, distance between the memory probe and the retrieved pattern (Fig. 2B). This occurs because for correct retrieval, the distance reflects only the magnitude of the memory noise ε_i, whereas for incorrect retrieval, the distance additionally includes the interpattern distance to the memory trace of another item. These results indicate that correct retrieval in both stages of retrieval is relevant for the subsequent memory performance. Memory traces of incorrectly retrieved items are likely to deviate more from the test probes and therefore more likely to miss the recognition threshold.

Given the importance of reliable memory traces, we next examined precision of the retrieval in each stage. We analyzed the distances between the cue and the retrieved memory pattern (Fig. 2C) and found that the retrieval precision decreases in the second stage, because the cue is a corrupted version of the originally presented item. Together with the accuracy findings in Figure 2A, this result suggests that associative retrieval is more difficult and error-prone as compared with item retrieval. If a threshold is applied to reject less reliable retrieval instances (γ_d in Fig. 2C), the second stage of the retrieval (i.e. associative retrieval) is more likely to fail.

Next, we compared associative retrieval for intact, recombined, and novel pairs (Fig. 2D). We defined associative failure as cases, where the distance between the cue and the retrieved pair is larger than the predefined threshold γ for both strands of the retrieval. For recombined pairs, the probability of associative failure is similar to that of the target pairs because both items in a recombined pair are stored in a memory trace at item level and are part of a memory trace at the associative level. Nevertheless, a noticeable difference between recombined and target pairs emerges as memory robustness increases, because, in the second stage, the cues in the recombined pairs lead to two different memory traces (Fig. 1C), while the cues in the target pairs are likely to lead to the same trace that was stored before ($A''B''$ in Fig. 1C). Thus, when the memory robustness increases, the retrieval of a well-preserved memory trace is less likely to fail when two rather than one cue is linked to that trace. In contrast, associative failure is more likely for novel pairs, because no memory traces are stored for the items, and the cues are less likely to closely match another memory trace. As a result, we expect to see a difference in associative recognition performance between recombined and novel pairs in our model, similar to those observed experimentally.

We therefore simulated experiments testing intact pairs against either recombined or novel pairs. To assess the role of the associative failure on the associative recognition ROC curve, we examined three settings. In the first setting, a low threshold value for associative retrieval success ($\gamma = 0$) ensured that memory retrieval always fails in the second stage. Since no association is retrieved, the recognition judgment is made solely based on the outcome of the first retrieval stage; that is, item recognition. Since there is no difference in item memory between intact and recombined pairs, this condition results in random recognition for recombined pairs (Fig. 3A, left). In contrast, since no memory traces were stored for items in novel pairs, associative recognition based on item information alone yielded a good performance for novel pairs (Fig. 3A, right). This result confirms that informed guessing based on item information can be useful for associative judgements.

The second setting examined the other extreme, when associative retrieval always succeeds ($\gamma = 10$). In this setting, associative recognition is based exclusively on associative information and, therefore the ROC curves are significantly above chance performance for both recombined and novel pairs (Fig. 3B). The two are similar because associative information was stored for neither type of lures.

The ROC curves in the first two settings are both inconsistent with experimental results, which show that recognition performance is above chance for both recombined and novel pairs and typically lower for recombined pairs due to higher false alarm rate. This pattern of results is predicted by our model when the two extremes are balanced with an intermediate setting of the associative failure threshold ($\gamma = 0.03$) (Fig. 2C,D). This setting ensures that retrieval in the second stage fails sometimes and succeeds sometimes, so that associative recognition is based on both item and associative information. This setting reproduces the aforementioned experimental results quite well (Fig. 3C).

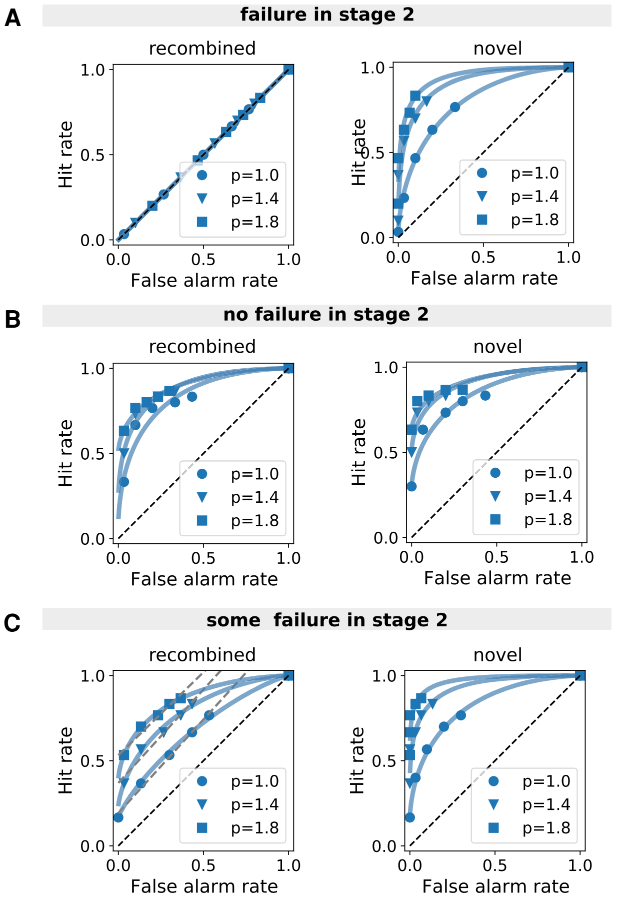

Figure 3. Associative recognition ROC curves (*A*) When the decision criterion is set such that associative retrieval fails in every case ($\gamma = 0$), associative recognition is solely based on item recognition; that is, correct retrieval in stage 1. In that case, the performance is at chance level for recombined (*left*), but not for novel lures (*right*). (*B*) When associative retrieval succeeds every time ($\gamma = 10$), which means that associative recognition is based solely on correct retrieval in stage 2, both recombined and novel pairs lead to good performance. (*C*) If the decision criterion is chosen such that associative retrieval fails sometimes ($\gamma = 0.03$), then associative recognition relies on both item and associative information. As a result, recombined lures lead to poorer performance compared with novel lures and the ROC curve becomes more linear for some conditions (gray dashed lines indicate linear fit).

Intriguingly, our model reproduces results on the curvilinearity of the ROC curve without additional fine-tuning. For recombined pairs, the ROC curve was close to linear in some conditions (that is, at low memory robustness; $P = 1.0$, adjusted R^2 linear = 0.99, adjusted R^2 quadratic = 0.99), whereas in others (that is, at high memory robustness; $P = 1.8$), it was more curvilinear (adjusted R^2 linear = 0.94, adjusted R^2 quadratic = 0.98) (Fig. 3C, left). This parallels observations that stronger associative memory increases the curvilinearity of associative ROC curves (Kelley and Wixted 2001).

Based on our modeling results, we propose that successful item recognition is a necessary, but not sufficient, condition for associative recognition. Although the correlation between item and associative recognition had been discussed in the literature (Slotnick and Dodson 2005; Onyper et al. 2010), we went one step further by suggesting that item retrieval is the first step in a multistage retrieval process that leads to the retrieval of further details, such as associations. The first reason behind this assumption is that the quality of item memory can aid or disrupt relational memory (Glanzer et al. 2004; Slotnick and Dodson 2005). Second, it appears that associative information becomes available in the later stages of memory retrieval (Cox and Criss 2020). Finally, it seems intuitive that recognition of individual items would precede and guide the retrieval of further details, similar to the "butcher on the bus" phenomenon, when the mere recognition of a person precedes the recollection of the context, in which the person was encountered. Nevertheless, the item and associative representations are stored independently to account for cases where associative information might fail despite successful item recognition (Hockley and Cristi 1996).

Our model predicts that if item recognition is poor, the retrieval of associative information is likely to fail, too, consistent with findings attributing noise in relational tasks to failed item recognition, which results in linear ROC curves (Slotnick and Dodson 2005; Mickes et al. 2010). Linear ROC curves in our model are linked to associative failure due to either the poor quality of stored associations or the noisy item cue. Especially the latter is consistent with the aforementioned studies. However, the former scenario still allows for good item and noisy associative recognition. We predict improved associative recognition even when item information is strengthened independently of associative information, as it would benefit from item cues with higher quality. Obviously, establishing the necessary conditions is a challenge for experiments (e.g., avoiding recall of pairs when constituent items are restudied), but perhaps items could be presented in isolation before they are paired with other items. A similar strengthening of the target item in a cued-retrieval task improved subsequent recall, even when the associated, but nonstrengthened, item was used as the cue (Humphreys and Bowyer 1981).

The variable shape of the associative ROC curves has generated much speculation regarding the involvement of different memory processes and their statistical properties, such as familiarity, continuous recollection, unitization, and the role of noise. In contrast, we moved away from postulating specific hypothetical distributions (Gaussian or otherwise) a priori and proposed a more mechanistic model, which operates upon distributions derived from sensory input. Using generic memory storage and retrieval mechanisms, we showed that qualitatively different ROC curves can arise without the involvement of distinct memory processes. It is possible that such generic memory mechanisms coupled with an attribution system (Bastin et al. 2019) might lead to certain phenomenological experiences, such as familiarity and recollection. However, the phenomenology plays no functional role in our model. The failure threshold for retrieval in our model may resemble the threshold process proposed for recollection in the dual-process model. However, the processes at different stages did not differ qualitatively here, as they do in the dual-process model. Importantly, our model exhibits a behavior very similar to the widely studied recall-to-reject mechanism without explicitly implementing it. Consider probing the memory in our model with the recombined pair AD. If the cues in this pair (*A*,*D*) lead to the retrieval of the original associations (A″B″, C″D″) and the quality of the associative memory trace is relatively high, the recognition judgment (based on the distance between B″ and *D*, or C″ and *A*) is likely to lead to a rejection.

In sum, our work shows that several conflicting findings in the associative recognition literature may be explained using only a generic memory model and multistage retrieval paired with decision-making processes.

Acknowledgments

This work was funded by the German Research Foundation (Deutsche Forschungsgemeinschaft [DFG])—project number 122679504-SFB 874, B2.

References

Bastin C, Besson G, Simon J, Delhaye E, Geurten M, Willems S, Salmon E. 2019. An integrative memory model of recollection and familiarity to understand memory deficits. *Behav Brain Sci* **42:** e281. doi:10.1017/S0140525X19000621

Clark SE, Gronlund SD. 1996. Global matching models of recognition memory: how the models match the data. *Psychonom Bull Rev* **3:** 37–60. doi:10.3758/BF03210740

Cox GE, Criss AH. 2020. Similarity leads to correlated processing: a dynamic model of encoding and recognition of episodic associations. *Psychol Rev* **127:** 792–828. doi:10.1037/rev0000195

DeCarlo LT. 2003. An application of signal detection theory with finite mixture distributions to source discrimination. *J Exp Psychol Learn Mem Cogn* **29:** 767–778. doi:10.1037/0278-7393.29.5.767

Eichenbaum H, Yonelinas A, Ranganath C. 2007. The medial temporal lobe and recognition memory. *Annu Rev Neurosci* **30:** 123–152. doi:10.1146/annurev.neuro.30.051606.094328

Glanzer M, Hilford A, Kim K. 2004. Six regularities of source recognition. *J Exp Psychol Learn Mem Cogn* **30:** 1176–1195. doi:10.1037/0278-7393.30.6.1176

Gold JJ, Hopkins RO, Squire LR. 2006. Single-item memory, associative memory, and the human hippocampus. *Learn Mem* **13:** 644–649. doi:10.1101/lm.258406

Hakobyan O, Cheng S. 2019. How do memory modules differentially contribute to familiarity and recollection? *Behav Brain Sci* **42:** e288. doi:10.1017/S0140525X19001833

Hakobyan O, Cheng S. 2021. Recognition receiver operating characteristic curves: the complex influence of input statistics, memory, and decision-making. *J Cogn Neurosci* 1–24. doi: 10.1162/jocn_a_01697

Haskins AL, Yonelinas AP, Quamme JR, Ranganath C. 2008. Perirhinal cortex supports encoding and familiarity-based recognition of novel associations. *Neuron* **59:** 554–560. doi:10.1016/j.neuron.2008.07.035

Healy MR, Light LL, Chung C. 2005. Dual-process models of associative recognition in young and older adults: evidence from receiver operating characteristics. *J Exp Psychol Learn Mem Cogn* **31:** 768–788. doi:10.1037/0278-7393.31.4.768

Hirshman E. 1995. Decision processes in recognition memory: criterion shifts and the list-strength paradigm. *J Exp Psychol Learn Mem Cogn* **21:** 302–313. doi:10.1037/0278-7393.21.2.302

Hockley WE, Cristi C. 1996. Tests of encoding tradeoffs between item and associative information. *Mem Cognit* **24:** 202–216. doi:10.3758/BF03200881

Humphreys MS. 1976. Relational information and the context effect in recognition memory. *Mem Cognit* **4:** 221–232. doi:10.3758/BF03213167

Humphreys MS, Bowyer PA. 1981. Cued recall and cue recognition: how similar are the processes? *Am J Psychol* **94:** 65–84. doi:10.2307/1422343

Kelley R, Wixted JT. 2001. On the nature of associative information in recognition memory. *J Exp Psychol Learn Mem Cogn* **27:** 701–722. doi:10.1037/0278-7393.27.3.701

Mandler G. 1980. Recognizing: the judgment of previous occurrence. *Psychol Rev* **87:** 252–271. doi:10.1037/0033-295X.87.3.252

Mickes L, Johnson EM, Wixted JT. 2010. Continuous recollection versus unitized familiarity in associative recognition. *J Exp Psychol Learn Mem Cogn* **36:** 843–863. doi:10.1037/a0019755

Onyper SV, Zhang Y, Howard MW. 2010. Some-or-none recollection: evidence from item and source memory. *J Exp Psychol Gen* **139:** 341–364. doi:10.1037/a0018926

Quamme JR, Yonelinas AP, Norman KA. 2007. Effect of unitization on associative recognition in amnesia. *Hippocampus* **17:** 192–200. doi:10.1002/hipo.20257

Serre T, Riesenhuber M. 2004. *Realistic modeling of simple and complex cell tuning in the HMAX model, and implications for invariant object recognition in cortex*. Defense Technical Information Center, Fort Belvoir, VA.

Slotnick SD, Dodson CS. 2005. Support for a continuous (single-process) model of recognition memory and source memory. *Mem Cognit* **33:** 151–170. doi:10.3758/BF03195305

Squire LR, Wixted JT, Clark RE. 2007. Recognition memory and the medial temporal lobe: a new perspective. *Nat Rev Neurosci* **8:** 872–883. doi:10.1038/nrn2154

Tulving E, Thomson DM. 1971. Retrieval processes in recognition memory: effects of associative context. *J Exp Psychol* **87:** 116–124. doi:10.1037/h0030186

Wixted JT. 2007. Dual-process theory and signal-detection theory of recognition memory. *Psychol Rev* **114:** 152–176. doi:10.1037/0033-295X.114.1.152

Yonelinas AP. 1997. Recognition memory ROCs for item and associative information: the contribution of recollection and familiarity. *Mem Cognit* **25:** 747–763. doi:10.3758/BF03211318

Yonelinas AP, Parks CM. 2007. Receiver operating characteristics (ROCs) in recognition memory: a review. *Psychol Bull* **133:** 800–832. doi:10.1037/0033-2909.133.5.800

Yonelinas AP, Kroll NE, Dobbins IG, Soltani M. 1999. Recognition memory for faces: when familiarity supports associative recognition judgments. *Psychonom Bull Rev* **6:** 654–661. doi:10.3758/BF03212975

Yonelinas AP, Aly M, Wang W-C, Koen JD. 2010. Recollection and familiarity: examining controversial assumptions and new directions. *Hippocampus* **20:** 1178–1194. doi:10.1002/hipo.20864

Received April 21, 2021; accepted in revised form August 16, 2021.