Item Distinctiveness Modulates the Relationship Between Hits and False Alarms in Old-New Recognition

Daniil Azarov[1,2] (dazarov@iu.edu) and Robert M. Nosofsky[1]
[1]Department of Psychological and Brain Sciences and [2]Department of Statistics, Indiana University Bloomington

Abstract

Recognition memory is strongly shaped by similarity structure, yet how item distinctiveness reorganizes the relationship between hits and false alarms is not well understood. Using distances in a multidimensional feature space to quantify item similarity, we examined item-level recognition performance for real-world objects. Low-distinctive items showed positive correlations between hit rates and false-alarm rates, reflecting variation in inter-item similarity as predicted by standard global-familiarity models. In contrast, highly distinctive items exhibited high hits but low false alarms, reducing or reversing the overall hit–false-alarm correlation. Split-half analyses showed that both hits and false alarms were reliable across participants, but their relationship depended on an item's position within the surrounding similarity space. The Hybrid-Similarity Exemplar model combining inter-item similarity with a distinctiveness-dependent self-match signal captured this reorganization. These findings show that memorability is reliable yet relational, emerging from similarity structure rather than intrinsic stimulus properties.

Keywords: recognition memory; distinctiveness; similarity; memorability; computational modeling

Introduction

Recognition memory allows people to determine whether a current experience has occurred before. A central determinant of recognition performance is similarity: targets that are highly similar to other studied items and to lures presented at time of test tend to be harder to discriminate from the lures, whereas more dissimilar targets are typically remembered more accurately (Brady & Störmer, 2024; Shen et al., 2023). This emphasis on similarity is shared across a wide range of theoretical approaches, including global familiarity models, which assume that recognition judgments depend on the structure of similarity relations among items (e.g., Gillund & Shiffrin, 1984; Hintzman, 1988; Nosofsky, 1991).

Within similarity-based frameworks such as standard global-familiarty models, recognition decisions are often assumed to depend on aggregated similarity between a probe and stored memory representations. For example, Schurgin et al. (2020) demonstrated that both perception and memory rely on similarity representations, and showed that once psychophysical similarity is taken into account, a wide range of memory phenomena can be explained within a unitary signal detection framework. These findings highlight the importance of item-level similarity structure in shaping recognition behavior across different memory systems.

The Hybrid-Similarity Exemplar model (Nosofsky & Zaki, 2003) provides a useful conceptual framework for understanding how different sources of similarity contribute to recognition judgments. As is the case in standard global-familarity models, familiarity arises from a summing of inter-item similarity (similarity between a probe and other studied items) and self-similarity (similarity between a probe and its own stored representation). Self-similarity selectively boosts familiarity for studied items compared to unstudied ones, allowing the model to account for the general finding that hit rates to targets tend to exceed false-alarm rates to lures. Likewise, lures that occupy high-density regions of a similarity space are predicted to have high false-alarm rates, because their summed inter-item similarity will be high. The crucial component of the Hybrid-Similarity Exemplar model that extends earlier versions of global-familiarity models is its assumption that old target items with highly distinctive features provide stronger matches to their own memory traces than do typical targets without highly distinctive features: This mechanism allows the model to account for enhanced recognition of targets that occupy relatively isolated regions of a multidimensional feature space (cf. Valentine, 1991).

Recent work demonstrates the importance of this mechanism for explaining item-level recognition performance in realistic stimulus domains. Meagher and Nosofsky (2023) and Nosofsky et al. (2025) examined old–new recognition for high-dimensional, real-world stimuli and showed that a standard exemplar model offered a weaker account of item-level variability in hit rates than did the hybrid-similarity version. The extended hybrid-similarity model that allowed for boosts in self-similarity due to matching distinctive features provided a substantially improved account of item-level recognition data. These findings suggest that distinctiveness-related self-similarity plays a critical role in shaping recognition judgments beyond what can be explained by inter-item similarity alone.

This distinction between inter-item similarity and self-similarity has direct implications for the predicted relationship between hit rates and false-alarm rates at the item level. For items that are not especially distinctive, self-similarity is relatively uniform, and differences in recognition performance are driven primarily by inter-item similarity. Under these conditions, items that are highly similar to other

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

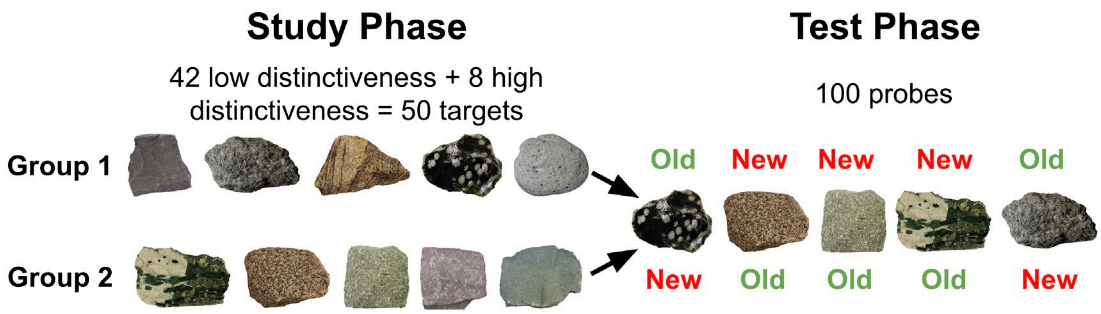

Figure 1: The study phase consisted of 42 rock images with low distinctiveness and 8 rocks with high distinctiveness, resulting in 50 targets total. Each participant saw only one set of targets depending on the group they were assigned to (set A or B). After the study phase both sets (A and B) were included in the test phase. In this example, the fourth study image in Group 1 and the first study image in Group 2 provide examples of rocks that received high distinctive feature ratings.

studied items should tend to produce both higher hit rates and higher false-alarm rates, whereas items that are less similar should tend to produce both lower hit rates and lower false-alarm rates. Thus, for low-distinctive items, the hybrid-similarity model predicts a positive correlation between item-wise hit rates and false-alarm rates.

Highly distinctive items, however, introduce opposing forces on these measures. Because distinctive items have greater self-similarity, they may yield high hit rates when they serve as old studied targets. At the same time, their isolation in feature space implies low inter-item similarity, leading to low false-alarm rates when they serve as new unstudied lures. As a result, including highly distinctive items in an item-level analysis should attenuate the positive hit – false-alarm correlation observed for low-distinctive items and may even produce a negative relationship in an overall analysis across all items. This predicted qualitative pattern follows naturally from the hybrid-similarity model as well as other modern global-familiarity models that distinguish between self-similarity and inter-item similarity and that predict high self-similarity for distinctive items (Cox, 2026; Shiffrin & Steyvers, 1997).

Recent work provides evidence consistent with this logic. For example, Utochkin et al. (2025a) reported a study that found that items with greater average distance from other studied items showed lower false-alarm rates and higher hit rates, suggesting that item-level similarity structure plays a critical role in shaping recognition outcomes. However, the relationship between hits and false alarms as a function of item distinctiveness has not been systematically examined.

In the present study, we investigate how item-level distinctiveness modulates the relationship between hit rates and false-alarm rates in recognition memory. Using distances in a multidimensional feature space to quantify item similarity, we test the prediction that low-distinctive items exhibit a positive correlation between hits and false alarms, reflecting variation in inter-item similarity alone. Crucially, we further test whether this relationship is reduced or reversed when highly distinctive items are included, consistent with the contribution of self-similarity. By focusing on item-level performance, this work clarifies how distinctiveness shapes recognition judgments and provides a qualitative test of core assumptions shared by modern global familiarity models of memory.

Method

The experiment was conducted online through SONA and implemented using jsPsych (de Leeuw et al., 2023). The study was approved by the Indiana University Institutional Review Board.

Participants

All participants were students who completed the online study for course credit. Before beginning the experiment, participants provided informed consent. A total of 220 participants were recruited. Each participant was randomly assigned to one of two groups (see Stimuli and Procedure).

In this report, our central focus is on participants who scored in the upper-median of hit – false alarm differences (although we consider patterns of performance for the lower-median participants as well). There is a wide variety of reasons why lower-median participants may have performed poorly, including lack of motivation, failure to understand instructions, and unknown events taking place in the online experiment itself. Such occurrences go outside the ability of extant formal models to explain, so our main focus is on the good-performing participants. After excluding the lower-median participants, 112 participants remained (61 in Group 1, 51 in Group 2).

Stimuli

We selected items from the set of rock-image stimuli used by Nosofsky et al. (2018) and Meagher and Nosofsky (2023), which was composed of 30 categories of 16 members each (480 items). A multidimensional scaling (MDS) solution for the 480 images was derived in those studies as well. Item-level distinctiveness ratings for all 480 images were taken from Meagher and Nosofsky (2023). We constructed two stimulus sets (Set A and Set B) from the complete collection of 480 rocks described above. For each set, we selected 8 target rocks with high distinctiveness ratings and 42 target

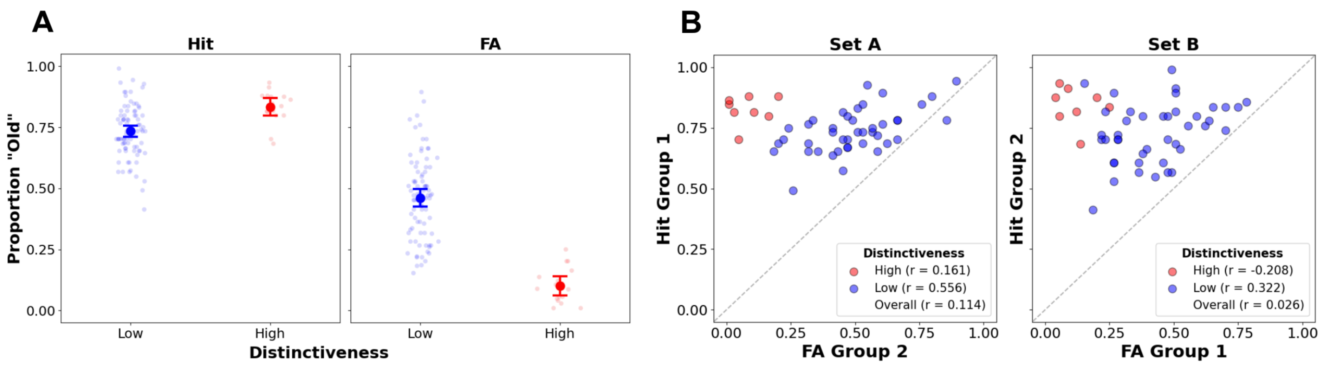

Figure 2: (A) Jitter-plots of mean individual-item hit and false-alarm rates as a function of distinctiveness. Under the present experimental conditions, high-distinctive items produce lower false alarms, but higher hit rates compared to low-distinctive items. Error bars represent 95% CIs. **(B)** These panels show that, regardless of stimulus set (A or B, i.e. Group 1 or Group 2 targets), low-distinctive items exhibit positive correlations between hits and false alarms, whereas high-distinctive items tend to cluster in the upper-left corner, reducing the overall correlation across all items.

rocks with low distinctiveness ratings. None of the target images were the same across the two sets. Items that served as targets in Set A served as lures in Set B, and vice-versa. The 16 high-distinctive rocks were chosen from those in the 480-item set that had the 40 highest distinctive-feature ratings; in addition, we sought to choose the items such that no two distinctive rocks resembled one another. The 84 low-distinctive rocks were chosen from among those that had the 100 lowest distinctive-feature ratings; an additional constraint is that we chose 6 rocks from each of 14 low-distinctive categories, with 3 rocks from each category assigned as targets in each set. Thus, the low-distinctive rocks tended to resemble subsets of rocks from their own set and from the contrast set. In other words, we sought to choose distinctive items such that they occupied isolated regions of the multidimensional similarity space, whereas low-distinctive items tended to occupy denser regions of the similarity space, both within and across sets.

Procedure

Participants were randomly assigned to one of two groups and studied targets of only one set (A or B). They then engaged in a recognition test phase in which all items from both sets were tested and they judged whether each item was old or new. Thus, each participant studied 50 targets (either Set A or Set B) and was tested with 100 probes (all items from Sets A and B; see Figure 1). Items that served as targets for Group 1 served as lures for Group 2, and vice-versa.

All items were presented in a new random order for each participant in both the study and test phases. The study phase was repeated three times, with a new random order on each repetition. During the study phase, each image was presented for 2,500 ms along with its category label, followed by a 1,000 ms inter-stimulus interval (ISI). The category labels were included to keep participants engaged in the task by giving them a chance to learn the geologic names of different categories of rocks.

In the test phase, each probe was presented until the participant made a response. Participants made an old/new judgment, followed by 500 ms of corrective feedback.

Results

"Old"-Judgments as a Function of Distinctiveness

The mean probabilities with which old items were judged to be old (Hits) and new items were judged to be old (FA) are shown in Figure 2A. Using items as the unit of analysis, we conducted a 2 x 2 ANOVA on the "Old" response data with stimulus status (old vs new) and distinctiveness (high vs low) as factors. There was: a main effect of distinctiveness, $F(1, 196) = 23.16$, $p < .001$, driven by the very low FA rates of the high-distinctive items; a main effect of stimulus status, $F(1, 196) = 778.55$, $p < .001$, confirming simply that hit rates were higher than FA rates; and a significant interaction between stimulus status and distinctiveness, $F(1, 196) = 184.26$, $p < .001$. Most important, as anticipated by the hybrid-similarity model, planned comparisons showed that hit rates were significantly higher for high-distinctive than for low-distinctive items, $p < .001$; whereas false-alarm rates were significantly lower for high-distinctive than for low-distinctive items, $p < .001$.

Discussion The finding that lures with highly distinctive features had low FA rates is unsurprising: the summed inter-item similarity for such lures will be low, so essentially all global-familiarity models predict they will have low FA rates. The logic involving the high hit rates for the high-distinctive items is more subtle. For old target items, there are two opposing forces. On the one hand, the presence of a highly distinctive feature leads to a greater degree of self-match between the test probe and its own representation in memory. On the other hand, the highly distinctive feature tends to lower its inter-item similarity to the remaining study items. The interpretation from the hybrid-similarity model is that the elevated hit rates for the high distinctive targets arise because, under the present conditions, the boost in self-similarity outweighs the reductions in inter-item similarity in determining the overall summed similarity. We discuss this issue in greater depth in our General Discussion.

Hits – False Alarms Correlations

We now address the central question of interest, namely the item-level relationship between hits and false alarms. We computed correlations separately for high-distinctive items, low-distinctive items, and across all items.Because each object in the two sets was a target in one group and a lure in the other, we correlated hits (H) from Group 1 with false alarms (FA) from Group 2, and vice versa (Figure 2B).

As predicted, for the low-distinctive items (blue dots in Figure 2B), there was indeed a positive correlation between the individual-item H and FA rates in both sets: r = .56, p < .001 in Set A; and r = .32, p =.038 in Set B. By contrast, when the high-distinctive (red dots) and low-distinctive (blue dots) items are analyzed collectively, the correlations between the H and FA rates are vastly reduced and do not approach statistical significance: r = .11, p = .43 in Set A; and r = .03, p = .86 in Set B. The basis for the vastly reduced correlations is easily seen in the Figure 2B scatterplots: in both sets, the red dots are clustered in the upper-left corner of the plots, reflecting the high H rates and low FA rates of the high-distinctive items – and causing any positive overall correlation between H and FA rates across all items to disappear. (A technical question involves the correlations within the small group of high-distinctive (red-dot) items considered on their own: Because of reduced range of H and FA rates in these small clusters and because of the opposing forces of self- and inter-item similarity, we did not have clear predictions regarding any H-FA correlations that might be observed *within* the red-dot clusters.)

Discussion In sum, here we report significant positive correlations between hits and false alarms for typical (low-distinctive) objects that occupy dense regions of similarity space; but vastly reduced correlations when high- and low-distinctive items are analyzed collectively. Therefore, we demonstrate that the distinctiveness of an object can act as a moderator of the relationship between hits and false alarms. As we will show, this qualitative pattern of results is as predicted by the hybrid-similarity exemplar model.

Split-half Correlations

Finally, we assessed the reliability of item-level effects by computing split-half correlations (10,000 iterations for each condition) across two independent groups of participants, corrected using the Spearman–Brown formula. Hit rates showed high reliability (mean ρ_{SB} = .74, SD = 0.04), as did false-alarm rates (mean ρ_{SB} = .91, SD = 0.01). In contrast, correlations computed after random permutation of item labels were near zero for both hits (mean ρ_{SB} = −.02, SD = 0.21) and false alarms (mean ρ_{SB} = −.03, SD = 0.21). These large split-half correlations indicate that both hit and false-alarm rates are highly consistent across participants at the item level, rather than the results reflecting noisy averages. This reliability supports the interpretation that the observed item-wise relationships reflect stable properties within this experimental context.

Modelling

To account for these patterns, we applied the hybrid-similarity exemplar model (Nosofsky & Zaki, 2003; Meagher & Nosofsky, 2023), a global-familiarity model that incorporates item distinctiveness in its machinery. According to this model, the probability of responding "Old" to probe-item i in the test phase is:

$$P(\text{Old}|i) = \frac{F_i}{F_i + k} \tag{1}$$

where F_i is the global-familiarity signal produced by item i, and k is a freely estimated criterion parameter. Familiarity is defined as:

$$F_i = \sum_{j=1}^{J} s'_{ij} \tag{2}$$

where s'_{ij} is the *hybrid*-similarity between the probe i and the individual study items j stored in memory. The hybrid similarity is computed as a joint function of the distance between items in a continuous-dimension similarity space and on the basis of matching versus mismatching distinctive features. Here, *continuous*-dimension similarity (s_{ij}) is computed as an exponential decay function of the Euclidean distance (D_{ij}) between items in an 8-dimensional MDS solution for the rock images derived by Nosofsky et al. (2018) and Meagher and Nosofsky (2023):

$$s_{ij} = e^{-cD_{ij}} \tag{3}$$

where c is a free sensitivity parameter that describes the rate at which similarity decreases with distance. The hybrid *self*-similarity of item i to itself (s'_{ii}) is then computed as

$$s'_{ii} = e^{\beta\delta_i} \cdot s_{ii} \tag{4}$$

where δ_i is the distinctive-feature rating for item i (Meagher & Nosofsky, 2023) and β is a freely estimated scaling parameter. The factor $e^{\beta\delta_i}$ in Equation 4 measures the *boost* in self-similarity that arises for test-probe i because of its distinctive-feature *match* to its representation in memory. By contrast, the *inter-item* hybrid-similarity of test item i to a mismatching study item j is given by

$$s'_{ij} = e^{-\alpha\delta_i} \cdot s_{ij} \tag{5}$$

The factor $e^{-\alpha\delta_i}$ in Equation 5 (where α is a free parameter) measures the *reduction* in inter-item similarity of

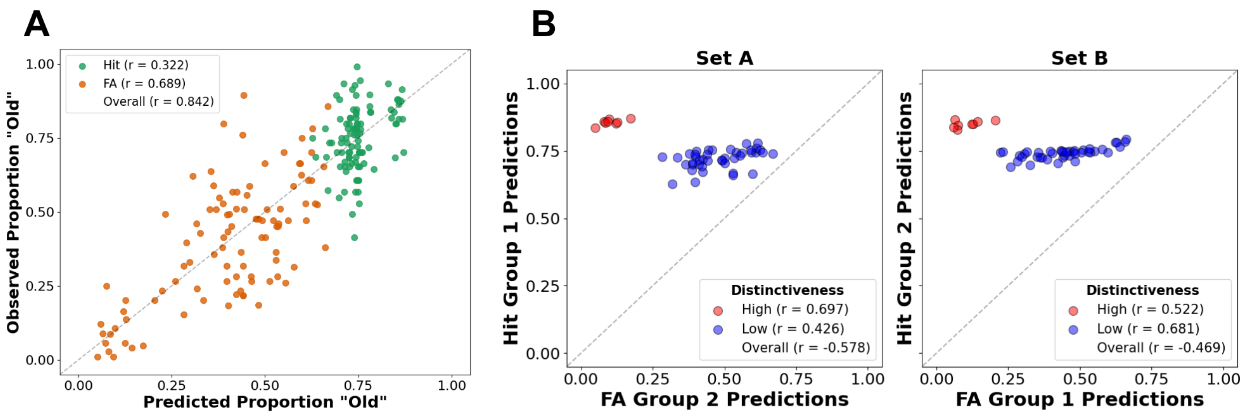

Figure 3: (A) Observed individual-item "old" response probabilities plotted against the hybrid-similarity model predictions. **(B)** Distinctiveness-dependent correlations between hits and false alarms. The model predicts positive H – FA correlations for low-distinctive items, but a cluster of high H and low FA rates for high-distinctive objects in the upper-left corner, resulting in a reduction of the positive correlation when computed across all items collectively.

item i to all other study items due to the *mismatching* distinctive feature (cf. Tversky, 1977).

Model Predictions

We fitted the model (Equations 1 – 5) to the individual-trial data using a maximum-likelihood estimation, but for simplicity assumed that the four free parameters (k, c, β and α) were fixed across subjects. (In future work we plan to conduct Bayesian hierarchical model fitting to account for individual-subject differences.) As shown in Figure 3A, this low-parameter model captured the item-level recognition probabilities reasonably well, although there is room for improvement in predicting the individual-item hit-rate data.

Most important, as shown in Figure 3B, the model also captures the qualitative pattern of results regarding the H-FA correlations. For the low-distinctive items (blue dots), it does indeed predict a positive correlation between the hit and false-alarm rates of the individual items ($r = .43$, $p < .001$ for Set A; and $r = .68$, $p < .001$ for Set B). But when the high-distinctive items and low-distinctive items are analyzed collectively, the positive correlation between H and FA rates is vastly reduced, turning negative overall ($r = -.58$ for Set A and $r = -.47$ for Set B). The model does indeed place the H and FA rates of the high-distinctive items toward the upper-left of the overall scatterplot. Importantly, a special case of the model in which the distinctive-feature scaling parameters (β and α) were held fixed at zero failed to account for the data. It incorrectly predicted *lower* hit rates for high-distinctive targets than for low-distinctive ones; and incorrectly predicted a strong positive correlation between H and FA rates when all items were analyzed collectively.

In future work, to achieve improved quantitative fits, we plan to test higher-parameter extended versions of the hybrid-similarity model that make allowance for differential attention-weighting of the MDS dimensions, nonlinear relations between psychological distinctiveness and rated distinctiveness, and so forth. The important take-home message here, however, is that with a minimum of parameter estimation, the model does a good job of capturing the main qualitative pattern of results regarding relations between H and FA rates of individual items and how these relations are modulated by the structure of the similarity space in which the items are embedded.

General Discussion

The goal of the present study was to examine how item distinctiveness shapes recognition memory, with particular emphasis on the joint behavior of hits and false alarms at the item level. Behavioral analyses and computational modeling converge on the conclusion that distinctiveness robustly reduces false alarms and increases hit rates, reflecting contributions from opposing memory forces. These findings support a theoretical account in which a global-familiarity signal arises from a summing of inter-item similarity between a test probe and study items, plus a self-similarity term in the case of old target probes. The presence of highly distinctive features boosts the self-similarity term in targets, but tends to reduce the inter-item similarities of both targets and lures to non-matching study items. When the self-match signal is sufficiently large, high-distinctive targets will tend to have higher hit rates than low-distinctive ones.

A consistent result was the strong negative relationship between distinctiveness and false alarms. Highly distinctive items, when presented as lures, were rarely endorsed as "Old." This pattern follows naturally from similarity-based accounts: distinctive items tend to occupy isolated regions of psychological space and therefore produce weak summed similarity to stored memory traces.

In contrast, the effect of distinctiveness on target-item hit rates was more modest. As explained above, this more modest effect reflects the opposing forces of distinctiveness on the global-familiarity signal: a boost in a target's self-similarity to its trace in memory but a reduction in its inter-item similarity to non-matching study items. Under the present conditions of testing, the self-similarity boost apparently exceeded the inter-item-similarity reduction.

Clearly, however, the outcome in any given experimental paradigm will depend on factors such as the detailed structure of the similarity space in which the items are embedded, the extent to which experimental conditions promote the coding of distinctive features, and individual-subject differences in memory capabilities. Regarding individual-subject differences, our focus in this report was on participants with upper-median performance. We should note that participants performing in the lower-median but with positive H-minus-FA rates (roughly 25% of the full sample) showed a different pattern of results: their hit rates on both high- and low-distinctive targets barely exceeded their FA rates on low-distinctive lures. Correct rejections were high for only the high-distinctive lures. This pattern of results can be captured by the model in cases in which the memory-sensitivity parameter (c in Equation 3) is set at a low value. In future work we plan to conduct Bayesian hierarchical modeling to capture the full gamut of individual-subject differences observed in our paradigm.

Returning to the upper-median participants, the opposing influences of distinctive features on self- and inter-item similarity become especially clear when considering the relationship between individual-item hits and false alarms. For low-distinctive items, we observed significant positive correlations between hit rates and false-alarm rates: items that were generally more confusable tended to elicit both more hits and more false alarms. High-distinctive items, however, occupied a qualitatively different region of the response space, characterized by high hits and low false alarms, and their inclusion substantially reduced the overall H-FA correlation when all items were analyzed collectively. Distinctiveness therefore acts as a modulator of the H-FA relationship, reorganizing the structure of H-FA correlations in accord with predictions from the hybrid-similarity exemplar model that served as a conceptual guide for the present work.

This interpretation sharpens recent findings that showed that distances in psychological space can be differently associated with hits and false alarms (Utochkin et al., 2025a), even when overall recognition performance remains stable across retrieval contexts (Utochkin et al., 2025b). Our results specify when and why these dissociations emerge. Specifically, a positive correlation between hits and false alarms arises primarily for low-distinctive objects, for which self-similarity is relatively uniform and variation in recognition is driven mainly by inter-item similarity. Under these conditions, items that are more globally confusable tend to elicit both higher hit rates and higher false-alarm rates. In contrast, high-distinctive items introduce an additional source of evidence – enhanced self-similarity – that selectively boosts hits while leaving inter-item similarity low. This added signal decouples hits from false alarms, breaking the positive relationship observed for low-distinctive items and, in some cases, reversing it.

Recent memorability research argues that memory performance reflects an intrinsic property of stimuli that is stable across observers and contexts (e.g., Bainbridge, 2019; Davis & Bainbridge, 2023; Isola et al., 2014). From this perspective, an image's tendency to elicit both hits and false alarms is treated as a fixed characteristic of the stimulus itself. Our split-half analyses replicate the first part of this claim: both hit rates and false-alarm rates were highly reliable across independent groups of participants. Critically, however, this reliability does not imply that these effects are intrinsic to the stimulus in isolation. Instead, our results show that the relationship between hits and false alarms depends systematically on an item's location within the surrounding similarity space (for related ideas and findings, see Nosofsky et al., 2025; Simpson et al., 2025). Distinctiveness is not a property of an item alone, but of its position relative to other items, and changes in the stimulus context can reorganize how stable item-level tendencies map onto recognition outcomes. In this sense, memorability is reliable – but relational rather than intrinsic.

The Hybrid-Similarity Exemplar model (Nosofsky & Zaki, 2003) provides a coherent computational account of these patterns, as would other modern global-familiarity models of memory that provide alternative mechanistic accounts of enhanced self-similarity for distinctive items (e.g., Cox, 2026; Shiffrin & Steyvers, 1997). By combining summed inter-item similarity with a distinctiveness-dependent self-match boost for targets and a similarity reduction for lures, the model captured not only the joint effects of distinctiveness on hits and false alarms, but also the reduction in their correlation when highly distinctive items are included in the analysis. Notably, the model reproduced the positive H – FA correlation for low-distinctive items and its attenuation or reversal in the full item set. Traditional global-familiarity models that place emphasis on only the role of inter-item similarity fail to account for such effects.

Acknowledgements

This work was supported in part by Australian Research Council Discovery Project Grant DP240101264.

References

Bainbridge, W. A. (2019). Memorability: How what we see influences what we remember. In *Psychology of Learning and Motivation* (Vol. 70, pp. 1-27). Academic Press.

Brady, T. F., & Störmer, V. S. (2024). Comparing memory capacity across stimuli requires maximally dissimilar foils: Using deep convolutional neural networks to understand visual working memory capacity for real-world objects. *Memory & Cognition*, *52*(3), 595–609. https://doi.org/10.3758/s13421-023-01485-5

Cox, G. E. (2026). Similarity as likelihood ratio: Coupling representations from machine learning (and other sources) with cognitive models. *Psychonomic Bulletin & Review*, *33*(1), 6.

Davis, T. M., & Bainbridge, W. A. (2023). Memory for artwork is predictable. *Proceedings of the National*

Academy of Sciences of the United States of America, *120*(28), e2302389120. https://doi.org/10.1073/pnas.2302389120

De Leeuw, J. R., Gilbert, R. A., & Luchterhandt, B. (2023). jsPsych: Enabling an open-source collaborative ecosystem of behavioral experiments. *Journal of Open Source Software*, *8*(85), 5351.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, *91*(1), 1.

Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, *95*(4), 528.

Isola, P., Jianxiong Xiao, Parikh, D., Torralba, A., & Oliva, A. (2014). What Makes a Photograph Memorable?. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *36*(7), 1469–1482. https://doi.org/10.1109/TPAMI.2013.200.

Meagher, B. J., & Nosofsky, R. M. (2023). Testing formal cognitive models of classification and old-new recognition in a real-world high-dimensional category domain. *Cognitive Psychology*, *145*, 101596.

Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance*, *17*(1), 3.

Nosofsky, R. M., Ehinger, K. A., & Osth, A. F. (2025). Tests of a hybrid-similarity exemplar model of context-dependent memorability in a high-dimensional real-world category domain. *Journal of Experimental Psychology: General*, *154*(12), 3450.

Nosofsky, R. M., Sanders, C. A., Meagher, B. J., & Douglas, B. J. (2018). Toward the development of a feature-space representation for a complex natural category domain. *Behavior Research Methods*, *50*(2), 530–556. https://doi.org/10.3758/s13428-017-0884-8

Nosofsky, R. M., & Zaki, S. R. (2003). A hybrid-similarity exemplar model for predicting distinctiveness effects in perceptual old-new recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(6), 1194.

Schurgin, M. W., Wixted, J. T., & Brady, T. F. (2020). Psychophysical scaling reveals a unified theory of visual memory strength. *Nature Human Behaviour*, *4*(11), 1156–1172. https://doi.org/10.1038/s41562-020-00938-0

Shen, K. J., Colloff, M. F., Vul, E., Wilson, B. M., & Wixted, J. T. (2023). Modeling face similarity in police lineups. *Psychological Review*, *130*(2), 432–461. https://doi.org/10.1037/rev0000408

Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*, *4*(2), 145-166.

Simpson, D., Johnson, B., & Brady, T. (2025). Are images truly intrinsically memorable? Memorability results are explained by between-image similarity structure, not features of individual images. *Journal of Vision*, *25*(9), 2833-2833.

Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*(4), 327.

Utochkin, I., Azarov, D., & Grigorev, D. (2025b). Invariant Recognition Memory Spaces for Real-World Objects Revealed With Signal-Detection Analysis. *Psychological Science*, *36*(11), 831-845.

Utochkin, I., Chiang, N., & Bainbridge, W. (2025a). Recognition-memory asymmetries predicted by individual item memorability. Preprint at *PsyArXiv* [*https://doi.org/10.31234/osf.io/6n74d_v2*](https://doi.org/10.31234/osf.io/6n74d_v2)

Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *The Quarterly Journal of Experimental Psychology Section A*, *43*(2), 161-204.