Nothing Intrinsically Memorable: The Relational Nature of Animacy Effects in Visual Memory

Dyllan Simpson (dysimpson@ucsd.edu)
Department of Psychology, University of California, San Diego
La Jolla, CA 92093 USA

Benjamin Johnson (bej001@ucsd.edu)
Department of Computer Science and Engineering, University of California, San Diego
La Jolla, CA 92093 USA

Timothy Brady (tfbrady@ucsd.edu)
Department of Psychology, University of California, San Diego
La Jolla, CA 92093 USA

Abstract

Are some images intrinsically more memorable than others? Previous research has demonstrated that animate objects are remembered better, on average, than inanimate objects, leading to claims that animate features are intrinsically more memorable. We challenge this view by showing that the same animate objects are either highly memorable or highly forgettable depending solely on their representational distinctiveness to other objects in the study set. By manipulating the proportion of animate objects (10% vs. 90%, N=1600 objects) in our dataset, we created a reversal in how distinctive animate items were, on average, from all items according to neural data from IT cortex in macaques, and in a recognition memory experiment (N=93), this led to a complete reversal of traditional memory advantages for animate items. Overall, these findings demonstrate that animate stimuli do not have intrinsically higher memorability — instead, their higher memorability emerges from their relation to other items in memory.

Keywords: memorability; long-term memory; electrophysiology

Introduction

What makes an image more or less likely to be remembered? The notion that certain images possess intrinsic memorability has gained significant traction in memory research (Rust & Mehrpour, 2020). Animate objects, in particular, have been consistently shown to be better remembered than inanimate objects (Nairne et al., 2013; Bonin et al., 2014; Kramer et al., 2023). The prevailing theoretical explanation for this empirical memorability effect posits that animate objects possess inherent properties that make them intrinsically more memorable, regardless of context or distinctiveness. In the case of animacy, this is often attributed to evolutionary adaptations that prioritize the processing of potentially relevant social stimuli (Nairne et al., 2013). This view aligns with broader claims in the memorability literature that certain stimuli possess stable, observer-independent memorability signatures that can be predicted from visual features of individual items alone (Bainbridge et al., 2013; Isola et al., 2014; Khosla et al., 2015).

Yet the idea of intrinsic memorability conflicts with long-established principles of recognition memory research. A century of work demonstrates that recognition decisions depend on purely relational features: how well the test item matches the representations of all of the items in memory, and how well we would expect an old vs. new item to match in this context (Shiffrin & Steyvers, 1997). Consistent with this view, memory is reliably heavily dependent on context, including study list composition (Hirshman, 1995), target-foil similarity (Benjamin & Bawa, 2004), and global matching processes that integrate information across all studied items (Shiffrin & Steyvers, 1997). In this framework, where decisions are made solely based on the relationships between items, arguing animate items are inherently memorable is like arguing red is inherently salient: it seems to mistake a relational feature for an intrinsically item-based feature. Instead we argue that, on average, red may be unlike its neighbors in many visual scenes; and on average, animate items may tend to stand out more from other items in their study lists, but red is not 'intrinsically' salient and animate items are not 'intrinsically' memorable.

Previous work has argued that context cannot account for intrinsic memorability because items are tested among a random set of items sampled from a larger database. However, on average, distinctiveness in a study list is dependent on distinctiveness in the larger database. Thus, here, we directly test these competing accounts by manipulating the prevalence of animate objects in the study set, which affects these items distinctiveness relative to other items that are seen.

If animate objects are intrinsically more memorable, then they should maintain their memory advantage regardless of their prevalence in the larger database. In contrast, if memorability emerges from relative distinctiveness, as predicted by all major memory models, then animate objects should lose or even reverse their advantage when they become the majority category in the dataset, even though the study lists are just small random samples from the dataset. To validate that our prevalence manipulation affects the relative distinctiveness of animate objects within the dataset, we use representational similarity analysis (RSA) of neural data from macaque IT as a proxy for representational distinctiveness.

Recognition Memory Experiment

We created two datasets by randomly sampling 1600 images from the THINGS dataset (Hebart et al., 2019) with varying proportions of animate objects: 10% or 90%. 93 undergraduates (aged 18-25) were then randomly assigned to one of the two dataset conditions (18 in the animate-biased dataset and 75 in the inanimate-biased dataset).

The experiment consisted of 16 blocks, where each block had both a study phase followed by a test phase. During study, participants viewed 40 images randomly selected from their assigned dataset, presented sequentially for 500ms each. ROC analysis is necessary to accurately measure recognition memory (Brady et al., 2023). Thus, during test, participants saw 40 images (20 old, 20 new) and rated their recognition confidence on a 6-point scale (1="definitely saw" to 6="definitely did not see"). Images were never repeated across blocks.

Neural dataset To assess the distinctiveness of objects in both stimulus sets, we leveraged the THINGS ventral stream spiking dataset (TVSD) published by (Papale et al., 2025). This neurophysiological dataset provides recordings of multi-unit spiking activity in response to over 25,000 images from the THINGS database across three key regions of the macaque ventral visual stream: primary visual cortex (V1), area V4, and inferior temporal cortex (IT). From this data set, we used recordings in the IT cortex (N = 2 macaques) to all the images used in our experiments as a proxy for human representational similarity. We averaged across the two monkeys, computing representational similarity matrices (RSMs)

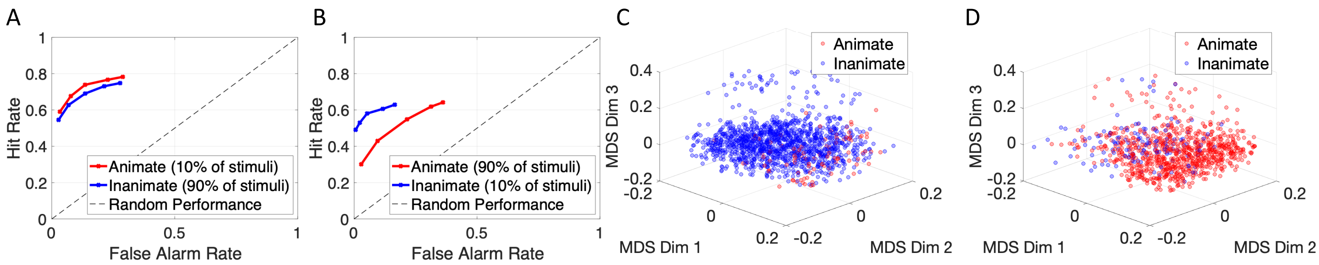

Figure 1: Panel A and B demonstrate the complete reversal of the animacy advantage when manipulating category prevalence. Panel A shows ROC curves when animate objects comprise the minority (10%) of the dataset, with animate objects (red) showing significantly better memory performance than inanimate objects (blue). Panel B demonstrates the reversal when animate objects comprise the majority (90%), with inanimate objects now showing superior memory performance. Panels C and D visualize the geometry of the representational space after multidimensional scaling (MDS). Neural activity patterns from IT cortex are projected into three dimensions, with animate objects shown in red and inanimate objects in blue. Panel C illustrates how animate objects primarily cluster to the right in the 10% condition, leading to them being distinct from the average item, and mirroring their enhanced memorability. Panel D shows how this pattern reverses in the 90% animate condition as now the animate items form the main cluster with inanimate objects primarily cluster to the left. Thus, the same objects can vary in their representational distinctiveness depending solely on the distributional properties of the dataset, and this distinctiveness pattern mirrors the memorability pattern observed in human participants.

with an exponential transform of distance in neural space, $S_{ij} = e^{-0.3d_{ij}}$.

Results To assess memory performance across our experimental conditions, we fit an unequal variance signal detection model and estimated d_a values for both animate and inanimate objects.

In the inanimate-biased dataset (10% animate, 90% inanimate), animate objects demonstrated better performance ($d_a = 1.45$) compared to inanimate objects ($d_a = 1.25$). This pattern is consistent with previous literature suggesting an intrinsic memorability advantage for animate objects. However, when we manipulated the prevalence of categories in the animate-biased dataset (90% animate, 10% inanimate), this pattern completely reversed: inanimate objects($d_a = 0.80$) were more memorable compared to animate objects ($d_a = 0.71$). The ROC curves in Figure 1A and 1B illustrate this reversal, showing how recognition performance systematically varies with category prevalence rather than category type. These behavioral findings suggest that memorability is not an intrinsic property of animate stimuli but emerges from their distinctiveness within the encoding context.

Next, we analyzed IT cortex activity patterns to understand the representational basis of our behavioral findings. A permutation test confirmed that animate objects in the inanimate-biased dataset (10% animate) were significantly more dissimilar from the dataset average than would be expected by chance ($p < .01$). This dissimilarity was observed both when analyzing the entire THINGS database and when restricting analysis to just our inanimate-biased subset. Critically, animate objects in the animate-biased dataset (90% animate) were significantly more similar to the dataset average than in the inanimate-biased dataset ($p < .01$), resulting in no such distinctiveness advantage, and mirroring our behavioral results. Thus, animacy advantages in the standard dataset and or 10% animate condition may arise solely from animate items being more distinct from other items in the dataset. Flipping prevalence also flips both this distinctiveness and memory performance.

Discussion

In line with standard recognition memory models, and in contrast to claims that sampling from a large dataset allows the assessment of ”intrinsic” memorability, our results demonstrate that the memorability of animate objects depends on their prevalence in the dataset and thus their representational distinctiveness. This complete reversal of the animacy advantage when manipulating prevalence (10% vs. 90%) directly challenges claims of intrinsic memorability: What has previously been interpreted as evidence for inherent memorability of animate objects is better explained by their typical distinctiveness in standard experimental paradigms, where they are underrepresented relative to macaques, and likely people's, representational space. There is nothing special about animacy per se; rather, memorability emerges from the relationship between a stimulus and its encoding context.

References

Bainbridge, W. A., Isola, P., & Oliva, A. (2013). The intrinsic memorability of face photographs. *Journal of Experimental Psychology: General*, *142*(4), 1323–1334. doi: 10.1037/a0033872

Benjamin, A. S., & Bawa, S. (2004). Distractor plausibility and criterion placement in recognition. *Journal of Memory and Language*, *51*(2), 159–172. doi: 10.1016/j.jml.2004.04.001

Bonin, P., Gelin, M., & Bugaiska, A. (2014). Animates are better remembered than inanimates: Further evidence from word and picture stimuli. *Memory & Cognition*, *42*(3), 370–382. doi: 10.3758/s13421-013-0368-8

Brady, T. F., Schurgin, M. W., & Wixted, J. T. (2023). Measuring recognition memory: A tutorial on the performance of roc analysis. *Journal of Experimental Psychology: General*, *152*(1), 93–114. doi: 10.1037/xge0001191

Hirshman, E. (1995). Decision processes in recognition memory: Criterion shifts and the list-strength paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(2), 302–313. doi: 10.1037/0278-7393.21.2.302

Isola, P., Xiao, J., Parikh, D., Torralba, A., & Oliva, A. (2014). What makes a photograph memorable? In *Ieee transactions on pattern analysis and machine intelligence* (Vol. 36, pp. 1469–1482). doi: 10.1109/TPAMI.2013.200

Khosla, A., Raju, A. S., Torralba, A., & Oliva, A. (2015). Understanding and predicting image memorability at a large scale. In *Proceedings of the ieee international conference on computer vision* (pp. 2390–2398). doi: 10.1109/ICCV.2015.275

Kramer, M. A., Hebart, M. N., Baker, C. I., & Bainbridge, W. A. (2023). The features underlying the memorability of objects. *Science Advances*, *9*(17), eadd2981. doi: 10.1126/sciadv.add2981

Nairne, J. S., VanArsdall, J. E., Pandeirada, J. N. S., Blunt, J. R., & Cogdill, M. (2013). Adaptive memory: The mnemonic value of animacy. *Psychological Science*, *24*(10), 2099–2105. doi: 10.1177/0956797613480803

Papale, P., Wang, F., Self, M. W., & Roelfsema, P. R. (2025). An extensive dataset of spiking activity to reveal the syntax of the ventral stream. *Neuron*, *113*(4), 539-553.e5. doi: 10.1016/j.neuron.2024.12.003

Rust, N. C., & Mehrpour, V. (2020). Understanding image memorability. *Trends in cognitive sciences*, *24*(7), 557–568.

Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*, *4*(2), 145–166. doi: 10.3758/BF03209391