One extra viewing leaves no detectable tuning-proportional trace in human medial temporal lobe neurons

Ryuzo Kijima

November 18, 2026

Preregistration · Code

Abstract

Responses of human medial temporal lobe neurons decline when an image is repeated within a session, but it is not known whether a single extra viewing leaves a trace in a neuron's later response, or whether such a trace depends on how strongly the neuron responds to the image. We tested a preregistered model in which one extra viewing changes the log response at a later test by $k \cdot x$, where x is the neuron's relative response to the image's category, measured on images that never appeared at test. In a public dataset of 950 hippocampal and amygdala neurons from 39 participants, we compared images seen once with images seen twice, the second time as a working-memory probe, about 35 minutes before a recognition test. The analysis code was fixed, and its false-positive rate and power were measured on synthetic data, before any test-phase spike was counted. In 130 category-tuned neurons, k was +0.055 (bootstrap SD 0.101, $z = 0.55$), and no robustness check met the decision rule. On synthetic data, the analysis detected $k = -0.3$, a change of 7–11% at the quartiles of x, in 28 of 30 datasets, and $k = -0.2$ in 17 of 30. In an exploratory analysis, the same kind of model found only a weak, inconclusive change at the second viewing itself ($k = -0.09$, $z = -1.28$), so these data cannot tell whether a tuning-proportional effect is absent or too small for this measure of tuning.

Introduction

Seeing something once can change how neurons respond to it later. In the human medial temporal lobe (MTL), responses decline when an image is repeated within a session[1], and some neurons respond differently to novel and familiar images after a single exposure, about 30 minutes later[2]. Activity during working-memory maintenance[3] and during encoding[4] predicts whether an image will be remembered. In macaque inferotemporal cortex, the form of the change, that is, how it depends on the neuron's initial response, has been inferred from responses to novel and familiar images[5,6].

We asked a narrow question about that form in humans: does one extra viewing of a novel image change a neuron's response half an hour later in proportion to how strongly the neuron responds to the image's category? We predicted a larger decline where the response is larger. Because we had already examined the encoding-phase data, we preregistered the test before counting any test-phase spike.

Methods

Data and task

We used DANDI 001187[7], recorded for a study of working and long-term memory[3], which describes the recordings, spike sorting and consent procedures: 46 sessions from 39 patients, with 950 single neurons (570 in the amygdala, 380 in the hippocampus). In the encoding phase, each of 140 working-memory trials showed 1 or 3 novel images from five categories, followed by a probe image; the participant answered whether the probe had been shown in the trial. A probe was either an image from the same trial (a match) or an image from an earlier trial (a familiar non-match), so probed images were seen twice. Of the 140 probed images, 60 were tested; the rest are among the never-tested images that define x, which uses only their first viewing. In the test phase, a median 11 minutes later (range 2–35), up to 400 images (200 new, 140 seen once, 60 seen twice) were shown, and the participant answered old or new with a confidence of 1–6 (4–6 = old; Fig. 1). An image was tested a median 35 minutes after its last viewing (5th–95th percentile, 15–63). Which images were probed was fixed by the task and was the same for every participant with the same image set (three sets).

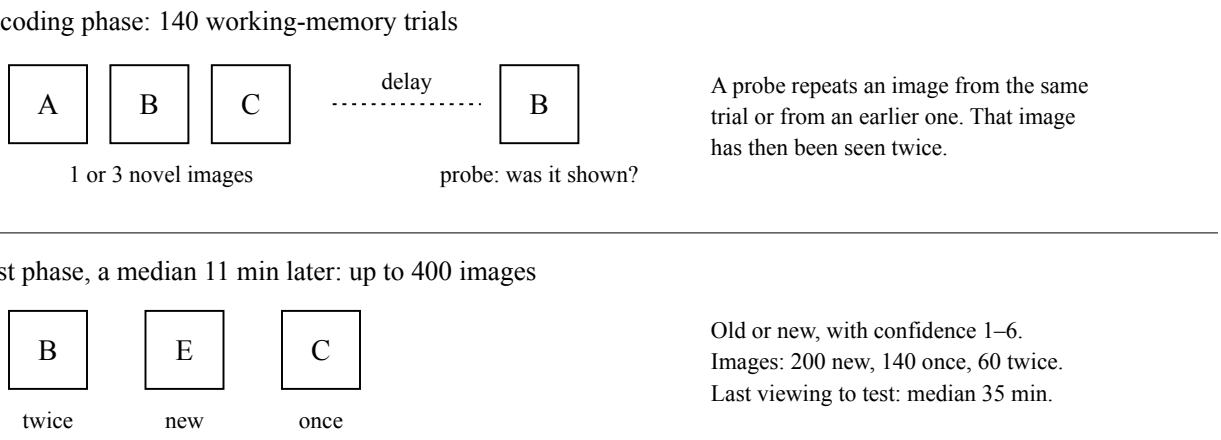


Fig. 1 | Task. An image shown again as a probe has been seen twice by the test, and other tested encoding images have been seen once. Encoding images that never appear at test (not drawn) define x.

Neurons

Neurons were excluded if, between the two phases, the correlation of their mean waveforms was below 0.98, the ratio of waveform amplitudes was outside 0.86–1.12, or the ratio of firing rates was outside 0.33–3.3, or if either phase had more than 1% of inter-spike intervals under 3 ms or fewer than 40 spikes. These rules used whole-phase spikes, not image or answer labels; 243 of 950 neurons were excluded. A neuron was category-tuned if it fired at least 20 spikes in the 0.2–1.2 s windows after image onset, had at least three never-tested encoding images per category, and its responses to those images differed across categories (one-way ANOVA, $p < 0.05$). In sessions with tuned neurons, the least-represented category had a median of 14 never-tested images (range 14–36).

Model

For each tuned neuron and category, x was the neuron's mean response (0.2–1.2 s) to that category's encoding images that never appeared at test, divided by its mean response to all such images, minus 1. In the analysed neurons, x had a median of −0.10 and an interquartile range of −0.34 to 0.24 (−0.08 and −0.32 to 0.21 before exclusion, as registered). At $k = -0.3$, one extra viewing lowers the response by 7% at $x = 0.24$ and raises it by 11% at $x = -0.34$.

For each tuned neuron and each image seen once or twice, spikes 0.2–0.7 s after onset were counted at the first viewing and at test. Test trials answered faster than 0.7 s were excluded, so that the window never contained the answer. With $T = 1$ for test rows and $D = 1$ for twice-seen images, the counts entered a Poisson regression

$$\log \mathrm{E}[\mathrm{count}] = \alpha_{\mathrm{neuron} \times \mathrm{category} \times \mathrm{phase}} + (\beta_1 + \beta_2 x + \gamma_{\mathrm{category}})\, D + (\beta_3 + k\, x + \delta_{\mathrm{category}})\, TD + \mathrm{covariates},$$

with a separate intercept α for each neuron × category × phase. The D terms describe how twice-seen images already differed at the first viewing and are subtracted from the test; k is the tested coefficient. First-viewing rows had three covariates, each also × x: time in the phase, one-image trial, and first image of a trial, because twice-seen images had been shown earlier in the phase and more often in one-image trials. Test rows had position in the test (cubic; its linear term also × x), whether the previous test image was old and whether it was called old, image typicality, log time since the last viewing and log reaction time, the last three also × x. Typicality was the similarity of the image's CLIP ViT-B/32 embedding[8] to the other test images of its category and image set, z-scored within category. Because twice-seen images were remembered better (Results), each neuron's response to the answer itself was subtracted from every test row as a log offset. It was estimated on new images only: the neuron's log ratio of responses to new images called old versus new, shrunk toward the population mean by a split-half estimate of its noise, plus a pooled effect of confidence within each answer.

Inference

The standard deviation of k was estimated from 100 bootstrap draws that resampled participants and images independently, and $z = k/\mathrm{SD}$. The decision rule was $|z| > 1.6$, with the same sign when the model was fitted on the difference scale (least squares on counts divided by their mean within neuron × category × phase) and when said-old × x and confidence 5–6 × x (the memory terms) were added.

We calibrated this rule on synthetic datasets built on the real task, answers and noise correlations, without neuron exclusion and with at most each session's number of tuned neurons (131–160 in total, against 130 in the real analysis), so the power below is, if anything, slightly overstated. With $k = 0$, the rule was met in 1 of 310 datasets (0.3%; 95% CI 0–1.8%), including datasets in which responses to the preferred category drifted across the encoding phase, and datasets in which each participant's own meaning of each image raised all of that participant's responses to it. With $k < 0$, it was met in 3 of 15 datasets at $k = -0.1$, 17 of 30 at −0.2, 28 of 30 at −0.3 and 25 of 25 at −0.5, and in 9 of 15 when one extra viewing lowered responses by 15% only in categories above the neuron's mean. Estimates were about 70% of the true k, because x is measured with noise. With $k = 0$, z had a standard deviation of 0.60, as expected when

participants and images are resampled independently[9]; the bootstrap SD therefore overstates the spread of k, and the threshold of 1.6 corresponds to about 2.7 standard deviations of the null.

Preregistration and blinding

Before the test, we had examined all encoding-phase spikes, the test-phase behaviour and the images, but not test-phase spikes. One examined quantity bears on the test: at the first viewing, the coefficient on x differed between twice- and once-seen images by +0.036 (SD 0.038 over 40 resamples of participants), or +0.027 (SD 0.045) with the first-viewing covariates; the model subtracts this difference. The analysis was registered at osf.io/z56gw with the code at commit 7d1f9b0 of github.com/ryuzo-k/second-viewing, which reads test-phase spikes only once the registration exists. There were no deviations from the registered analysis; analyses marked exploratory were added after the results.

The registration also specified an exploratory test of neuron pairs, motivated by Hebbian plasticity: if neurons A and B, active together at the first viewing, change their connection, B's test response should depend on the product of A's and B's first-viewing responses. For each neuron B, the coefficient of that product was computed for every same-participant neuron A and compared with equally sized sets of neurons from other participants who saw the same images; the statistic is the mean rank of B's strongest real partner among those sets, minus 0.5, and its z is the statistic divided by its standard deviation over 1,000 draws in which another participant's set plays the real one. It was to be tested in DANDI 000004[10,11] only if $|z| > 2.5$ here and two controls (spikes before image onset; images swapped within category) gave $|z| \leq 2.5$. Before registration, a pilot of this analysis on encoding-phase probes gave $z = -0.25$ and +0.15, and five variants then tried gave $|z| < 2$. The SHA-256 of the 001187 results was posted to the registration before any 000004 test-phase spike was counted.

Results

After exclusion, 130 neurons were category-tuned (91 in the amygdala, 39 in the hippocampus), from 36 sessions of 30 participants (median 2.5 per participant, range 1–15). The model used 16,517 once-seen and 6,997 twice-seen test rows (neuron × image). Trials answered before 0.7 s were 187 of 6,166 once-seen, 89 of 2,626 twice-seen and 283 of 8,809 new-image trials. Recognition was better for twice-seen images. Averaged over sessions, with 0.5 added to each count, hit rates were 0.64 for once-seen and 0.72 for twice-seen images and the false-alarm rate was 0.26 (d′ 1.10 and 1.38); d′ was higher for twice-seen images in 38 of 46 sessions.

The estimate of k was +0.055 (bootstrap SD 0.101, $z = 0.55$), opposite in sign to the prediction. It was +0.079 on the difference scale ($z = 0.52$) and +0.050 with the memory terms ($z = 0.58$). The decision rule was not met. The difference already present at the first viewing, which the model subtracts, was +0.022 in the analysed neurons with all covariates (+0.027 before exclusion; Methods).

No robustness check met the decision rule (Fig. 2, Table 1). Leaving out each participant in turn gave estimates of +0.017 to +0.077. Shuffling the twice/once labels 200 times within image set and category, which keeps participants fixed and leaves time since the last viewing at its true labels, gave a null

standard deviation of 0.053 and a two-sided p of 0.30. One check exceeded 1.6 on the difference scale alone, in the positive direction (0.2–0.5 s window, $z = 1.67$).

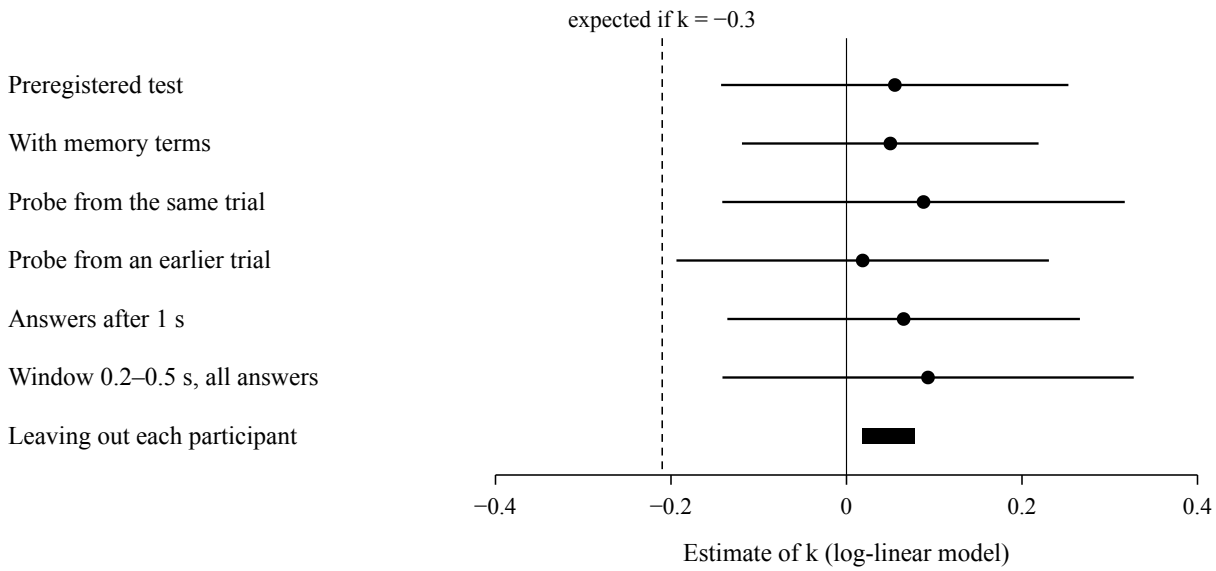


Fig. 2 | Estimates of k, with conservative intervals of ±1.96 bootstrap SD; the bar for leaving out each participant is the range of estimates. The dashed line is the estimate expected at $k = -0.3$ (−0.21), from the attenuation of the full analysis in synthetic data.

Table 1 | Estimates of k with bootstrap SD and z, on the log scale and the difference scale. The 0.2–0.5 s window was applied at the first viewing and at test. – : not computed.

Analysis	k	SD	z	k, difference	z
Preregistered test	+0.055	0.101	0.55	+0.079	0.52
With memory terms	+0.050	0.086	0.58	–	–
Probe from the same trial	+0.088	0.117	0.75	−0.031	−0.17
Probe from an earlier trial	+0.018	0.108	0.17	+0.170	0.88
Answers after 1 s	+0.065	0.102	0.64	+0.158	0.99
Window 0.2–0.5 s, all answers	+0.093	0.120	0.78	+0.244	1.67

The remaining analyses are exploratory. The uniform effect of the extra viewing at test (β_3, at $x = 0$) was −0.021 on the log scale (bootstrap SD 0.085). We also applied the same kind of model to the second viewing itself, comparing the probes of tested images with the first viewings of tested images (one intercept per neuron × category, not matched by image). The never-tested images that define x were left out of both, since sharing images with x biases the estimate; two earlier versions of this analysis that shared them gave −0.108 and −0.031. Probes answered before 0.7 s were excluded, and a term for match probes was added. At the second viewing, the decline in proportion to x was weak and indistinguishable from zero ($k = -0.088$, SD 0.069, $z = -1.28$), and there was no uniform change (−0.013, SD 0.048).

The pair analysis was negative ($z = -0.83$; controls −0.12 and +1.56), so DANDI 000004 was analysed only descriptively ($z = -1.20$). Of three registered secondary pair analyses, one had $|z| > 2.5$ (amygdala neurons as B, $z = -2.64$; hold period, +0.26; hippocampus, +0.52). On synthetic data without any effect, the primary pair statistic had $|z| > 2.5$ in 7 of 280 datasets; the subsets were not calibrated, and if each behaves alike, three looks give a false-positive rate of up to about 7.5%. As the primary pair test failed, we do not interpret it.

Discussion

We found no evidence that one extra viewing of a novel image changes the response of human MTL neurons, about 35 minutes later, in proportion to their category tuning. Divided by the attenuation found in synthetic data (0.66–0.78 across the simulated values of k), a conservative interval for the true k (±1.96 bootstrap SD) is about −0.2 to +0.4.

The exploratory analysis of the second viewing limits what this null can mean. Within-session repetition suppression[1] predicts a decline at the probe, seconds to minutes after the first viewing, yet we found no uniform change there and only a weak, inconclusive tuning-proportional one. The probe was a task event that the participant judged, which may dominate its response, and it was a single repetition, whereas suppression has been measured over several. Either way, the data offer no positive control: they cannot tell whether a tuning-proportional trace is absent or below what this measure of tuning can detect. Single-exposure familiarity signals[2] are not contradicted, since they distinguish seen from unseen images and both groups here had been seen.

Noise in x shrinks estimates toward zero and lowers power; the calibration includes this shrinkage. Which images were probed was fixed within each of three image sets, so the twice-seen images are a small fixed set; the image bootstrap accounts for this. A better measure of tuning, or more neurons, would test smaller effects.

Data and code

Data: DANDI 001187, version 0.240904.2347, and DANDI 000004, version 0.220126.1852. Code, synthetic-data calibration, results and the registered text: github.com/ryuzo-k/second-viewing; exploratory analyses are in its `extra` folder. Preregistration: osf.io/z56gw, registered on November 17, 2026.

References

1. Pedreira, C. et al. Responses of human medial temporal lobe neurons are modulated by stimulus repetition. *J. Neurophysiol.* 103, 97–107 (2010).
2. Rutishauser, U., Mamelak, A. N. & Schuman, E. M. Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. *Neuron* 49, 805–813 (2006).
3. Daume, J. et al. Persistent activity during working memory maintenance predicts long-term memory formation in the human hippocampus. *Neuron* 112, 3957–3968.e3 (2024).
4. Xiao, Z. et al. Stimulus-specific recruitment of human amygdala neurons predicts episodic memory encoding success. *Nat. Neurosci.* https://doi.org/10.1038/s41593-026-02493-9 (2026).

5. Lim, S. et al. Inferring learning rules from distributions of firing rates in cortical neurons. *Nat. Neurosci.* 18, 1804–1810 (2015).

6. Mohan, K. et al. Visual learning at fast and slow timescales is driven by distinct plasticity rules in primate inferotemporal cortex. *Nat. Commun.* 17, 8305 (2026).

7. Daume, J., Cheng, S. & Rutishauser, U. Data for: Persistent activity during working memory maintenance predicts long-term memory formation in the human hippocampus. DANDI archive, version 0.240904.2347. https://doi.org/10.48324/dandi.001187/0.240904.2347 (2024).

8. Radford, A. et al. Learning transferable visual models from natural language supervision. *Proc. Mach. Learn. Res.* 139, 8748–8763 (2021).

9. Owen, A. B. The pigeonhole bootstrap. *Ann. Appl. Stat.* 1, 386–411 (2007).

10. Chandravadia, N. et al. A NWB-based dataset and processing pipeline of human single-neuron activity during a declarative memory task. *Sci. Data* 7, 78 (2020).

11. Chandravadia, N. et al. A NWB-based dataset and processing pipeline of human single-neuron activity during a declarative memory task. DANDI archive, version 0.220126.1852. https://doi.org/10.48324/dandi.000004/0.220126.1852 (2022).