

Dynamics of brain activity reveal a unitary recognition memory signal

Christoph T. Weidemann^{1,2*}, Michael J. Kahana²

1 Department of Psychology, Swansea University, Swansea, Wales, UK

2 Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

* ctw@cogsci.info

Abstract

A repeated encounter with a person or object frequently elicits feelings of familiarity and recollections of previous interactions. Dual-process theories of recognition memory assume that these independently contribute to recognition decisions. We ask whether the memory system integrates available evidence into a unitary evidence signal, or if, instead, different types of mnemonic evidence separately drive distinct routes to recognition. To address this question, we quantified neural evidence for recognition decisions as a function of time using multivariate classifiers trained on spectral EEG features. Assuming distinct temporal profiles for different types of mnemonic evidence, a classifier trained on neural features from multiple time bins during the recognition period should be able to capitalize on independent signals from these evidence sources, if these are indeed kept separate. Instead, we found practically identical classifier performance for classifiers trained on a small portion of the recognition period compared with those also trained on neural features from previous time bins. Recollection has been frequently linked to recall suggesting that a comparison between previously recalled and not previously recalled targets should highlight differences between familiarity- and recollection-based recognition decisions assuming such a distinction is meaningful. Separating classifier performance for targets that were previously recalled and those that were not revealed qualitatively similar increases of classifier performance with time, consistent with different strengths (rather than different kinds) of evidence driving the recognition decisions. These results, along with a strong correspondence between classifier outputs and task performance, firmly link recognition decisions to other types of decisions under uncertainty, which are commonly assumed to rely on a unitary evidence signal differentiating between the response options.

Author summary

Dual-process models of recognition memory propose that two independent sources of evidence (familiarity and recollection) drive recognition decisions. Usually these are assumed to give rise to different kinds of recognition decisions, with individual “old” responses reflecting either a familiarity or a recollection signal. Alternative accounts posit that all available evidence is integrated into a unitary “memory strength” signal. Using multivariate (“machine learning”) classifiers, we quantified the neural evidence for recognition decisions across various partitions of the recognition period to distinguish between these single- and dual-route accounts of recognition memory. The performance of classifiers that are trained on neural features from a small portion of the recognition

period should only match the performance of similar classifiers that also incorporate information from previous time periods to the extent that neural evidence indexes a unitary evidence signal integrating over all available evidence. Compatible with this prediction, we found performance for these two types of classifiers to be practically identical, supporting the view that each recognition decision makes use of all available evidence.

Introduction

A repeated exposure to people or objects sometimes evokes only a vague sense of familiarity; at others, it elicits vivid recollections of contextual details from previous encounters. This distinction is formalized in dual-process models of recognition memory that posit two independent types of evidence subserving recognition decisions (with recollection commonly, but not always, conceptualized as a threshold process) [1–4]. In apparent support of these models, neuroscientific studies of recognition memory have identified patterns of brain activity with distinct time courses thought to reflect an early familiarity signal (peaking around 400 ms after onset of a memory probe) and a later recollection signal (peaking around 600 ms after probe onset) [5,6].

Most dual-process models assume that familiarity and recollection signals can each separately lead to recognition [7]. In some models, however, the memory system integrates evidence from different sources into a unitary evidence signal [8,9]. This results in a single route to recognition despite the contributions from different types of evidence. From this perspective, such models are conceptually similar to single-process models which assume only a single evidence source [4]. One indication that two separate routes to recognition may not be necessary to account for recognition performance is the fact that single-process models have been highly successful at accounting for intricate relationships between response time distributions, accuracy, and confidence ratings across a wide range of experimental manipulations (e.g., [10–19]).

Because the single- vs. dual-process labels do not reliably differentiate between the number of routes to recognition, we will refer to models as single- or dual-*route* models to make this distinction explicit. Specifically, we label models that assume that different types of evidence can give rise to different kinds of recognition decisions (e.g., [1, 7, 20, 21]) as dual-route models. Single-route models are those that assume a single type of evidence source and those that assume that evidence from multiple sources/processes is integrated into a unitary evidence signal (e.g., [8,9]). Within the framework of dual-route models, it makes sense to label individual recognition decisions with respect to the type of evidence (e.g., “familiarity” vs. “recollection”) that gave rise to them, whereas such a categorization of individual recognition memory decisions is not meaningful within the framework of single-route models, because information from all available sources contributes to recognition decisions. We propose that conflating the question about the number of recognition signals (i.e., the distinction between single- vs. dual-process models) with the question about the number of different routes to recognition may have contributed to the apparent disconnect between the evidence for separate familiarity and recollection signals and the success of single-process models.

Capitalizing on the presumed temporal separation of familiarity and recollection signals [1], we quantify the neural evidence distinguishing targets from lures in various partitions of the period leading up to the recognition decision. Specifically, we ask if combining neural evidence from multiple time bins during the recognition decision tells us more about whether an item has been studied than just the latest considered time bin by itself. If we are picking up on independent signals at different points in the recognition decision, then combining information from both should boost our ability to use neural activity to distinguish between old and new items. If, however, the neural

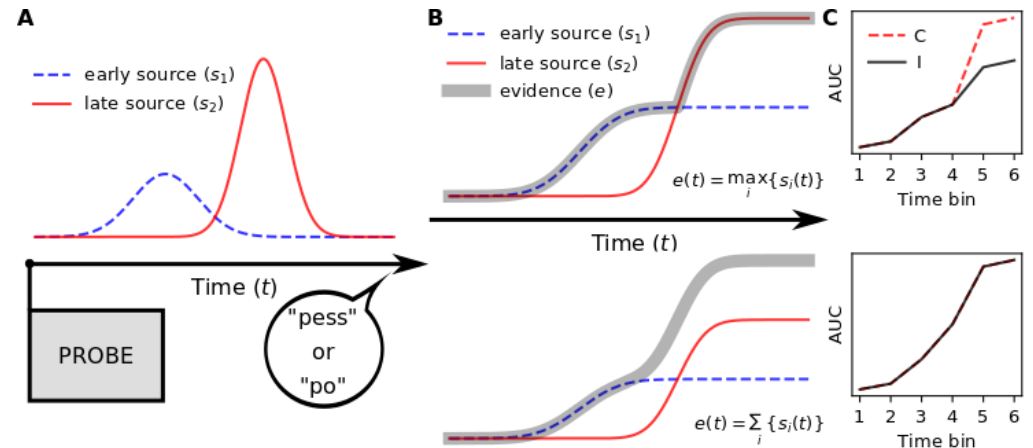


Fig 1. Toy models of evidence in recognition memory. (A) Probability density functions (PDFs) illustrating two sources of evidence for recognition memory decisions. (B) Cumulative density functions (CDFs) for the PDFs shown in (A) along with CDFs for the evidence on which the recognition memory decision is based. The top panel illustrates a case where the evidence is determined by a single source of evidence leading to different routes to recognition memory depending on which source determines the evidence signal at the time of response (in the case of two sources, we label this class of models “dual-route models”). The bottom panel illustrates a case where the evidence signal integrates information from all sources (we label this class of models “single-route models” regardless of the number of sources contributing to the evidence). (C) Expected patterns of performance for classifiers trained on features from individual (I) or cumulative (C) time bins partitioning the time between probe onset and recognition response (see text for details). Assuming the sources contribute independent information with distinct dynamics, dual-route models predict diverging performance for classifiers trained on individual and cumulative time bins (top panel), whereas single-route models predict identical performance (bottom panel; C and I lines are overlapping). AUC indicates area under the receiver operating characteristic curve, a measure of classifier performance.

signal corresponds to an integrated/unitary evidence signal, information from previous time points should not contain information that is not also present in the neural activity at later points.

Fig 1 illustrates our approach with the help of two toy models of evidence in recognition memory. Fig 1A shows activation for two sources of evidence containing information about the old/new status of an item as a function of time, and Fig 1B shows two alternative ways these sources could give rise to an evidence signal for the recognition decision (in this toy example we assume an “early” and a “late” source, analogous to the presumed dynamics of familiarity and recollection signals). The top panel of Fig 1B illustrates a dual-route model: the recognition decision is based exclusively on whichever source has accumulated more evidence at the time of response. Thus any information from the non-dominant evidence source is lost. Assuming sources with different temporal signatures, the evidence signal will initially be determined by activity from the early source, which sometimes will be exceeded by activity from the later source by the time the response is initiated. The bottom panel of Fig 1B illustrates a single process model: here the evidence for the recognition memory decision at any given time reflects the information accrued across all sources so far. Even when the relative contributions of the different sources change, no information is lost, because all relevant information contributes to the evidence signal.

It is difficult to distinguish between these alternative accounts on the basis of recognition decisions alone, because these presumably only reflect a snapshot of the evidence signal from around the time when the response was initiated. Recordings of brain activity, however, allow us to assess the evolution of a neural evidence signal in the lead-up to a recognition response. We used multivariate (“machine learning”) classifiers to quantify the neural evidence distinguishing between targets and lures during the processing of the probe (i.e., between probe onset and just prior to the execution of a response). By comparing performance for classifiers trained on neural features from various partitions of this time period, we can make inferences about whether relevant information is integrated into a single evidence signal or whether evidence from an earlier signal is sometimes lost.

Fig 1C illustrates the logic of the main analyses. As explained in the Methods section, we partition each recognition decision into time bins and train classifiers either on individual time bins or on a cumulatively increasing number of time bins. If brain activity reflects different evidence signals that contribute independent information at different time points, then performance of a classifier trained on neural features from multiple time bins should exceed that of a classifier trained on features from a single time point, since it is able to capitalize on the information from distinct evidence signals (top panel of Fig 1C). If, on the other hand, the neural evidence signal integrates information from all sources, then the signals from previous time points do not contain additional information. Thus, we would expect no benefit for classifiers trained on neural features from multiple time bins in that case (lower panel of Fig 1C).

Materials and methods

Participants

The current data set of 132 participants is a subset of the data set for which we previously presented analyses of overt responses [22] (basic analyses of recognition accuracy and response times are repeated here for this subset). Each participant provided informed consent and all procedures were approved by the Institutional Review Board of the University of Pennsylvania. We selected those participants who completed 20 sessions of various free recall tasks. This enabled us to train statistical classifiers on individual participants’ data from 19 sessions (holding out data from one session for cross-validation of classifier performance). This yielded enough data to train non-linear classifiers even in cases where not all trials contributed to the classification (as detailed below, some of our analyses placed restrictions on response times).

Experimental task

As part of a large-scale study of episodic memory, we asked participants at the end of each of 20 sessions to make recognition memory decisions and confidence ratings about words that had been presented earlier in the session for study in various free recall tasks. Throughout the experiment, we obtained high-density EEG recordings, allowing us to investigate brain activity as it unfolds during processing of a memory probe. Each session included a final recognition memory test that probed target words, which had been previously (i.e., in the same session) studied for free recall, and previously unstudied lure words. Each recognition memory trial consisted of the presentation of a probe word, which required a verbal response to the question of whether the given item had been previously studied. We asked participants to substitute “pess” and “po” for “yes” and “no” when answering this question in order to facilitate determination of response times on the basis of the onset of the verbal response (we excluded trials with

response times below 300 ms and above 3000 ms from further analyses). Following each binary recognition memory decision, we asked participants to indicate their confidence in the response on a scale from 1 to 5 with 5 indicating the highest level of confidence and 1 indicating low confidence. Most participants indicated confidence ratings verbally; any reference to response times in this manuscript is with respect to the binary recognition decision and not for the confidence ratings. Further details were as described previously [22].

Data availability

De-identified data and analysis code used in this study may be freely downloaded from the authors' websites (<http://cogsci.info> and http://memory.psych.upenn.edu/Electrophysiological_Data).

EEG data collection and processing

EEG data were recorded with 129 channel Geodesic Sensor Nets using the Netstation acquisition environment (Electrical Geodesics, Inc.). Cz was used as a reference during recording, but all recordings were converted to an average reference offline. Twenty-six electrodes that were placed on the face (rather than the scalp) were excluded from further analyses.

EEG data were partitioned into events starting 500 ms before the onset of a test item and ending 100 ms before the onset of the verbal recognition response. We applied a time-frequency decomposition using Morlet wavelets with 5 cycles for 15 log-spaced frequencies between 2 and 200 Hz, log-transformed the resulting power values, and z -transformed these within session. We used a 1500 ms buffer at the beginning of the events and mirrored 1500 ms at the end of each event to avoid edge artifacts and to prevent EEG activity from periods during the verbal recognition memory response from bleeding into the analyzed time bins [23]. Data were initially sampled at 500 Hz and down-sampled to 100 Hz after wavelet transformation. We then discarded samples before the onset of the test items, resampled power values for each event to 360 samples, and averaged these samples into 36 equal-time bins for the univariate analyses (Fig 2) and into six equal-time bins for the multivariate classifiers. The lengths of the individual time bins were identical within each trial, but because response times varied across trials, so did the lengths of the ("vincentized" [24]) time bins. We chose to fix the number of time bins to allow us to compare the neural signals across trials as a function of the proportion of each trial's response time, but we also present some complementary analyses using fixed-length (100 ms) time bins below. To aid with interpretation, whenever reasonable, figures show mean times associated with time bins rather than indicating the corresponding ordinal time bin numbers.

Classification of EEG data

We used the scikit-learn library [25] to train support vector machine classifiers with a radial basis function kernel for each participant using a leave-one-session-out cross-validation procedure (all reported classifier results are from left-out sessions). Features were all z -transformed log-power values (the z -transformation was within each session and thus completely separate for training and testing data) either from an individual time bin or from varying ranges of time bins starting with the first time bin. We used the default regularization parameter ($C = 1.0$) and set up the classifier such that the weights were adjusted inversely proportional to class frequencies (using "balanced" as input to the "class_weight" keyword, to take into account unequal numbers of target and lure trials).

For the classification of features from 100 ms time bins, we only included trials where responses occurred 750 ms or more after probe onset and only included the 100 (out of 132) participants with at least 30 such trials in each session. Because we considered the time bin starting at probe onset, as well as 11 additional time bins that each had a 50 ms overlap with the previous time bin, this ensured that the last time bin (ending 650 ms after the probe onset) was separated from the response by at least 100 ms.

Results

Traditionally, researchers have averaged voltage time series from EEG recordings to obtain event-related potentials (ERPs) whose peaks and troughs can be compared across conditions [26]. ERPs mainly reflect phase-locked low-frequency power of the underlying EEG activity and are less sensitive to other spectral features that have been shown to reflect cognitive processes involved in episodic memory [27,28]. For our analyses, we therefore decomposed the EEG signal into power across a wide range of frequencies. As described above, we were careful to exclude any brain activity overlapping with the verbal response by analyzing brain activity only up to 100 ms before the execution of the recognition decision and by using a mirrored buffer that prevented any later brain activity from leaking into the analyzed time period [23].

Power contrasts

Studies of testing effects in recognition memory have suggested that previous recall of an item selectively enhances recollection in a recognition memory test [29]. We aimed to identify any signals reflecting recollective processes by contrasting spectral power for correctly recognized targets (“hits”) that were previously recalled with that for hits that were not. Fig 2A shows the dynamic patterns of these contrasts for sensors in two regions of interest (ROIs) that have been frequently used in EEG investigations of familiarity and recollection [30]. The contrasts in the two ROIs were broadly similar, although there was a pattern of larger θ (4–7 Hz) power for previously recalled hits in later time bins at the posterior ROI, which may be related to higher ERP amplitudes often found at the posterior ROI for hits with recollective experiences [30].

To directly track neural evidence distinguishing old from new items, we also calculated contrasts between spectral power for targets and lures irrespective of the subsequent response. Fig 2B shows that the pattern of these contrasts was remarkably similar to those for contrasts between previously recalled and not previously recalled hits (shown in Fig 2A). Under the assumption that memory is strongest for previously recalled targets, weaker for not previously recalled targets, and weakest/absent for lures, the patterns in Fig 2 could reflect a single memory strength signal that falls out of any contrast between two conditions that vary in memory strength [31]. The fact that these patterns changed quite drastically in the lead-up to the memory decisions, however, might also reflect independent sources of evidence with distinct time courses. An assessment of the relative merits of these alternative accounts, therefore, requires us to quantify the neural evidence in the trial-by-trial variability of EEG activity that distinguishes between targets and lures in the lead-up to a recognition memory decision.

Quantifying neural evidence

Despite previous efforts to relate the familiarity and recollection components of dual-process models to different (temporally distinct) neural signals [5,6], little is known about the actual dynamics of information accumulation in recognition memory decisions and how they relate to accuracy, response times (RTs), and response confidence.

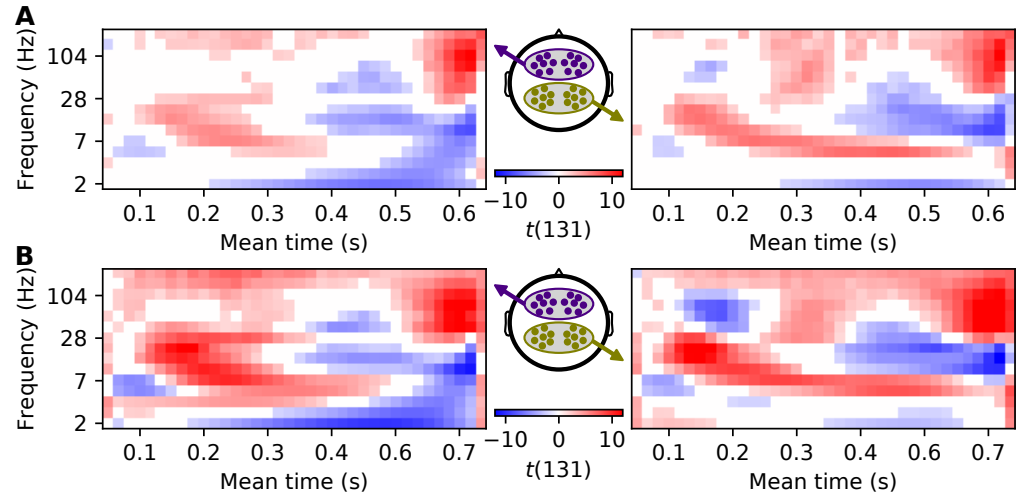


Fig 2. Power contrasts. Differences in power for previously recalled vs. not previously recalled hits (**A**) and for all targets vs. lures irrespective of response (**B**). Left and right panels show these differences for anterior and posterior ROIs (illustrated in the middle panels) respectively across frequencies and vincentized time bins (mean times associated with some of the time bins are indicated on the abscissas). We used t -values for the differences between trial categories for each participant to calculate t -values across participants. Red shades indicate higher power for previously recalled hits (**A**) or for all targets (**B**) and blue shades indicate higher power for not previously recalled hits (**A**) or for lures (**B**; within each panel, values that did not reach statistical significance with a false discovery rate of .05 are set to white).

Building on the success of machine learning techniques in neural data analyses that have provided unique insights into the dynamics of cognitive processes [32–35], we trained statistical classifiers on spectral EEG features to track the neural dynamics of evidence accumulation during recognition memory decisions. A classifier’s ability to distinguish targets from lures can be directly compared to an individual’s recognition memory performance through the use of receiver operating characteristic (ROC) functions relating hits to “false alarms” (incorrect classifications of lures as “old”). The area under an ROC curve (AUC) serves as a convenient index of classification performance, with an AUC of .5 indicating chance performance and an AUC of 1.0 indicating perfect classification [36]. We previously used confidence ratings and latencies for binary recognition memory decisions to generate ROC functions and showed a strong correspondence between the respective AUCs in the dataset from which the current dataset was derived [22]. Here we assess the evolution of a neural signal indexing evidence for the recognition memory decision by generating ROC functions from the outputs of classifiers that were trained to distinguish targets from lures using spectral EEG features from various intervals during the recognition period. To reduce computational complexity and generate more reliable features for the classifiers, we aggregated the time bins shown in Fig 2 by averaging them in groups of six, partitioning each recognition memory decision into six equal-time bins (see Methods for details).

Neural evidence across the entire recognition period

For each participant, we trained a classifier on spectral EEG features from all six time bins to confirm (in held out sessions) that the neural signal in individual trials reliably distinguished between targets and lures (AUC = .71, $t(131) = 34.59$, $SE = 0.021$, $p < .001$; Fig 3A). Single-process models of recognition memory commonly assume that

evidence for targets is more variable than that for lures [12], with converging evidence for this assumption coming from fits of detailed models of evidence accumulation [18,37,38]. Larger target variability can result in increased AUCs that are based only on “old” responses (or corresponding classifier output) compared to those reflecting overt responses or classifier output for “new” decisions only [22]. These conditional AUCs indicate how much signal the measure of interest contains for each response class (beyond the signal contained in the binary classification of test items as “old” or “new” [22]) and were consistently larger for “old” classifications across all measures ($t(131) = 8.22-33.22$, $SE = 0.005-0.008$, $ps < .001$).

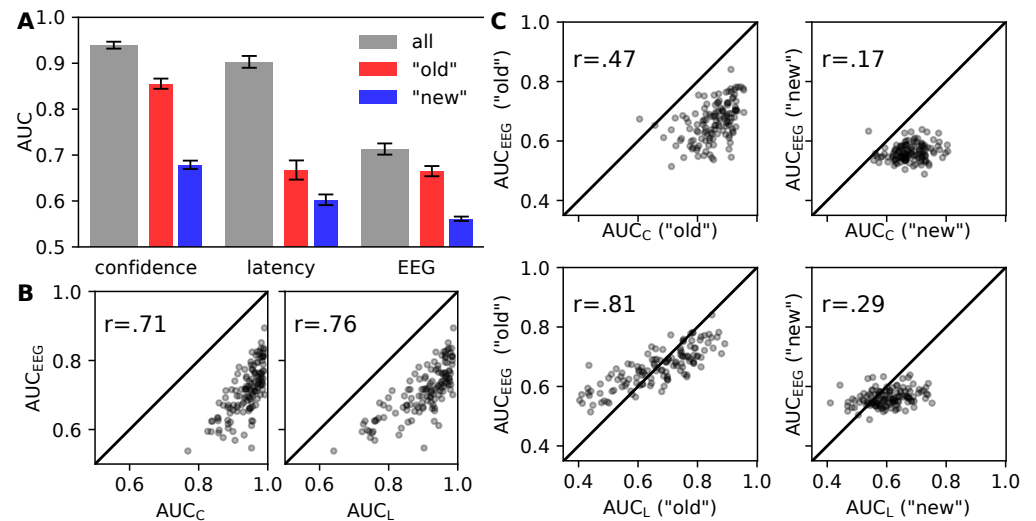


Fig 3. Overall performance for statistical and human classifiers. (A) Areas under the ROC functions (AUCs) for confidence ratings, response latencies, and EEG activity. AUCs for conventional ROC functions are shown in gray and those for conditional ROC functions based on only “old” or “new” responses (or corresponding classifier output) are shown in red and blue respectively. Error bars correspond to 95% confidence intervals. (B) Scatter plots illustrating the relationships between AUCs for either confidence ratings (AUC_C; left panel) or response latencies (AUC_L; right panel) and EEG activity (AUC_{EEG}). (C) Scatter plots illustrating the relationships between conditional AUCs for either confidence ratings (AUC_C; top panels) or response latencies (AUC_L; bottom panels) corresponding to “old” (left panels) or “new” (right panels) recognition decisions and EEG activity (AUC_{EEG}) for “old” and “new” classifications. Corresponding correlation coefficients are indicated in the top left of each scatter plot and the main diagonals are shown for convenience.

In principle, a classifier trained on neural data to distinguish targets from lures may use different signals than those which are most important for the individual’s recognition memory decision. Indeed, it is unlikely that the coarse measure of scalp EEG activity (compressed into power for a small number of frequencies) could reflect the neural signals leading to the recognition memory decision with high fidelity. In that light, it is of particular interest to what extent the qualitative pattern of (conditional) AUCs is similar across measures. Fig 3B illustrates similarly close relationships between AUCs based on overt responses (AUC_C and AUC_L, with the subscripts denoting confidence ratings and response latency, respectively) and AUCs based on EEG-classifier output (AUC_{EEG}; $r = .71$ & $.76$, $t(130) = 11.53$ & 13.36 , $ps < .001$) We also observed strong correlations between conditional AUCs based on overt responses and EEG activity (Fig 3C; $r = .17-.81$, $t(130) = 1.97-15.58$, $p \leq .05$).

Tracking neural evidence across time

Having established a neural signal that reliably distinguishes between targets and lures and that strongly correlates with recognition memory decisions, we next asked how this evidence accrues over time. If different types of evidence accrue with distinct time courses, this would lead to distinct relative contributions at different points in the decision process. In dual-route models, these signals are assumed to reflect independent sources of evidence and thus should give a distinct advantage to any classifier trained on features from multiple time bins (which would thus be posed to capitalize on both types of evidence to boost performance), relative to classifiers trained on features from a smaller portion of the recognition period. To test this prediction, we compared performance for classifiers trained on features from a cumulatively increasing number of time bins to performance from classifiers trained on features from individual time bins only. Fig 4A shows the corresponding AUCs for these two types of classifiers as a function of time, as well as respective differences (Δ AUC). The AUCs for both types of classifiers were very similar and increased gradually with time (regression slopes and intercepts for both were 0.02 and 0.59, respectively; corresponding r values for classifiers trained on individual and cumulative features were 0.52 and 0.53 respectively, $SEs = 0.001$, $ps < .001$). The two types of classifiers are identical for the first time bin and, in the absence of over-fitting, the additional features used by the cumulative classifier for later time bins cannot decrease classification performance relative to the classifiers using features from only one time bin. At the last time bin, the number of features differ by a factor of six for the two classifiers and, even though the features from earlier time bins clearly contained relevant signal (as shown by reliable classification performance for previous time bins), the difference in classification performance between the two classifiers was minuscule (Δ AUC = .003 at the last time bin and $< .01$ throughout). This pattern of results is what would be expected if the classifier output reflected an integrated evidence signal as it accumulates in the lead-up to a decision.

To properly assess this evidence against dual-route accounts of recognition memory, it is important to link our neural classifiers to overt responses. Above we compared (conditional) AUCs for memory decisions to those from classifiers trained on neural features from all time bins (Fig 3). Here we track the correspondence between brain activity and overt responses across time by correlating the trial-by-trial output of classifiers trained on individual time bins with subsequent confidence ratings. If classifiers trained on features from different time bins were picking up on different types of relevant signals, we would expect the correlation between classifier output and confidence ratings to peak whenever each signal type provides maximal evidence (e.g., a peak in correlation reflecting an early familiarity signal, followed by another peak reflecting contributions from later recollective processes). From the perspective of a single, continuously accumulating, memory strength signal, however, we would expect gradually increasing correlations as a function of time. Fig 4B shows positive and increasing correlations (measured by Spearman's ρ) for outputs from classifiers and confidence ratings as a function of time (regression slope and intercept were 0.43 and 0.22, respectively, $r = .56$, $SE = 0.002$, $p < .001$).

Neural evidence as function of prior recall

Given the converging evidence against dual-route accounts, we set out to maximize our ability to detect any contributions from recollective processes by assessing the classifiers' performance conditional on previous recall of targets [29]. Fig 4C shows that AUCs from recalled trials exceeded those for unrecalled trials across all time bins ($ts(131) = 21.08-31.19$, $SEs = 0.003$, $ps < .001$) and that both types of AUCs increased with time (regression slopes for both were 0.02 with intercepts of 0.54 and 0.61 for




Fig 4. Classifier performance as a function of time. (A) AUCs for the EEG classifier as a function of time bin for features from individual time bins (I) and for cumulatively adding features from each time bin (C; top panel). The bottom panel shows the differences between AUCs (Δ AUCs). (B) Mean Spearman's ρ between confidence ratings and output of the EEG classifier using features from individual time bins as a function of time bin. (C, D) AUCs for the EEG classifier as a function of vincentized time bin (C) or fixed-length (100 ms) time bin (D) when tested on targets that were previously recalled (R) or previously unrecalled (U) and corresponding differences (Δ AUC). Mean times corresponding to vincentized time bins are indicated on the abscissas in (A) and (B); because RTs differed as a function of previous recall only ordinal time bin labels are indicated in (C). Shaded regions correspond to 95% confidence intervals.

AUCs based on unrecalled and recalled targets respectively). Furthermore, the differences between AUCs based on recalled and unrecalled targets also increased with time (regression slope and intercept for Δ AUC were 0.006 and 0.073, respectively, $r = .27$, $SE < 0.001$, $p < .001$). Regardless of previous recall status, classifier performance increased gradually with time as one would expect if evidence accumulated continuously (but at different rates) for both types of trials.

Because differential RTs for targets as a function of previous recall could have contributed to the differences shown in Fig 4C, we also trained classifiers on 100 ms bins of spectral EEG features that we moved from probe onset in steps of 50 ms until the time window that ended at 650 ms after probe onset. Again, AUCs from recalled trials exceeded those for unrecalled trials across all time bins ($ts(99) = 3.6-19.76$, $SEs = 0.002-0.003$, $ps < .001$) and both types of AUCs increased with time (with regression slopes of 0.004 and 0.009 and corresponding intercepts of 0.5 for the AUCs based on unrecalled and recalled targets respectively; corresponding $rs = .45$ & $.64$, $SEs < 0.001$, and $ps < .001$). Also as above, the differences between AUCs based on

299
300
301
302
303
304
305
306
307
308
309
310
311
312
313

recalled and unrecalled targets increased with time (regression slope and intercept for Δ AUC were 0.004 and 0.003, respectively, $r = .53$, $SE < 0.001$, $p < .001$). For both types of analyses, evidence accrual for recalled and unrecalled targets appears most consistent with the continuous accumulation of different amounts (rather than different kinds) of evidence.

Discussion

Despite a long history of research on recognition memory, there is considerable disagreement about the nature of the evidence that allows us to distinguish repeated encounters from novel experiences [1–4, 9, 12–14, 31, 39–41]. By training multivariate classifiers to distinguish between previously studied and novel items based on spectral EEG features recorded prior to the execution of a recognition response, we were able to track the accrual of this evidence. We found performance of classifiers trained on a small fraction of the recognition period to be nearly identical to that of classifiers that were able to also capitalize on features from all previous time bins (Fig 4A). This suggests that the classifiers directly tracked an evidence signal, rather than a signal over which a decision process integrates to calculate the accumulated evidence—an interpretation also supported by the strong correspondence between classifier output and overt responses (Fig 3 & 4B) and by the increasing classifier performance as a function of time (Fig 4A, C–D). These findings, as well as the strong qualitative similarities between classifier performance as a function of time for previously recalled vs. not previously recalled targets (Fig 4C–D), are difficult to reconcile with dual-route models that posit different kinds of recognition decisions based on independent and temporally distinct familiarity and recollection signals [1–3, 7].

A recent study pursued similar goals to the present work by training a linear classifier on a sliding window of EEG voltages recorded during recognition memory decisions and by relating classifier output at two time points to drift rates in a drift diffusion model (DDM) [38]. Despite several differences in methodology (e.g., the use of linear vs. non-linear classifiers, the use of fixed vs. vincentized time bins, and the use of voltage vs. spectral EEG activity as classifier features), the conclusions from both studies largely complement each other. Specifically, both studies make a case in favor of a unitary evidence signal underlying recognition memory decisions that is more variable for targets than for lures (as described above, our finding that areas under conditional ROC functions were larger for “old” compared to “new” classifications is compatible with the latter assertion [22]). That previous study’s conclusions, however, relied on strong assumptions about the relationship between EEG voltage in two fixed time bins and the drift rates in a DDM [38]. None of our conclusions depend on a specific model of recognition memory or binary choice. Our inferences, however, do rely on a set of important assumptions, to which we now turn.

Does classifier output reflect recognition evidence?

We have used outputs of multivariate classifiers trained on spectral EEG activity to track evidence for recognition decisions as it evolved in the lead-up to a response. Our conclusions are conditional upon this approach’s ability to faithfully reflect information that is relevant for the recognition decision. Alternatively, our conclusion that a unitary evidence signal drives recognition decisions could also be due to our approach’s inability to detect neural activity associated with a separate evidence signal. The strong correlations of classifier output with overt responses offer some reassurance by limiting the variance that could be explained by unobserved evidence signals. Additionally, ERP analyses during recognition decisions are often interpreted as reflecting contributions of

independent familiarity and recollection signals [5, 6], which should render these signals observable in our approach.

Might different evidence signals overlap?

Our analyses depend on different evidence signals exhibiting distinct temporal profiles. Whereas many dual-process models do not specify the dynamics of familiarity and recollection, as discussed above, the near universal assumption is that these signals are temporally distinct. To the extent that evidence signals for different routes to recognition overlap in time, we would not be able to distinguish them with our approach. Furthermore, to detect evidence signals corresponding to multiple routes to recognition, our analyses assume that combining information from the different corresponding evidence signals allows for a better discrimination between targets and lures, compared to the use of the evidence signal for one route only. To the extent that different routes to recognition are indeed distinct, and thus rely on independent evidence, the corresponding evidence signals should meet this requirement.

Novelty detection

Whereas we have considered different ways in which a previously studied item might be recognized as “old”, some evidence suggests that the detection of novelty can also support recognition memory [42–45]. Our approach quantifies evidence distinguishing between targets and lures and thus is agnostic with respect to whether the relevant signals index familiarity or novelty. It is likely that any familiarity and novelty signals would be strongly (negatively) correlated in standard recognition memory tasks, and some evidence suggests similar temporal profiles for familiarity and novelty signals [45]. Our approach does not distinguish between strongly correlated and/or temporally overlapping signals and thus is unable to differentiate familiarity and novelty signals with these properties.

Conclusion

Debates about the relative merits of single- vs. dual-process models often presuppose a false dichotomy between a single route to recognition and contributions from multiple sources of evidence (such as recollection and familiarity) to recognition decisions. Notable exceptions are formal models positing that recognition decisions are driven by an evidence signal that combines contributions from familiarity and recollection signals [8, 9]. Rather than asking what different sources of evidence might lead to recognition, we focused on the question to what extent available evidence is integrated into a unitary evidence signal that drives recognition decisions. There are inherent trade-offs between integrating all available evidence, and thus maximizing the ability to distinguish old from new items, and separately considering different sources of evidence, and thus maximizing the ability to qualify recognition decisions (e.g., with remember-know judgments). It is therefore possible that the extent to which evidence from different sources is integrated into a unitary signal is sensitive to task demands [8, 9]. At least for the standard old-new discrimination task considered here, however, our results indicate that the memory system integrates available evidence into a unitary evidence signal that drives recognition decisions. These findings thus firmly link recognition decisions to other types of decisions under uncertainty, which are commonly assumed to rely on a unitary evidence signal differentiating between the response options [34, 35, 46].

Acknowledgments

The authors would like to thank the members of the Computational Memory Laboratory at the University of Pennsylvania for their assistance with data collection and preprocessing as well as Youssef Ezzyat, James Kragel, Nora Herweg, Ethan Solomon, and Rivka Cohen, for helpful comments on a draft of this paper. This work was supported by grant NIMH RO1 55687 to MJK. Both authors contributed to the study design, analysis plan, and write-up. M.J.K. supervised data collection. C.T.W. analysed the data and wrote the manuscript. The authors declare that they have no competing financial interests.

References

1. Diana RA, Reder LM, Arndt J, Park H. Models of recognition: A review of arguments in favor of a dual-process account. *Psychonomic Bulletin & Review*. 2006;13(1):1–21.
2. Yonelinas AP. The Nature of Recollection and Familiarity: A Review of 30 Years of Research. *Journal of Memory and Language*. 2002;46(3):441–517. doi:10.1006/jmla.2002.2864.
3. Yonelinas AP, Aly M, Wang WC, Koen JD. Recollection and Familiarity: Examining Controversial Assumptions and New Directions. *Hippocampus*. 2010;20:1178–1194. doi:10.1002/hipo.20864.
4. Malmberg KJ. Recognition memory: A review of the critical findings and an integrated theory for relating them. *Cognitive Psychology*. 2008;57:335–384. doi:10.1016/j.cogpsych.2008.02.004.
5. Curran T. The electrophysiology of incidental and intentional retrieval: ERP old/new effects in lexical decision and recognition memory. *Neuropsychologia*. 1999;37:771–785.
6. Rugg MD, Curran T. Event-related potentials and recognition memory. *Trends in Cognitive Sciences*. 2007;11(6):251–257. doi:10.1016/j.tics.2007.04.004.
7. Reder LM, Nhouyvanisvong A, Schunn CD, Ayers MS, Angstadt R, Hiraki KA. A mechanistic account of the mirror effect for word frequency: A computational model of remember-know judgments in a continuous recognition paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2000;26:294–320.
8. Rotello CM, Macmillan NA, Reeder JA. Sum-Difference Theory of Remembering and Knowing: A Two-Dimensional Signal-Detection Model. *Psychological Review*. 2004;111:588–616.
9. Wixted JT, Mickes L. A continuous dual-process model of remember/know judgments. *Psychological Review*. 2010;117(4):1025–1054. doi:10.1037/a0020874.
10. Ratcliff R. A theory of memory retrieval. *Psychological Review*. 1978;85(2):59–108. doi:10.1037/0033-295x.85.2.59.
11. Ratcliff R, Starns JJ. Modeling confidence and response time in recognition memory. *Psychological Review*. 2009;116(1):59–83. doi:10.1037/a0014086.

12. Wixted JT. Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*. 2007;114(1):152–176. doi:10.1037/0033-295x.114.1.152.
13. Dunn JC. Remember-Know: A Matter of Confidence. *Psychological Review*. 2004;111(2):524–542. doi:10.1037/0033-295x.111.2.524.
14. Dunn JC. The dimensionality of the remember-know task: A state-trace analysis. *Psychological Review*. 2008;115(2):426–446. doi:10.1037/0033-295x.115.2.426.
15. Cox GE, Shiffrin RM. Criterion Setting and the Dynamics of Recognition Memory. *Topics in Cognitive Science*. 2012;4(1):135–150. doi:10.1111/j.1756-8765.2011.01177.x.
16. Diller DE, Nobel PA, Shiffrin RM. An ARC–REM model for accuracy and response time in recognition and recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2001;27(2):414–435. doi:10.1037/0278-7393.27.2.414.
17. Starns JJ, White CN, Ratcliff R. The strength-based mirror effect in subjective strength ratings: The evidence for differentiation can be produced without differentiation. *Memory & Cognition*. 2012;40(8):1189–1199. doi:10.3758/s13421-012-0225-1.
18. Starns JJ, Ratcliff R. Validating the unequal-variance assumption in recognition memory using response time distributions instead of ROC functions: A diffusion model analysis. *Journal of Memory and Language*. 2014;70:36–52. doi:10.1016/j.jml.2013.09.005.
19. Shiffrin RM, Steyvers M. A Model For Recognition Memory: REM—Retrieving Effectively From Memory. *Psychonomic Bulletin & Review*. 1997;4(2):145–166. doi:10.3758/bf03209391.
20. Yonelinas AP. Receiver-operating characteristics in recognition memory: Evidence for a dual-process model. *Journal of Experimental Psychology: Learning Memory, and Cognition*. 1994;20(6):1341–54.
21. Yonelinas AP. Recognition memory ROCs for item and associative information: The contribution of recollection and familiarity. *Memory & Cognition*. 1997;25:747–763.
22. Weidemann CT, Kahana MJ. Assessing recognition memory using confidence ratings and response times. *Royal Society Open Science*. 2016;3(4):150670. doi:10.1098/rsos.150670.
23. Cohen MX. *Analyzing Neural Time Series Data*. MIT University Press Group Ltd; 2014. Available from: http://www.ebook.de/de/product/21273084/mike_x_cohen_analyzing_neural_time_series_data.html.
24. Ratcliff R. Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin*. 1979;86(3):446–461. doi:10.1037/0033-2909.86.3.446.
25. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*. 2011;12:2825–2830.

26. Luck SJ. An Introduction to the Event-Related Potential Technique. The MIT Press; 2005.
27. Nyhus E, Curran T. Functional role of gamma and theta oscillations in episodic memory. *Neuroscience & Biobehavioral Reviews*. 2010;34(7):1023–1035. doi:10.1016/j.neubiorev.2009.12.014.
28. Jacobs J, Hwang G, Curran T, Kahana MJ. EEG oscillations and recognition memory: Theta correlates of memory retrieval and decision making. *NeuroImage*. 2006;32(2):978–987. doi:10.1016/j.neuroimage.2006.02.018.
29. Chan JCK, McDermott KB. The testing effect in recognition memory: A dual process account. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2007;33(2):431–437. doi:10.1037/0278-7393.33.2.431.
30. Schwikert SR, Curran T. Familiarity and recollection in heuristic decision making. *Journal of Experimental Psychology: General*. 2014;143(6):2341–2365. doi:10.1037/xge0000024.
31. Squire LR, Wixted JT, Clark RE. Recognition memory and the medial temporal lobe: A new perspective. *Nature Reviews Neuroscience*. 2007;8(11):872–883. doi:10.1038/nrn2154.
32. Polyn SM, Natu VS, Cohen JD, Norman KA. Category-Specific Cortical Activity Precedes Retrieval During Memory Search. *Science*. 2005;310(5756):1963–1966. doi:10.1126/science.1117645.
33. Norman KA, Polyn SM, Detre GJ, Haxby JV. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*. 2006;10(9):424–430. doi:10.1016/j.tics.2006.07.005.
34. Ratcliff R, Philiastides MG, Sajda P. Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the EEG. *Proceedings of the National Academy of Sciences*. 2009;106(16):6539–6544. doi:10.1073/pnas.0812589106.
35. Philiastides MG, Sajda P. Temporal Characterization of the Neural Correlates of Perceptual Decision Making in the Human Brain. *Cerebral Cortex*. 2006;16(4):509–518. doi:10.1093/cercor/bhi130.
36. Fawcett T. An introduction to ROC analysis. *Pattern Recognition Letters*. 2006;27(8):861–874. doi:10.1016/j.patrec.2005.10.010.
37. Starns JJ. Using response time modeling to distinguish memory and decision processes in recognition and source tasks. *Memory & Cognition*. 2014;42(8):1357–1372. doi:10.3758/s13421-014-0432-z.
38. Ratcliff R, Sederberg PB, Smith TA, Childers R. A single trial analysis of EEG in recognition memory: Tracking the neural correlates of memory strength. *Neuropsychologia*. 2016;93:128–141. doi:10.1016/j.neuropsychologia.2016.09.026.
39. Kirwan CB, Wixted JT, Squire LR. A demonstration that the hippocampus supports both recollection and familiarity. *Proceedings of the National Academy of Sciences*. 2010;107(1):344–348. doi:10.1073/pnas.0912543107.
40. Dede AJO, Wixted JT, Hopkins RO, Squire LR. Hippocampal damage impairs recognition memory broadly, affecting both parameters in two prominent models of memory. *Proceedings of the National Academy of Sciences*. 2013;110(16):6577–6582. doi:10.1073/pnas.1304739110.

41. Merkow MB, Burke JF, Kahana MJ. The human hippocampus contributes to both recollection and familiarity components of recognition memory. *Proceedings of the National Academy of Science of the United States of America*. 2015;112(46):14378–14383. doi:10.1073/pnas.1513145112.
42. Daselaar SM, Fleck MS, Prince SE, Cabeza R. The medial temporal lobe distinguishes old from new independently of consciousness. *J Neurosci*. 2006;26(21):5835–5839. doi:10.1523/JNEUROSCI.0258-06.2006.
43. Davelaar EJ, Tian X, Weidemann CT, Huber DE. A Habituation Account of Change Detection in same/different judgments. *Cognitive, Affective, and Behavioral Neuroscience*. 2011;11:608–626. doi:10.3758/s13415-011-0056-8.
44. Kafkas A, Montaldi D. Two Separate, But Interacting, Neural Systems for Familiarity and Novelty Detection: A Dual-Route Mechanism. *Hippocampus*. 2014;24:516–527. doi:10.1002/hipo.22241.
45. Bunzeck N, Doeller CF, Fuentemilla L, Dolan RJ, Duzel E. Reward Motivation Accelerates the Onset of Neural Novelty Signals in Humans to 85 Milliseconds. *Current Biology*. 2009;19:1294–1300. doi:10.1016/j.cub.2009.06.021.
46. Nosofsky RM, Little DR, James TW. Activation in the neural network responsible for categorization and recognition reflects parameter changes. *Proceedings of the National Academy of Sciences*. 2012;109(1):333–338. doi:10.1073/pnas.1111304109.