Published in final edited form as:

Neuropsychologia. 2016 December ; 93(Pt A): 128–141. doi:10.1016/j.neuropsychologia.2016.09.026.

A Single Trial Analysis of EEG in Recognition Memory: Tracking the Neural Correlates of Memory Strength

Roger Ratcliff^{1,*}, Per Sederberg¹, Troy Smith¹, and Russ Childers¹

¹The Ohio State University

Abstract

Recent work in perceptual decision-making has shown that although two distinct neural components differentiate experimental conditions (e.g., did you see a face or a car), only one tracked the evidence guiding the decision process. In the memory literature, there is a distinction between a fronto-central evoked potential measured with EEG beginning at 350 ms that seems to track familiarity and a late parietal evoked potential that peaks at 600 ms that tracks recollection. Here, we applied single-trial regressor analysis (similar to multivariate pattern analysis MVPA) and diffusion decision modeling to EEG and behavioral data from two recognition memory experiments to test whether these two components contribute to the recognition decision process. The regressor analysis only involved whether an item was studied or not and did not involve any use of the behavioral data. Only the EEG activity that distinguishes studied from not studied items at about 600 ms following each test item onset predicted the diffusion model drift rate derived from the behavioral choice and reaction times but only for studied items. When drift rate was made a linear function of the trial-level regressor values, the estimate for studied items was different than zero. This showed that the later EEG activity indexed the trial-to-trial variability in drift rate for studied items. Our results provide strong evidence that only a single EEG component reflects evidence being used in the decision process.

Keywords

Diffusion model; recognition memory; EEG; single-trial regressor; reaction time

Introduction

In studies examining the neural processes used in decision making, it is often assumed that if a neural response varies as a function of some perceptual or memory variable, it tracks the evidence being used to make a decision. However, this view does not discriminate between early representations of stimuli and evidence extracted from them that is used in the decision process.

*Corresponding author. Roger Ratcliff, Department of Psychology, The Ohio State University, Columbus, OH, 43210. Fax.: +614 688 3984; ratcliff.22@osu.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

In previous work, single-trial analysis of EEG in concert with diffusion decision modeling of choice and reaction time (RT) has been applied to examine the neural components of perceptual decisions. In a face/car discrimination task, Philiastides, Ratcliff & Sadja (2006) used a single-trial analysis, also known as multivariate pattern analysis (e.g., Blankertz et al. 2011; Pereira et al. 2009) that weighted signals from an array of electrodes to produce a single component value that represented how car-like or face-like was the EEG signal. Two components were obtained that tracked stimulus quality, one at around 180 ms and one at around 380 ms (by component, it is meant a significant single trial regressor signal as in Philiastides et al., 2006). Ratcliff, Philiastides & Sadja (2009) sorted the behavioral data for each condition into halves based solely on the EEG component value to which they fit the diffusion model. Drift rate (quality of evidence used in the decision process) differed for the two halves, but only for the late component. This suggested that the later component tracks information being used in the decision, but the earlier component represents the quality of the stimulus encoding from which decision-relevant information is extracted.

In the behavioral memory literature, there is considerable debate about whether one or two processes are involved in making recognition decisions. The single-process view, reflected in most computational modeling approaches, proposes that recognition decisions are made based on either a single source of information or on information from multiple sources that is combined into a single continuous measure of memory strength (Cohen, Rotello & Macmillan, 2008; Dennis & Humphreys, 2001; Dunn, 2004; Gillund & Shiffrin, 1984; Hintzman, 1984; Shiffrin & Steyvers, 1997; Starns & Ratcliff, 2008; Wixted, 2007). By contrast, the dual-process view proposes that there are two distinct decision processes that are used in recognition-- a sense of familiarity that is a continuous variable and an all-or-none recollection component based on details about the encoding event (Rotello, Macmillan & Reeder, 2004; Yonelinas, 1994, 1997; see also Buchler, Light & Reder, 2008).

Key evidence for the dual process view comes from studies examining event-related potential (ERP) components when subjects are making recognition decisions (Rugg, 1995; Rugg & Curran, 2007). Specifically, it has been argued that a frontal component occurring in the range of 300 to 500 ms after onset of a test item can be identified with familiarity while a parietal component occurring in the range of 400 to 600 ms after onset can be identified with recollection (see Eichenbaum, Yonelinas & Ranganath, 2007; Rugg & Yonelinas, 2003; but see Paller, Voss, & Boehm, 2007).

In this article, we apply this single trial regressor and diffusion modeling approach to data from a recognition memory task. This allows us to determine whether these proposed neural correlates of the memory processes guide the recognition decision.

Methods

We report the results of two recognition memory experiments, the second experiment serving as an independent replication of the first experiment, with a simplified design. The two experiments both used a standard item recognition procedure in which a list of words is presented followed by a test list. For every test word, subjects have to decide whether it was studied or not. In many prior experiments, response time and accuracy data are well-

explained by the diffusion decision model (Ratcliff, Thapar, & McKoon, 2004, 2010, 2011). In Experiment 1 we manipulated number of repetitions of study words and word frequency of studied and new words and in Experiment 2 we manipulated number of repetitions of study words. These manipulations were done to provide constraints on modeling and to demonstrate that the experiments provided standard results from this task.

Subjects

All subjects were recruited from the university community at The Ohio State University and spoke and read English fluently. Subjects provided written consent in accordance with requirements of the local institutional review board and were paid \$20 for their time. Experiment 1 comprised twenty-five right-handed subjects (14 female) between 18 and 38 years of age ($M = 21.6$). Data from subjects for whom there were excessive motion artifacts or recording noise ($n = 5$) or who failed to follow instructions ($n = 2$) are not included in the analyses. Experiment 2 comprised fourteen right-handed subjects, however three were discarded due to experimental error, excessive motion artifacts, and/or recording noise, leaving 11 subjects (3 female) between 18 and 24 years of age ($M = 20.6$).

Stimuli

Stimuli were drawn from the SUBTLEXUS database (Brysbaert & New, 2009). A total of 549 high-frequency words (between 40 and 400 occurrences per million, $M=122.31$, $SD=84.52$) and 553 low-frequency words (between 1 and 7 occurrences per million, $M=1.44$, $SD = 0.39$) formed the stimulus pool. All words in the pool had between 4 and 7 letters, and word length was equated across the frequencies ($M=5.03$ and 4.99 , respectively).

Design

The first experiment used a 2 (Word Frequency: high vs. low) $\times$ 3 (Item Strength: strong vs. weak vs. new) factorial within-subject and within-list design. Each subject studied and was tested on 13 lists of words. Each study list was constructed by randomly selecting 9 words of each frequency from the pool to be studied one time (the “weak” condition) and 9 words of each frequency to be studied three times (the “strong” condition). Thus, there were 36 unique words and 72 item presentations in each study list, with the order of the presentations pseudo-randomized such that there were no immediate repetitions of the strong items. A matching test list consisting of all 36 studied words along with a set of 36 matching lures was constructed for each study list. The order of words on the test lists was randomized in 2 blocks with the first 18 unique studied words tested in the first block and the second 18 tested in the second block. This was done to ensure that the end of list items were never tested immediately, thereby reducing the probability of recency effects. After discounting the first list, which was used for practice, this design yielded a total of 108 target trials and 108 lure trials for each experimental condition.

Experiment 2 had only the item strength manipulation, giving rise to a 3 factorial (strong vs. weak vs. new), within-subject and within-list design. Each study list was constructed from a pool excluding low-frequency words, but without the word length restriction applied in Experiment 1. Subjects studied and were tested on 11 lists, each comprising 15 strong (presented three times), 15 weak (presented one time), and two once-presented buffer words

at the end of the list, giving rise to 62 item presentations. Each study list had a matching test list, where each of the 30 unique target words was presented in random order along with 30 lure words. The two buffer target items, along with two buffer lure items were tested at the start of the list to account for recency effects. After discounting the first list, which was used for practice, the design of Experiment 2 yielded 150 trials in each of the strong, weak, and lure conditions.

Equipment

For Experiment 1, a desktop computer with a 17" LCD display was used to present the stimuli and collect subject responses. A custom program written using the Python experiment programming library (PyEPL; Geller et al., 2007) was used to generate the study lists for each subject, control the timing of the tasks, and record subject responses. For Experiment 2, a desktop computer with a 24" LCD display, running at 120 Hz was used to present the stimuli and collect subject responses. A custom Python program written using the State Machine Interface Library for Experiments (SMILE; <https://github.com/compmem/smile>) generated the study lists, controlled the timing of the tasks, and recorded subject responses.

For both experiments, scalp EEG data were sampled and recorded reference-free at 1000 Hz using a DC-powered actiCHamp amplifier/analog-to-digital converter connected to a desktop PC equipped with PyCorder software (BrainProducts GmbH, Munich, Germany). Prior to receiving instructions, each subject was fitted with an elastic cap containing 64 active electrodes arrayed in an extended 10-10 layout. Electrode impedances were reduced to less than 25K ohms in accordance with operating instructions for the actiCAP system (BrainProducts GmbH, Munich, Germany).

Procedure

After informed consent was obtained, subjects were fitted with the EEG cap, seated in front of a computer, and given instructions by the experimenter. Subjects were told that they would be studying lists of words and would be given a recognition memory test after each study list. They were informed that some of the study words would be repeated, but were not told that this was an experimental manipulation. Subjects were also instructed to try to avoid thinking back to previous words on the list during the study session. Subjects were then given the first study and test list as a practice list. The experimenter answered any questions and then started the experiment proper.

The precise timing of study and test item presentations was similar, but not identical for the two experiments. In Experiment 1, prior to each study list, an orientation cross was displayed in the center of the screen for 2600–3000 ms. Then the list of words were displayed, one at a time, in the center of the screen for 1600 ms followed by a blank screen for 300–700 ms. Following the study list, an orientation cross was displayed for 1000–1200 ms to signify the test was about to begin. On each test trial, the probe word was displayed at the center of the screen and a prompt “Old or New?” was displayed below it. Subjects made an old judgment by pressing the “J” key on a standard keyboard with the index finger of their right hand or a new judgment by pressing the “K” key with the middle finger of their

right hand. Once the subject made a response, the prompt was removed from the screen. Subjects were allowed up to 1800 ms to make a response, then the stimulus was removed from the screen. Test trials were separated by a jittered inter-trial interval of 400–800 ms.

In Experiment 2, prior to each study list an orientation cross was displayed in the center of the screen for 2000 ms, followed by a 1000–1400 ms blank pause before the words. The list of study words was presented one at a time in the center of the screen for 1600 ms, followed by a blank screen for 300–700 ms before the next word. Following the study list, an instruction reminder was provided for 4000 ms to refresh the subject of their Old/New response mapping, which was counterbalanced across subjects and did not remain on the screen during the memory test (see below). After the response-mapping reminder an orienting stimulus was presented at the center of the screen for 4000 ms, followed by a 1000–1400 ms blank pause before the test words. For each test word, the subject had as much time as they would like to make a response and once a response was made, the word was removed from the screen 500 ms later, followed by a 300–700 ms blank inter-trial interval before the next test word.

EEG Data Processing

All preprocessing was done using the Python Time Series Analysis (ptsa) library (<https://github.com/compmem/ptsa>). Voltage measurements from each electrode were first re-referenced to the average of the right and left mathematically-linked mastoid electrodes at each time-point then high-pass filtered at 0.5 Hz. Eye-blinks and minor motion artifacts were corrected using the wavelet-enhanced independent components analysis algorithm (Castellanos & Makarov, 2006). Following eye blink correction, the voltage data from the test phase were segmented into events using a 1.25 s window starting 0.25 s prior to the onset of each stimulus. Any epoch that exhibited a kurtosis value greater than 5.0 or a voltage range exceeding 100 micro volts was removed from the analysis. Event-related potentials (ERPs) were calculated separately for each channel and down-sampled to 100 Hz for subsequent analyses.

The eyeblink correction algorithm cleans eyeblinks from the EEG signal quite well, but it does not remove artifacts due to horizontal eye-movements. Given that we provided a response mapping cue to participants throughout the testing period of Experiment 1, to ensure that systematic eye-movements did not contaminate our old/new classification we performed a within-subject counterbalancing. Specifically, we doubled the number of trials used to train the logistic regression classifier to discriminate old and new items, including first the actual data and second the same event data with the neural activity mirrored horizontally. Thus, the classifier would be unable to harness any lateral differences in voltage to guide the old/new classification, with the side-effect that only laterally symmetric signals will remain. This side-effect is acceptable in the present application because the standard evoked potentials of interest (the FN400 and late positive component) fall primarily on the mid-line electrodes.

One of the primary motivations for Experiment 2 was to replicate the results of the first experiment while significantly reducing the effect of eye-movements. Here we did not provide a response-mapping reminder during the test lists and we further counterbalanced

the response mapping across subjects. Thus, eye movements were minimized and not systematic across participants and it was not necessary to perform the within-subject counterbalancing we performed for the first experiment.

Diffusion Model

The diffusion model is a model of the decision process. The model is fit to accuracy and response time (RT) distributions for correct and error responses yielding measures of distinct components of processing including estimates of the strength of evidence entering the decision process, the amount of evidence necessary to make a decision, and the duration of other components of processing (encoding, memory access, and response output).

Figure 1A illustrates the decision process: Evidence is accumulated from a starting point z toward one of two criteria, or boundaries, a and 0 . A response is initiated when a boundary is reached. RTs and accuracy are naturally integrated by the model: RTs are determined by the time it takes for accumulated evidence to reach one of the boundaries plus nondecision time, and the choice is determined by which boundary is reached.

The values of drift rate, the boundaries, and the nondecision parameter are assumed to vary from trial to trial. This means that subjects do not set identical parameter values from trial to trial, and the assumption is required to fit the difference in RTs between correct and error responses (e.g., Laming, 1968; Ratcliff, 1978). Across-trial variability in drift rate is assumed to be normally distributed with standard deviation η , across-trial variability in the starting point (equivalent to across-trial variability in the boundary positions) is assumed to be uniformly distributed with range s_z , and across-trial variability in the nondecision component is assumed to be uniformly distributed with range s_t . These distributional assumptions are not critical: if modestly different distributions are assumed, the predictions of the model are similar (Ratcliff, 2013).

Across-trial variability in drift rate represents variability in stimulus information across trials that come from the same condition. It is exactly equivalent to variability in stimulus strength in signal detection theory. In some diffusion-process models (Palmer, Huk & Shadlen, 2005; Shadlen & Kiani, 2013; Usher & McClelland, 2001), it is argued that there is no across trial variability in drift rate. Our analysis will test the assumption that drift rate varies on a trial-by-trial basis and determine whether the single-trial EEG component is related to trial-level changes in drift rate.

Diffusion Model Fitting

For the chi-square method (Ratcliff & Childers, 2015; Ratcliff & Tuerlinckx, 2002), the chi-square value is minimized using the SIMPLEX minimization routine (Nelder & Mead, 1965), typically with RTs divided into 5 quantiles. The data entered into the minimization routine for each experimental condition are the .1, .3, .5, .7, .9 quantile RTs for correct and error responses and the corresponding accuracy values. The quantile RTs and parameter values of the model are used to generate the predicted cumulative probabilities of a response by that quantile RT. Subtracting the cumulative probabilities for each successive quantile from the next higher quantile gives the proportion of responses between adjacent quantiles. For the chi-square computation, these are the expected values, to be compared to the

observed proportions of responses between the quantiles (i.e., the proportions between 0, .1, .3, .5, .7, .9, and 1.0 are .1, .2, .2, .2, .2, and .1). These proportions are multiplied by the number of observations to give the expected frequencies, and summing over $(\text{Observed} - \text{Expected})^2 / \text{Expected}$ for all conditions gives a single chi-square value to be minimized. The number of degrees of freedom are the number of experimental conditions multiplied by 11 (12 bins between and outside the quantiles for correct and error responses minus 1 because the total probability sums to 1) minus the number of degrees of freedom.

For the maximum likelihood method (MLH method, Ratcliff & Tuerlinckx, 2002), the predicted probability density, $f(t_i)$ for each RT, t_i , for each correct and error response is computed and the product over all densities for all the RTs t_i is the likelihood, $L = \prod f(t_i)$. To obtain the maximum likelihood parameter estimates, the value of the likelihood is maximized by adjusting parameter values using a function minimization routine. However, because products of densities can become very large or very small, numerical problems occur and so the standard approach is to maximize the log likelihood, i.e., the sum of the logs of the densities (summing logs of the values is the same as the log of the product of the values, $\log(ab) = \log(a) + \log(b)$). Summing the logs of the predicted probability densities for all the RTs gives the log likelihood and minimizing minus the log likelihood produces the same parameter values as maximizing the likelihood.

Minus the log likelihood or chi-square can be minimized using a variety of software routines and we use the robust SIMPLEX routine (Nelder & Mead, 1965). This routine takes starting values for each parameter, calculates the value of the function to be minimized, then changes the values of the parameters (usually one at a time) to reduce minus the log likelihood. This process is repeated until either the parameters do not change from one iteration to the next by more than some small amount or the value to be minimized does not change by more than some small amount.

Single Trial Regressor Analysis

Philiastides et al. (2006) used a single-trial regressor analysis applied to a face/car discrimination task with degraded, briefly presented stimuli to examine the time course of perceptual processing. They used the method to find an optimal weighting of signals from 60 electrodes for discriminating between the two conditions for signals averaged over a small temporal window (Parra, Alvino, et al., 2002; Parra, Spence, Gerson, & Sajda, 2005). They performed this analysis over a series of 60 ms wide time windows following a poststimulus onset time of 200 ms. The resulting EEG component, $y(t)$, provided a single value that represented how strongly the stimulus represented a face or car.

A single trial regressor, $y(t)$, computed at a time window, t , is:

$$y(t) = \mathbf{w}^T \mathbf{x}(t) = \sum_{i=1}^D w_i x_i(t)$$

Where D =number of sensors. $\mathbf{w}$ =vector of weights for the sensors. $\mathbf{x}$ =vector of sensor values.

The key is to learn the vector of weights that best discriminates the two types of stimuli.

The method assumes that the data are distributed as a logistic function:

$$p(c=\text{old}|x) = \frac{1}{1+e^{-w^T x}} = f(y)$$

and $p(c = \text{new} | \mathbf{x}) = 1 - f(y)$.

Then, the weight vector, $\mathbf{w}$, is adjusted to minimize minus log likelihood of the data:

$$-\sum_j \log(p(c_j|x))$$

Philiastides et al. (2006) found different weightings of the electrodes that significantly discriminated between face and car stimuli, one that peaked at around 180 ms and another that peaked at around 380 ms. A critical feature of this analysis for all the later analyses and conclusions is that the regressor is derived from the electrode voltages as a function of the nature of the stimulus (i.e., whether the test item was a car or a face) and is not based on behavioral data (i.e., the choice made or response time). Below we present a similar analysis for our memory data.

Results

We begin with presentation of behavioral results followed by the results from the regressor analysis. Then we divide the behavioral data in halves based on the size of the single-trial regressor and fit the diffusion model as in Ratcliff et al. (2009). We then present a new diffusion model analysis in which drift rate is assumed to be a linear function of the single-trial regressor. In this analysis, we estimate how much drift rate changes as a function of the EEG regressor value. Finally, we present an analysis of the size of drift rates, regressor, and memory strength derived from signal detection theory for hits (H), false alarms (FA), misses (M), and correct rejections (CR).

For Experiment 1, accuracy for studied old items was 0.72 and for new items was 0.85 and for Experiment 2, accuracy for studied old items was 0.78 and for new items was 0.84. For Experiment 1, mean correct RT for old items was 811 ms and for new items was 852 ms and for Experiment 2, mean correct RT for old items was 819 ms and for new items was 902 ms. Full results are presented in Table 1. These are typical of results in the memory literature. Figure 1, Panels B and C, show plots of RT quantiles against response proportions for the conditions of the experiment averaged over subjects. The positions on the x-axis represent the proportion of responses and the values on the y-axis represent quantiles (the middle value of each column is the median RT). The lines are fits of the diffusion model that we discuss later.

Figures 2A and 3A show nine sample ERP plots for parietal, central, and frontal electrodes for median activity averaged over subjects for the two experiments. There are signals in the

range from 350–400 ms and up that discriminate between old and new test items, encompassing both the traditional FN400 and late-positive components often observed in EEG studies of recognition memory. Figure 4A shows the differences in voltages between “old” and “new” test items for frontal electrodes (Fz, F1, F2, AFz, FCz) and parietal electrodes (Pz, P1, P2, CPz, POz) as a function of time. The error bars are the same for frontal and parietal values (1 SE values are 0.28 and 0.29 respectively). A primary question is whether the frontal signal is significant at the 375 ms time point. A t-test on the data from both experiments combined shows it is ($t(27)=3.4$, $p<.05$).

An analysis of variance on the data from both experiments combined showed a significant difference between frontal and parietal electrodes ($F(1,27)=38.4$, $p<.05$), and significant effect of time ($F(3,81)=6.2$, $p<.05$), and a significant interaction ($F(3,81)=9.6$, $p<.05$). The pattern showed little difference between frontal and parietal electrodes at 375 ms (the range was 350–400 ms) with the parietal difference increasing with time. This shows a robust frontal voltage difference, which means that any lack of frontal effects in later analyses is not due to lack of a frontal signal.

Figures 2B and 3B show scalp maps of the average activity for old items minus the activity for new items from 200 to 800 ms in steps of 100 ms following presentation of the test item (averaged over trials and subjects). Both experiments show parietal activation most strongly in the 500–600 ms range. Experiment 1 shows frontal suppression of old relative to new items in the 500 ms window.

Note that the data are symmetric in both Figure 2A and 2B in Experiment 1 because of the within-subject counterbalancing we performed to ensure no unwanted effects of horizontal eye-movements. To verify that this within-subject counterbalancing did not greatly affect the underlying signal used to classify old vs. new we compared the ERPs with the symmetric data to data with horizontal eye-movement removed via the Gratton et al. (1983) regression method. The ERPs were nearly identical to the symmetric counterbalancing approach presented here and the classification and split-half model results presented below were also the same. Experiment 2 did not have the reminder on the screen and so there was no cue that might have induced eye-movements, thus the counterbalancing was not performed on Experiment 2.

In the next stage of the analyses, single-trial regressors are computed and these are related to the diffusion model analysis as in Ratcliff et al. (2009). The results of this analysis will, first, extend the methodology into the memory domain. If by sorting data into halves based on the old/new regressor, it is found that drift rates differ across the halves, then this is strong evidence that there is systematic variation in drift rate across trials and strong evidence that this EEG component represents evidence being used in the decision process. On the other hand, if an EEG component does not sort data, then the signal does not represent information used in the decision process. This has serious implications for the interpretation of EEG results in the memory domain because it means that the EEG signal does not directly track information used in the decision process.

For both experiments we performed a single-trial analysis of the EEG signal from all the electrodes to distinguish whether a test word was old or new (studied or not studied). Logistic regression was used to find an optimal weighting of signals from electrodes for a series of training windows in 50 ms intervals with 50 ms widths following stimulus onset time. Then for each trial, we constructed a histogram of the regressor value for old and new items and used these to generate a receiver operator characteristic (ROC) curve. The area under this curve, A_z , was used as a measure of “old”/“new” discriminability from the EEG data as in Figure 5A. Note that it is difficult to produce a scalp map of the weighting on electrodes that discriminate between old and new because the weighting is based on the size of the signal and the reliability of the signal (including the variability).

Plots of the A_z value of the regressor for the experiment are shown in Figure 5B. The 95% significance level (the horizontal lines at $A_z=0.62$ and $A_z=0.71$ and the top dotted lines) were obtained by randomizing the assignment of conditions to EEG activity then performing logistic regression. This is necessary because the method will find some weighting on electrodes to produce discrimination between old and new test items because of random variability even if there is no significant signal. The randomization allows estimation of the size of this effect, for example, the bottom dotted lines are the median A_z values. The 5% significance levels for A_z were the average values for which 95% of the maximum values of A_z were lower. The regressor values are significant over the range 400–800 ms (e.g., the regressor at 375 ms was significant by t-test, for Experiment 1, $t(16)=5.7$ and for Experiment 2, $t(10)=9.7$, both $p<.05$), but the more convincing result is that the regressor is significant for Experiment 1 by the leave-one-out analysis). The significant regressor, even at the earliest range, is further indication that there is a significant difference in neural activity between old and new items, but as we will see below, this difference in neural activity does not impact the decision process.

For the first diffusion model analysis, we fit the model to each of the six conditions for Experiment 1 (high and low frequency words studied 3 times, once, and not studied) and the three conditions for Experiment 2 (words studied 3 times, once, and not studied) for each subject using the chi-square method (Ratcliff & Tuerlinckx, 2002). Mean parameter values are shown in Tables 2 and 3. The model fit the data as well as in most other applications (Ratcliff & Childers, 2015, p. 253) with the mean value of chi-square a little larger than the critical value of 71.0 ($df=53$) for Experiment 1 (82.1) and a little smaller than the critical value of 35.2 ($df=23$) for Experiment 2 (28.3).

For the second analysis, for each condition of the experiment and each subject, the data were divided into halves based on the value of the regressor at about 600 ms (at the time 10% of responses had terminated - the window was from 575–625 ms). We took the regressor values for each of the six conditions for Experiment 1 and three conditions for Experiment 2 and divided the behavioral data for each condition into a more “old” and a less “old” half.

We first fit the diffusion model to each condition for the more “old” half of the data and to each condition for the less “old” half of the data with separate parameters for each half of the data. The only parameters that were significantly different were drift rates shown in Table 3. For the next analysis, we fit the model to the data from each half simultaneously

with only drift rates allowed to differ between halves. Tables 2 and 3 show the parameter values and Figure 5C shows plots of the drift rates for 17 individual subjects for the 6 conditions for the old-half versus the new-half of the data for Experiment 1 and the 11 individual subjects for the 3 conditions for the old-half versus the new-half of the data for Experiment 2. Results for both experiments show most of the drift rates for the old-half are more positive than the drift rates for the new-half so that the points lie above the diagonal line. For Experiment 1, an ANOVA showed that drift rates were significantly different across the six conditions ($F(5,80)=69.7$, $p<.05$), drift rates for the old-half were significantly different from those for the new-half ($F(1,16)=22.1$, $p<.05$), and the interaction was significant ($F(5,80)=4.0$, $p<.05$). For Experiment 2, an ANOVA showed that drift rates were significantly different across the three conditions ($F(2,20)=61.7$, $p<.05$), drift rates for the old-half were significantly different from those for the new-half ($F(1,10)=25.8$, $p<.05$), and the interaction was significant ($F(2,20)=34.0$, $p<.05$). The interactions show that the difference between the old and new-half were different across the different conditions (reflecting a smaller difference for new items than old items as shown in Figure 5C) and post hoc tests showed that the differences between drift rates for the new items were not significant (values of 0.015 and -0.025 with $t(80)=0.07$ and $t(20)=0.03$ for Experiments 1 and 2 respectively). Also, the differences in drift rates for old items across halves was quite large (0.115 and 0.138 for Experiments 1 and 2 respectively) relative to the absolute values of the drift rates in Table 3.

For comparison, we performed a similar analysis for the component at 375 ms. Figures 2 and 3 show that there was separation of the ERPs for old versus new items at this time for the electrodes shown. Fitting the data to the two halves separated on the basis of the regressor value produced no difference in drift rates. Figure 5C (right panels) shows that the drift rates for the old-half versus the new-half of the data lie close to the diagonal. In an ANOVA for Experiment 1, drift rates were significantly different across the three conditions ($F(5,80)=70.0$, $p<.05$), but the drift rates for the old-half and the new-half were not significantly different ($F(1,16)=0.6$) and neither was the interaction ($F(5,80)=1.0$). In an ANOVA for Experiment 2, drift rates were significantly different across the three conditions ($F(2,20)=45.4$, $p<.05$), but the drift rates for the old-half and the new-half were not significantly different ($F(1,10)=0.2$) and neither was the interaction ($F(2,20)=0.5$). Because there were no significant effects at 375 ms, this analysis provides a control for the analysis at the 600 ms time point and shows it is not spurious.

For completeness, we also performed these analyses at two more intermediate points, at 450 ms and 525 ms. Figure 4B shows plots of the difference in drift rates for the data sorted by the regressors over time. The plot shows increasing functions over time to peaks at 600 ms.

In addition to the analyses above (which followed those in Philiastides et al. 2006 and Ratcliff et al. 2009), we also performed regressor analyses on correct versus error responses, fast versus slow responses, and beta and low gamma frequency bands at the peak of the regressor in Figure 5B for Experiment 1. None of these showed significant sorting of the behavioral responses.

We can now extend our analysis by using the single-trial regressor value at 600 ms in the diffusion model analysis directly. The assumption is that drift rate for each individual trial is a constant (baseline value, one for each condition) plus a coefficient (slope) multiplied by the regressor value (normalized to a mean of 0 and SD of 1 for each condition separately), that is a linear model with the drift rate a linear function of the regressor. It is also assumed that there is random trial to trial variability in drift rate because the regressor value is noisy and not an exact measure of evidence). If the regressor had no influence on drift rate (through the RT and choice), then the slope would be zero. In the model that we fit, we allowed drift rate, across trial variability in drift rate, and the regressor coefficient to differ between the old and new conditions.

In order to perform model fitting, a maximum likelihood method (Ratcliff & Tuerlinckx, 2002) was used (as opposed to the quantile method) in which each individual choice and RT is used in fitting. In this method, the likelihood or probability density of each response choice and RT is computed and the model parameters adjusted until the product of the likelihoods is maximized (or, equivalently, the sum of minus the log likelihoods minimized). Tables 2 and 3 show the parameter values recovered from this method and we present t-values for Experiments 1 and 2 (respectively). The results showed that, first the coefficients of the regressor for old items are significantly different from zero ($t(16)=5.29$, $p<.05$ and $t(10)=2.86$, $p<.05$), and for new items they are not ($t(16)=1.14$ and $t(10)=0.44$). Also, the two coefficients differed from each other ($t(16)=4.78$, $p<.05$ and $t(10)=2.72$, $p<.05$) and across trial variability in drift rate differed between old and new items for Experiment 1 but not Experiment 2 ($t(16)=4.52$, $p<.05$ and $t(10)=1.58$), see Starns and Ratcliff (2014). These results show that the regressor captures a proportion of trial-by-trial variability in drift rate.

If the value of the regressor on each trial is a (noisy) measure of the drift rate on that trial, then it is possible to generate simulated data from the diffusion model while recording the drift rate to each trial and examine how well the mean drift rates for hits, H, misses, M, false alarms, FA, and correct rejections, CR for each of the conditions of the experiment match the regressor values. We do this for cases with predictions from the diffusion model and signal detection theory. For the diffusion model there are different possibilities with regard to how much of the variability in drift rate is measured by the regressor (we term this “systematic” variability). We consider cases in which none and all are systematic and cases in which part of the variability is systematic.

For signal detection theory (Figure 6A), for predictions for hits, the mean of the normal distribution above the criterion can be computed by integration over the distribution from the criterion up (similar integrations can be carried out to find the means for M, FA, and CR). For the diffusion model for the example in Figure 6B, 2000 simulated RTs and choices were generated using the random walk method in Tuerlinckx, Maris, Ratcliff & De Boeck (2001). The parameters used were from the top row of Table 2 with drift rate 0.101, i.e., for once presented high frequency words. For every simulated decision, the choice, the RT, and the drift rate for that trial (randomly selected from the drift rate distribution with mean v and SD η) were recorded. If the EEG component was an accurate estimate of drift rate on each trial, then the drift rate would match the EEG component value (of course, in practice, there is variability in the EEG component value). Figure 6B shows distributions of hits and misses as

a function of drift rate. From the “old” and “new” responses in this histogram, the hit rate and miss rate can be computed. Using this analysis for distributions for each condition, the mean drift rate was computed for H, M, FA, and CR. For the diffusion model with no systematic across trial variability in drift rate, we used the values of the mean drift rates (but simulated data were generated with trial to trial variability in drift rate). In the cases in which across trial variability is composed of two parts, one is systematic and the other is random. In this case (Figure 6E), the SD in drift across trials estimated from fits to the data is divided into two values and one of these is used to generate drift rates and is saved for the plot in Figure 6E. To each of these drift rates, another random value is added using the other SD and this new drift rate is used to generate simulated data that is used to generate predictions in Figure 6E. The parameter values for all the diffusion model predictions are from fits to all the data combined in Tables 2 and 3.

Figure 6C–F shows plot of the regressor value on the y-axis versus the predicted values from SDT and three diffusion models with different assumptions about across trial variability in drift rate. The regressor values show an ordering of CR, FA, M and H with values for FAs and M close to each other for Experiment 1 and with the values for FA closer to those for CR for Experiment 2. For Figure 6C, SDT produces the ordering in Figure 6A, namely, CR, M, FA, and H. The plot shows that FA are predicted to have a higher value than the regressor values. The diagonal line is a linear regression.

For the diffusion model with no across trial variability in drift rate, the drift rates for H and M are identical (because they both arise from random variability in the decision process) and likewise for FA and CR. This can be seen in Figure 6D because FA and CR line up vertically as do H and M. This produces a mismatch between the predictions and the regressor values.

For the diffusion model with across trial variability in drift rate all systematic, the match between the regressor and drift rates are good, but for Experiment 2, there is a mismatch for FA. Figure 6E shows plots for 2/3 of across trial variability in drift rate systematic and this produces the best fits for between drift rates and the regressor values for both experiments (we also generated values for 1/3 and 1/2 of the variability systematic and the value 2/3 fit best).

This analysis provides additional evidence for the mapping between the single trial regressor from the single-trial analysis of EEG data and between the diffusion model with the assumption that drift rate varies from trial to trial.

Discussion

Our results provide strong evidence for an EEG component that reflects evidence used in the decision process. This component is centered on a parietal location and peaks around 600 ms following stimulus presentation. The single-trial regressor computed at this time influences drift rate in a diffusion model analysis for old items. It can be used to sort data from each old item condition into groups representing more “old” and less “old” and diffusion model fits show that drift rates differ for those groups. Another analysis assumes that drift rate on each trial is a linear function of the regressor value and estimates of the slope are different from

zero. A comparison between the regressor values for hits, misses, false alarms and correct rejections for diffusion model and for signal detection theory predictions shows that diffusion model predictions match best.

The EEG signals in an ERP analysis showed initial separation between old and new items a little before 400 ms (Figure 2A). The observed frontal signal around 400 ms and the later parietal signal at around 600 ms have been interpreted in prior studies as evidence for different components used in making recognition decisions, an early familiarity component indexed by the frontal signal and a later recollection component indexed by the parietal signal. If the two components were involved in decisions, then single-trial regressors at both time points should sort data and produce differences in drift rate. However, our analyses showed that the signal at 375 ms produces a significant regressor (both regressors are significant by t-test, but the regressor is significant by the leave-one-out analysis only for Experiment 1) and sorting data based on the best regressor computed at this time does not affect drift rate (in contrast to the signal at 600 ms). This is evidence that EEG activity at the 375 ms time point does not represent information being used in the decision process (cf., the early perceptual component studied in Ratcliff et al., 2009), and only the later component represents information used in the decision process. The early component reflects global old versus new discrimination because of the training, but it does not map into whether an individual item is more “old” or more “new” in the decision process. The late component result is also consistent with a “mnemonic accumulator” hypothesis that proposes a link between activity in the posterior parietal cortex and memory retrieval (Cabeza, Ciaramelli, Olson & Moscovitch, 2008; Konishi, Wheeler, Donaldson & Buckner, 2000; Wagner, Shannon, Kahn & Buckner, 2005).

The results also show that the regressor begins to peak around 550 ms. This places it near the beginning of the diffusion decision process as indicated by the estimated non-decision time from the diffusion model fits (around 600 ms). This is consistent with the view that the regressor indexes drift rate, that is, evidence driving the decision process.

These results have two major implications for recognition memory. First, they show that early EEG signals that have been associated with familiarity do not index the decision process, that is, the signal does not indicate or predict how strong or weak an item is relative to other items in that condition. In contrast, the signal at 600 ms does indicate how strong or weak an item is. This suggests that results that show differences in EEG signals as a function of experimental variables do not provide evidence that these signals reflect the decision process. This is especially important in recognition memory because EEG data have been used as one of the major pieces of empirical support for dual process theory.

The second major implication for recognition memory is that there is greater across trial variability in drift rate for old items than for new items. The usual way to examine the relative variability in the memory strength of old and new items is to use ROC functions. Starns and Ratcliff (2014) analyzed a large number of subjects (376) from a number of experiments and found that across trial variability in drift rate was larger for old items than for new items. This shows that a diffusion model analysis can replace ROC analyses and use simple two choice tasks to perform the same measurements. If we combine the results for

the two experiments, values are 0.258 for old items and 0.165 for new items and the difference is significant ($t=5.9$, $p<.05$). This result replicates and is consistent with the results from Starns and Ratcliff (2014).

There is an important methodological implication of this analysis and that is that a linear model for drift rates can be embedded in the diffusion model. This provides a way of measuring the effect of the EEG regressor on the decision process.

The single process view assumes a signal detection representation in which the old item distribution has larger standard deviation than the new item distribution. In the diffusion model, this corresponds to a larger value of across trial variability in drift rate for old items relative to new items (see Starns & Ratcliff, 2014). In our analysis, a larger slope of drift rate as a function of the regressor for old relative to new items further supports this single process view.

Similar results distinguishing two potentially distinct processes have been obtained in the perceptual domain. Philastides et al. (2006) showed that there are two EEG components that are significant, one at around 180 ms and one at 380 ms in a face/car perceptual discrimination task. Dividing the data into halves based on the late component produced different values of drift rate, but dividing the data based on the early component produced no difference in drift rates or any model parameter (Ratcliff et al., 2009). This was interpreted as showing that the early component tracks encoding during which a stimulus representation is constructed, and the late component represents decision-related information extracted from the stimulus representation. Whereas the stimulus representation has many attributes that are not necessarily relevant to the decision, for example, color, size, shape, and so on, only some of these are relevant to the decision. Consequently, part of the nondecision time involves extracting the decision-related information from the stimulus before the decision process can begin.

The analyses presented here show that the regressor predicts trial-level fluctuations in drift rate. Originally, across trial variability was introduced into the diffusion model because in application to memory, it is difficult to see how items in different lists are encoded to exactly the same degree (Ratcliff, 1978). This also produces slow errors and later it was found to produce good fits to experimental data (Ratcliff, Van Zandt, & McKoon, 1999). However, it has been claimed that this is not correct, that there is no such across-trial variability in drift rates and models have been developed to account for data using collapsing bounds without trial to trial variability in drift rate (Churchland et al., 2008; Ditterich, 2006a, 2006b; Drugowitsch et al., 2012; Kiani et al., 2014; Palmer et al., 2005; Zhang et al., 2014). Hawkins et al. (2015) performed a comprehensive study fitting a flexible collapsing boundary model and the standard model and found that for human subjects, the standard model with across trial variability in drift rate was preferred. The results presented here and in Ratcliff et al. (2009) provide strong independent evidence for the hypothesis that drift rate varies from trial to trial.

There is also a recent claim in the psychological literature by Jones and Dhazafarov (2014) who argued that if the forms of the across-trial variability distributions were unconstrained,

the model could exactly match any data, that is, any accuracy values and any RT distributions, hence rendering the models unfalsifiable. However, Smith et al. (2014) pointed out (among other things) that their argument is predicated on eliminating within-trial noise from the accumulation process. Within-trial noise in the accumulation process is a fundamental assumption of the model, and it is within-trial noise that leads to most errors and most variability in RTs. The Jones and Dhazafarov argument boils down to a mapping between velocity, distance, and time in which each RT is transformed to a constant times $1/RT$ which serves as a drift rate. Our analysis shows a consistent relationship between across trial variability in drift rate, an independent EEG measure of trial to trial variability, and the time course of this effect in decision making in recognition memory.

The combination of the single-trial regressor and behavioral data avoids the multiple comparison problems that can affect analyses of neural data. The single-trial regressor is computed based on electrical signals and whether the test item was studied or not and this analysis is completely independent of the behavioral data. The single-trial regressors are applied to behavioral data with no modification based on the results of the behavioral analysis. The fact that the analysis produces a robust effect on drift rates at around 600 ms but no effect at 375 ms shows that the effect at 600 ms is not spurious or caused in some way by overfitting.

Experimental approaches in the human neuroscience of memory make the strong assumption that EEG (and fMRI) signals index evidence that is being used to make a decision. However, as Philiastides et al. (2006 and Ratcliff et al., 2009) demonstrated and as our results show, simply finding differential activity does not guarantee that it represents information being used in making the decision. Philiastides et al. colored the stimuli red and green and made the task red/green discrimination and found that the early 180ms component was unaffected but the late component disappeared. Ratcliff et al. showed that the late 380 ms component sorted the data but the 180 ms component did not. These results were interpreted as showing that the 180 ms component reflected perceptual encoding of all aspects of the stimulus and the 380 ms component reflected the decision relevant aspects of the stimulus extracted from the encoded representation.

More generally, these examples serve to illustrate a pervasive problem in studies of cognition: neural and behavioral responses are not direct measures of cognitive processes. It could be that a neural response represents the results of earlier processing that is then transformed to later processing. But it is also possible that processing produces several possible output streams, only some of which are used in any particular task. The outputs not used might produce a strong neural response (face vs. car stimuli in Philiastides et al.), but that encoding might have little to do with the decision process (red vs. green discrimination). We believe that methods like this one that allows item level effects to be examined provide more convincing evidence that the neural response is involved in the processing stream giving rise to the decision.

Cognitive processes are latent variables that give rise to both the observed neural activity and behavior. While the goal of cognitive neuroscience is to map neural activity to these underlying mechanisms, we believe that traditional means of analyzing neural data that map

neural activity to behavior are greatly enhanced by explicit theory of this mapping. Recently, there has been a push to incorporate computational models that explicitly represent the cognitive mechanisms into the analysis of neural data. Such models provide a theory of the link between the observed neural activity and the cognitive processes used in the task (e.g., Gold & Shadlen, 2001; Purcell, Schall, Logan & Palmeri, 2012; Ratcliff, Cherian & Segraves, 2003; Ratcliff, Philiastides & Sadja, 2006; Turner, Forstmann, et al., 2013).

In our view, the present memory results complement the perceptual results from Ratcliff et al. (2009) by showing a (nonsignificant) component at about 375 ms that does not map into the decision process and a later component around 600 ms does map into the decision process. The strongest conclusion from this research, in both the perceptual and memory domains, is that it is necessary to show that an EEG component has behavioral consequences on a trial-to-trial basis to conclude that the component represents processing in making a decision.

Acknowledgments

This article was supported by NIA Grant R01-AG041176, AFOSR Grant FA9550-11-1-0130, and NIMH grant R01-MH085092.

References

- Blankertz B, Lemm S, Treder M, Haufe S, Muller K-R. Single-trial analysis and classification of ERP components - A tutorial. *NeuroImage*. 2011; 56:814–825. [PubMed: 20600976]
- Brysbaert M, New B. Moving beyond Kucera and Francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior Research Methods*. 2009; 41:977–990. [PubMed: 19897807]
- Buchler NG, Light LL, Reder LM. Memory for items and associations: Distinct representations and processes in associative recognition. *Journal of Memory and Language*. 2008; 59:183–199. [PubMed: 18677424]
- Cabeza R, Ciaramelli E, Olson IR, Moscovitch M. The parietal cortex and episodic memory: An attentional account. *Nature Reviews Neuroscience*. 2008; 9:613–625. [PubMed: 18641668]
- Castellanos NP, Makarov VA. Recovering EEG brain signals: Artifact suppression with wavelet enhanced independent component analysis. *Journal of Neuroscience Methods*. 2006; 158:300–312. [PubMed: 16828877]
- Churchland AK, Kiani R, Shadlen MN. Decision-making with multiple alternatives. *Nature Neuroscience*. 2008; 11:693–702. [PubMed: 18488024]
- Cohen AL, Rotello CM, Macmillan NA. Evaluating models of remember-Know judgments: Complexity, mimicry, and discriminability. *Psychonomic Bulletin & Review*. 2008; 15:906–926. [PubMed: 18926982]
- Dennis S, Humphreys MS. A context noise model of episodic word recognition. *Psychological Review*. 2001; 108:452–477. [PubMed: 11381837]
- Ditterich J. Computational approaches to visual decision making. In: Chadwick, DJ.; Diamond, M.; Goode, J., editors. *Percept, decision, action: Bridging the gaps*. Chichester, UK: Wiley; 2006a. p. 114
- Ditterich J. Stochastic models of decisions about motion direction: Behavior and physiology. *Neural Networks*. 2006b; 19:981–1012. [PubMed: 16952441]
- Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, Pouget A. The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience*. 2012; 32:3612–3628. [PubMed: 22423085]

- Dunn JC. Remember-Know: A Matter of Confidence. *Psychological Review*. 2004; 111:524–542. [PubMed: 15065921]
- Eichenbaum H, Yonelinas AP, Ranganath C. The medial temporal lobe and recognition memory. *Annual Review of Neuroscience*. 2007; 30:123–152.
- Geller AS, Schleifer IK, Sederberg PB, Jacobs J, Kahana MJ. PyEPL: A cross-platform experiment-programming library. *Behavior Research Methods*. 2007; 39:950–958. [PubMed: 18183912]
- Gillund G, Shiffrin RM. A retrieval model for both recognition and recall. *Psychological Review*. 1984; 91:1–67. [PubMed: 6571421]
- Gold JI, Shadlen MN. Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Science*. 2001; 5:10–16.
- Gratton G, Coles MGH, Donchin E. A new method for off-line removal of ocular artifact. *Electroencephalography and Clinical Neurophysiology*. 1983; 55:468–484. [PubMed: 6187540]
- Hawkins GE, Forstmann BU, Wagenmakers E-J, Ratcliff R, Brown SD. Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience*. 2015; 35:2476–2484. [PubMed: 25673842]
- Hintzman DL. MINERVA 2: A simulation model of human memory. *Behavior Research Methods, Instruments, & Computers*. 1984; 16:96–101.
- Jones M, Dzhafarov EN. Unfalsifiability and mutual translatability of major modeling schemes for choice reaction time. *Psychological Review*. 2014; 121:1–32. [PubMed: 24079307]
- Kiani R, Corthell L, Shadlen MN. Choice certainty is informed by both evidence and decision time. *Neuron*. 2014; 84:1329–1342. [PubMed: 25521381]
- Konishi S, Wheeler ME, Donaldson DI, Buckner RL. Neural correlates of episodic retrieval success. *Neuroimage*. 2000; 12:276–286. [PubMed: 10944410]
- Laming, DRJ. *Information theory of choice reaction time*. New York: Wiley; 1968.
- Nelder JA, Mead R. A simplex method for function minimization. *Computer Journal*. 1965; 7:308–313.
- Paller KA, Voss JL, Boehm SG. Validating neural correlates of familiarity. *Trends in Cognitive Sciences*. 2007; 11:243–250. [PubMed: 17475539]
- Palmer J, Huk AC, Shadlen MN. The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*. 2005; 5:376–404. [PubMed: 16097871]
- Parra L, Alvino C, Tang A, Pearlmutter B, Young N, Osman A, Sajda P. Linear spatial integration for single-trial detection in encephalography. *NeuroImage*. 2002; 17:223–230. [PubMed: 12482079]
- Parra L, Spence C, Gerson A, Sajda P. Recipes for the linear analysis of EEG. *NeuroImage*. 2005; 28:326–341. [PubMed: 16084117]
- Pereira F, Mitchell T, Botvinick M. Machine learning classifiers and fMRI: a tutorial overview. *NeuroImage*. 2009; 45:S199–S209. [PubMed: 19070668]
- Philiastides MG, Ratcliff R, Sajda P. Neural representation of task difficulty and decision making during perceptual categorization: A timing diagram. *Journal of Neuroscience*. 2006; 26:8965–8975. [PubMed: 16943552]
- Purcell BA, Heitz RP, Cohen JY, Schall JD, Logan GD, Palmeri TJ. Neurally-constrained modeling of perceptual decision making. *Psychological Review*. 2010; 117:1113–1143. [PubMed: 20822291]
- Purcell BA, Schall JD, Logan GD, Palmeri TJ. From salience to saccades: multiple-alternative gated stochastic accumulator mode of visual search. *Journal of Neuroscience*. 2012; 32:3433–3446. [PubMed: 22399766]
- Ratcliff R. A theory of memory retrieval. *Psychological Review*. 1978; 85:59–108.
- Ratcliff R. Parameter variability and distributional assumptions in the diffusion model. *Psychological Review*. 2013; 120:281–292. [PubMed: 23148742]
- Ratcliff R, Cherian A, Segraves M. A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of simple two-choice decisions. *Journal of Neurophysiology*. 2003; 90:1392–1407. [PubMed: 12761282]
- Ratcliff R, Childers R. Individual differences and fitting methods for the two-choice diffusion model. *Decision*. 2015; 2:237–279.

- Ratcliff R, Philastides 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; 106:6539–6544.
- Ratcliff R, Thapar A, McKoon G. A diffusion model analysis of the effects of aging on recognition memory. *Journal of Memory and Language*. 2004; 50:408–424.
- Ratcliff R, Thapar A, McKoon G. Individual differences, aging, and IQ in two-choice tasks. *Cognitive Psychology*. 2010; 60:127–157. [PubMed: 19962693]
- Ratcliff R, Thapar A, McKoon G. Effects of aging and IQ on item and associative memory. *Journal of Experimental Psychology: General*. 2011; 140:46–487.
- Ratcliff R, Tuerlinckx F. Estimating the parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin and Review*. 2002; 9:438–481. [PubMed: 12412886]
- Ratcliff R, Van Zandt T, McKoon G. Connectionist and diffusion models of reaction time. *Psychological Review*. 1999; 106:261–300. [PubMed: 10378014]
- 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. [PubMed: 15250777]
- Rugg MD. Memory and consciousness: a selective review of issues and data. *Neuropsychologia*. 1995; 33:1131–1141. [PubMed: 7501134]
- Rugg MD, Curran T. Event-related potentials and recognition memory. *Trends in Cognitive Sciences*. 2007; 11:251–257. [PubMed: 17481940]
- Rugg MD, Yonelinas AP. Human recognition memory: A cognitive neuroscience perspective. *Trends in Cognitive Neuroscience*. 2003; 7:313–319.
- Shadlen MN, Kiani R. Decision making as a window on cognition. *Neuron*. 2013; 80:791–806. [PubMed: 24183028]
- Shiffrin RM, Steyvers M. A model for recognition memory: REM: Retrieving effectively from memory. *Psychonomic Bulletin and Review*. 1997; 4:145–166. [PubMed: 21331823]
- Smith PL, Ratcliff R, McKoon G. The diffusion model is not a deterministic growth model: Comment on Jones and Dhafarov (2013). *Psychological Review*. 2014; 121:679–688. [PubMed: 25347314]
- Starns JJ, Ratcliff R. Two dimensions are not better than one: STREAK and the univariate signal detection model of remember/know performance. *Journal of Memory and Language*. 2008; 59:169–182. [PubMed: 19649141]
- 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. [PubMed: 24459327]
- Starns JJ, Ratcliff R, McKoon G. Evaluating the unequal-variability and dual process explanations of zROC slopes with response time data and the diffusion model. *Cognitive Psychology*. 2012; 64:1–34. [PubMed: 22079870]
- Tuerlinckx F, Maris E, Ratcliff R, De Boeck P. A comparison of four methods for simulating the diffusion process. *Behavior, Research, Instruments, and Computers*. 2001; 33:443–456.
- Turner BM, Forstmann BU, Wagenmakers E-J, Brown SD, Sederberg PB, Steyvers M. A Bayesian framework for simultaneously modeling neural and behavioral data. *NeuroImage*. 2013; 72:193–206. [PubMed: 23370060]
- Usher M, McClelland JL. The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*. 2001; 108:550–592. [PubMed: 11488378]
- Wagner AD, Shannon BJ, Kahn I, Buckner RL. Parietal lobe contributions to episodic memory retrieval. *Trends in Cognitive Science*. 2005; 9:445–453.
- Wixted JT. Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*. 2007; 114:152–176. [PubMed: 17227185]
- Yonelinas AP. Receiver-operating characteristics in recognition memory: Evidence for a dual-process model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1994; 20:1341–1354.

- Yonelinas AP. Recognition memory ROCs for item and associative information: The contribution of recollection and familiarity. *Memory & Cognition*. 1997; 25:747–763. [PubMed: 9421560]
- Zhang S, Lee MD, Vandekerckhove J, Gunter M, Wagenmakers E-J. Time-varying boundaries for diffusion models of decision making and response time. *Frontiers in Quantitative Psychology and Measurement*. 2014; 5(1364)

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Highlights

- A single trial regressor analysis was conducted on EEG data from recognition memory
- The EEG regressor was trained by whether the test item was studied or not
- A diffusion decision model with drift rate a function of the regressor fit the data
- A regressor trained at 600 ms (but not one at 375 ms) indexed evidence accumulation rate
- An EEG difference does not guarantee that it reflects evidence used in the decision

Figure 1.

Panel A shows an illustration of the diffusion model. Panels B and C show quantile probability plots for "old" and "new" responses for Experiments 1 and 2 respectively. The o's represent the experimental data and the x's joined by lines are the model predictions. The conditions are shown on the x-axis in terms of proportions of responses (the orders can be derived from Tables 1 and 2). Proportions on the right are for correct responses and proportions on the left for error responses. The RT quantiles are, in order from bottom to

top, the .1, .3, .5, .7, and .9 quantiles. Predictions were generated for each subject and then averaged over subjects. The data were averaged over subjects in the same way.

Figure 2.

Panel A shows sample ERP's for old (solid line) and new (dashed line) test items averaged over subjects for Experiment 1. Panel B shows ERP differences between old and new test items, ERPs for old items, and ERPs for new items as a function of scalp location and as a function of time.

Figure 3.

Panel A shows sample ERP's for old (solid line) and new (dashed line) test items averaged over subjects for Experiment 2. Panel B shows ERP differences between old and new test items, ERPs for old items, and ERPs for new items as a function of scalp location and as a function of time.

Figure 4.

Panel A shows differences in voltage between "old" and "new" test words for frontal electrodes and for parietal electrodes as a function of time from stimulus onset. Panel B shows a plot of the difference in drift rates for data sorted by the regressor value as a function of time from the stimulus onset.

Figure 5.

Panel A shows the regressor values for old and new test items used to generate the area under the ROC curve, Az. Panel B shows plots of the single trial regressor Az values for Experiments 1 and 2. The solid line is the regressor, the bottom dotted line is median regressor value obtained from scrambling the trial labels and reassigning them to the EEG signals, and the top dotted line is the highest 5% of the regressors obtained from scrambling the trial labels. This latter value represents the 95% confidence limit on the regressor. The top left of Panel C shows plots of the drift rates for 17 individual subjects for each of the 6

conditions for the old-half versus the new-half of the data with the data divided into halves based on the regressor value at 600 ms for Experiment 1. L=low frequency words, H=high frequency words, 3=three times presented, 1=one time presented, and N=new items. The top right of Panel C shows plots of the drift rates for 11 individual subjects for each of the 3 conditions for the old-half versus the new-half of the data with the data divided into halves based on the regressor value at 600 ms for Experiment 2. The bottom plots of Panel the drift rates as in top panels but divided into halves based on a regressor at 375 ms.

Figure 6.

Panel A shows an illustration of signal detection theory and the mean values for hits, misses, false alarms, and correct rejections. Panel B shows a plot of drift rates for the two choices ("old" and "new") for 2000 simulated trials from the diffusion model (across trial variability in drift rate produces the different values of drift rate). Panel C shows plots of the regressor value against signal detection theory predictions for H, M, FA, and CR. Panel D shows plots of the regressor value against drift rates for H, M, FA, and CR from the diffusion model without across trial variability in drift rate. Panel E shows plots of the regressor value against

drift rates for H, M, FA, and CR from the diffusion model with 2/3 of the variability in drift rate systematically related to the regressor value. Panel F shows plots of the regressor value against drift rates for H, M, FA, and CR from the diffusion model variability in drift rate systematically related to the regressor value.

Table 1

Accuracy and Mean RT

Word frequency	Study	Old Pr	Old RT	New RT
Experiment 1 Low frequency	3 presentations	0.836	778	928
	1 presentation	0.670	825	886
	New	0.108	921	830
Experiment 1 High frequency	3 presentations	0.777	804	935
	1 presentation	0.584	855	916
	New	0.191	943	877
Experiment 2	3 presentations	0.872	781	1056
	1 presentation	0.690	856	986
	New	0.160	1018	902

Table 2

Parameter values for diffusion model fits

Expt	Analysis	EEG signal	a	z	T _{er} (ms)	η old	η new	s _z	s _t (ms)	a _v old	a _v new	χ ²
1	Fits to all data		0.143	0.067	570	0.232	0.161	0.052	211			82.1
	Separate fits for the 600ms regressor	newer	0.148	0.071	569	0.234	0.173	0.049	223			59.0
		older	0.144	0.067	564	0.218	0.150	0.045	187			52.7
	Separate drifts: 600ms regressor		0.152	0.070	568	0.222	0.195	0.067	214			158.7
	Separate drifts: 375ms regressor		0.144	0.067	567	0.203	0.185	0.056	229			157.2
	Separate drifts: 450ms regressor		0.147	0.069	574	0.220	0.066	0.180	227	0.180		161.7
	Separate drifts: 525ms regressor		0.145	0.068	572	0.222	0.053	0.195	229	0.195		174.2
	MLH: 600ms regressor		0.153	0.070	575	0.207	0.143	0.111	254	0.067	0.009	
	Fits to all data		0.170	0.088	518	0.299	0.170	0.034	152			28.3
	Separate fits for the 600ms regressor	newer	0.166	0.080	530	0.290	0.213	0.033	162			26.3
2		older	0.166	0.084	521	0.310	0.167	0.046	172			22.5
	Separate drifts: 600ms regressor		0.169	0.084	519	0.269	0.183	0.039	172			59.9
	Separate drifts: 375ms regressor		0.172	0.088	518	0.276	0.189	0.042	171			64.1
	Separate drifts: 450ms regressor		0.172	0.090	520	0.272	0.207	0.046	159			55.6
	Separate drifts: 525ms regressor		0.171	0.086	521	0.290	0.201	0.044	167			62.9
	MLH: 600ms regressor		0.149	0.069	512	0.160	0.125	0.046	201	0.033	0.003	

Note. a=boundary separation, z=starting point, T_{er}=nondecision component of response time, η=standard deviation in drift across trials, s_z=range of the distribution of starting point (z), s_t= range of the distribution of nondecision times, a_v=slope of drift rate as a function of the regressor values, and χ² is the chi-square goodness of fit statistic.

Table 3

Drift rates for diffusion model fits

Analysis: Experiment 1	EEG signal	v _{3L}	v _{1L}	v _{NL}	v _{3H}	v _{1H}	v _{NH}
Fits to all data							
Separate fits for the 600ms regressor	newer	0.291	0.113	-0.292	0.204	0.029	-0.183
	older	0.263	0.059	-0.307	0.187	0.031	-0.187
Separate drifts: 600ms regressor	newer	0.358	0.161	-0.254	0.292	0.102	-0.185
	older	0.237	0.059	-0.333	0.147	0.016	-0.196
Separate drifts: 375ms regressor	older	0.367	0.182	-0.294	0.270	0.097	-0.205
	newer	0.255	0.096	-0.285	0.158	0.037	-0.191
Separate drifts: 450ms regressor.	older	0.274	0.090	-0.294	0.206	0.034	-0.189
	newer	0.270	0.065	-0.313	0.169	0.055	-0.188
Separate drifts: 525ms regressor.	older	0.297	0.142	-0.276	0.250	0.053	-0.182
	newer	0.237	0.066	-0.302	0.130	0.016	-0.180
MLH: 600ms regressor	older	0.337	0.151	-0.285	0.274	0.074	-0.179
		0.300	0.157	-0.294	0.219	0.116	-0.213
Analysis: Experiment 2							
Fits to all data							
Separate fits for the 600ms regressor	newer	v ₃	v ₁	v _N			
	older	0.428	0.181	-0.227			
Separate drifts: 600ms regressor	newer	0.360	0.145	-0.206			
	older	0.549	0.234	-0.224			
Separate drifts: 375ms regressor	newer	0.314	0.119	-0.195			
	older	0.510	0.221	-0.216			
Separate drifts: 450ms regressor	newer	0.382	0.164	-0.219			
	older	0.405	0.162	-0.217			
Separate drifts: 525ms regressor.	newer	0.337	0.139	-0.248			
	older	0.450	0.162	-0.216			
MLH: 600ms regressor	newer	0.362	0.130	-0.208			
	older	0.451	0.211	-0.222			
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
MLH: 600ms regressor							
 							