

A Bayesian Model for the Time Course of Lexical Processing

Mark Steyvers

(msteyver@psych.stanford.edu)

Department of Psychology, Stanford University
Stanford, CA 94305-2130

Richard Shiffrin

(shiffrin@indiana.edu)

Department of Psychology
Indiana University, Bloomington, IN 47405-7007

Eric-Jan Wagenmakers

(pn_wagenmakers@macmail.psy.uva.nl)

Department of Psychology, University of Amsterdam
Amsterdam, The Netherlands.

René Zeelenberg

(pn_zeelenberg@macmail.psy.uva.nl)

Department of Psychology, University of Amsterdam
Amsterdam, The Netherlands.

Jeroen Raaijmakers

(raaijmakers@psy.uva.nl)

Abstract

A Bayesian-based model for lexical decision, REM-LD, is fit to data from a novel version of a signal-to-respond paradigm. REM-LD calculates the odds that a test item is a word, by accumulating likelihood ratios for each lexical entry in a small neighborhood of similar words. The new model predicts the time course of observed effects of nonword lexicality, word frequency and repetition priming. It can also make qualitative predictions for the response time distributions in tasks with subject paced responding.

Introduction

It is generally assumed that the understanding of the skill of reading should be based in part on an understanding of the storage and retrieval of words. These processes are often studied through the use of the lexical decision task, requiring participants to distinguish words (e.g., CHAIR and FUME) from nonwords (e.g., GREACH and ANSU). In tasks in which accuracy is near ceiling, three critical findings are seen in the response times: (1) The *word frequency effect*. Words that occur regularly in natural language (high frequency or HF words such as CHAIR) are classified correctly faster than words that occur relatively rarely (low frequency or LF words such as FUME). (2) The *repetition priming effect*. Prior exposure to a word leads to faster correct classifications for that word on a second presentation. This increase in performance is particularly pronounced for LF words (e.g., FUME benefits more from prior exposure than CHAIR). (3) The *nonword lexicality effect*. Nonwords that look like words (e.g. GREACH) take longer to be classified correctly than nonwords that are relatively dissimilar to words (e.g. ANSU). In this article¹, we use a new variant of a signal-to-respond procedure that produces findings in the accuracy domain that mimic those listed above for response times. We will fit a new Bayesian model, REM-LD, to the data. The advantage

of the signal-to-respond technique is that it allows one to track the time course of processing, obtaining multiple data points for each stimulus category while reducing concerns about higher order task strategies and speed-accuracy trade-off's. We will also present some preliminary work on predicting response time distributions in the lexical decision task with subject paced responding.

Experimental Data

The signal-to-respond paradigm has occasionally been applied to lexical decision (Antos, 1979; Hintzman & Curran, 1997). We used our new version of the signal-to-respond paradigm to replicate and extend Experiment 2 from Hintzman and Curran (1997).

Method

We used four types of stimuli: (1) 168 HF words, each occurring more than 30 times per million according to the CELEX lexical database (Burnage, 1998) (2) 168 LF words, each occurring 1 or 2 times per million (3) 168 pronounceable nonwords created by replacing one letter of an existing word (e.g., GREACH created from PREAMCH) (4) 168 pronounceable nonwords differing by at least two letters from any word (e.g., ANSU; this condition was absent in the Hintzman and Curran study). The first three stimulus categories were matched on neighborhood structure (i.e., a neighbor is a word differing from another word in one letter, so TIED is a neighbor of LIED); These categories had the same summed logarithmic word frequency of the neighbors. Stimuli were presented twice to study how prior exposure affects performance. To control for practice effects and shifts in response criteria, we presented stimuli in blocks of 48 trials, half of which were stimuli that were encountered in the previous block, half of which were new. For the first block, half of the items were filler items and half were stimuli that would be repeated in the next block. Each block contained 24 words and 24 nonwords. Subjects were required to respond at six different lags: 350, 400, 450, 500, 550, and 600 ms. The appropriate lag was indicated to the

¹ More details and related research can be found in Wagenmakers (2001).

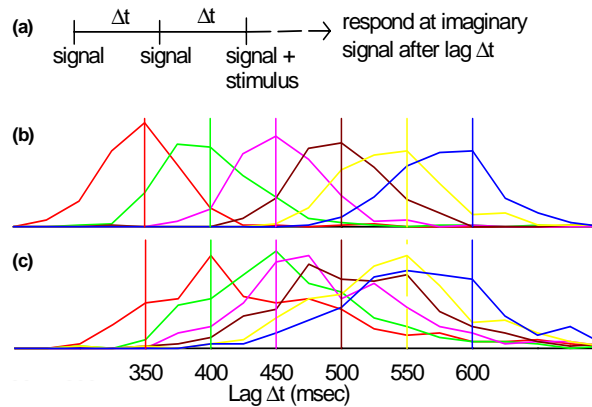


Figure 1. (a) the signal-to-respond procedure. (b) response time distributions for a participant with good timing and (c) for a participant with bad timing. Matching line-colors indicate correspondence between lag (vertical line) and response distribution.

subject by means of three tones (see Figure 1a). The tones were equidistant in time, and the onset of the third and last tone coincided with the onset of the stimulus. The subject had to respond at the fourth *imaginary* tone. We adopted this procedure in the hope that it would produce less interference than the presentation of a tone during processing. After each trial, subjects received feedback concerning the accuracy and latency of their response relative to the desired latency.

Results and Discussion

Forty-three students at Indiana University participated. we excluded 14 participants from the analyses because of extremely bad performance or bad timing. More specifically, subjects were excluded when their overall logarithmic d' was less than 1.0 and/or their average response latency was off by more than 50 ms from the required response latency. Figure 1b and 1c show the distribution of response times for a subject with good and bad timing, respectively. All response latencies were grouped into six bins for each subject separately, the first bin containing the 16.7% slowest responses, and so forth. Next, the accuracy data from each bin were averaged over subjects. Other analyses such as binning by actual response latency or analyzing accuracy data by lag yielded similar results. The results can be seen in Figure 2. Performance for HF words is better than for LF words, and performance for nonwords that differ from any word in two letters (i.e., NW2) is better than for nonwords that differ from a word in one letter (i.e., NW1). Repeated stimuli (indicated by open symbols) are more likely to be classified as 'word' than new stimuli (indicated by the filled symbols), an effect larger for LF words than for

HF words. As expected, performance increases dramatically with processing time, except perhaps for new LF words. This lack of increase could be due to a very slow retrieval process for LF words, or to the possibility that some subjects might be uncertain concerning the lexicality of some LF words. One might argue that the gain in performance for repeated LF words reflects a retrieval of the feedback given on the earlier presentation ('I remember this stimulus is supposed to be a word'). However such a memory process would lead to improved performance for repeated nonwords ('I remember this stimulus is a nonword'), whereas the data show a *decrease* in performance for repeated nonwords. The hypothesis that repetition priming involves two distinct processes (i.e., familiarity and recollection) will be elaborated upon in the Discussion. Overall, the data consistently show effects of processing time, nonword lexicality, word frequency and repetition priming.

The REM-LD Model

The REM-LD model is similar to the REM model for episodic recognition (Shiffrin & Steyvers, 1997). In episodic recognition, participants have to distinguish 'old' words (i.e., words that were presented in a previous study list) from 'new' words. The REM-LD model is an application of the REM model to the lexical decision task. In the REM-LD model, we will make the following assumptions.

(1) Words and nonwords can be represented by vectors of feature values. We assume that these features arbitrarily represent attributes such as orthography and




Figure 2. Results of the signal-to-respond lexical decision experiment. The observed data is indicated by the symbols while the REM-LD model fit is indicated by the solid lines (stimuli presented once) and dashed lines (stimuli presented twice).

phonology. Here we represent each word by a collection of 30 features with values 1 to 10 randomly drawn from a uniform distribution.

(2) Words have lexical entries (i.e., representations) in memory whereas nonwords do not. The presentation of the probe (i.e., the stimulus) leads to activation of n lexical entries that are orthographically *similar* to the probe (see Figure 3a). In a more complete model the value of n would probably be smaller for tests of dissimilar nonwords (i.e., nonwords that differ in two letters from any word), but for simplicity we set $n=10$ for all test items and instead vary the feature similarity for dissimilar nonwords. In case the probe is a word, one of the activated lexical entries is the probe (denoted s-entry for ‘same’, e.g. BEG in Figure 3a). The other activated entries are similar but different from the probe (denoted d-entries for ‘different’). Note that a nonword can only activate lexical entries that are *similar* to it, since nonwords do not have lexical representations.

The degree of similarity between the probe and the entries is determined by the probability that the same feature value is present in probe features as in features from the lexical entries. Stored s-entry features will match the probe features with probability β_1 ; stored d-entry features will match the probe features with probability β_2 , where $\beta_2 < \beta_1$. With probability $1-\beta$, feature values in the entries can differ from those in the probe. These feature values are obtained by sampling randomly from the uniform feature distribution, enabling matches to occur by chance. In the experiment, the similarity between the probe and the lexical neighbors was equated for NW1, LF and HF stimuli. However, the NW2 stimuli were more dissimilar from their lexical neighbors because they differed from any word in at least two letters. Therefore, we set $\beta_2(NW2) < \beta_2(NW1) = \beta_2(LF) = \beta_2(HF)$.

In REM-LD we assume that word frequency is represented by parameter β_f ; the probability of a probe feature matching a trace feature is assumed to be higher for HF word probes than for LF word probes. The increased matching probability for HF words might involve various mechanisms such better matching context features because HF words occur in many different contexts. Therefore, we set $\beta_f(LF) < \beta_f(HF)$. In sum, the β_f parameters determine the similarity or degree of overlap between a word probe and its corresponding lexical entry whereas the β_2 parameters determine the amount of overlap between the probe and lexical entries that are most similar to it.

(3) Mistakes in this task are made because the features in the lexical entries become only gradually available over time. The function $\alpha(t)$ gives the probability that a feature from a lexical entry is available at time t , in order to be compared with a probe

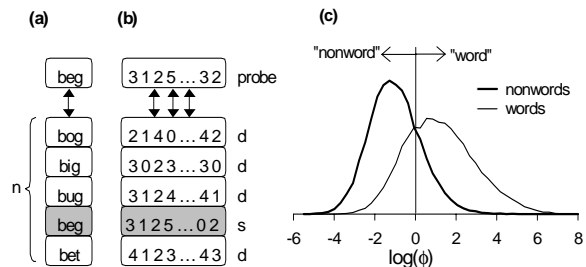


Figure 3. (a) a presented letter string activates n orthographically similar lexical entries. (b) letter strings and lexical traces are represented by vectors. (c) the distribution of $\log(\phi)$ for two conditions.

feature. To account for the improvement in the lexical decision task as a result of more processing time, we assume that the function $\alpha(t)$ is monotonically increasing over time according to:

$$\alpha(t) = 1 - e^{-b(t-t_0)} \quad (1)$$

where b and t_0 are parameters of rate and starting point. Based on Equation (1), the probability that r features are available in the comparison process at time t is given by a binomial distribution:

$$P(r) = \binom{k}{r} \alpha(t)^r (1-\alpha(t))^{k-r} \quad (2)$$

where k is the total number of features in each trace (always 30 here).

(4) The decision to respond ‘word’ or ‘nonword’ is based on an assessment of the evidence that the activated set of lexical entries contains an s-entry. An optimal decision is based on the odds ϕ that the probe is a word rather than a nonword, given the available data D : $\phi = P(w|D) / P(nw|D)$ where the data consists of the number of matches and mismatches between the probe features and the features of all activated lexical entries. Figure 3c illustrates typical $\log \phi$ distributions generated by word and nonword probes, also illustrating the fact that $\log \phi = 0$ is the optimal response criterion. By Bayes’ rule, $\phi = [P(D|w)P(w)] / [P(D|nw)P(nw)]$. Because in our experiment the prior probability of the probe being a word, $P(w)$, equaled the probability of the probe being a nonword, $P(nw)$, we have:

$$\phi = \frac{P(D|w)}{P(D|nw)} \quad (3)$$

When the probe is a word, there is an equal probability that any activated lexical entry is a s-entry. This can be used in a simple derivation (Shiffrin & Steyvers, 1997) that leads to:

$$\frac{P(D|w)}{P(D|nw)} = \frac{1}{n} \sum_{j=1}^n \frac{P(m_j, q_j | s_j)}{P(m_j, q_j | d_j)} \quad (4)$$

where m_j and q_j are the number of matching and mismatching features respectively in the comparison of the lexical entry j to the probe. Note that the total number of features that participate in the comparison process is given by Equation (2). The terms s_j and d_j represent the assumptions that the lexical entry j is a s -entry and d -entry respectively. Therefore, the odds for *word* is an average of the likelihood ratios for each of the lexical entries in the activated set. We can calculate each likelihood ratio in the following way:

$$\frac{P(m_j, q_j | s_j)}{P(m_j, q_j | d_j)} = \left(\frac{\hat{\beta}_1 + (1 - \hat{\beta}_1) \frac{1}{v}}{\hat{\beta}_2 + (1 - \hat{\beta}_2) \frac{1}{v}} \right)^{m_j} \left(\frac{1 - \hat{\beta}_1}{1 - \hat{\beta}_2} \right)^{q_j} \quad (5)$$

In this equation, v is the number of distinct values from the uniform feature distribution (always 10 here).. The calculations of the system also involves the estimates $\hat{\beta}_1$ and $\hat{\beta}_2$. These system estimates are based on an arithmetic average of the different values that β_1 and β_2 can take on in the different experimental conditions.

(5) Prior exposure to a word primes the corresponding lexical entry. Therefore, the features of a repeated word probe will better match the features in the corresponding lexical entry. We model this by assuming that β_1 is increased by a small amount $\Delta\beta$ for repeated word probes. Similarly, prior exposure to a nonword primes the lexical entry of the word that is most similar to it. Therefore, the second occurrence of the nonword string will lead to more matching features in the comparison of the repeated nonword probe and the most similar lexical neighbor. We model this by increasing β_2 for one lexical entry by the amount $\Delta\beta$ for repeated nonword probes.

Simulation results

Figure 2 shows the results of a quantitative model fit of the REM-LD model to the observed data involving seven free parameters. The mean squared error (MSE) of the fit is 1.02e-003. The values of the seven parameters values found to produce the predictions were: $\beta_1(LF)=.66$, $\beta_1(HF)=.77$, $\beta_2(NW1)=.44$, $\beta_2(NW2)=.40$, $\Delta\beta=.070$, $t_0=339$, $b=0.0050$. The qualitative predictions were found to be relatively robust against variations in these parameter values. Because accuracy for HF words is higher than for LF words, $\beta_1(HF)$ was set higher than $\beta_1(LF)$ so that lexical probes would match their lexical entries better for HF words than LF words. Because NW1 nonwords are more often mistakenly judged to be words than NW2 nonwords, $\beta_2(NW1)$ was set higher than $\beta_2(NW2)$ so that a NW1 probe would activate its similar lexical neighbors to a greater extend than a NW2 probe. For both word and nonword conditions, probes that are

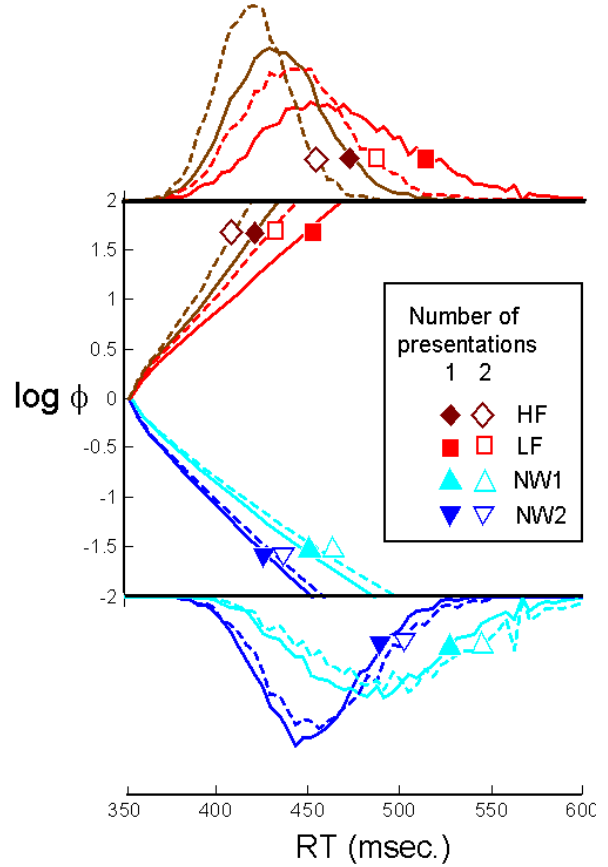


Figure 4. Predicted response time distributions for the lexical decision experiment. A word response is generated by the model at time t when the log odds ($\log \phi$: the current likelihood of responding “word” vs. “nonword”) reaches the upper boundary set arbitrarily at 2. Similarly, a nonword response is generated when $\log \phi$ reaches the lower boundary set at -2 . The middle plot shows the mean $\log \phi$ as a function of time; it displays the growth of evidence over time. The upper and lower plots show the response time distributions for correct responses in the eight conditions of the experiment. Note that the response time results mirror the accuracy results: repetitions speed up word responses but slow down nonword responses.

repeated are classified as ‘word’ more often than probes encountered for the first time. The model predicts this because a repeated word probe primes the corresponding lexical entry while a repeated nonword probe leads primes the lexical entry of the word that it is most similar to. The model also predicts that the repetition priming effect is more pronounced for LF words than for HF words. This is because the average value of $\log \phi$ is closer to zero for the LF words than for the HF words. Hence, an identical increase in $\log \phi$ due to priming of a

lexical entry will have a greater impact on performance for LF words than for HF words.

Predicting Response Time Distributions

There are various ways to derive predictions for the response time distributions in the REM-LD model. Because the main aim of the model was to fit the accuracy results in the signal to respond task, we will only show some qualitative predictions for the response time distributions in the lexical decision task. We explored a method here that is based on diffusion models for response time distributions (e.g., Ratcliff, 1978; Ratcliff, Van Zandt, & McKoon, 1999). Instead of forcing the model to respond at specified lags, the model is simulated at a fine grained time scale. As time progresses, more and more features in the lexicon are activated and participate in the comparison process. For words, the calculated log odds ($\log \phi$) increases on average to larger positive values. For nonwords, the calculated $\log \phi$ will increase to larger negative values. We predict response times by recording the times at which the $\log \phi$ reaches an upper boundary a (for a word response), or a lower boundary $-b$ (for a nonword response). The boundaries a and b were arbitrarily set at 2. In Figure 4, the predicted response time distributions are shown using the same parameters as in the previous simulation. Only the distributions for correct responses are shown. The results mirror the accuracy results: the conditions that lead to more accurate responses are also the conditions that lead to faster responses. For example, words that are repeated are responded to more accurately and faster than non-repeated words. Similarly, nonwords that are repeated are responded to less accurately and slower than non-repeated nonwords. The model also predicts that the response time distributions are more skewed for the slower conditions. Such results have been observed several times in the literature (e.g. Balota & Spieler 1999).

Discussion

We have shown that a Bayesian-based model, REM-LD, can predict lexical decision effects such as word frequency, repetition priming, and nonword lexicality. This model takes into account the similarity of nonwords to words, thereby keeping the system 'centered' around the optimal criterion of $\log \phi$ of zero. REM-LD can also handle the observed improvement in performance with processing time. In contrast to most extant models and empirical work in lexical decision, we focused on changes in accuracy over time, as seen in a variant of a signal-to-respond procedure. A Bayesian model is particularly suited toward explaining data from the signal-to-respond paradigm, since the

system bases its decisions on the diagnosticity of the evidence, simultaneously considering the evidence for and against the 'word' response. When, early in processing, the evidence is noisy and supports neither the 'word' response nor the 'nonword' response, performance is at chance. Empirically, the most interesting finding is the decrease in performance for repeated nonwords. The current model assumes prior exposure of a nonword primes the most similar activated lexical entry, predicting the observed decrement in performance. However, Logan (1988) observed an *increase* in performance for repeated nonwords with subject-paced responding. This result might be attributed to the very short lag between repetitions in Logan's experiment. Therefore, it is possible that the net result of repetition priming for nonwords is the sum of two opposing effects: (1) An *implicit priming* component such as modeled by REM-LD, leading subjects to give the *erroneous* 'word' response, and (2) A *recollection* component that leads subjects to remember the *correct* 'nonword' response. This recollection process might be operative when subjects are under less pressure to give speeded responses, such as in experiments in which responding is subject-paced (e.g., Wagenmakers, 2001).

References

- Antos, S. J. (1979). Processing facilitation in a lexical decision task. *Journal of Experimental Psychology: Human perception and performance*, 5, 527-545.
- Balota, D.A., & Spieler, D.H. (1999). Word frequency effects, repetition, and lexicality effects in word recognition tasks: beyond measures of central tendency. *Journal of Experimental Psychology: General*, 128, 32-55.
- Burnage, D. (1998). *CELEX: a guide for users*. Nijmegen: Centre for Lexical Information.
- Hintzman, D. L., & Curran, T. (1997). Comparing retrieval dynamics in recognition memory and lexical decision. *Journal of Experimental Psychology: General*, 126, 228-247.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492-527.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59-108.
- Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, 106, 261-300.
- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4, 145-166.
- Wagenmakers, E.J. (2001). *Priming in Visual Word Recognition: Empirical Studies and Computational Models*. Unpublished doctoral dissertation.