Published in final edited form as:

J Math Psychol. 2007 June ; 51(3): 198–204. doi:10.1016/j.jmp.2007.02.001.

Analytic Expressions for the BCDMEM Model of Recognition Memory

Jay I. Myung^{a,c}, Maximiliano Montenegro^b, and Mark A. Pitt^a

^aDepartment of Psychology, Ohio State University, 225 Psychology Building, 1835 Neil Avenue, Columbus, Ohio 43210-1287, {myung.1@osu.edu, pitt.2@osu.edu}, 614-292-1862 (voice), 614-688-3984 (fax)

^bDepartment of Mathematics, Science and Technology Education, Ohio State University, Columbus, Ohio 43210, montenegro.3@osu.edu

Abstract

We introduce a Fourier Transformation technique that enables one to derive closed-form expressions of performance measures (e.g., hit and false alarm rates) of simulation-based models of recognition memory. Application of the technique is demonstrated using the bind cue decide model of episodic memory (BCDMEM; Dennis & Humphreys, 2001). In addition to reducing the time required to test the model, which for models like BCDMEM can be excessive, asymptotic expressions of the measures reveal heretofore unknown properties of the model, such as model predictions being dependent on vector length.

Keywords

Recognition Memory; Cognitive Modeling; Fourier Transformation; Signal Detection Theory

1 Introduction

A number of simulation-based models of recognition memory have been proposed in recent years. Examples include BCDMEM (Dennis & Humphreys, 2001), REM (Shiffrin & Steyvers, 1997), SAM (Shiffrin, Ratcliff, & Clark, 1990), MINERVA (Hintzman, 1988), and SLiM (McClelland & Chappell, 1998). The popularity of this style of modeling is due in part to the fact that computing technologies permit both flexibility in model design and precision in implementation. A model can be built that incorporates realistic assumptions about the processes underlying encoding, storage, and retrieval. With the aid of random number generators, a model's predictions are obtained as sample averages over a large number of simulation replications.

^cCorresponding Author

³Interestingly, the same asymptotic expressions as in (15) and (16) can also be obtained from the multivariate normal approximation of the multinomial distribution. That is, the normal approximation of the multinomial distribution implies that the sampling distribution of the loglikelihood ratio $\log(L(\mathbf{n}|\theta, v))$ is approximately normally distributed for large v . The derivation of the asymptotic approximations for hit and FA rates in this manner is shown in the Appendix.

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 this paper, we introduce a technique for deriving closed-form solutions of these models. It involves performing a Fourier transform on a model, which in essence requires one to rewrite the model in another form. The value of the technique is at least two-fold. Most importantly, the asymptotic expressions provide insights into properties of the model that are difficult to glimpse in the model's original formulation. In addition, the excessive computing time required to derive model predictions and estimate model parameters in simulation-based models is eliminated, thereby facilitating model evaluation. We demonstrate the application of the technique and what can be learned with it using BCDMEM.

BCDMEM is a model of episodic recognition memory that, among other things, simulates human performance in tasks in which participants have to discriminate old from new words after a study phase in which only the old words were presented. The contextual information accompanying each encounter with a word (i.e., episodes) is fundamental to the model and its behavior. Context is encoded with the word and used to discriminate old from new ones during recognition. The essence of the context-based recognition process is briefly reviewed.

Associated with each word is a context vector made of multiple feature units, each taking on the value of either 1 (active, or presence of the feature) or 0 (inactive, or absence of the feature). At study, presentation of a word activates a noisy context vector representing the current experimental context. As a result, some of the feature units in the context vector are "learned", meaning that the corresponding feature values become stabilized or permanent and are stored in memory along with the word. Context vectors can undergo spontaneous forgetting, with the value of each feature unit of the vector decaying from one to zero. During recognition, when a test word is presented, the memory retrieval process is activated and the context vector corresponding to the test word is retrieved from memory. Simultaneously, the current experimental context vector is reinstated in memory. If the test word is old (i.e., presented during study phase), the overlap between the retrieved context and the reinstated context should be maximal. If the test word is new (i.e., a distracter), then the retrieved context vector would randomly overlap with the reinstated context vector, which in this case would consist of all the contexts in which the test word has been encountered. On both types of trials, the two context vectors, retrieved and reinstated, are compared against each other, feature by feature, to determine whether the test word should be judged old or new.

The model has five parameters: d , p , r , s , v . The parameter d is the contextual reinstatement parameter, and represents the probability that an originally active feature unit in the experimental context vector fails to get reinstated so becomes inactive in the recognition process. The parameter p is the context noise parameter, and denotes the probability that any feature element of a context vector is active as a result of previous learning. The learning rate parameter, r , stands for the probability that each new feature in the current context vector is copied successfully (i.e., learned) to the memory image. The sparsity parameter, s , represents the probability that each feature in the context vector is turned on spontaneously, independent of prior learning. Finally, the vector length parameter, v , represents the number of distinctive binary features that makes up the context vector.

In the model, the recognition decision of 'old' or 'new' is based on the evaluation of a likelihood ratio defined as the probability of observing the data (i.e., test word) under the hypothesis that it is old to the probability of observing the data under the hypothesis that it is new:

$$\text{Decide "old" if } L = \frac{P(\text{data}|\text{old})}{P(\text{data}|\text{new})} > 1 \text{ and "new" otherwise.} \quad (1)$$

Given that both the retrieved context vector $\mathbf{m} = (m_1, \dots, m_v)$ and the reinstated context vector $\mathbf{c} = (c_1, \dots, c_v)$ are binary (i.e., $m_i = \{0,1\}$, $c_i = \{0,1\}$), there are four types of match that arise in the comparison between the two vectors. They are denoted by the symbols '00', '01', '10', and '11' where, for example, the '01' match represents the observation of $m_i = 0$ and $c_i = 1$ for the i -th feature unit. The likelihood ratio, L , in terms of match types and model parameters is given by (Dennis & Humphreys, 2001, eq. 6)

$$L(\mathbf{n} = (n_{00}, n_{01}, n_{10}, n_{11}) | \theta = (d, p, r, s), v) = \left[\frac{1 - s + ds(1 - r)}{1 - s + ds} \right]^{n_{00}} \left[\frac{p(1 - s) + ds(r + p - rp)}{p(1 - s) + dsp} \right]^{n_{01}} (1 - r)^{n_{10}} \left[\frac{r + p - rp}{p} \right]^{n_{11}} \quad (2)$$

where n_{00} denotes the total number of 00 matches, and so on. By definition, the sum of the four match counts must be equal to the vector length parameter v , i.e., $n_{00} + n_{01} + n_{10} + n_{11} = v$. Note that $\mathbf{n} = (n_{00}, n_{01}, n_{10}, n_{11})$ is a random variable, and so is the likelihood ratio $L(\mathbf{n} | \theta, v)$.

2 Statistical Formulation of BCDMEM

In this section, we provide a statistical treatment of the context noise process assumed in BCDMEM. This allows us to rewrite the model as a statistical one defined as a parametric family of probability distributions, so as to facilitate application of the Fourier Transformation we carry out next.

2.1 Memory Retrieval as a Multinomial Process

Given the assumption that features are independently sampled during recognition, the matching counts, n_{00} , n_{01} , n_{10} and n_{11} , follow a multinomial probability distribution because they are discrete random variables. Specifically, for a test item that is old (i.e., presented during study), the probability distribution of the matching counts is given by

$$f(\mathbf{n} | \mathbf{p}(\theta)) = v! \prod_{\lambda} \frac{p_{\lambda}(\theta)^{n_{\lambda}}}{n_{\lambda}!} \quad (3)$$

In the equation, $\lambda = \{00, 01, 10, 11\}$ is an index variable representing the four types of pattern matches, and $\mathbf{p}(\theta) = (p_{00}(\theta), p_{01}(\theta), p_{10}(\theta), p_{11}(\theta))$ is a vector consisting of four multinomial probability parameters, one for each matching type, defined as: $p_{00}(\theta) = (1 - s)(1 - p) + sd(1 - r)(1 - p)$; $p_{01}(\theta) = (1 - s)p + sd(r + p - rp)$; $p_{10}(\theta) = s(1 - d)(1 - r)(1 - p)$; $p_{11}(\theta) = s(1 - d)(r + p - rp)$. Note that $p_{00}(\theta) + p_{01}(\theta) + p_{10}(\theta) + p_{11}(\theta) = 1$ for all θ , by definition, and that $p_{\lambda}(\theta) = (d, p, r, s)$ does not depend upon the vector length parameter v .

Similarly, for a new item that was not presented during the study phase, the matching counts follow another multinomial probability distribution of sample size v ,

$$f(\mathbf{n} | \mathbf{q}(\theta)) = v! \prod_{\lambda} \frac{q_{\lambda}(\theta)^{n_{\lambda}}}{n_{\lambda}!} \quad (4)$$

In this equation, the multinomial probability vector $\mathbf{q}(\theta) = (q_{00}(\theta), q_{01}(\theta), q_{10}(\theta), q_{11}(\theta))$ is defined as: $q_{00}(\theta) = (1 - s(1 - d))(1 - p)$; $q_{01}(\theta) = (1 - s(1 - d))p$; $q_{10}(\theta) = s(1 - d)(1 - p)$; $q_{11}(\theta) = s(1 - d)p$.

Using the multinomial probabilities $\mathbf{p}(\theta)$ and $\mathbf{q}(\theta)$, the logarithm of the likelihood ratio, called the *loglikelihood ratio*, is re-expressed as

$$\log(L(\mathbf{n}|\theta, \nu)) = \sum_{\lambda} \beta_{\lambda}(\theta) n_{\lambda} \quad (5)$$

where $\beta_{\lambda}(\theta) = \log(p_{\lambda}(\theta)/q_{\lambda}(\theta))$. Note that the loglikelihood ratio is a linear combination of multinomial counts n_{λ} weighted by $\beta_{\lambda}(\theta)$.

2.2 Hit and False Alarm Rates as Expectations

The model makes predictions using two measures of recognition accuracy, hit and false alarm (FA) rates. The hit rate is defined as the proportion of the time that an old item is correctly recognized as old, whereas the FA rate as the proportion of the time that a new item is *incorrectly* recognized as old. Statistically speaking, $P(\text{Hit}|\theta, \nu)$, the probability of recognizing an old item as old, is an expectation under the multinomial distribution in equation (3),

$$P(\text{Hit}|\theta, \nu) = \text{Prob.} [\log(L(\mathbf{n}|\theta, \nu)) > 0 | f(\mathbf{n}|\mathbf{p}(\theta))] = \sum_{n_{00}+n_{01}+n_{10}+n_{11}=\nu} S(\log(L(\mathbf{n}|\theta, \nu))) f(\mathbf{n}|\mathbf{p}(\theta)) \quad (6)$$

where the summation is over all values of n_{00} , n_{01} , n_{10} and n_{11} satisfying $n_{00}+n_{01}+n_{10}+n_{11} = \nu$ and $S(x)$ is the step function defined as $S(x) = 1$ if $x > 0$ and 0 otherwise. In the above equation, the likelihood-based decision rule in (1) has, for mathematical convenience, been translated into an equivalent, log-likelihood based rule since $x > 1 \Leftrightarrow \log(x) > 0$ for all $x > 0$. Similarly, $P(\text{FA}|\theta, \nu)$, the probability of recognizing a new item as old, can be written as another expectation under the multinomial distribution in equation (4),

$$P(\text{FA}|\theta, \nu) = \sum_{n_{00}+n_{01}+n_{10}+n_{11}=\nu} S(\log(L(\mathbf{n}|\theta, \nu))) f(\mathbf{n}|\mathbf{q}(\theta)) \quad (7)$$

To derive the model's predictions on hit and FA rates given a specific set of parameter values, all that we need to do is to evaluate the two equations (6) and (7). This can, however, be computationally very expensive to calculate because of the combinatorial explosion problem. The number of terms that must be summed rapidly increases as ν increases (in proportion to ν^3), and thus becomes prohibitively large even for a reasonable value of ν .¹ For instance, for $\nu = 200$, there are over one million terms to sum! In short, a direct evaluation of the multinomial sums is generally not possible in practice unless ν is trivially small (e.g., 5 - 10).

Faced with this obstacle, the modeler must rely on a less-than-satisfactory shortcut: Brute force simulations of the model are run in which for each chosen set of parameter values, the entire process of the model is simulated on computer using random number generators. This is obviously an extremely time-consuming procedure such that finding parameter values that best fit observed data can be next to impossible.

¹The total number of terms to sum in (6) is obtained by evaluating the sum $\left(\sum_{n_{11}=0}^{\nu} \sum_{n_{10}=0}^{\nu-n_{11}} \sum_{n_{01}=0}^{\nu-n_{11}-n_{10}} 1 \right)$ with the help of an algebraic computing package such as *Mathematica* or using the Faulhaber's Formula (Gradshteyn & Ryzhik, 2000). The number obtained is equal to $(\nu^3 + 6\nu^2 + 11\nu + 6)/6$, that can be approximated to $\nu^3/6$ for large ν .

3 Integral Expressions for Recognition Probabilities

With the statistical formulation of BCDMEM in hand, we are ready to apply the Fourier transformation technique and show how it enables one to derive easy-to-compute integral expressions for hit and FA rates, thus freeing the modeler from cumbersome and restrictive simulation methods.

The Fourier transformation (FT), a mathematical tool widely used in many scientific and engineering applications, is an integral transformation technique that decomposes a function into an infinite sum (an integral) of its complex-valued frequency components with coefficients determined specific to the function (see, e.g., Arfken & Weber, 2001; Boas, 1983).

3.1 Fourier Integral Expressions for Hit and FA Rates

Using the FT technique, the step function $S(x)$ defined in equation (6) can be equivalently written as

$$S(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{e^{ikx}}{ik} dk \quad (8)$$

where $e^{ikx} \equiv \cos(kx) + i\sin(kx)$ is a complex number consisting of the real and imaginary parts, with the latter containing the mathematical symbol $i = \sqrt{-1}$ or equivalently $i^2 = -1$.

By replacing $S(x)$ in (6) with its Fourier transformation in (8) and the loglikelihood ratio with the expression in (5), we rewrite $P(\text{Hit}|\theta, v)$ as

$$P(\text{Hit}|\theta, v) = \sum_{n_{00}+n_{01}+n_{10}+n_{11}=v} \left[\frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{e^{ik \log(L(n|\theta, v))}}{ik} dk \right] f(n|p(\theta)) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{dk}{ik} \left[\sum_{n_{00}+n_{01}+n_{10}+n_{11}=v} e^{ik \sum_{\lambda} \beta_{\lambda}(\theta) n_{\lambda}} v! \prod_{\lambda} \frac{p_{\lambda}(\theta)^{n_{\lambda}}}{n_{\lambda}!} \right] \quad (9)$$

The sum in the hard bracket [...] reduces to

$$\text{Sum} = \sum_{n_{00}+n_{01}+n_{10}+n_{11}=v} v! \prod_{\lambda} \frac{(e^{ik \beta_{\lambda}(\theta)} p_{\lambda}(\theta))^{n_{\lambda}}}{n_{\lambda}!}$$

using the relation $e^{ik \sum_{\lambda} \beta_{\lambda}(\theta) n_{\lambda}} = \prod_{\lambda} (e^{ik \beta_{\lambda}(\theta)} p_{\lambda}(\theta))^{n_{\lambda}}$. We then notice that this sum corresponds to the multinomial expansion,

$$\text{sum} = \left(\sum_{\lambda} e^{ik \beta_{\lambda}(\theta)} p_{\lambda}(\theta) \right)^v$$

Finally, plugging this result into equation (9), we obtain the following one-dimensional integral expressions for hit rate, and a similar expression for FA rate, as follows:

$$\begin{aligned}
 P(\text{Hit}|\theta, v) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \left(\sum_{\lambda} e^{ik\beta_{\lambda}(\theta)} p_{\lambda}(\theta) \right)^v \frac{dk}{ik} \\
 P(\text{FA}|\theta, v) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \left(\sum_{\lambda} e^{ik\beta_{\lambda}(\theta)} q_{\lambda}(\theta) \right)^v \frac{dk}{ik}
 \end{aligned}
 \tag{10}$$

The above FT integrals can be easily evaluated numerically. To understand why, note that the integral for $P(\text{Hit}|\theta, v)$ is a line integral along the real axis in the complex plane that contains the singularity point at $k = 0$. From the Residue Theorem of complex analysis (Arfken & Weber, 2001), the integration around the singularity yields the value of $1/2$, and the integration along the rest of the real axis can be calculated just for the real part of the integrand function because the real part of the function is even around $k = 0$ (i.e., $f(k) = f(-k)$). Also, because the imaginary part of the function is odd, its integral along the real axis cancels out. In short, the FT integral for the hit rate is rewritten as

$$P(\text{Hit}|\theta, v) = \frac{1}{2} + \frac{1}{\pi} \int_{0+}^{\infty} \Re \left[\frac{1}{ik} \left(\sum_{\lambda} e^{ik\beta_{\lambda}(\theta)} p_{\lambda}(\theta) \right)^v \right] dk
 \tag{11}$$

where $\Re[Z]$ is the real part of a complex number z and $0+$ is an infinitesimally small positive number.² It is straightforward to evaluate numerically the above expression using *Mathematica* or *Matlab*.

We verified the correctness of the FT results by comparing them with simulation data of Dennis and Humphreys (2001). In one simulation study, they evaluated BCDMEM's ability to reproduce the data of Glanzer and Adams (1990), who first had participants perform a lexical decision task using low- and high-frequency words presented auditorily or visually. This was followed by a standard old-new recognition memory task. The empirical data from the four conditions are in the second column of Table 1. To the right are the best-fit predictions from Dennis and Humphreys' (2001) simulations, which are quite good. To validate that our implementation of BCDMEM was correct, we ran the same simulation using the same parameter values as Dennis and Humphreys (2001). The data, shown in the fourth column, are very similar. Next we evaluated the FT integral expression in (10) to determine if equivalent results would be obtained. As shown in the rightmost column, this is in fact the case.

A second test was performed to verify further the close correspondence between the two forms of the model (simulation versus FT integration). Three parameters d , p , s were fixed at values different from the preceding simulation. The learning rate parameter, r , was varied across its range and model predictions were compared across this range. As r increases, discriminability of old from new items should increase, causing the hit and FA rates to diverge from one another. The results are shown in Figure 1, with the circles representing the simulation predictions and the lines the FT predictions. As is clear, the values from the analytic expression are in almost perfect agreement with the simulation values. The results from this and the preceding test provide demonstrate the accuracy and correctness of the FT integral.

²The second term of the right hand side of the equation is evaluated in practice by approximating the term as

$\lim_{\epsilon \rightarrow 0} \frac{1}{\pi} \int_{\epsilon}^{\infty} \Re[\dots] dk$, provided that the result converges to a finite value.

3.2 Asymptotic Expressions of Fourier Integrals

Although the Fourier-transformed integrals in (10) are readily computable, it is useful to seek their asymptotic expressions to identify fundamental properties of BCDMEM. To derive an asymptotic expression for $P(\text{Hit}|\theta, \nu)$, we expand the second term on the right-hand side of equation (11) for large values of the vector length parameter, ν . By plotting the integrand function in the term as a function of k , we observe that the function concentrates its mass around $k = 0$ for large ν . The first step of asymptotic expansion is therefore to expand this function around $k = 0$. To do this, we successively apply the two Taylor series expansions of the

exponential and logarithmic functions, $e^x \approx 1+x+\frac{x^2}{2}$ and $\log(1+x) \approx x - \frac{x^2}{2}$, to obtain the following approximation:

$$\begin{aligned} \log \left(\sum_{\lambda} e^{ik\beta_{\lambda}(\theta)} p_{\lambda}(\theta) \right) &\approx \log \left(1 + ik \sum_{\lambda} \beta_{\lambda}(\theta) p_{\lambda}(\theta) - \frac{k^2}{2} \sum_{\lambda} \beta_{\lambda}^2(\theta) p_{\lambda}(\theta) \right) \approx \left(ik \sum_{\lambda} \beta_{\lambda}(\theta) p_{\lambda}(\theta) - \frac{k^2}{2} \sum_{\lambda} \beta_{\lambda}^2(\theta) p_{\lambda}(\theta) \right) \\ &\quad - \frac{1}{2} \left(ik \sum_{\lambda} \beta_{\lambda}(\theta) p_{\lambda}(\theta) - \frac{k^2}{2} \sum_{\lambda} \beta_{\lambda}^2(\theta) p_{\lambda}(\theta) \right)^2 \\ &= i\mu_p(\theta)k \\ &\quad - \frac{\sigma_p^2(\theta)}{2}k^2 + O(k^3) \end{aligned} \quad (12)$$

where $\mu_p(\theta)$ and $\sigma_p^2(\theta)$ are the mean and variance of the single-trial “payoff” $\beta(\theta)$ defined as $\mu_p(\theta) = \sum_{\lambda} \beta_{\lambda}(\theta) p_{\lambda}(\theta)$ and $\sigma_p^2(\theta) = \sum_{\lambda} \beta_{\lambda}^2(\theta) p_{\lambda}(\theta) - \mu_p^2(\theta)$. From the above approximation and remembering the identity $z^{\nu} = \exp(\nu \log(z))$, we obtain

$$\left(\sum_{\lambda} e^{ik\beta_{\lambda}} p_{\lambda} \right)^{\nu} = \exp \left(\nu \log \left(\sum_{\lambda} e^{ik\beta_{\lambda}} p_{\lambda} \right) \right) \approx \exp \left(i\nu\mu_p k - \frac{\nu\sigma_p^2}{2}k^2 \right) \quad (13)$$

where, for simplicity, we employed the shorthand symbols $\beta_{\lambda}, p_{\lambda}, \mu_p$ and σ_p^2 by dropping the explicit θ in the argument. This convention will be used in the remainder of this section, whenever convenient.

Now, using the approximation in (13), the the second term on the right hand side of (11) is rewritten as

$$\begin{aligned} \frac{1}{\pi} \int_{0+}^{\infty} \Re \left[\frac{1}{ik} \left(\sum_{\lambda} e^{ik\beta_{\lambda}(\theta)} p_{\lambda}(\theta) \right)^{\nu} \right] dk &\approx \frac{1}{\pi} \int_{0+}^{\infty} \Re \left[\frac{1}{ik} \exp \left(i\nu\mu_p k - \frac{\nu\sigma_p^2}{2}k^2 \right) \right] dk \\ &= \frac{1}{\pi} \int_{0+}^{\infty} \left(\frac{1}{2} \int_{-\nu\mu_p}^{\nu\mu_p} \cos(xk) dx \right) \exp \left(-\frac{\nu\sigma_p^2}{2}k^2 \right) dk \\ &= \text{ncdf} \left(\frac{\mu_p}{\sigma_p/\sqrt{\nu}} \right) - \frac{1}{2} \end{aligned} \quad (14)$$

where $\text{ncdf}(x) \equiv \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp \left(-\frac{y^2}{2} \right) dy$ is the unit normal cumulative distribution function. By plugging the above result into (11), we obtain the desired asymptotic hit rate as

$$P(\text{Hit}|\theta, v) \approx \text{ncdf}\left(\frac{\mu_p(\theta)}{\sigma_p(\theta) / \sqrt{v}}\right) \quad (15)$$

for large v . Similarly, the asymptotic FA rate for large v is obtained as

$$P(\text{FA}|\theta, v) \approx \text{ncdf}\left(\frac{\mu_q(\theta)}{\sigma_q(\theta) / \sqrt{v}}\right) \quad (16)$$

where $\mu_q(\theta) = \sum_{\lambda} \beta_{\lambda}(\theta) q_{\lambda}(\theta)$ and $\sigma_q^2(\theta) = \sum_{\lambda} \beta_{\lambda}^2(\theta) q_{\lambda}(\theta) - \mu_q^2(\theta)$.³

Three properties of BCDMEM directly fall out of these asymptotic expressions. First, the mean $\mu_p(\theta)$ turns out to be equal to the Kullback-Leibler distance measure between two distributions in information theory (Cover & Thomas, 1991)

$$\mu_p(\theta) = \sum_{\lambda} p_{\lambda}(\theta) \log\left(\frac{p_{\lambda}(\theta)}{q_{\lambda}(\theta)}\right) = D(p \parallel q) \geq 0$$

and similarly, $\mu_q(\theta) = -D(q \parallel p) \leq 0$. The Kullback-Leibler distance is not symmetric, i.e., $D(p \parallel q) \neq D(q \parallel p)$, unless $p = q$. One can also easily show the inequality $\mu_q(\theta) \leq 0 \leq \mu_p(\theta)$. These observations together imply that the hit rate will always be equal to or greater than the FA rate; that is, $P(\text{FA}|\theta, v) \leq 0.5 \leq P(\text{Hit}|\theta, v)$ for all θ and v , at least asymptotically.

Second, the hit and FA rates that the model predicts depend on the value of the vector length parameter, v . In particular, as v is increased to infinity, hit and FA rates converge to 1 and 0, respectively, regardless of the values of the other four parameters $\theta = (d, p, r, s)$. Formally, we observe

$$P(\text{FA}|\theta, v) \rightarrow 0 \quad \& \quad P(\text{Hit}|\theta, v) \rightarrow 1 \quad \text{as} \quad v \rightarrow \infty$$

for all θ .

Third, related to the second point, inspection of the asymptotic expressions reveals a trade-off between the parameter vector $\theta = (d, p, r, s)$ and the vector length parameter, v . The same hit rate is predicted for two different choices of model parameters, (θ_1, v_1) and (θ_2, v_2) , as long as they are chosen such that the value of the argument of the cumulative normal distribution function $\text{ncdf}(\cdot)$ remains unchanged,

$$\frac{\mu_p(\theta_1)}{\sigma_p(\theta_1) / \sqrt{v_1}} = \frac{\mu_p(\theta_2)}{\sigma_p(\theta_2) / \sqrt{v_2}}.$$

The same can be said about the model's prediction of FA rate. An implication of this behavior for BCDMEM is that the model is unidentifiable to the extent that different sets of parameter values for which the above equality holds can be found. When the model becomes unidentifiable, one cannot meaningfully interpret the parameter values, as there exist multiple sets of parameter values that fit the data equally well.

Another benefit of the asymptotic expressions for $P(\text{Hit}|\theta, \nu)$ and $P(\text{FA}|\theta, \nu)$ is that one can derive an explicit equation describing the receiver operating curve (ROC) in signal detection theory (Green & Swets, 1974). To show this, first note that the asymptotic expressions are obtained under a fixed decision criterion: Decide “old” if $\log L(\mathbf{n}|\theta, \nu) > 0$ and “new” otherwise. Under a more general decision rule assuming a variable criterion c instead of $c = 0$, the asymptotic expressions in (15) and (16) are then modified to

$$P(\text{Hit}|\theta, \nu, c) \approx \text{ncdf}\left(\frac{\mu_p(\theta) - c}{\sigma_p(\theta) / \sqrt{\nu}}\right); \quad P(\text{FA}|\theta, \nu, c) \approx \text{ncdf}\left(\frac{\mu_q(\theta) - c}{\sigma_q(\theta) / \sqrt{\nu}}\right)$$

for $-\infty < c < \infty$. From this, the ROC plotted on Gaussian axes (z_{Hit} vs z_{FA}) is obtained as

$$z_{\text{Hit}}(\theta, \nu, c) = \left(\frac{\sigma_q(\theta)}{\sigma_p(\theta)}\right) z_{\text{FA}}(\theta, \nu, c) + \left(\frac{\mu_p(\theta) - \mu_q(\theta)}{\sigma_p(\theta) / \sqrt{\nu}}\right) \quad (17)$$

which is known as the z -ROC.

4 Discussion

Application of the FT technique makes BCDMEM not only more user-friendly but also more understandable. No longer are extended simulations required to generate reliable predictions. Rather, the analytically exact integral expression enables one to quickly compute model predictions in the form of hit and FA rates. To perform this analysis, we first recast BCDMEM as a multinomial process, characterized by a multinomial probability distribution with hit and FA rates as expectations under that distribution. To our knowledge, this is the first such statistical treatment of the model.

By taking the analysis one step further and deriving the asymptotic expressions of hit and FA rates for large ν , we were able to glean previously unknown properties of the model. Most importantly, the context vector length, ν , turns out to be a non-ignorable parameter that can significantly affect model predictions. The analysis also showed that the model is unidentifiable, in the sense that multiple sets of parameter values can provide the same fit to a data set. At least one parameter such as ν must be fixed to avoid this.⁴

An obvious next question to ask in this line of work is whether one can utilize the asymptotic expressions to address other substantive issues of current interest to memory researchers. For example, it would be interesting if one could show that BCDMEM can predict a reversal or elimination of the low frequency hit-rate advantage that recent experiments have found (e.g., Criss & Shiffrin, 2004). Another possible analysis, suggested by a reviewer, is to determine whether the model can exhibit the list strength effect (Ratcliff, Clark, & Shiffrin, 1990; Ratcliff, Shen, & Gronlund, 1992). More broadly, one can also ask whether the FT technique can be applied to other simulation-based models (e.g., REM, SAM, MINERVA), and of course, whether asymptotically they exhibit similar properties as BCDMEM. Unique challenges are faced in applying the technique to each of these models. Work is currently underway that will address many of the above issues.

⁴Because Dennis and Humphreys (2001) always fixed the vector length parameter ν in their simulation analyses, they may well have suspected that BCDMEM is unidentifiability when ν varies.

Acknowledgments

This work was supported by National Institute of Health grant R01-MH57472 to JIM and MAP. We wish to thank Simon Dennis, Richard Shiffrin, Roger Ratcliff, Mark Steyvers, and an anonymous reviewer for valuable feedback on earlier versions of this paper. Correspondence concerning this article should be addressed to Jay Myung, Department of Psychology, Ohio State University, 1835 Neil Avenue, Columbus, OH 43210. E-mail: myung.1@osu.edu.

Appendix: Sampling Distribution of Loglikelihood Ratio

In this appendix we show, using the normal approximation of the multinomial distribution, that the sampling distribution of the loglikelihood ratio $\log(L(\mathbf{n}|\theta, v))$ is approximately normal for large v . From this, we then derive the asymptotic expressions for the Fourier integral in (15).

According to equation (5), the loglikelihood ratio, which itself is a random variable, is written as

$$\log(L(\mathbf{n}|\theta, v)) = \sum_{\lambda} \beta_{\lambda}(\theta) n_{\lambda}$$

where $\lambda = \{00, 01, 10, 11\}$. Suppose that n_{λ} represents the pattern matching count given a target item that is old, and thus it follows the multinomial distribution $f(\mathbf{n}|\mathbf{p}(\theta))$. The mean of the loglikelihood ratio under this distribution is given by

$$E[\log(L(\mathbf{n}|\theta, v))] = \sum_{\lambda} \beta_{\lambda}(\theta) E(n_{\lambda}) = \sum_{\lambda} \beta_{\lambda}(\theta) v p_{\lambda}(\theta) = v \mu_p(\theta) \quad (18)$$

where $\mu_p(\theta)$ is defined in the main text. Next, the variance of the loglikelihood ratio is obtained as

$$\begin{aligned} \text{Var}[\log(L(\mathbf{n}|\theta, v))] &= \text{Var}\left[\sum_{\lambda} \beta_{\lambda} n_{\lambda}\right] \\ &= \sum_{\lambda} \beta_{\lambda}^2 \text{Var}(n_{\lambda}) \\ &\quad + 2 \sum_{\lambda \neq \kappa} \beta_{\lambda} \beta_{\kappa} \text{Cov}(n_{\lambda}, n_{\kappa}) \\ &= \sum_{\lambda} \beta_{\lambda}^2 v p_{\lambda} (1 - p_{\lambda}) \\ &\quad + 2 \sum_{\lambda \neq \kappa} \beta_{\lambda} \beta_{\kappa} (-v) p_{\lambda} p_{\kappa} \\ &= v \sum_{\lambda} \beta_{\lambda}^2 p_{\lambda} \\ &\quad - v \left(\sum_{\lambda} \beta_{\lambda} p_{\lambda} \right)^2 = v \sigma_p^2(\theta) \end{aligned} \quad (19)$$

where $\sigma_p^2(\theta)$ is defined in the main text. In the above equation, for simplicity, we employed the shorthand symbols β_{λ} and p_{λ} for $\beta_{\lambda}(\theta)$ and $p_{\lambda}(\theta)$, respectively.

Next, the multivariate normal approximation of the multinomial distribution implies that each n_{λ} would be approximately normally distributed for large v . Given that the loglikelihood ratio is a linear sum of normals (i.e., n_{λ} 's) as shown in (5), its sampling distribution would also be approximately normal. The *linearity* of the loglikelihood equation in n_{λ} is crucial for the normal approximation to work. Otherwise, the approximation is not justified.

Finally, putting together the above results and observation, we can now rewrite the probability of the loglikelihood ratio exceeding 0 (i.e., $P(\text{Hit}|\theta, v)$) in terms of a unit normal random variable z as

$$\text{Prob.} [\log (L (n|\theta, v)) > 0 | f (n|p (\theta))] = P \left(z > \frac{0 - v\mu_p (\theta)}{\sqrt{v\sigma_p^2 (\theta)}} \right) = P \left(z < \frac{\mu_p (\theta)}{\sigma_p (\theta) / \sqrt{v}} \right) = \text{ncdf} \left(\frac{\mu_p (\theta)}{\sigma_p (\theta) / \sqrt{v}} \right) \quad (20)$$

yielding the desired result. The asymptotic expression for $P(\text{FA}|\theta, v)$ can be derived in a similar fashion.

5 References

- Arfken, GB.; Weber, HJ. Mathematical methods for physicists. 5th ed.. Weber, HJ., editor. Academic Press; San Diego: 2001.
- Boas, ML. Mathematical methods in the physical sciences. 2nd ed.. Wiley; New York: 1983.
- Cover, T.; Thomas, J. Elements of information theory. Wiley; New York: 1991.
- Criss AH, Shiffrin RM. Interactions between study task, study time, and the low-frequency hit rate advantage in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2004;30(4):778–786.
- Dennis S, Humphreys MS. A context noise model of episodic word recognition. *Psychological Review* 2001;108(2):452–478. [PubMed: 11381837]
- Glanzer M, Adams JK. The mirror effect in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1990;16(1):5–16.
- Gradshteyn, IS.; Ryzhik, IM. *Tablitsy integralov, summ, riadov i proizvedenii* [Table of integrals, series, and products]. 6th ed.. Jeffrey, A., editor. Academic Press; San Diego: 2000.
- Green, D.; Swets, J. Signal detection theory and psychophysics. Krieger; New York: 1974.
- Hintzman DL. Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review* 1988;95(4):528–551.
- McClelland JL, Chappell M. Familiarity breeds differentiation: A subjective-likelihood approach to the effects of experience in recognition memory. *Psychological Review* 1998;105(4):724–760. [PubMed: 9830377]
- Ratcliff R, Clark SE, Shiffrin RM. The list-strength effect: I. data and discussion. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1990;16:163–178.
- Ratcliff R, Shen C-F, Gronlund SD. Testing global memory models using roc curves. *Psychological Review* 1992;99:518–535. [PubMed: 1502275]
- Shiffrin M, Steyvers M. A model for recognition memory: REM-retrieving effectively from memory. *Psychonomic Bulletin & Review* 1997;4:145–166.
- Shiffrin RM, Ratcliff R, Clark SE. List-strength effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1990;16(2):179–195.

Fig. 1.

Comparison between simulated and exact FT results. Solid and open circles represent simulation values of hit and FA rates, respectively, averages based on 5,000 independent replications. Solid and dotted curves are exact values obtained by analytically evaluating the FT integral expressions in (10). The parameter values used in both cases were $d = 0.20$, $p = 0.40$, $s = 0.20$ and $v = 200$.

Table 1

The second column contains observed data from Glanzer and Adams (G & A, 1990), obtained under high (H) and low (L) word-frequency conditions. The third column shows best-fit predictions of BCDMEM as reported in Dennis and Humphreys (D & H, 2001). Parameter values for their simulation were $d = 0.590$; $p_L = 0.094$; $p_H = 0.336$; $r = 0.602$; $s = 0.02$, and $v = 200$. The probability values shown in the third column are obtained by digitally scanning the figure image in Dennis and Humphreys (2001, Figure 9A). The next two columns show our own results obtained by direct simulation and exact evaluation of the Fourier transformation (FT) integral expressions in (10).

	G & A (1990) Data	D & H (2001) Simulation	Present Simulation	FT Integration
Hit _L	0.7437	0.744±0.024	0.7371	0.7338
Hit _H	0.6785	0.687±0.022	0.6793	0.6792
FA _L	0.2172	0.189±0.022	0.1818	0.1825
FA _H	0.2804	0.299±0.025	0.2935	0.2943