# Mechanisms of skill refinement: A model of long-term repetition priming

**Michael D. Colagrosso**[+*], **Michael C. Mozer**[+*], **David E. Huber**[#*]
[+] Department of Computer Science
[#] Department of Psychology
[*] Institute of Cognitive Science
University of Colorado, Boulder, CO 80309
{mozer,colagrom,dhuber}@colorado.edu

## Abstract

We address an omnipresent and pervasive form of human learning—*skill refinement*, the improvement in performance of a cognitive or motor skill with practice. A simple example of skill refinement is the psychological phenomenon of *long-term repetition priming*: Participants asked to read briefly presented words are more accurate if they viewed the word earlier in the experiment. We simulate various phenomena of repetition priming using a probabilistic model that describes information flow along cortical processing pathways. The model suggests two distinct mechanisms of adaptation with experience, one that updates prior probabilities of pathway outputs, and one that improves information transmission through a pathway. These two mechanisms loosely correspond to bias and sensitivity changes that have been observed in experimental studies of priming. Both mechanisms are extremely sensible from a rational perspective, and serve as the foundation of skill acquisition and skilled performance.

Computational modeling has focused primarily on two aspects of human learning—the induction of new concepts and categories, and the acquisition of new skills. Another aspect of human learning has received little attention—the refinement of existing skill. Skill refinement, also called skill practice, is an omnipresent and pervasive form of learning. As we type, drive, read, or play video games, our behavior becomes less error prone and more fluent, rapid, and robust to distraction and irrelevant aspects of the task. Skill refinement is sometimes explicit, such as the rehearsal of a piano sonata, but is often implicit, such as entering one's personal identification number at an automated teller machine. Understanding skill refinement is fundamentally about discovering the mechanisms by which one trial or performance of the skill leads to improvements on the next trial.

## 1 Long-term repetition priming

Perhaps the most direct and easily studied manifestation of skill refinement in the psychological literature is the phenomenon of *long-term repetition priming*. In the priming paradigm, participants engage in a series of experimental trials, and experience with a stimulus or response on one trial results in more efficient processing on subsequent trials.

Efficiency is defined in terms of shorter shorter response times, lower error rates, or both. A typical long-term perceptual priming experiment consists of a *study phase* in which participants are asked to read a list of words one at a time, and a *test phase*, during which they must respond a series of brief, masked target words. The time between target onset and mask onset is is called the *flash duration*. Typical response paradigms include speaking the target aloud (*naming*) and a forced choice between two alternatives (*2AFC*). Repetition priming occurs when a word from the study phase influences performance during the test phase. Priming is an implicit memory phenomenon: participants are not told the study and test phases are related, and they do not try to recall study words during the test phase as a deliberate strategy for performing the task. Thus, priming is incidental and not task related; it comes about as a result of experience and is thus a form of skill refinement, where the "skill" here is perceptual processing of a letter string.

Priming in this paradigm is long term, in that it persists over a period of many minutes and many intervening trials. Priming can also be short term, in that it persists only from one trial to the next. Priming can occur based not only on repetitions of an item, but based on semantic or orthographic similarity. Repetition priming is an easy case to study because it is the case of maximal similarity between prime and target. Models have been proposed for other forms of priming (e.g., Huber et al., 2001).

A key question concerning repetition priming is whether—to use the language of signal detection theory—priming is due to increased *bias* or increased *sensitivity*. Bias means that participants are more likely to report studied items regardless of what word is presented for identification. Sensitivity means that participants become better at perceptual discrimination of the studied items. From signal detection theory, it is well known that an increased bias toward a studied word can either benefit (by increasing the correct detection rate) or hinder (by increasing the false detection rate) overall performance. In contrast, increased sensitivity to a studied word has the specific effect of improving the ability to perceive that word during the test phase. A key finding in long-term repetition priming research has been that priming reflects both increased bias and increased sensitivity, although the sensitivity increase is robust only for low-frequency words or novel items.

The goal of this paper is to introduce a model of skill rehearsal and performance. The model has two distinct learning mechanisms which contribute to skill improvement with practice. The model explains various data from psychological studies of long-term repetition priming. In this paper, we model two experiments isolating bias and sensitivity effects in priming, and show that our two learning mechanisms correspond to these two effects.

## 2 Modeling long-term repetition priming

Our theory posits that cortical computation is performed by a set of functionally specialized *pathways*. Each pathway performs a primitive cognitive operation, e.g., visual word-form recognition, identification of semantic features of visual objects, computation of spatial relationships, or construction of motor plans. To model the effects of long-term repetition priming, we propose a model with two pathways in cascade. A *perceptual* pathway maps visual features to word identities. A *response* pathway takes the output of the perceptual pathway and maps it to a task-appropriate response. We assume the pathways communicate continuously during processing and that communication is unidirectional.

### 2.1 Implementing a pathway as a dynamic belief network

We present a probabilistic model of a pathway, which characterizes the transformation of pathway input to output, the time course of information processing during a single trial, and tuning of the pathway behavior over many trials. The inputs and outputs of a pathway

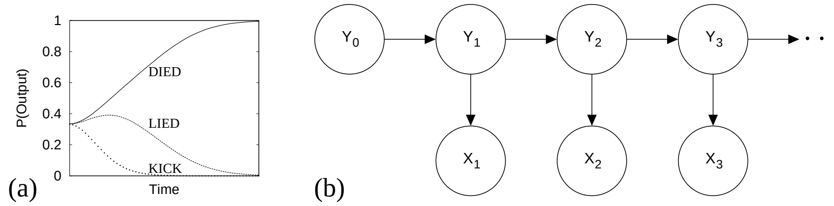


Figure 1: (a) Illustration of a perceptual pathway when the static visual input corresponding to the word DIED is presented. The three curves show the probabilities of different output alternatives as a function of processing time. (b) An HMM implementation of a pathway

are represented as probability distributions over distinct alternatives. Formally, the input and output states of a pathway at a particular time $t$, denoted $X_t$ and $Y_t$, respectively, are discrete random variables. Each variable can take on one of a finite set of values selected from a multinomial distribution, with set size $N_X$ and $N_Y$ for $X_t$ and $Y_t$, respectively. We wish to model the temporal dynamics of a pathway, i.e., how $X_t$ and $Y_{t-1}$ combine to determine $Y_t$. To link this notation to the repetition priming paradigm, consider a perceptual pathway. To model the processing of some word $x$ for a brief duration $d$, we would set $X_1 = X_2 = \ldots = X_d = x$ (i.e., assigning the random variables a particular value $x$); to model the masking of the word, $X_t$ for $t > d$ might be set to a uniform distribution over alternatives. Given this input sequence corresponding to a single trial, we can then observe the temporal evolution of the pathway output (Figure 1a).

The relationship among the input and output variables is specified by the graphical model in Figure 1b, known as a *dynamic belief network* (DBN) (Dean & Kanazawa, 1989; Kanazawa, Koller, & Russell, 1995). Each arrow corresponds to a conditional probability distribution (CPD) specifying the relationship between two dependent variables. For the reader unfamiliar with graphical models, one should not be concerned with the direction of the arrows; Figure 1b is cast as a generative process—depicting the flow from outputs to inputs—but inference can be carried out in either direction. The graphical model allows us to infer the probability distribution over $Y_t$, $\mathrm{P}(Y_t)$, given $X_1, X_2, \ldots, X_t$. This computation is performed via iterative Bayesian belief revision. Figure 1b is simply a hidden Markov model (HMM), used in a novel way. In typical usage, an HMM is presented with a sequence of distinct inputs, whereas we maintain the same input for many successive time steps. Further, in typical usage, an HMM transitions through a sequence of distinct hidden states, whereas we attempt to converge with increasing confidence on a single state.

In Figure 1b, the set of arrows from $X_t$ to $Y_t$ corresponds to $\mathrm{P}(X_t|Y_t)$, and can be thought of as a *strength of association* between $X_t$ and $Y_t$. The set of arrows from $Y_{t-1}$ to $Y_t$ corresponds to $\mathrm{P}(Y_t|Y_{t-1})$, and can be thought of as a *short-term memory* in the pathway output. In dynamic belief networks, it is typical to assume temporal invariance of the conditional distributions, i.e., $\mathrm{P}(X_t|Y_t) = \mathrm{P}(X|Y)$ and $\mathrm{P}(Y_t|Y_{t-1}) = \mathrm{P}(Y|Y_{prev})$ for all $t$. This assumption is equivalent to stating that the parameters of these distributions are *homogeneous*—the relationship between pathway inputs and outputs does not change on the brief time scale of information processing modeled. The two CPDs, $\mathrm{P}(X|Y)$ and $\mathrm{P}(Y|Y_{prev})$, embody the knowledge in a pathway. In the following two sections, we discuss these forms of knowledge, which are the central claims of the model.

#### 2.1.1 Association strengths

The association strength between some $X = i$ (the random variable $X$ taking value $i$) and some $Y = j$ is formulated as $\mathrm{P}(X = i|Y = j) \sim \epsilon + \alpha_{ij}$ where $\epsilon$ is a constant representing the intrinsic difficulty of the task and $\alpha_{ij}$ indicates the familiarity with the association between states $i$ in $X$ and $j$ in $Y$. The greater the association strength, the more rapidly that information about $X$ will be communicated to $Y$.

Although the input representation is localist, in that there is one value of $X$ for each possible input, one can design in the similarity structure inherent in a distributed representation using explicit terms, $\gamma_{ik}$, that specify the similarity between input states $i$ and $k$:

$$\mathrm{P}(X = i | Y = j) \sim \epsilon + \sum_k \gamma_{ik} \alpha_{kj}$$

### 2.1.2 Short-term memory

We assume that the transition probability matrix from $Y_{prev}$ to $Y$ acts as a memory with diffusion. That is, with probability $\beta$, $Y$ is reset to its prior state and with probability $(1 - \beta)$, $Y$ remains in the same state as $Y_{prev}$:

$$\mathrm{P}(Y = i | Y_{prev} = j) = \begin{cases} (1-\beta) + \beta \mathrm{P}(Y = i) & \text{if } i = j \\ \beta \mathrm{P}(Y = i) & \text{otherwise} \end{cases}$$

where $\beta$ is the diffusion constant and $\mathrm{P}(Y)$ is the prior distribution (described later). If $\beta = 0$, the transition matrix acts as a perfect memory.

### 2.1.3 Processing dynamics

The distribution over $Y_t$ can be derived from Bayes' theorem, given the input sequence, $\mathbf{X}_t \equiv \{X_1, X_2, \ldots, X_t\}$, the association strengths encoded in $\mathrm{P}(X|Y)$, the pathway output memory in $\mathrm{P}(Y|Y_{prev})$, and the prior distribution $\mathrm{P}(Y_0 = k|\mathbf{X}_0)$, which we also write as $\mathrm{P}(Y = k)$:

$$\mathrm{P}(Y_t = k | \mathbf{X}_t) \sim \left( \sum_{i=1}^{N_Y} \mathrm{P}(Y_{t-1} = i | \mathbf{X}_t) \mathrm{P}(Y_t = k | Y_{t-1} = i) \right) \left( \sum_{j=1}^{N_X} \mathrm{P}(X_t = j) \mathrm{P}(X_t = j | Y_t = k) \right)$$

To model two pathways in cascade, such as the perceptual and response pathways, the output of the perceptual pathway is provided as input to the response pathway. Although the two pathways could be coupled into a single graphical model, inference in this model is intractable. Consequently, we approximate inference by assuming that at each time step the perceptual pathway output is copied to the response pathway input. This decoupling corresponds to the assumption of limited communication between pathways.

## 2.2 Learning mechanisms

We postulate that transmission of signals becomes more efficient with experience. By *more efficient*, we mean that a pathway produces the appropriate response with higher accuracy and more rapidly. By *experience*, we mean practice or repeated performance of a cognitive or motor skill. Efficiency is reflected in a leftward shift of the curve relating processing time to output probability (e.g., Figure 1a). The curve can shift in two ways, either by increasing the prior (Figure 2a), which raises the initial probability of the response, or by increasing the association strength between input and output (Figure 2b), which results in more rapid integration of the output probability. As a result of either change, the pathway is more accurate for a fixed amount of processing time, and the pathway is faster to attain a fixed level of accuracy.

We incorporated both mechanisms into our model. Following an experience in which input $i$ leads to activation of output $j$, the association strength $\alpha_{ij}$ is updated by the constant $\Delta\alpha$.

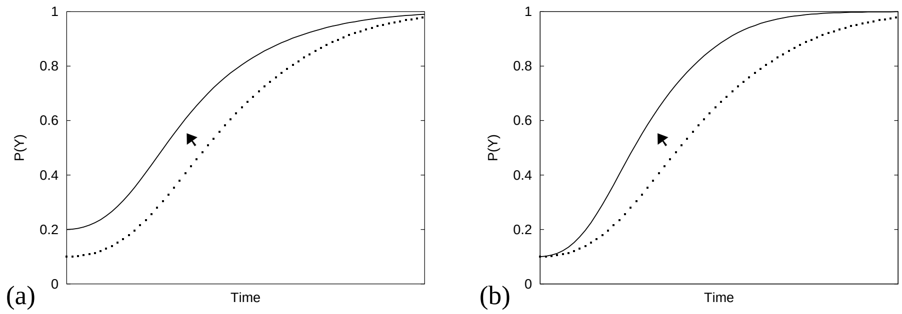


Figure 2: Change in the time course of activation of a pathway resulting from (a) an adjustment to the priors and (b) an adjustment to the association strengths

To form an analogous rule for the prior, we define the prior $\mathrm{P}(Y=j)$ in terms of secondary parameters, $\mathrm{P}(Y=j) \sim \kappa + \rho_j$, and update $\rho_j$ by the constant $\Delta\rho$.

# 3 Simulations

We model data from two psychological experiments on long-term repetition priming. Our simulations of these data utilized two pathways which were identical, except that the forgetting rates in the perceptual and response pathways were $\beta^p = 0.05$ and $\beta^r = 0.01$. The pathways were designed to produce 1-1 mappings, with $N_X = N_Y = 20$ input and output states for each pathway. We assume a similarity structure in which each pathway input is similar to two others with a uniform similarity coefficient $\gamma = 0.8$. Rather than independent $\alpha_{ij}$ parameters for each association, we used two values, $\alpha_{\text{high}} = 60$ and $\alpha_{\text{low}} = 10$, corresponding to high- and low-frequency words, respectively. Although there are two other parameters related to the associations, $\epsilon = 8$ and $\Delta\alpha = 5$, there is only one additional degree of freedom due to probability renormalization. Finally, the prior update rule has two free parameters, where we selected $\kappa = 1$ and $\Delta\rho = 3.3$, although one degree of freedom is also lost here due to probability renormalization. In total, the model had seven independent parameters, although the model's behavior was insensitive to the exact parameter values. One additional constraint was that we chose parameters such that one simulation time step corresponds to one millisecond in the experimental studies.

## 3.1 Experiment 1: bias effect

One explanation for the facilitatory effect of repetition priming is that study of the prime introduces a response bias that increases the probability of reporting the prime in the future. Ratcliff and McKoon (1997, Experiment 3) explored the bias account of priming in a 2AFC paradigm. During the test phase, masked target words were briefly presented, followed by a two alternative forced choice between the target and a distractor alternative. The target and distractor were orthographically similar, making the discrimination more difficult.

Three experimental conditions were contrasted: In the *congruent* condition, the target was presented during the study phase. In an *incongruent* condition, the distractor was presented during the study phase. In a neutral condition, neither was previously studied. For example, if DIED was studied, then target DIED with distractor LIED would be a congruent trial; target LIED with distractor DIED would be an incongruent trial; and target KICK with distractor SICK would be a neutral trial. The experiment also manipulated the *flash duration*, the asynchrony between target and mask onset.

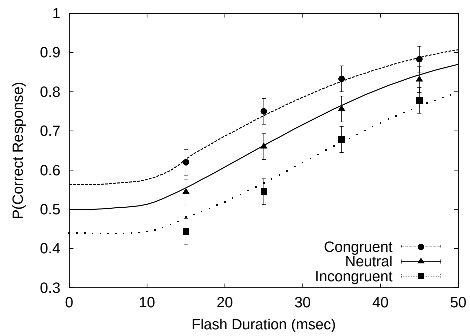


Figure 3: Accuracy of response for congruent, incongruent, and neutral conditions of the Ratcliff & McKoon (1997, Experiment 3) study of bias effects in priming. The points are results from human subjects, and the curves are produced by our model.

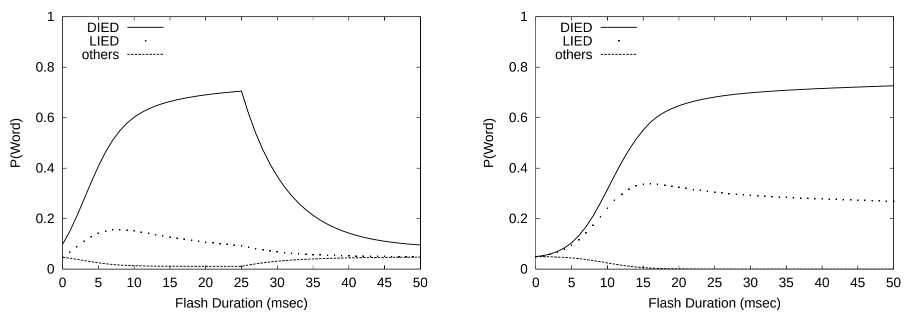


Figure 4: Output of the perceptual and response pathways (left and right panels) for a 25 msec presentation of the target DIED on a congruent trial

Human performance in the experiment is indicated by the data points in Figure 3 for flash durations of 15, 25, 35, and 45 msec. Across flash durations, the benefit on the congruent trials is balanced by the cost in accuracy on incongruent trials, diagnostic of a bias effect.

Our model produces an excellent fit to the data, as shown by the curves in Figure 3. The stimulus presentation durations were modeled by fixing the perceptual pathway input distribution, $\mathrm{P}(X_t)$, such that $\mathrm{P}(X_t = i) = 1$ for flashed word $i$. Following the flash duration, say at $t = 25$ msec, the pathway input was reset to the uniform distribution, causing the perceptual pathway to decay back to its prior distribution at a rate proportional to $\beta$, as shown in Figure 4a. The figure shows a model of a 25 msec presentation of the word DIED when DIED has been studied before, and the plot of the perceptual pathway shows that $\mathrm{P}(Y = \text{DIED})$ is higher at $t = 0$ than any other word. The response pathway, shown in Figure 4b, accumulates evidence from the perceptual pathway, reaching an asymptote as the perceptual pathway decays. To produce a 2AFC response, we renormalize the response pathway outputs conditional on the output being one of the two response alternatives.

Although two mechanisms of adaptation are built into the model, the prior update rule is almost entirely responsible for the differences in performance among conditions. Setting the association strength adjustment, $\Delta\alpha$, to zero has little impact on the simulation results. Thus, the prior update rule roughly corresponds to the notion of bias. However, the correspondence is only rough because as the flash duration increases, the differences among conditions asymptotically disappear. A simple rule that adjusted response probabilities independent of flash duration could not account for the data.

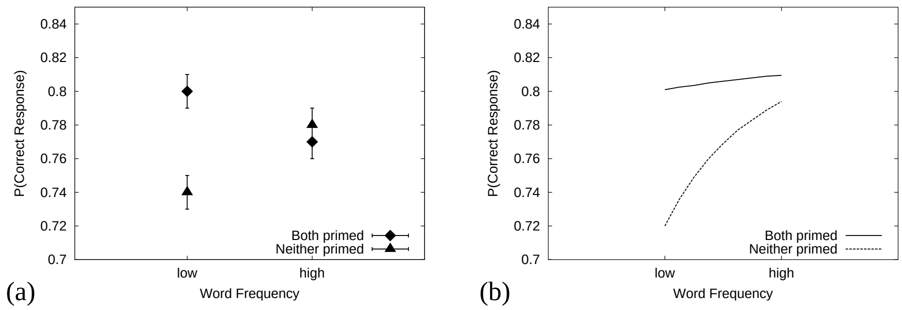


Figure 5: 2AFC accuracy for both- and neither-primed conditions, for low- and high-frequency words. (a) human data from McKoon and Ratcliff (2001); (b) simulation results from our model.

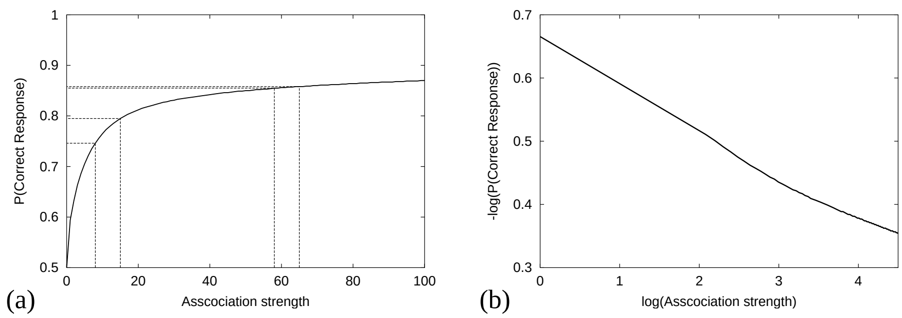


Figure 6: (a) The accuracy of the model as a function of association strength, $\alpha$, for a fixed flash duration. (b) A log-log plot of accuracy versus accociation strength.

### 3.2 Experiment 2: sensitivity effect

As a complement to the bias effect, the sensitivity effect is an improvement in perceptual discrimination of an item as a result of previous study. To determine if a sensitivity effect contributes to priming, McKoon and Ratcliff (2001, Experiment 2) studied a 2AFC task in which a comparison is performed between a condition in which both response alternatives are primed (*both primed*) and a condition in which neither response alternative is primed (*neither primed*). Any difference between these conditions could not be attributed to a bias effect, because the bias effect should cancel when both alternatives are primed. A reliable benefit in the both-primed condition relative to the neither-primed condition is diagnostic of a sensitivity effect. This experiment explicitly manipulated the frequency of the word stimuli. In one condition, uncommon, low-frequency words were used (e.g., WOMB and TWIG), and in another, high-frequency words were used (e.g., BEEN and THAN).

The human data showed a significant benefit of study for low-frequency words, but no statistically reliable effect for high-frequency words (Figure 5a). Thus, priming does produce a sensitivity boost for novel or low frequency items, but not for high frequency items. In our model, we assume that each experience with a word results in a constant increase in $\alpha$. Although $\alpha$ is proportional to word frequency, the benefit in performance diminishes with $\alpha$, as reflected by the convergence of the both-studied and neither-studied curves in Figure 5b. Thus, our model produces the same qualitative pattern as the human subjects—a large benefit of study for low-frequency words but little benefit for high-frequency words. This result is due to the adjustment of association strengths: if we turn off the adjustment of priors by setting $\Delta\rho = 0$, the qualitative pattern of results in Figure 5b is unaffected.

Figure 6a explains the reason for the diminishing benefit with word frequency. The graph shows the model's probability of correct response as a function of the association strength, $\alpha$. For equal $\Delta\alpha$ increments, the gain in accuracy is smaller for a high-frequency word than for a low-frequency word. Figure 6b is a log-log plot of accuracy versus $\alpha$, indicating a power law of practice. (A plot of response time versus $\alpha$ yields the same result.)

## 4 Discussion

We proposed a probabilistic model that characterizes the temporal dynamics of information transmission along cortical processing pathways, and explained key phenomena in the long-term repetition priming literature, including: the bias and sensitivity effects, the dependence of the sensitivity effect on word frequency, and the time course of priming within a trial. We have used exactly the same model to address other priming phenomena, including: the effects of target-distractor similarity, the decay of bias effects over time, alternative response paradigms including naming and matching, and response priming effects (Colagrosso, in preparation; Mozer, Colagrosso, & Huber, 2002). The model is compact and has few free parameters, yet it can explain a wide array of data. The elegance of the model stems in part from the Bayesian framework, which dictates the mechanisms of inference within a pathway, and in part from parameters that correspond directly to quantities of psychological interest, such as interitem similarity ($\gamma$) and degree of experience ($\alpha$).

A central claim of our work is that skill practice and refinement is subserved by two distinct mechanisms. In terms of our model, one mechanism adjust the pathway output prior probabilities and the other mechanism adjusts association strengths within a pathway. Both are extremely sensible mechanisms for an adaptive system. The priors can be viewed as a simple model of the environment, and updating this model is appropriate if encountering an object in one's environment implies that one is more likely to encounter the object in the future. The association strengths can be viewed as a limited-capacity resource, and allocating this resource to recently performed cognitive operations is judicious assuming that they are likely to be required again. Although these mechanisms are primitive forms of learning, they are the foundation of skill acquisition and skilled performance.


### Acknowledgments

This research was supported by Grant 97–18 from the McDonnell-Pew Program in Cognitive Neuroscience, NSF award IBN–9873492, and NIH/IFOPAL R01 MH61549–01A1.


### References


Colagrosso, M. D. (in preparation). *A rational theory of skilled performance and practice: Modeling long-term repetition priming.* Unpublished doctoral dissertation, University of Colorado.

Dean, T., & Kanazawa, K. (1989). A model for reasoning about persistence and causation. *Artificial Intelligence*, *93*, 1–27.

Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Ruys, K. I. (2001). Perception and preference in short-term word priming. *Psychological Review*, *108*, 149–182.

Kanazawa, K., Koller, D., & Russell, S. J. (1995). Stochastic simulation algorithms for dynamic probabilistic networks. In *Proceedings of the 11th Annual Conference on Uncertainty in AI (UAI)* (pp. 346–351). Montreal, Canada.

McKoon, G., & Ratcliff, R. (2001). Counter model for word identification: Reply to Bowers (1999). *Psychological Review*, *108*, 674–681.

Mozer, M. C., Colagrosso, M. D., & Huber, D. E. (2002). A rational analysis of cognitive control in a speeded discrimination task. In T. G. Dietterich, S. Becker, & Z. Ghahramani (Eds.), *Advances in Neural Information Processing Systems 14.* Cambridge, MA: MIT Press.

Ratcliff, R., & McKoon, G. (1997). A counter model for implicit priming in perceptual word identification. *Psychological Review*, *104*, 319–343.