A Race Model of Perceptual Forced Choice Reaction Time

David E. Huber (dhuber@psyc.umd.edu)
Department of Psychology, 1147 Biology/Psychology Building
College Park, MD 20742 USA

Denis Cousineau (Denis.Cousineau@UMontreal.CA)
Département de psychologie, C. P. 6128, succ. Centre-ville
Montréal (Québec), H3C 3J7 CANADA

Abstract

We present a race model for forced-choice data that provides a unified account of both latency and accuracy. The model is applied in the domain of short-term priming, but could characterize many other response tasks. A series of perceptual identification experiments found systematic bias changes as a function of prime duration. Notably, reaction times (RTs) were observed to change along with response bias. Furthermore, correct RTs changed in an opposite manner to error RTs. These results are explained by assuming a race between choice alternatives. The theory provides an alternative to signal detection theory, with faster finish times, rather than greater signal strength, determining both accuracy and RT.

Short-term priming

For more than 30 years it has been known that presenting related words results in facilitated processing. For instance, Meyer and Schvaneveldt (1971) observed that lexical decisions (deciding whether a letter string is a valid word) to target words were faster in the presence of an associated prime word (e.g., "doctor - nurse"), as compared to an unrelated prime word. Many additional experiments using lexical decision and other tasks, have observed similar facilitations for semantic priming, orthographic-phonemic priming, and repetition priming.

Prime identification and response bias

A question of theoretical interest is whether priming facilitations result from enhanced perceptual processing (i.e., more information extracted more quickly from the target) or from an item-specific bias such that participants are more likely to respond with primed words. Such a bias can be assessed using a forced-choice procedure in which incorrect foil words are also primed. For the perceptual identification task shown in Figure 1, where the task is to identify which of the final two words appeared as the target flash, a bias for primed words results in higher accuracy when the target is primed, and lower accuracy when the foil is primed, revealing a pattern of costs and benefits with priming.

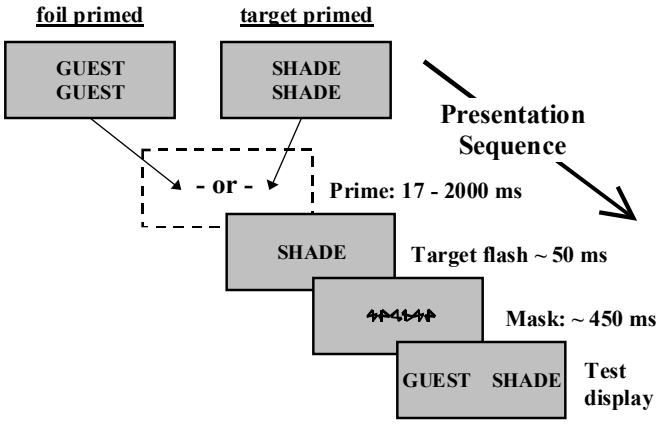

Figure 1: sequence of events on a trial in the forced-choice perceptual identification task, with repetition priming. Half the trials were target primed and half foil primed. The task was to choose which of the words in the test display appeared as the briefly flashed target. Target flash durations were determined separately for each participant, such that performance was near 75%. Mask durations were set equal to 500 ms minus the target flash duration. Participants were instructed that the prime was irrelevant to their task, considering that the incorrect foil and correct target were primed just as often. In addition, feedback at the end of each trial provided further evidence there was no effective strategy in relation to which choice word was primed.

Mean RT, accuracy, and a neural race model

Using the paradigm seen in Figure 1, Huber (in preparation) found that prime duration had a dramatic effect on bias, as indexed by accuracy (see Figure 2). After reaching a maximum bias to choose repeated words at 50 ms (i.e., target primed performance was better, and foil primed worse), further increases in prime duration caused a switch to a bias against repeated words (i.e., target primed performance was worse, and foil primed better).

In addition, there were equally dramatic reaction time (RT) effects as a function of prime duration. The prime duration at which accuracy changed was also the prime duration at which RT changed. However, the pattern was opposite for correct versus error trials (the middle and lower panels of Figure 2).

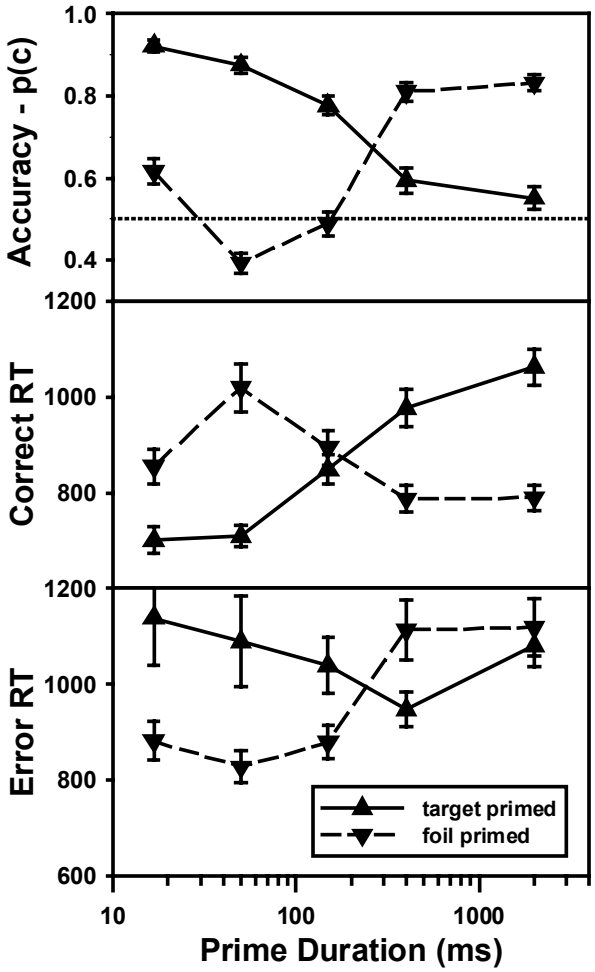

Figure 2: short-term repetition priming experiment (Huber, in preparation), using the paradigm seen in Figure 1. Prime durations occurred at 17, 50, 150, 400, and 2000 ms.

The neural network theory proposed by Huber and O'Reilly (2003) quantitatively captures these accuracy and mean RT data. In that model, processing of the choice words occurs in parallel, and the decision results from a race between the target and foil. The choice word identified first is chosen (i.e., accuracy), and the finish time of the winner determines RT. Residual activation from the target flash speeds target identification, which is the basis of accurate performance. However, residual activation from presentation of the prime likewise speeds identification of primed words, resulting in a bias. When the target is primed, this speeds target processing which boosts accuracy and results in faster correct RTs (i.e., trials where the target wins the race). When the foil is primed, this speeds foil processing which harms accuracy and results in faster error RTs (i.e., trials where the foil wins the race). Due to an accommodative process, long prime durations affect the primed racer in an opposite manner, causing sluggish, rather than speeded processing. In their model, accommodation results from transient synaptic depression, although the authors also consider an abstract version of the theory, relating accommodation to the Bayesian concept of "explaining away" observances in the face of known causes.

The race model of Huber and O'Reilly (2003) is deterministic, and they left specification of the noise process to future work. Critical to the development of appropriate forms of noise is the shape of the RT distribution. In this paper we model RT distributions with a generic version of the race model. Analysis of the best-fitting parameters will guide future development of appropriate noise mechanisms.

Generic race model

As outlined above, it is assumed that the target and foil race in parallel with the winner determining both the chosen response as well as the RT of that response. In order to quantitatively specify the theory, it is necessary to consider position and variability in the finish times of the target and foil; assuming specific finish time distributions (i.e., assuming a particular form of variability) allows calculation of the correct and error RT distributions.

Inverse signal detection theory

In signal detection theory, there is a distribution of signal strength for target present trials and another distribution for target absent trials (i.e., foil trials). The distribution for target present trials is typically shifted to the right, by some fixed amount, due to the extra evidence from the target. Assuming specific distributions (e.g., identical independent Gaussians, with shifted means), forced-choice accuracy can be calculated and is related to the area of overlap between the distributions.

A race between two racers is similar to forced-choice signal detection theory except that the target present distribution is shifted to the left (faster), and the distributions are finish times rather than signal strength (see Figure 3a). Slowing the target finish time, or speeding the foil finish time, results in greater overlap and therefore accuracy decreases (i.e., the gray area in Figure 3a becomes larger).

Figure 3a also provides an indication of RT distributions. Correct RTs occur when the target wins the race. Therefore, the vertical hatched area is indicative of correct RTs (i.e., RTs at which it is more likely that the target will win). Error RTs occur when the foil wins the race. Therefore, the horizontal hatched area is indicative of error RTs (i.e., RTs at which it is more likely that the foil will win). In order to precisely determine correct and error RTs, the probability of one racer finishing at a given time is multiplied by the probability of the other racer not finishing by that time (assuming processing independence between the racers). Equations 1 and 2 provide the correct (p_C) and error (p_E) RT probability functions (which equal the correct and error rates when integrated).

$$p_C(t) = f_T(t)\,[1 - F_F(t)] \qquad (1)$$

$$p_E(t) = f_F(t)\,[1 - F_T(t)] \qquad (2)$$

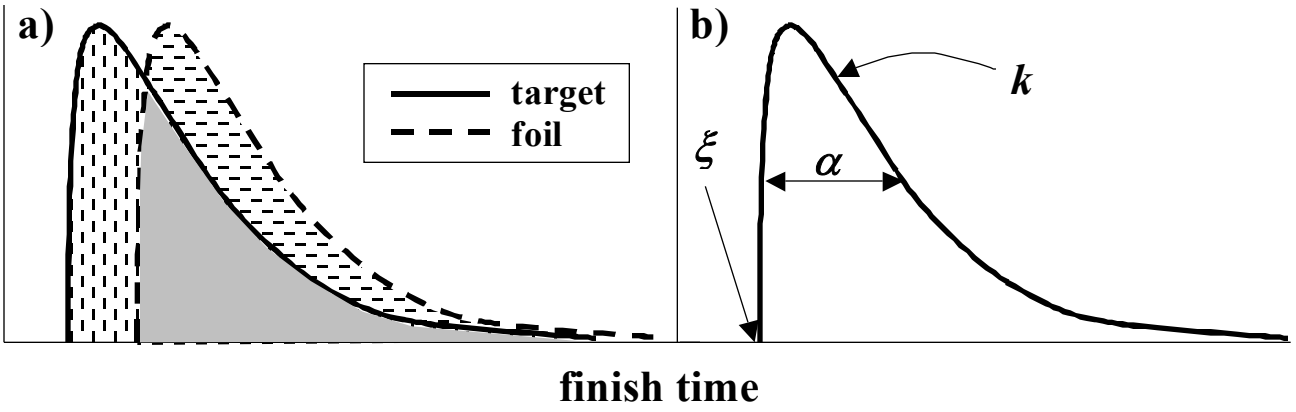

Figure 3: a) hypothetical finish time density functions for the target (f_T) and foil (f_F) racers, which have the corresponding target (F_T) and foil (F_F) cumulative distributions (see Equations 3 and 4 for the Weibull density function and cumulative distribution, and Equations 1 and 2 for the RT probability functions necessary to determine accuracy and RT distributions). The vertically hatched area is indicative of correct RTs and the horizontally area is indicative of error RTs. As in standard signal detection theory, the overlapping gray area is inversely related to accuracy. b) graphical representation of the Weibull parameters (onset-ξ, scale-α, and shape-k).

The Weibull distribution

There are many positively skewed finish time distributions that could be substituted into Equations 1 and 2. We selected the Weibull because it is the distribution of the winning time when there is a large number of similarly behaved racers.[1] In the current context, the finish time of each of the two choice word racers can be viewed as resulting from a race between a large number of perceptual features (Cousineau, submitted). For instance, assume a pool of features that uniquely identifies the target and another pool that uniquely identifies the foil. If any one of these features is accessed, then the identity of the corresponding word is known. Furthermore, since the two pools of features identify different words, the properties of the racers in each pool could be different, as modeled by different Weibull distributions for each pool. These properties relate to the Weibull's three parameters: onset (ξ), scale (α), and shape (k), which provide a readily interpretable description of the finish time distribution for each choice word (see Figure 3b). Equations 3 and 4 provide the Weibull density function and cumulative distribution.

$$f(t) = \frac{k}{\alpha}\left(\frac{t-\xi}{\alpha}\right)^{k-1} e^{-\left(\frac{t-\xi}{\alpha}\right)^{k}} \qquad (3)$$

$$F(t) = 1 - e^{-\left(\frac{t-\xi}{\alpha}\right)^{k}} \qquad (4)$$

A similar model was proposed by Logan (1988), in which racers were assumed to be memorized instances of presented items. Because the racers are internal in that model, it was not immediately obvious how it should be applied to the effect of externally presented choice alternatives, such as in the forced-choice procedure.

Fitting reaction time data

Huber, Curran, O'Reilly, and Woroch (submitted) ran a follow up study using the same paradigm as Huber (in preparation). They selected the 150 ms (short) and 2000 ms (long) prime durations, and recorded event-related potentials (ERPs) while participants performed the task. Because ERPs require a large number of trials per condition for each participant, these data are adequate for RT distribution modeling.

Empirical data

Figure 4 shows the observed correct and error RT distributions, averaged over participants, for the four conditions contained in Huber et al.'s (submitted) study. Consistent with a race model, the correct and error distributions are highly overlapping in the low accuracy conditions (i.e., the short foil-primed and long target-primed conditions). Also consistent with a race model, the correct distribution is faster than the error distributions in the high accuracy conditions (i.e., the short target-primed and long foil-primed conditions).

[1] The original theorem assumed independent and identically distributed racers. Cousineau, Goodman, and Shiffrin (2002) recently demonstrated that non-identical racers nevertheless produce Weibull finish times and avoid degeneration (Colonius, 1995).

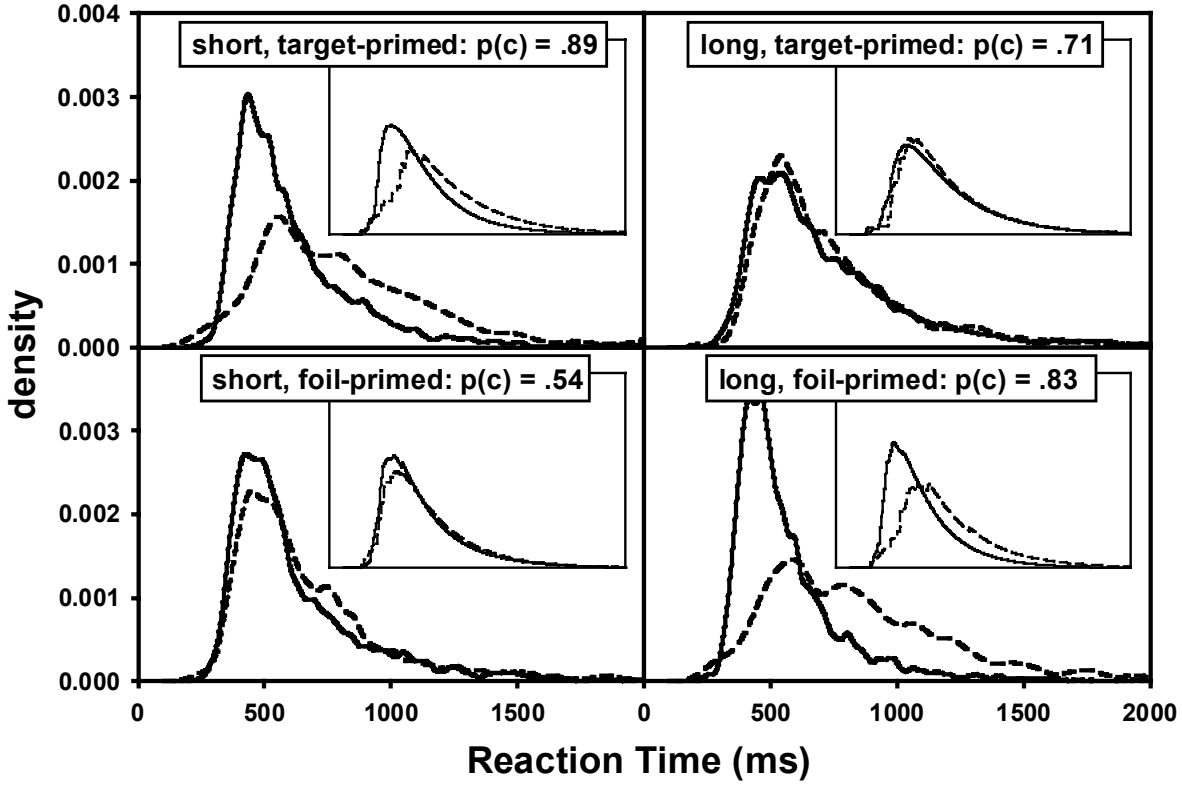

Figure 4: Huber, Curran, O'Reilly, and Woroch's (submitted) RT distributions. Solid lines indicate correct responses and dashed lines indicate error responses. Results with the best-fitting race model parameters are shown in the upper right-hand inset for each of the four conditions. The race model was applied separately to the data of each participant, although the results of these separate fits are combined in producing these aggregate plots.

Application of the race model

Assuming separate Weibull finish times for the target and foil racers, we modeled the data of Figure 4. This was done at the level of individual participants using the maximum likelihood method (Cousineau & Larochelle, 1997). This method attempts to maximize the probability of observing the particular sample of RTs given candidate correct and error RT probability functions. With 100 data points per condition for 33 participants, it is difficult to visualize the distributions of a typical participant. Instead, Figure 4 aggregates the separate individual best-fitting race model distributions. Because the race model was fit to individual data this allowed application of inferential statistics on the separate parameters for each participant (see Figure 5).

Using separate target onset, target scale, foil onset, and foil scale parameters for each condition and each participant, none of the individual predicted probability correct values was greater than .025 different than the observed values and only a few were greater than .01 different. Little was gained by allowing the shape parameter to freely vary and it was therefore fixed at 1.2 for all distributions. To produce the correct and error distributions in Figure 4, the best-fitting correct and error RT probability functions of each participant were averaged and then normalized by the predicted group accuracy values. The apparent noise in the model fits of Figure 4 result from individual differences and does not indicate variability in the simulation process.

Figure 5 depicts the best-fitting onset and scale parameters for the target and foil finish time distributions in each of the four conditions, as averaged across participants. For both the onset and scale parameters, in both the target- and foil-primed conditions, there was a highly significant interaction between prime duration and target versus foil racer (i.e., the lines in each panel are not parallel). In addition, all the parameters significantly changed between the short and long prime duration, except for the foil onset parameter in the target-primed condition.

As seen in Figure 5, the target and foil finish time parameters were conversely affected by prime duration. This occurred for both the onset (left panels) and scale (right panels) parameters. Increasing prime duration caused the finish time distributions to converge in the target-primed conditions, yet diverge in the foil-primed conditions. In other words, prime-induced changes in the finish time distributions were opposite, and roughly equal, for the target versus the foil finish times, and, furthermore, these effects were opposite for target versus foil priming.

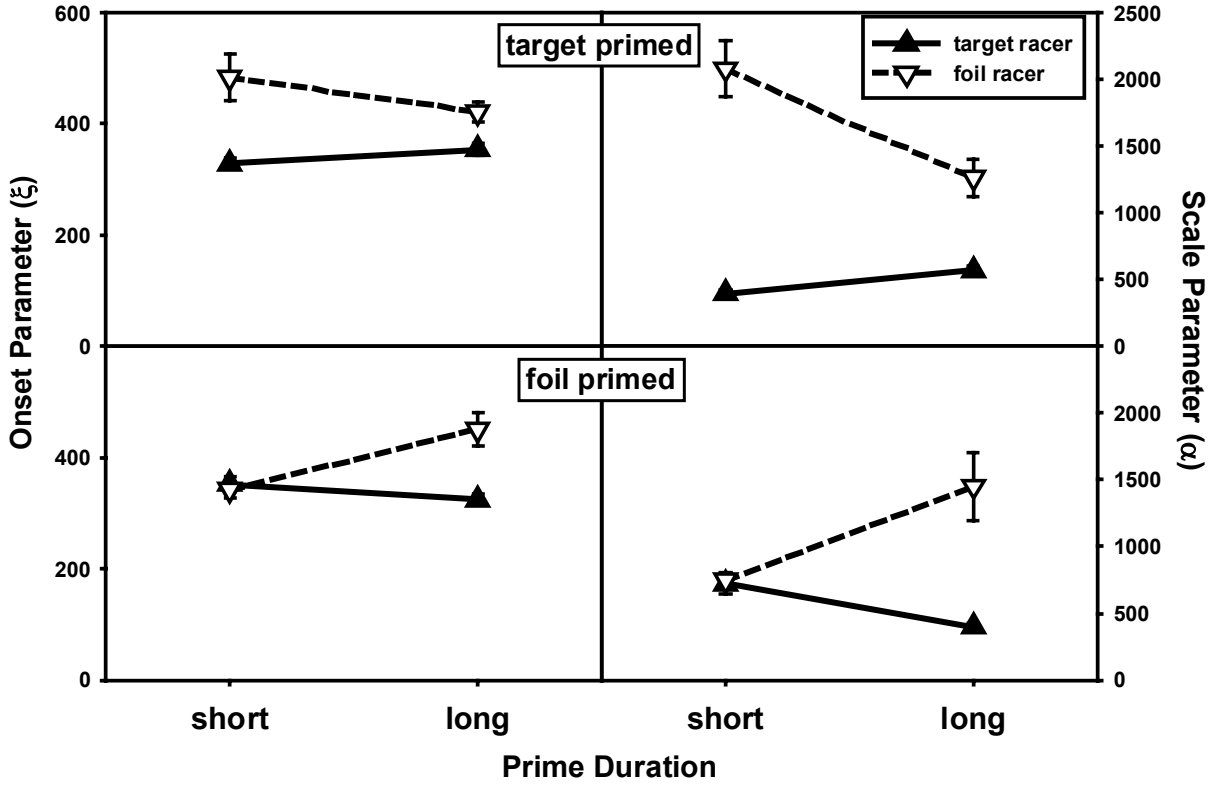

Figure 5: average best-fitting onset and scale parameters for the separate target and foil finish time distributions, for each of the four conditions shown in Figure 4. Error bars show the standard error of the mean, for variability of the parameters across participants. In moving from the short to the long prime duration, the primed racer is slowed down and the unprimed racer sped up, producing an opposite pattern of results for the target primed condition compared to the foil primed condition.

In summary, the short prime duration resulted in a speed up of the primed alternative and a slow down of the unprimed alternative. In contrast, the long prime duration resulted in a slow down of the primed alternative and a speed up of the unprimed alternative. These changes affected both the onset and scale in a similar manner (i.e., later finish time onsets corresponded to greater finish time variability).

Discussion

In this paper we present a generic, descriptive form of the race model, as applied to forced choice data from a perceptual identification task. Similar to applications of signal detection theory, we use the race model as a tool for describing accuracy and latency data under the assumption that the response alternatives accrue information in parallel. Other researchers have used race models to describe RT distributions, but such applications are few. More specifically, the Poisson race model has been applied in several domains, including perceptual identification (e.g., Van Zandt, Colonius, & Proctor, 2000). However, the Poisson race model assumes that the finish times for each racer are distributed as Gammas, as results from a Poisson accumulation process. Our assumption that finish times are distributed as Weibulls provides an interpretation in terms of a race between separate pools of racers for each response alternative. Because the Weibull itself is the distribution that results from a race process, we are essentially assuming a "race-race" process, placing into a final competition the fastest racers from separate pools of racers.

Future work will contrast this interpretation of these data with that provided by more traditional single accumulator models, such as a random walk or diffusion process (e.g., Ratcliff, 1978). A random walk does not describe the data in terms of the offsetting effects of the target versus the foil, but might, for instance, describe the data in terms of changes in the decision boundaries or starting point bias. In order to distinguish between these closely related explanations, we are currently employing other experimental techniques, such as Receiver Operating Characteristic (ROC) and signal to respond analyses.

Consideration of the best-fitting parameters demonstrates that priming one of the alternatives affected the other alternative in an opposite manner. This suggests a "rich get richer" interaction that could be realized through various mechanisms such as lateral inhibition or capacity limitations. For example, if the primed word is identified more rapidly, it may attract attention, removing processing capacities from the other word. It is not clear at this time whether such interactions imply processing dependence between the finish time distributions, but it is crucial to answer this question since it may be the main distinction

between a race process and a random walk process. Indeed, a race between perfectly negatively correlated racers is identical to a random walk. However, it is important to realize that observing opposite effects upon the race parameters does not imply a negative processing dependency. For instance, it may be that priming affects the properties of the two racers in an opposite manner, but, nevertheless, the race proceeds in an independent fashion once those properties have been established.

Discerning the nature of the competitive interaction observed in Figure 5 is an ongoing subject of research and understanding the relationship between accuracy and latency in these tasks will guide the development of mechanistic process models. In particular, Huber and O'Reilly (2003) proposed a detailed race model, based upon neural mechanisms. The current results with a descriptive race model will aid the extension of their model to include appropriate forms of variability.

References

Colonius, H. (1995). The instance theory of automaticity: Why the Weibull? *Psychological Review*, *102*, 744-750.

Cousineau, D. (submitted). Merging parallel race model and adaptive network: A parallel race network.

Cousineau. D., Goodman, V. W., & Shiffrin, R. M. (2002). Extending statistics of extremes to distributions varying on position and scale, and implications for race models. *Journal of Mathematical Psychology, 46,* 431-454.

Cousineau, D., & Larochelle, S. (1997). PASTIS: A Program for Curve and Distribution Analyses. *Behavior Research Methods, Instruments, & Computers*, *29*, 542-548.

Huber, D. E. (in preparation). The time course of perceptual processing for words and faces.

Huber, D. E., Curran, T., O'Reilly, R. C., & Woroch, B. (submitted). Testing a hierarchical perceptual network: Early brain potentials to immediately repeated words

Huber, D. E., & O'Reilly, R. C. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: Temporal segregation through synaptic depression. *Cognitive Science, 27*, 403-430.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492-527.

Meyer, D., & Schvaneveldt, R. W. (1971). Facilitation in recognizing parts of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, *90*, 227-234.

Ratcliff, R. (1978). Theory of memory retrieval & application over range of experimental paradigms. Psychological Review, *85*, 59-108.

Van Zandt, T., Colonius, H., & Proctor, R. W. (2000). A comparison of two reaction-time models applied to perceptual matching. *Psychonomic Bulletin and Review*, *7*, 208-256