

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Short-Term Word Priming Across Eye Movements

Permalink

<https://escholarship.org/uc/item/3g26h7ck>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 32(32)

ISSN

1069-7977

Authors

Denton, Stephen
Shiffrin, Richard M.

Publication Date

2010

Peer reviewed

Short-Term Word Priming Across Eye Movements

Stephen E. Denton (sedenton@uemail.iu.edu) and Richard M. Shiffrin (shiffrin@indiana.edu)

Department of Psychological and Brain Sciences, and Cognitive Science Program, Indiana University
1101 E. 10th Street, Bloomington, IN 47405 USA

Abstract

The authors conducted a short-term repetition priming experiment (using a visual, forced-choice word identification task) that compared a standard priming condition, where prime and target words appeared in the same spatial location, with an experimental condition in which prime and target words were spatially separated enough to necessitate an eye movement. Prime presentation duration was manipulated and, within both eye movement conditions, it was found that short primes produced a preference to choose a primed alternative, whereas for longer duration primes this preference was absent. Based on the similarity between eye movement conditions, it is argued that prime and target features from separate fixations are still confusable and that evidence regarding prime feature must still be discounted. A computation model that includes these offsetting components of source confusion and discounting provides an excellent account of our results.

Keywords: short-term priming; immediate priming; repetition priming; perceptual identification

The term *priming* refers to a well-known information processing effect wherein one stimulus (a *prime*) influences a similar or related stimulus (a *target*) presented at a later time. The influence the prime has on the target is usually one of facilitation. In the typical priming task trial, the prime is presented first followed by a target that is briefly flashed and masked. This paradigm is referred to as *short-term* or *immediate* priming because primes are presented immediately prior to the targets, with inter-stimulus intervals generally less than a second. The first stimulus is called a prime because it is thought to “prime the pump” for a related target, yielding faster and more accurate responding. However, priming does not simply result in facilitation. Previous research has indicated that there is an intricate set priming effects that occur across various situations and experimental conditions (e.g., Huber, Shiffrin, Lyle, & Ruys, 2001; Weidemann, Huber, & Shiffrin, 2005). Some experiments find facilitation, while other fail to find facilitation or even find reliable deficits due to priming.

Huber et al. (2001) began to systematically examine the effects of priming using a two-alternative forced-choice (2-AFC) testing procedure with the goal of separating the perceptual and decisional components involved in priming. Observers were asked to correctly identify a previously flashed target word when given a choice between it and an incorrect foil word. This study indicated that priming largely arises from preferences to choose whatever has been primed. For example, in repetition priming (priming in which the prime can be the same word as either the target or foil), it was found that with short prime presentation durations, priming with the target increases accuracy and priming with the foil decreases accuracy, when both are compared to a control condition in which the prime is unrelated to either the target

or the foil. The corresponding prime conditions are termed *target-primed*, *foil-primed*, and *neither-primed*, respectively.

The preference effects revealed by Huber et al. (2001) proved to be readily manipulatable, changing in magnitude and direction as a function of prime saliency (e.g., Huber, Shiffrin, Quach, & Lyle, 2002; Weidemann et al., 2005; Weidemann, Huber, & Shiffrin, 2008). Thus, when the prime is made more salient, either through a longer presentation duration or through an active task that required responding to the prime, the conventional priming effect diminishes or in some cases reverses. Prime saliency manipulations result in reduced facilitation (or slight deficits) in target-primed conditions, while leading to increased accuracy in foil-primed conditions.

The ROUSE Model

To account for a range of findings within the 2-AFC identification paradigm, Huber et al. (2001) developed a feature-based Bayesian model of short-term priming. The *responding optimally to unknown sources of evidence* (ROUSE) model accounted for experimental priming data by incorporating the two offsetting components of feature source confusion and discounting. The source confusion portion of the model posits that features can be carried over from the prime to the target percept, without source information. Thus, when a choice word is presented, feature activations could be due to the prime presentation, the target presentation, and/or noise without the source of the activation being known to the system. This factor alone can produce the standard priming effect, as it causes a preference toward prime-related choice words. The discounting mechanism in the model can counteract this preference, because this component posits that perceived features are assigned evidence and feature evidence is discounted when known to have come from the prime. A Bayesian decision process then arrives at an optimal response given it is operating on this noisy and discounted evidence. The implication here is that making the prime more salient (e.g., long presentation duration) leads to increased discounting of prime feature evidence. Thus, discounting mechanisms can explain a lack of, or a reversal in, the typical priming preference.

In the ROUSE model, choice words are represented as a feature vector, typically consisting of 20 binary features (Huber et al., 2001). Features can be independently activated by the prime, with a probability α , by the target, with a probability β , or be activated due to noise, with a probability γ . The system is assumed to only have access to what features are active and not the source of their activation (i.e., source confusion), therefore the probabilistic effect of α , if

not counteracted, will result in a preference toward primed choice words. This preference is counteracted by a decision process that assigns lower levels of evidence to features that might have been activated by the prime word (i.e., discounting). To discount features optimally, the system needs to know the probabilities that features are active due to particular sources (i.e., α , β , and γ). However, it is assumed that the decisional system does not have access to the exact probabilities and therefore uses estimates, α' , β' , and γ' . These estimates, which reflect the amount of discounting applied to particular sources, are used to evaluate evidence. These estimates are theorized to be close to their true values, but the model's behavior depends critically on the magnitude and direction of the difference between α and its estimate. Under-discounting ($\alpha' < \alpha$) results in a preference for primed choice words and over-discounting ($\alpha' > \alpha$) produces a preference against primed choice words.

Given the estimated source probabilities, the optimal response among choice words can be computed by combining feature evidence. Feature evidence takes the form of a likelihood ratio that specifies the probability that a feature is from the target word over the probability the feature is part of the foil, given the feature's current activation state and whether or not the feature appeared in the prime. Assuming feature independence, these likelihood ratios can be multiplied together across all word features to produce an overall likelihood that the choice word is the target. The likelihoods of the two words can then be compared with the larger being identified as the target. In the case of ties, a random selection is made between the choice words.

The above description of the ROUSE model is only intended as a brief summary, which means a number of details, large and small, have been left out. For a more comprehensive presentation of the ROUSE model, the reader is referred to Huber et al. (2001) for the original stochastic version and Huber (2006) for a later analytic version.

Feature carry-over across eye movements

The ROUSE model has been useful in accounting for and predicting priming data, but there is still much uncertainty about how its key components, source confusion and discounting, actually perform their theorized functions. Many questions remain unanswered regarding what causes or allows both feature confusions and discounting. Source confusion could be the result of a number of candidate causes, e.g., spatial or temporal proximity between prime and target, or similarity between prime and target on any number of dimensions. Similarity relations (semantic, orthographic, etc.) between target and prime have received considerable attention elsewhere and will be set aside for present purposes. Temporal proximity does appear to be important as increasing the inter-stimulus intervals (ISIs) between prime and target to be greater than 250 ms has been shown to disrupt and diminish priming (Hochhaus & Marohn, 1991). In the following experiment, we chose ISIs to be less than 250 ms so that we can effectively ignore the temporal dynamics of priming for

the purposes of the present study.

The focus of this research is how spatial proximity effects priming. We are interested in knowing if feature confusions, discounting, or both are location dependent. In the present study, the location at which the prime and target were presented were sometimes spatially separated at a 10° visual angle. Visual acuity drops to roughly 25% (of central vision) at a visual eccentricity of 10° (Low, 1951). Given diminished visual acuity with this spatial separation, we are effectively enforcing an eye movement between viewing the prime word and the target word, if both are to be seen. Not only will the target appear in a different location than the prime, but it will also appear in a different eye fixation.

The goal of the present experiment is to see if features are carried over from a prime fixation to a target fixation when the fixation locations are relatively distant spatially. We would like to see if features from the prime join with the target percept at the new location (i.e., will there still be source confusion when an eye movement is made). Discounting could also be affected by our eye movement manipulation. If there is some source confusion across eye movements, it is conceivable that prime features could be discounted completely (i.e., little evidence from the prime would be used in making the target identification decision). Then again, given the presence of some source confusion, it is also possible that better estimates of feature activations due the prime can be made when the prime and target appear at separate and distinct locations. This would reduce or eliminate under-discounting and over-discounting of prime evidence. The ROUSE model will be used to provide an indication of the relative contribution of source confusion and discounting across eye movement conditions.

Priming Across Eye Movements Experiment

We designed the current experiment to investigate how spatial separation between prime and target affects priming. The experiment compared an eye movement condition to an appropriately matched control condition in which all words were presented in the center of the screen (i.e., the stereotypical short-term priming paradigm). For all trials, the prime word was presented in the center of the screen. As with our previous studies, prime salience was manipulated by adjusting the prime presentation duration, which was either short (50 ms) or long (1000 ms). On any particular trial, the prime word itself could be either the same as the target, the foil, or neither, corresponding to the prime conditions discussed earlier. The target word was flashed in the center of the display on half to the trials and on the other half it appeared equally often in one of four locations, directly above, below, left or right of the center. The participant's task was to identify the flashed target given a 2-AFC test at the end of each trial.

The experiment was a within-subject design that crossed the 2 levels of prime presentation duration (short and long) with 3 levels of prime condition (target, foil, or neither primed) and crossed both with 5 different target locations split




Figure 1: An example sequence of events in the present experiment. This figure shows both short and long prime duration conditions for one particular peripheral target location (left of center). The duration for each frame is presented on the right. The left sequence corresponds to a short prime duration condition, whereas the right sequence shows a long prime duration sequence where the total time the prime is presented is 1000 ms (850 ms+100 ms+50 ms). Both sequences were preceded by a 500 ms presentation of a fixation cross in the center of the screen (not shown). The positions of the target and foil word on the last frame were randomized.

into 2 general categories (center and peripheral [top, bottom, left, and right]). A schematic of two example trials is presented in Figure 1. A more detailed description of our exact experimental procedure is below.

Methods

Participants Fifty one undergraduate students volunteered to participate for partial credit in an introductory psychology course at Indiana University.

Materials and Apparatus We used two pools of 1,100 five-letter and 1,300 six-letter words with a written-language frequency of at least 4 per millions as defined by Kučera and Francis (1967). All words were presented in uppercase using the fixed-width “Courier Bold” 17-point font. Throughout the experiment, stimulus words were sampled randomly without replacement with the only constraint being that 5 and 6 letter word never appeared together in the same block of trials. All masking was done with two rows of six “@” signs presented in “Arial Narrow Bold” 13-point font. This ensured dense coverage of the prime and target. The stimulus words were presented in white against a black background.

All stimuli were displayed on 16-in. (40.6 cm) PC CRT monitors with vertical refresh rates of 120 Hz and a screen

resolution of 800×600 pixels. The experiment was programmed using the Vision Egg library for the python programming language (Straw, 2008). The display was synchronized to the vertical refresh of the monitor providing display increments of 8.33 ms.

Participants sat individually in a dimly lit, ventilated, sound-dampened booth. Participants were asked to sit up straight to keep the distance from their eyes to the monitor at approximately 50 cm., but no head restraint was used to enforce this viewing distance. This viewing distance ensured that peripheral targets would appear at least a 10° visual angle away from the center of the screen. Participant responses were collected using a standard computer keyboard. In a 2-AFC test, participants were asked to press the “F” key to choose the left alternative or the “J” key to choose the right.

Procedure The procedure used in the present experiment was carefully designed as to compare an experimental condition involving eye movements to an appropriate control that maintained important aspects of the experimental condition (namely timing and perceptual masking), while not requiring an eye movement. The experiment was designed based on some knowledge of the timing of eye movements or saccades. A 10° saccade would take less than 50 ms to complete

once initiated, but takes more than 150 ms to plan and initiate (Irwin, Brown, & Sun, 1988). As such, having the target flash appear in the periphery immediately after the prime would make it impossible to see. To remedy this, we preceded the target with a cue that indicated the correct location where the target would appear. This target indicator cue appeared 250 ms before the target flash. We left a 100 ms interval between the prime offset and the target onset. During this 100 ms interval the prime was post-masked while the target was pre-masked. This timing is reflected in the example sequences presented in Figure 1.

Each trial began with a fixation cross appearing in the center of the screen for 500 ms. In long prime duration conditions, the prime is presented (in duplicate) for 850 ms then the prime and target indicator appear together for 150 ms. In trials with short prime durations, the target indicator appears first by itself for 100 ms, then the prime and the target indicator appear together for 50 ms. Participants should not be able to plan and initiate a saccade in under 100 ms, therefore the prime will be viewed; the prime is ‘snuck-in’ before the eyes have a chance to move. After the prime and the 100 ms intervening mask, the target is flashed. The target is post-masked and then the 2-AFC options are presented to the immediate left and right of the location where the target appeared. A peripheral target word can appear at the top, bottom, left, or right of the screen. Center target location trials have the same timing as peripheral target trials. Besides the actual target location, the only differences between center and peripheral trials are that the target indicator appears in the center of the screen (surrounding the prime) and the prime post-mask and the target pre-mask are one in the same. On each trial, once a 2-AFC selection is made, feedback is provided.

Each participant went through 672 priming trials broken into 7 blocks of 96. The first 32 trials were neither-primed practice trials with long target durations (150 ms) to get participants used to the task. These practice trials were followed by 64 neither-primed calibration trials. Target word durations were individually adjusted for each subject such that accuracy was roughly 75% on neither primed conditions. This calibration was done separately for the center and peripheral target locations using a staircase method. As with previous studies (e.g., Huber et al., 2001; Weidemann et al., 2005), there were large individual differences. For center target trials, target flash times ranged from 25 ms to 91.7 ms with a median of 50 ms. For peripheral targets, flash times ranged from 33.3 ms to 200 ms (the maximum allowed) with a median of 91.7 ms. The increased variance in target flash times for peripheral locations is likely due to individual differences in saccade latency on top of individual differences in target processing time, i.e., some participants may have fixated on the target later and took longer to process the target.

Results

Data from all peripheral target locations were combined for the purposes of analysis, creating two target location conditions, central and peripheral, that initially had equal sample

sizes. Reaction times were collected and used to eliminate deviant trials in the data. Trials in which the participant responded in less than 100 ms or took more than 3-s to respond were eliminated. Approximately 1% of the data was thrown out by this criterion. The first block of trials, which included practice and calibration trials, were not included in the analysis. Remaining experimental data were analyzed with a $3 \times 2 \times 2$ (Priming Condition \times Prime Duration \times Target Location) repeated measures analysis of variance (ANOVA).

There were large main effects of prime condition, $F(2, 100) = 167.6, p < .001$, and prime duration, $F(1, 50) = 135.5, p < .001$. Also, these two variables interacted, $F(2, 100) = 121.37, p < .001$. The main effect of target location was not significant, however this variable had a significant interaction with both prime condition, $F(2, 100) = 20.43, p < .001$, and prime duration, $F(1, 50) = 8.637, p < .005$. Finally, there was a significant 3-way (Priming Condition \times Prime Duration \times Target Location) interaction, $F(2, 100) = 48.57, p < .001$.

Average accuracy across all conditions is shown graphically in Figure 2. The dots in the figure give ROUSE model predictions which will be discussed in the next section. As can be seen from Figure 2, the typical priming effect was found when the target and a short prime were both presented in the center of the screen. When the prime was presented for a longer duration this prime preference disappeared. When the target was presented in the periphery after a centrally presented prime, the trend remained largely the same, although the magnitude of the priming effect decreased.

Applying the ROUSE model

Fitting the ROUSE model to the current experiment involved estimating the following eight parameters:

1. the probability that a choice word feature is activated by the prime (α),
2. the estimated probability that a feature is activated by a short prime when the prime and target are located in the center ($\alpha'_{S,C}$),
3. the estimated probability that a feature is activated by a long prime located in the center ($\alpha'_{L,C}$),
4. the estimated probability that a feature is activated by a short prime from a previous fixation ($\alpha'_{S,P}$),
5. the estimated probability that a feature is activated by a long prime from a previous fixation ($\alpha'_{L,P}$),
6. the probability that a feature is activated due to the target flash (β),
7. the probability that a feature is activated due to noise given a short prime presentation (γ_S), and
8. the probability that a feature is activated due to noise given a long prime duration (γ_L)¹

¹Two separate noise parameters were included to account for trend in the data that performance is better on neither-primed trials when prime duration is long compared to when the prime duration is short. The γ values from the to-be-discussed model fit indicate there is higher noise in the short prime conditions when compared to long prime conditions. This is somewhat sensible given the erratic na-

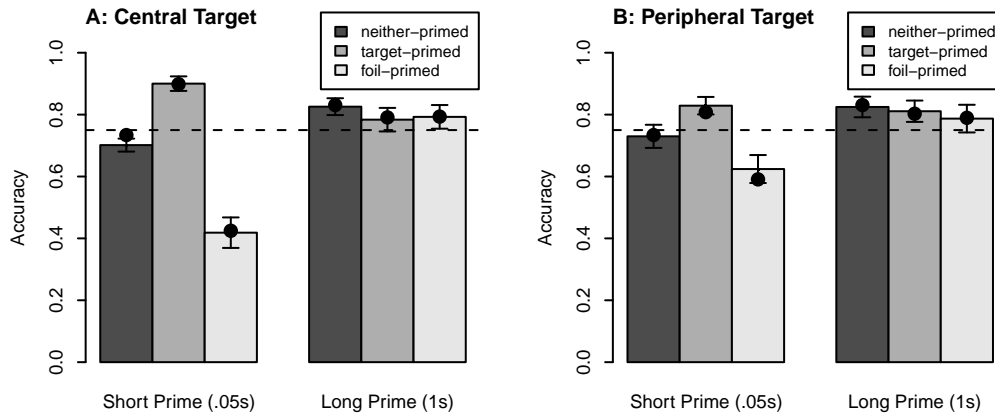


Figure 2: Forced choice performance and corresponding ROUSE predictions (represented by the dots). The bar heights show the mean proportions of correct target identification choices (error bars show $\pm 95\%$ confidence intervals) within each condition. Panel A shows accuracy for centrally located targets and Panel B shows accuracy when the target was presented in a peripheral location. Each panel is further subdivided by prime duration and priming condition. The dashed horizontal line indicates 75% performance; the accuracy participants should roughly achieve on neither-primed trials due to the target duration calibration procedure.

As in previous work (e.g., Huber et al., 2001; Weidemann et al., 2005), it is assumed that estimates of feature activations due to targets and noise are equal to the actual values (i.e., $\beta' \equiv \beta$ and $\gamma' \equiv \gamma$).

The parameters listed above were estimated to generate the ROUSE model fit that appears in Figure 2. The parameter estimates used were: $\alpha = .12$, $\alpha'_{S,C} = .032$, $\alpha'_{L,C} = .65$, $\alpha'_{S,P} = .10$, $\alpha'_{L,P} = .20$, $\beta = .064$, $\gamma_S = .067$, and $\gamma_L = 0.011$. It is important to stress that exact values of these parameters are relatively immaterial, especially the exact values of the α and α' parameters, as it is their relative magnitudes that dictate the behavior of the model. Multiple model fits were done with both more and less parameters free to vary, but this fit provided the most satisfying account to the data. The large number of parameters (8 parameters for 12 data points) and possible over-fitting should be somewhat of a concern here, but since the model is largely descriptive, these complexity concerns are left unaddressed in the present context.

As can be seen in Figure 2, the ROUSE model provides a very good qualitative and quantitative description of the priming data we collected. Interpretation of the fit parameter values can be a tricky, because attending to the exact magnitude can be deceptive. For example, the same α value was used to fit all conditions. When the 2 target location conditions

ture of short duration priming trial (i.e., the target indicator, prime, mask, and target all appear in rapid succession) which could induce noise in the system. This seems especially true in neither-primed trials where all prime features presented are effectively noise. Increasing noise (γ) in the model decreases accuracy and uniformly translates all performance predictions down. Separate β parameters could be used similarly to selectively increase and decrease performance across short and long prime duration conditions, but saying that there is an increased probability of target features being activated for long prime durations provides a much less sensible explanation of the results.

were allowed to have separate freely estimated α parameters it produced virtually no fit improvement. This could be taken as evidence that source confusion was the same at peripheral locations as at the center. However, α' was fit separately for each condition and changes in these values can compensate for possible underlying differences in source confusion (α).

Given the relationship between the fit α' values, it does appear that the magnitude of discounting decreased when the target was moved to a novel location. The estimates of primed feature intrusion contracted closer to optimal in the peripheral conditions (i.e., $\alpha'_{S,C} < \alpha'_{S,P} < \alpha$ (optimal) $< \alpha'_{L,P} < \alpha'_{L,C}$), which means there was less under-discounting for short primes and less over-discounting for long primes when the prime and the target did not appear in the same spatial location.

Discussion

In the present research, we investigated whether spatial proximity of prime and target was necessary to find priming effects, particularly when there was a large spatial distance between the two requiring an eye movement. Our findings indicated that even after a relatively large eye movement (a 10° visual angle), participants showed similar priming preferences as when they viewed all stimuli in one location. More specifically, in both eye movement and control conditions, participants showed a preference to choose the prime word when the prime duration was short. Further, this prime preference was undetectable when the prime was made more salient by increasing its presentation duration.

The ROUSE model provided an excellent account of our experimental data. The model's success, the resulting best fitting parameters, and data themselves all provide evidence that across eye movements (a) source confusion is still present,

and (b) discounting of evidence regarding prime features persists. The evidence that source confusion endures after an eye movement to a novel location implies that prime features appearing in one eye fixation are carried over to the next and join with the target precept at the new location. Further, source confusion from two spatially distinct locations does not appear to be substantially different from what it is when all feature sources are in the same spatial location.

The presence of discounting across eye movements implies that the decisional system is, in some sense, aware that features are not tied to particular locations and therefore must continue to estimate the likelihood that feature activations in choice words are the result of the prime in an effort to produce optimal responses. Even when the prime was in a previous eye fixation, the estimation process evidently succumbs to the same biases that are present without an eye movement, i.e., short, difficult to detect primes are under-discounted and long, salient primes are over-discounted. Since evidence for the prime is only discounted to the extent that features are known to exist in the prime, it is conceivable that the distinct spatial location of the prime provides some additional information to the decisional system that helps it better estimate the actual probability that a feature activation is due to the prime. Consequently, spatially separating the prime and target would make discounting slightly more optimal. Our modeling results hint that this is the case as the model estimates of feature intrusions from the prime better match the actual intrusion probabilities in the peripheral condition.

Our study suggests that the offsetting mechanisms of source confusion and discounting operate on features that are in general not tied to specific locations. Eye saccades are thought to effectively erase iconic memory, which is a visual store with high capacity but with a very limited duration (Irwin, 1992). There are a few visual items (3 to 4) that are retained from one eye fixation to the next in what is termed trans-saccadic memory, which has a limited capacity compared to iconic memory but has a longer duration (at least 750 ms). Nevertheless, an enforced eye movement will eliminate most low-level visual features from the previous fixation. After a saccade, the features that remain are presumably more high-level and location-independent. If source confusion and discounting operate with such features then an eye movement would not eliminate priming effects. This conforms to the findings of the present study. It is possible that source confusion and discounting operate on low-level features as well, and this may provide some explanation as to why eye movements slightly degrade the effectiveness of priming. The extent to which priming involves low-level features, and if it is indeed their suppression during saccades that produce the minor differences across eye movement conditions, are topics for future research.

Acknowledgments

This research was supported by NSF Grant 6804643 awarded to the second author.

References

- Hochhaus, L., & Marohn, K. M. (1991). Repetition blindness depends on perceptual capture and token individuation failure. *Journal of Experimental Psychology: Human Perception and Performance*, *17*(2), 422–432.
- Huber, D. E. (2006). Computer simulations of the ROUSE model: An analytic simulation technique and comparison between the error variance–covariance and bootstrap methods for estimating parameter confidence. *Behavior Research Methods*, *38*(4), 557–568.
- Huber, D. E. (2008). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology: General*, *137*(2), 324–347.
- Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Quach, R. (2002). Mechanisms of source confusion and discounting in short-term priming 2: Effects of prime similarity and target duration. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(6), 1120–1136.
- Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Ruys, K. I. (2001). Perception and preference in short-term word priming. *Psychological Review*, *108*(1), 149–182.
- Huber, D. E., Shiffrin, R. M., Quach, R., & Lyle, K. B. (2002). Mechanisms of source confusion and discounting in short-term priming: 1. effects of prime duration and prime recognition. *Memory & Cognition*, *30*(5), 745–757.
- Irwin, D. E. (1992). Memory for position and identity across eye movements. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(2), 307–317.
- Irwin, D. E., Brown, J. S., & Sun, J.-S. (1988). Visual masking and visual integration across saccadic eye movements. *Journal of Experimental Psychology: General*, *117*(3), 276–287.
- Kučera, H., & Francis, W. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.
- Low, F. N. (1951). Peripheral visual acuity. *Archives of Ophthalmology*, *45*, 577–593.
- Straw, A. D. (2008). Vision egg: An open-source library for realtime visual stimulus generation. *Frontiers in Neuroinformatics*, *2*(4), 1–10. (doi: 10.3389/neuro.11.004.2008)
- Weidemann, C. T., Huber, D. E., & Shiffrin, R. M. (2005). Confusion and compensation in visual perception: Effects of spatiotemporal proximity and selective attention. *Journal of Experimental Psychology: Human, Perception, and Performance*, *31*(1), 40–61.
- Weidemann, C. T., Huber, D. E., & Shiffrin, R. M. (2008). Prime diagnosticity in short-term repetition priming: Is primed evidence discounted, even when it reliably indicates the correct answer? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(2), 257–281.