Refixations gather new visual information rationally

Yunyan Duan (yduan@u.northwestern.edu) and **Klinton Bicknell (kbicknell@northwestern.edu)**
Department of Linguistics, Northwestern University
Evanston, IL 60201 USA

Abstract

The standard model is that word identification in reading is a holistic process, most efficient when words are centered in the visual field. In contrast, rational models of reading predict word identification to be a constructive process, where readers efficiently gather visual information about a word, and may combine visual information about different parts of the word obtained across multiple fixations to identify it. We tease apart these accounts by arguing that rational models should predict that the most efficient place in a word to make a second fixation (refixation) depends on the visual information the reader has already obtained. We present evidence supporting this prediction from an eye movement corpus. Computational model simulations confirm that a rational model predicts this finding, but a model implementing holistic identification does not. These results suggest that refixations can be well understood as rationally gathering visual information, and that word identification works constructively.

Keywords: eye movements; reading; word identification; rational analysis; refixation

Introduction

Reading is a complex information processing task with a goal usually related to comprehending the text. In general, accurate text comprehension requires the identification of many (if not most) of the words in a text. It is not surprising, then, that decades of research on eye movements in reading have established that word identification can be seen as the primary driver of eye movements (Rayner, 1998). A substantial body of work has studied the role in this process of many information sources relevant to word identification in reading, including especially word frequency and in-context predictability, among others. However, although visual information is the primary source of information used to ultimately identify a word, the fundamental way in which visual information is used in word identification remains unresolved.

In the standard model of word identification in reading, word identification is hypothesized to be a holistic process, during which visual information about the word as a whole constrains the efficiency of identification. Eye movement studies have shown that a word presented in isolation is most rapidly identified when fixating approximately at its center (O'Regan, 1990, 1992). It has also been found in natural reading that the fixation position that minimizes gaze duration (the total amount of time spent fixating a word in first pass) and refixation probability (the probability of making more than one fixation on a word in first pass) is on average at or slightly left of the center (Rayner, Sereno, & Raney, 1996). One explanation for these results is that when the word center is directly fixated, the largest possible part of the word falls in the central high-acuity portion of the visual field (the fovea), yielding the highest-quality visual input of the whole word; as the fixation deviates from the center, more letters of the word fall out of the fovea and suffer from a rapid drop in acuity, leading to poorer visual information about the overall word. Following this interpretation, it is hypothesized that visual processing efficiency of a word is maximized when fixating at word center, and decreases with increasing distance between word center and fixation position. This standard holistic account is incorporated in dominant eye movement models of eye movement control in reading (e.g. E-Z Reader, Reichle et al., 2009; and SWIFT, Engbert et al., 2005).

Alternatively, word identification may not utilize visual information holistically, especially in natural reading. Unlike in isolated word identification where information about a word comes only from visual input obtained by directly fixating it, in natural reading information about a word comes from more sources. These include contextual information from the preceding text and visual information obtained from fixations close to but not on the current word, which may still yield some visual preview of the word's initial letters. As a result, the most efficient positions from which to obtain useful new visual information about the word can vary from trial to trial, dependent on the information already obtained. Even in such an account, it is still possible that, on average, the most efficient positions are located near the center (as has been found in prior work). This account of word identification is implemented in rational models, which consider reading as a process of combining information from various sources to identify words and making eye movement decisions to maximize identification efficiency (Bicknell & Levy, 2010, 2012; Legge et al., 1997, 2002). For example, if a reader in this framework is working to identify a particular word, considering all the information that has already been gathered, there may be parts of the word that the reader has already identified relatively well and parts that are still relatively uncertain. It is intuitive in such a situation that identification efficiency will be maximized by moving the eyes next to the part of the word about which the reader is still relatively uncertain. This is because such fixations would obtain fine-grained visual information of a particular part of the word, which can be combined with visual information obtained from previous fixations (as well as linguistic contextual information), and identify the word in a constructive manner. Thus, contrary to the holistic account's view that any fixation landing on a non-central position slows identification efficiency relative to a central fixation, the view from rational models is that the position in the word to move the eyes next to maximize identification efficiency will vary from trial to trial and depend on information already obtained.

A phenomenon that can be used to tease apart these two accounts is that of refixations, cases in which a word is fixated more than once during first-pass reading. The goal of an intended refixation is assumed to be moving the eyes to a position that will maximize identification efficiency of the current word. Despite previous experiments showing that refixation rate varied on average as a quadratic function of the distance between word center and the fixation position (McConkie, Kerr, Reddix, Zola, & Jacobs, 1989) and was influenced by linguistic properties such as word frequency (Rayner et al., 1996), few studies shed light on where refixations go. The two accounts of word identification make different predictions for this question. The rational model predicts that refixations target the part of the word about which sufficient information has not yet been obtained. Which part of the word this is depends on the visual information already available.[1] In contrast, the standard model of word identification predicts that refixations should always target the center to maximize the holistic visual processing efficiency of the word, independent of information obtained about different parts of the word.

Naively, then, we could tease apart these two hypotheses by analyzing the relationship between the position of the initial fixation on a word (the 'landing position') and the refixation position. The rational account would predict that if the landing position is at the beginning of the word, a refixation should be at the end, and vice versa, whereas the standard model would predict that all refixations cluster around the center, regardless of the landing position. Empirically, this prediction of the rational models is borne out (Rayner et al., 1996), but the standard model explains this phenomenon in a different way. Specifically, there is a concept of *systematic error* (McConkie et al., 1988), which suggests that intended saccade sizes become biased toward the overall average saccade size. This means that refixation saccades intended to be short and target the center of the word in the standard model will tend to overshoot their target, landing on the opposite end of the word. Thus, both the standard model combined with systematic error and the rational model predict the effect of landing position on where refixations go.

Analyzing where refixations go as a function of the location of the *previous* fixation made before fixating a word (the 'launch site'), however, can tease apart these two accounts, when controlling for effects of landing site. If a reader's first eye movement to the word is launched from a position close to the word, then more visual information about the word's beginning should be available (relative to the launch site being further away), holding constant the landing site. Therefore, rational models predict that for closer launch sites, a refixation should be less likely to move the eyes back toward the beginning of the word (Fig. 1, right panel). In contrast, the standard model would not predict such an effect, but predict that an intentional refixation that follows a fixation on the left half of a word should always go forward, while one that follows a fixation on the right half should always go backward, always targeting the word center (Fig. 1, left panel).

[1] In general, the most efficient place to move the eyes next in a rational model depends not just on visual information already obtained but also contextual information. For the present paper, we ignore contextual information for simplicity.

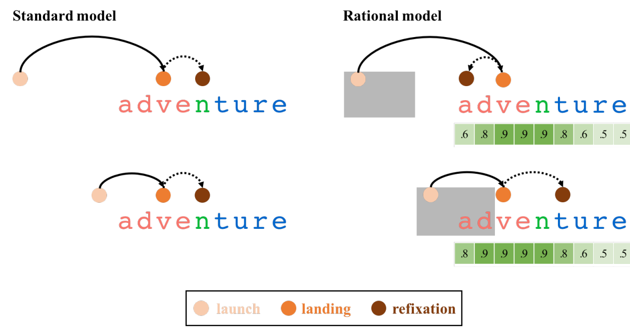

Figure 1: The standard model and the rational model make different predictions for where refixations go. For the standard model, refixations always target the center of the word, regardless of launch site. For the rational model, refixations target positions where character identity has low confidence (here represented by hypothetical $m(j)$ values). Therefore, closer launch sites, which provide more visual information about the word's initial letters (schematically represented here by grey rectangle) predict refixations are more likely to move forward. The refixation decisions here are based on eye movement policy parameters of $\alpha = .9$ and $\beta = .7$. (See Eye movement policy section for more details.)

In this paper, we empirically evaluate these two competing predictions by performing a statistical analysis of where refixations go in a large eye movement corpus, and we compare these results to simulations from computational models of both accounts. In the next section, we report the results of our statistical analysis of human refixation data, showing that it is as predicted by the rational account. We then confirm that an eye movement model that implements the standard model cannot accommodate this finding by performing simulations with E-Z Reader (Reichle et al., 2009). After that, we describe our rational model of refixations. Finally, we confirm that simulations using it show the same qualitative pattern as the human data, and then conclude.

Experiment 1: Human data in Dundee corpus

This analysis aims to tease apart the predictions of the rational model and the standard model on where refixations go. Specifically, we use the English part of the Dundee corpus (Kennedy, 2003) of eye movements during natural reading, and analyze the direction of refixation as a function of launch site, statistically controlling for landing site.

Methods

Data The English part of the Dundee corpus contained eye movement records from 10 native English-speaking participants as they read through newspaper editorials (see Kennedy & Pynte, 2005 for further details.) We first did a set of screening procedures, according to criteria that are generally applied to eye movement data, to remove fixations involving blinks, non-first-pass fixations, and the first/last two fixations of the line. After this procedure, the corpus contained 23,854 fixations that were followed by a refixation during first-pass reading (18.9% of first-pass fixations). These data then underwent screening procedures excluding: (a) extremely far launch sites (1%), leaving the launch sites of fixations in the range $[-16, -1]$ (in terms of number of characters from word beginning); (b) fixations that landed on the space right before the word (25.5%) or on the last character of the word (4.7%) to ensure the variability of refixation directions; and (c) fixations on words of which the previous word was skipped to eliminate possible overshootings of the previous word (20.9%), since these can be followed by corrective saccades. In the end, the data consisted of 7,667 fixations.

Statistical analysis A logistic generalized linear mixed-effects model (GLMM) was used to analyze the direction of refixations (forward vs. backward). Fixed effects included launch site and combinations of word length and landing site, which controlled for arbitrary effects of word length and landing site on refixation direction. Random effects included a random intercept and slope of launch site by subjects. Significance testing was via likelihood ratio test. All statistical analyses were implemented in the R environment, using the *glmer* function from the *lme4* package (Bates, Mächler, Bolker, & Walker, 2015) for GLMM implementation. In order to ensure model convergence, word length–landing site pairs for which all refixations (or all but 1) moved in the same direction were excluded, leaving 6714 fixations (87.6%).

Results and discussion

Fig. 2 shows the effect of launch site on the probability that refixations move forward for each word length–landing site pair. The GLMM showed that nearer launch sites predicted significantly more forward refixations, $\hat{\beta} = 0.15$, $SE = 0.03$, $\chi^2_1 = 13.98$, $p < 0.001$, 95% confidence interval (CI) $= [0.10, 0.20]$. As reported in the following section, the standard model can accommodate this effect only for landing sites on the right half of the word. To see whether this was also true of the human data, separate analyses were carried out for fixations with landing sites on the left and the right half of the word. For the left half (4790 fixations), launch site predicted more forward refixations, $\hat{\beta} = 0.16$, $SE = 0.03$, $\chi^2_1 = 10.91$, $p < 0.001$, $95\%CI = [0.09, 0.22]$, and the same was true for the right half (1362 fixations), $\hat{\beta} = 0.14$, $SE = 0.04$, $\chi^2_1 = 7.40$, $p < 0.01$, $95\%CI = [0.05, 0.22]$. These observations that closer launch sites predicted more forward-moving refixations confirm the rational model's predictions. The separate analyses of fixations on the left and right halves of the word indicated that this effect generalized across both.

Experiment 2: E-Z Reader

This section aims to show that the standard model does not predict the effect of launch site on direction of refixations. To this end, we carry out the same analyses as the previous section on simulation data from E-Z Reader, a computational model of eye movements in reading that incorporates the standard holistic model of word identification, and always targets refixations to the center of words. In principle, then, all intentional refixations following a fixation on the left half of the word should move forward and those following a fixation on the right half should move backward. Simulations with an implemented version of this model help to ensure that unintentional refixations – saccades intended for another word that happen to become a refixation due to motor error – do not in general change these predictions.

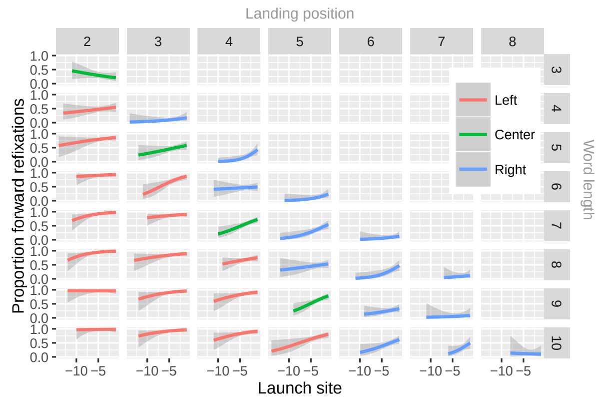

Figure 2: Effect of launch site on proportion of forward-moving refixations on data from Dundee corpus. Each panel contains data from a combination of word length and landing position, and shows a GAM smoother.

Methods

Data We used E-Z Reader 10 (Reichle et al., 2009) to generate eye movement data for 100,000 virtual readers reading sentences from the Schilling corpus (Schilling, Rayner, & Chumbley, 1998) of single English sentences typical of reading experiments. Each virtual reader was a simulation completed using a Monte Carlo run of the model.

The data cleansing procedure was the same as that in Expt. 1. Out of the 20,189,603 first-pass fixations, 3,417,999 (16.9%) of them were followed by a refixation. Excluding extreme launch sites, fixations landing on initial or final letters of a word, and skipping of the previous word left 1,029,801 fixations. Launch site ranged between $[-15, -1]$.

Statistical analysis A generalized linear model (GLM) with the same fixed effects as that in Expt. 1 was adopted to analyze the effect of launch sites on refixation direction. Random effects were removed from the GLMM used for Expt. 1 since the virtual readers were simply different Monte Carlo runs with no systematic differences. Excluding word length–landing position pairs where all refixations (or all but 1) moved in the same direction left 899,838 fixations (87.4%).

Results and discussion

Fig. 3 shows the effect of launch site on the probability for refixations moving forward. The GLM showed that nearer launch site predicted significantly more forward refixations, $\hat{\beta} = 0.08$, $SE = 0.004$, $\chi^2_1 = 386.66$, $p < 0.001$, $95\%CI = [0.07, 0.09]$. However, this effect was driven by fixations landing on the right half of the word, $\hat{\beta} = 0.10$, $SE = 0.004$, $\chi^2_1 = 542.99$, $p < 0.001$, $95\%CI = [0.09, 0.11]$, while fixations landing on the left half had 99% refixations moving forward and yielded an opposite effect, $\hat{\beta} = -0.33$, $SE = 0.03$, $\chi^2_1 = 147.37$, $p < 0.001$, $95\%CI = [-0.39, -0.27]$.

Therefore, E-Z Reader does not in general predict that closer launch sites should lead to refixations being more likely to go forward, contrary to our observations on the human data, although it can accommodate such a prediction for fixations on the right half of the word. Although this effect on the right half of the word may seem surprising, we note that the predictions we described above for this account only hold for *intentional* refixations. We believe that this effect on refixations on the right half of the word arises from unintentional refixations. Specifically, for a fixation position on the right half of a word, the E-Z Reader model will generally execute one of two behaviors: initiating a saccade to refixate the word or initiating a saccade to move on to the next word. In this case, an intended refixation will target a leftward position (since the center of the word is to the left of fixation) and an intended saccade to the next word will target a rightward position. Which of these two behaviors occurs depends on how quickly the identification (or more technically, L_1) is completed for the current word. Closer launch sites mean that identification of the word will be completed more quickly, which in turn will lead to a greater chance of making a forward saccade intended for the next word. Assuming some of these forward saccades become unintentional forward refixations, this creates exactly the predicted relationship between launch site and refixation direction. For the present purposes, however, the main conclusion here is that the standard model cannot reproduce a general effect of launch site on refixation direction.

Rational models of reading

In this section, we describe an implemented rational model of refixations, which we will use in the next section to confirm that the intuitively-derived predictions of the rational account for the relationship between launch site and refixations are actually produced by an implemented rational model. Rational models of reading use Bayesian inference to combine visual information with language knowledge (e.g., contextual information). Based on the posterior distribution, eye movements are selected to maximize identification efficiency. The rational model of refixations we describe in this paper also follows this idea, and can be viewed as an application of the more general-purpose rational models of eye movements in reading to the specific situation of refixations. This section introduces the framework of our model.

Word identification as Bayesian inference

Word identification consists of Bayesian inference, in which a prior distribution over possible identities of the text given by its language model is combined with a likelihood term given by 'noisy' visual input at the position of fixation to form a posterior distribution over the identity of the text given all information sources. Formalized with Bayes' theorem,

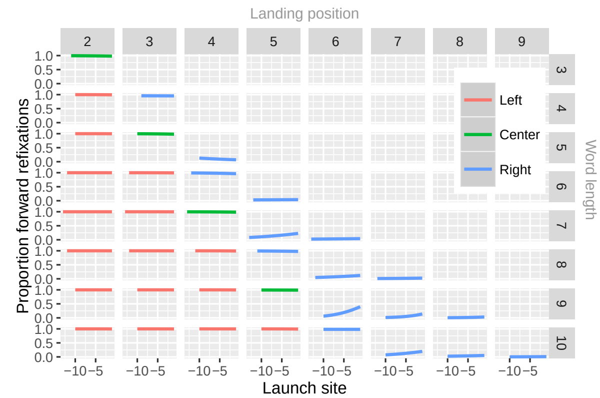

Figure 3: Effect of launch site on proportion of forward-moving refixations in data from E-Z Reader simulation. Each panel contains data from a combination of word length and landing position and shows a GAM smoother.

$$p(w|\mathcal{I}) \propto p(w)p(\mathcal{I}|w) \quad (1)$$

where the probability of the true identity of the word being w given uncertain visual input $\mathcal{I}$ is calculated by multiplying the language model prior $p(w)$ with the likelihood $p(\mathcal{I}|w)$ of obtaining this visual input from word w, and normalizing.

In general, the prior $p(w)$ represents reader expectations for words conditioned on the context, but for the present paper, we ignore context and use only a word frequency model for simplicity. The visual likelihood is computed similarly to in Bicknell and Levy (2010): each letter is represented as a 26-dimensional vector with a single element being 1 and the rest being 0s. Visual input about each letter is accumulated iteratively over time by sampling from a multivariate Gaussian distribution centered on that letter with a diagonal covariance matrix $\Sigma = \lambda^{-1}I$, where λ is the reader's visual acuity for that letter. Visual acuity depends on the location of the letter in relation to the point of fixation, which is a function of the letter's eccentricity ε. In our model, we assumed that acuity is a symmetric, exponential function of eccentricity:

$$\lambda(\varepsilon) = \int_{\varepsilon-.5}^{\varepsilon+.5} \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-\frac{x^2}{2\sigma^2})dx \quad (2)$$

with $\sigma = 3.075$, the average of two σ values for the asymmetric visual acuity function ($\sigma_L = 2.41$ for the left visual field, $\sigma_L = 3.74$ for the right visual field) used in Bicknell and Levy (2010). In order to scale the quality of visual information, we multiply each acuity λ by the overall visual input quality Λ, which is set to 12 in our simulation (see Expt. 3 below).

Eye movement policy

Based on the posterior distribution on possible identities of the word, eye movement decisions are selected to maximize reading efficiency. For example, the first rational model of reading, Mr. Chips, used this optimizing principle: the model reads input text sequentially, without error, in the minimum number of saccades (Legge et al., 1997, 2002). Specifically, saccades were made to minimize the expected entropy of the current word after the next fixation.

In a more recent rational model of eye movements in reading (Bicknell & Levy, 2010, 2012), eye movement decisions depend on the uncertainty of the posterior distribution about each letter position. Specifically, given a fixation landing on an unknown character c in position j, the marginal probability m of the most likely character under the posterior is

$$m(j) = \max_{c} p(w_j = c) \tag{3}$$

where w_j indicates the character in position j. A high value of $m(j)$ indicates relative confidence about the character's identity, and a low value relative uncertainty. The model then decided between four possible actions based on $m(j)$: continuing to fixate the current landing position, moving backward, moving forward, and ending the reading process.

We use a similar eye movement policy in our refixation model. If the value of the aforementioned statistic $m(j)$ is less than a parameter α, the model chooses to continue fixating the current position. Otherwise, if the value of $m(j)$ is less than the parameter β for some leftward position, the model initiates a saccade to the closest such position. If no such positions exist to the left, then the model initiates a saccade to the closest position to the right for which $m(j) < \alpha$. Once a refixation is executed, the simulation ends. If all $m(j)$ values to the right (left) are above α (β), we decide this word is identified with a satisfactory uncertainty level, and the identification of this word ends. In such a situation, we expect that the eyes move to the next word, which is beyond the current paper's scope of studying refixations.

The actual landing position is the intended fixation position with random motor error: the actual landing position ℓ_i is sampled from a Gaussian centered on the intended target t_i with standard deviation given by a linear function of the intended saccade distance

$$\ell_i \sim \mathcal{N}(t_i, (\sigma_0 + \sigma_1 |t_i - \ell_{i-1}|)^2) \tag{4}$$

for some linear coefficient σ_0 and σ_1.[2] In Expt. 3 in this paper, we follow the SWIFT model in using $\sigma_0 = 0.87$, $\sigma_1 = 0.084$. A refixation occurs if the actual landing site of the next fixation falls on the same word.

Experiment 3: Rational model

In this section, we analyze simulated data from our rational model of refixations to verify that it does indeed make the prediction that we derived from it intuitively: that refixations would be more likely to move forward for closer launch sites. As described in the previous section, the rational model of refixations we use combines information from previous fixations (including the launch site) to form a posterior distribution on the identity of a word through Bayesian inference. It then makes refixation eye movements to parts of the word about which it is uncertain.

Methods

Model parameters For the language model component of the word identification model (the prior), we used word frequency information (a unigram model) from the Corpus of Contemporary American English (COCA) (Davies, 2016). For this simulation, we did not optimize the behavior policy parameters to maximize reading efficiency as in Bicknell and Levy (2010), but set them manually to what we surmised might be reasonable values of $\alpha = 0.9$ and $\beta = 0.7$. Future work will optimize them, but we do not expect the qualitative predictions relevant to this analysis to change.

Data Eye movement data were generated to identify a word. All words were in the most frequent 5,000 words in COCA, and word lengths ranged between $[3, 10]$. Launch site had a range of $[-10, -1]$. For each word length, each possible landing position, and each launch site, 200 trials were run to model the word identification process as when a fixation landed on that landing position, preceded by a fixation on that launch site. In each trial, a word was randomly selected uniformly from words with the same length.

Procedure Each trial began with a fixation with a duration of 200 time steps on the launch site, in order to represent the visual information obtained prior to fixating the word. Then, the fixation at the landing site began. On each timestep of that fixation, visual information was obtained and integrated with prior information to update the posterior, and then a behavior decision was made: whether to continue fixating, make a refixation, or stop reading (see model description).

Statistical analysis A GLM with the same fixed effects as that in Expt. 2 was adopted to analyze the effect of launch site on refixation direction. Excluding word length–landing position pairs where all refixations (or all but 1) moved in the same direction left 25,636 fixations.

Results and discussion

Fig. 4 shows the effect of launch site on the probability for refixations moving forward. As expected, the GLM showed that nearer launch site predicted significantly more forward refixations, $\hat{\beta} = 0.07$, $SE = 0.005$, $\chi^2_1 = 187.62$, $p < 0.001$, $95\%CI = [0.06, 0.08]$. The same pattern held for both data with landing positions on the left half of the word, $\hat{\beta} = 0.04$, $SE = 0.008$, $\chi^2_1 = 28.85$, $p < 0.001$, $95\%CI = [0.02, 0.06]$, and the right half, $\hat{\beta} = 0.12$, $SE = 0.009$, $\chi^2_1 = 179.38$, $p < 0.001$, $95\%CI = [0.10, 0.14]$. These results confirm that an implemented rational model does indeed make this predic-

[2] Note that motor error in a rational model has only random error (variance), but not systematic error (bias).

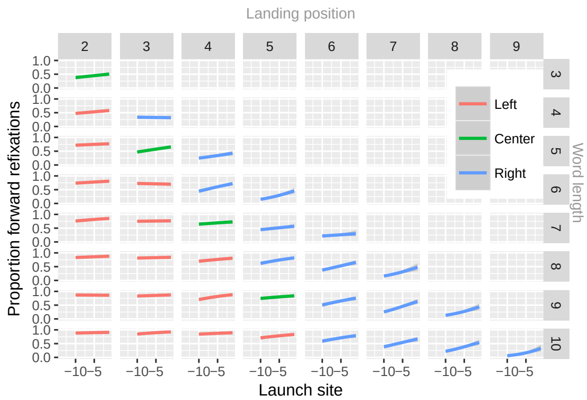

Figure 4: Effect of launch site on proportion of forward-moving refixations in data from rational model simulation. Each panel contains data from a combination of word length and landing position and shows a GAM smoother.

tion, which we observed in Expt. 1 to hold of human data.

General discussion

In this paper, we investigated how visual information is used for word identification during natural reading. We compared two accounts: (1) the standard holistic model, in which visual information about the word as a whole is used in word identification, and processing is always most efficient from the center; and (2) a rational model, in which readers combine information from many sources to identify a word constructively, and the fixation location that maximizes identification efficiency depends on what prior information has been obtained. We suggested that these two models make divergent predictions for the possible effects of launch site on where refixations go. Specifically, only the rational model should predict that refixations are more likely to go rightward for closer launch sites. An analysis of a large human eye movement corpus confirmed that this prediction of the rational account holds in human data. Model simulations confirmed that a rational model does indeed predict it, and that at least one of the implementations of the standard model (E-Z Reader) could not accommodate this finding.

These findings seem strongly inconsistent with models in which all intentional refixations target the center of a word, which in turn suggests that the standard holistic model of word identification in reading may be incorrect. However, it is possible to imagine that other refixation targeting schemes could be used even if the holistic model of word identification in reading is correct. For example, even under the standard model, it might be a useful strategy to target a refixation further forward in a word when that word is closer to being identified. Even if there is an efficiency penalty for being away from the center while that word is finished being identified, that penalty might be outweighed by the benefits of being closer to the next word when the reader's attention (soon) turns to it.

While it's possible that such eye movement models could be constructed while maintaining the standard model of word identification, our findings are completely consistent with the predictions of rational models of reading, and suggest that these models should be more fully explored. Here, we focused specifically on how visual information already obtained about a word influences where refixations should go, but rational models predict that the interaction of visual and linguistic information is what should ultimately matter. Future work should test these more complex predictions.

References

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48. doi: 10.18637/jss.v067.i01

Bicknell, K., & Levy, R. (2010). A rational model of eye movement control in reading. In *Proceedings of the 48th annual meeting of the Association for Computational Linguistics* (pp. 1168–1178).

Bicknell, K., & Levy, R. (2012). Word predictability and frequency effects in a rational model of reading. In *Proceedings of the 34th annual conference of the Cognitive Science Society.*

Davies, M. (2016). *Corpus of Contemporary American English (COCA).* Harvard Dataverse. Retrieved from `http://dx.doi.org/10.7910/DVN/AMUDUW`

Engbert, R., Nuthmann, A., Richter, E. M., & Kliegl, R. (2005). SWIFT: a dynamical model of saccade generation during reading. *Psychological Review*, *112*(4), 777–813.

Kennedy, A. (2003). The Dundee Corpus [cd-rom]. *Psychology Department, University of Dundee*.

Kennedy, A., & Pynte, J. (2005). Parafoveal-on-foveal effects in normal reading. *Vision Research*, *45*(2), 153–168.

Legge, G. E., Hooven, T. A., Klitz, T. S., Mansfield, J. S., & Tjan, B. S. (2002). Mr. chips 2002: New insights from an ideal-observer model of reading. *Vision Research*, *42*(18), 2219–2234.

Legge, G. E., Klitz, T. S., & Tjan, B. S. (1997). Mr. Chips: an ideal-observer model of reading. *Psychological Review*, *104*(3), 524–553.

McConkie, G. W., Kerr, P. W., Reddix, M. D., & Zola, D. (1988). Eye movement control during reading: I. the location of initial eye fixations on words. *Visual Research*, *28*, 1107–1118.

McConkie, G. W., Kerr, P. W., Reddix, M. D., Zola, D., & Jacobs, A. M. (1989). Eye movement control during reading: Ii. frequency of refixating a word. *Attention, Perception, & Psychophysics*, *46*(3), 245–253.

O'Regan, J. K. (1990). Eye movements and reading. *Reviews of Oculomotor Research*, *4*, 395–453.

O'Regan, J. K. (1992). Optimal viewing position in words and the strategy-tactics theory of eye movements in reading. In *Eye movements and visual cognition* (pp. 333–354). Springer.

Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, *124*(3), 372–422.

Rayner, K., Sereno, S. C., & Raney, G. E. (1996). Eye movement control in reading: a comparison of two types of models. *Journal of Experimental Psychology: Human Perception and Performance*, *22*(5), 1188–1200.

Reichle, E. D., Warren, T., & McConnell, K. (2009). Using E-Z Reader to model the effects of higher level language processing on eye movements during reading. *Psychonomic Bulletin & Review*, *16*(1), 1–21.

Schilling, H. E., Rayner, K., & Chumbley, J. I. (1998). Comparing naming, lexical decision, and eye fixation times: Word frequency effects and individual differences. *Memory & Cognition*, *26*(6), 1270–1281.