Prediction as a Basis for Skilled Reading:
Insights from Modern Language Models

Benedetta Cevoli,¹ Chris Watkins² & Kathleen Rastle¹

1. Department of Psychology, Royal Holloway, University of London
2. Department of Computer Science, Royal Holloway, University of London

RUNNING HEAD: PREDICTION IN NATURAL READING

Abstract

Reading is not an inborn human capability, and yet, English-speaking adults read with impressive speed. This study considered how predictions of upcoming words impact on this skilled behaviour. We used a powerful language model (GPT-2) to derive predictions of upcoming words in text passages. These predictions were highly accurate, and showed a tight relationship to fine-grained aspects of eye-movement behaviour when adults read those same passages, including whether to skip the next word and how long to spend on it. Strong predictions that did not materialise resulted in a prediction error cost on fixation durations. Our findings suggest that predictions for upcoming words can be made based on the analysis of text statistics, and that these predictions guide how our eyes interrogate text at very short timescales. These findings open new perspectives on reading and language comprehension more broadly, and illustrate the capability of modern language models to inform understanding of human language processing.

Reading is not an inborn human capability, and yet, literate English-speaking adults read at nearly 250 words per minute (Brysbaert, 2019). There has been substantial interest in understanding the extent to which *prediction* of upcoming words underpins this skilled language behaviour. Eye-tracking (e.g. Rayner et al., 2011) and electrophysiological studies (e.g. Kutas & Hillyard, 1984) have established that words that are highly predictable given the preceding context are easier to recognize than those that are less predictable (see Kuperberg & Jaeger, 2016 for a review). To illustrate, research suggests that the word *vegetarian* is easier to recognize in (a) than in (b) below (adapted from Rayner et al. 2011).

(a) *Megan decided to stop eating meat. She declared herself a vegetarian today.*

(b) *Megan went to lunch with colleagues. She declared herself a vegetarian today.*

However, despite decades of research investigating predictability effects in reading, there is considerable disagreement about whether readers routinely predict upcoming words. Our study offers a new perspective on this debate by using a modern language model to delineate the impact of predictions on eye-movement behaviour across the time-course of natural reading.

Several lines of argument suggest that skilled readers do not routinely predict or anticipate upcoming words (see Staub, 2015 for discussion). One set of arguments suggests that the predictability effects observed in the literature reflect *integration* processes rather than prediction. In the example above, the word *vegetarian* would be recognized more quickly in (a) than in (b) not because the reader has anticipated that word in advance of encountering it but because it is easier to integrate in the unfolding context (see Pickering & Gambi, 2018; Van Petten & Luka, 2012 for discussion). However, while prediction is difficult to disentangle from integration, the fact that predictability influences very early eye-movement measures including the probability of skipping a word (e.g. Rayner et al., 2011) seems unlikely to be consistent with a full integration account (Staub, 2015). Yet, even if there is some role for prediction, others have argued that natural language is too unconstrained to make prediction a viable processing strategy in most contexts (Jackendoff, 2007). Finally, while there is relative consensus that reading may involve pre-activation of general semantic or morpho-syntactic features (e.g. Federmeier & Kutas, 1999; Van Berkum et al., 2005; Wicha et al., 2004), evidence that readers predict specific, upcoming words is far more limited.

If readers engage in prediction of specific upcoming words, then there should be some evidence of a processing cost when those predictions do not materialise (Staub, 2015). The strongest demonstration of such a processing cost was reported by DeLong et al. (2005) in an ERP study. They observed a higher N400 amplitude when readers encountered *articles* that were inconsistent with a strongly predicted target as in the sentence “*The day was breezy so the boy went outside to fly an ...*”. This finding is important because it implicates pre-activation of a *specific phonological form* (see also Wang et al., 2018). However, this result was not replicated in a pre-registered, highly-powered study (Nieuwland et al., 2018). Evidence for a processing cost because of inaccurate predictions has also been sought using eye-tracking methodology. Frisson et al. (2017) monitored fixation duration on unpredictable target words embedded in highly-constraining or neutral sentence contexts, as in the below examples (target in bold).

- a) *The young nervous paratrooper jumped out of the **chair** when he heard the shots.*
- b) *The tired movie maker was sleeping in the **chair** when he was woken up by a scream.*

They reasoned that if readers generate specific predictions in highly-constraining contexts such as sentence (a) above (in which the word *plane* should be predicted), then there should be some evidence of a prediction error cost on fixations to unpredictable targets, relative to neutral contexts (b). Despite observing the usual predictability effect, no evidence for a prediction error cost was observed on any measure.

Drawing inferences regarding the use of prediction in reading has been hampered by a variety of methodological challenges. In addition to the challenge of disentangling prediction from integration, constraints on stimuli and stimulus presentation in different paradigms make it difficult to determine whether prediction is a phenomenon associated with natural text reading, as opposed to more artificial laboratory paradigms. In electrophysiology studies, words are typically presented one at a time with a fixed duration (Clifton & Staub, 2012; Rayner et al., 2004; Schotter et al., 2014). This unnatural presentation disrupts typical characteristics of natural reading including fixation durations, skips, and regressions. It also elicits predictive processes that may differ from those that occur in normal reading (Brothers et al., 2017; Dambacher et al., 2012; Huettig & Guerra, 2015; Huettig & Mani, 2016). Eye-tracking studies avoid these pitfalls but investigate predictability effects on single targets contained in artificial sentences deliberately designed to be predictable or unpredictable (see Staub, 2015 for review).

It is also difficult to quantify the predictability of words in different contexts. Most studies in this domain measure predictability as cloze probability (e.g. Rayner et al., 2011; Staub et al., 2015; Taylor, 1953). In cloze tasks, participants are presented with a sentence frame and asked to guess the next word. The distribution of participants' top predictions is then taken to reflect the probability distribution of the target in context. The underlying assumption is that participants' responses reflect their individual probability distributions for upcoming words; however, this assumption has been challenged (Smith & Levy, 2011, 2013). Most critically, cloze measures underestimate the probability of unpredictable words; such words are rarely indicated by participants and are therefore associated with a cloze probability of zero. The bias against responding with an unpredictable word may hide important variation in the lower range. These limitations are partly mitigated by use of corpus measures of predictability, which use n -gram models to compute the conditional probability of a word given a preceding local context (Hale, 2001; Ong & Kliegl, 2008). This approach allows more accurate estimation of the lower range of predictability; however, these estimates are usually based on local contexts of only a few words and thus ignore background information and long-range dependencies.

The present study advances the debate by using a newly-developed language model to characterise the predictability of words in a large eye-tracking corpus of natural passage reading (Luke & Christianson, 2018), and then relating those metrics to reading behaviour. The language model used is GPT-2, a deep neural network based on a transformer architecture that uses attention mechanisms to focus selectively on words in the text that are most relevant in a context (Radford et al., 2019). The model is trained on vast corpora to predict upcoming words accurately given the preceding context. The attention mechanisms allow for context-specific training of word representations (i.e. *contextualized embeddings*). This means that unlike predecessor language models such as LSA (Landauer & Dumais, 1997) or word2vec (Mikolov et al., 2013), the learned high-dimensional representation of a single word will vary as a function of its context; for example, representation of the word *bark* will differ if it is in the context of a tree or a dog (Cevoli et al., 2022).

In essence, GPT-2 provides a computational instantiation of a next-word prediction strategy that does not rely on deep understanding. Use of GPT-2 therefore offers opportunity to assess the extent to which prediction is *feasible* in natural reading, and it also provides a means of quantifying precisely the probability distribution of any word in any

context, including those words with very low probability. By presenting each word in the eye-tracking corpus sequentially to the model, we are able to extract novel measures that quantify (a) uncertainty *before* a target is encountered; and (b) surprisal *when* a target is encountered. This distinction is critically important because it allows us definitively to separate prediction from integration. Our analyses sought to determine whether there is unambiguous evidence that specific predictions made before a target is encountered drive fixations to that target and whether there is evidence that inaccurate predictions yield a prediction error cost.

Method

The study derived fine-grained predictability metrics from GPT-2 and applied these to eye-movement data derived from the Provo corpus (Luke & Christianson, 2018). The materials, data, and analysis scripts for this project can be found on [link provided to reviewers].

Provo Corpus

This study exploited the Provo Corpus (Luke & Christianson, 2018), a large and publicly available corpus of eye movements in reading. The corpus consists of eye-tracking data from 84 adult American English native speakers reading 55 short passages derived from a variety of sources, including online news articles, science magazines, and fiction. The eye-movement measures of interest to us include whether the target was fixated in the first pass (i.e. before the eyes move rightward from the target), first fixation duration (length of the first fixation on a word), gaze duration (sum of all fixation durations in the first pass), total reading time (sum of all fixation durations), and regressions back to the target. The corpus also includes predictability norms derived from an online word-by-word cloze task completed by 470 native American English speakers. These norms provide information about the percentage of participants that accurately predicted the target as the next word.

Language Model

Predictability metrics were derived from a pre-trained version of GPT-2 (Radford et al., 2019; model version: *base*; 12-layer, 768-hidden, 12-heads, 117M parameters; PyTorch Transformers library; Wolf et al., 2020) with no fine-tuning of hyper-parameters. This version of the model was trained for optimal probabilistic prediction according to the minimal cross-

entropy criterion and was accessed through the OpenAI GPT2 website (https://huggingface.co/transformers/v3.0.2/model_doc/gpt2.html). We used the python scientific stack to compute predictability metrics from the model.

The language model's predictions about an upcoming word in a passage are expressed in terms of the likelihood of each word in the model's vocabulary occurring as the next word. These predictions consist of a vector of logit prediction scores for each vocabulary token, hereafter referred to as the 'prediction-probability vector'. We derived likelihood values by normalizing these scores with a softmax function. An incremental moving window was used to calculate predictability information for each word in the corpus passages (for a total of 2,685 tokens and 1,191 types). The language model's predictions were analysed for each word given the preceding context. Starting from the very beginning of a text passage, the first word is provided to the model and predictability information about the subsequent word (the *target*) is extracted. Then, the first and the second word are provided to examine predictions of the third word (the new target), and so on (see Figure 1). In this way, predictability information based on the language model's predictions were computed for each of the 2,685 words in the corpus.

Figure 1. Illustration of language model next-word predictions (on the right in grey) given the preceding context (on the left in black or coloured text). The transparency of the grey predictions indicates their strength. The colour of the final word indicates the language model surprisal when the target word is encountered.

We derived two metrics from the prediction-probability vectors. The first metric quantified the uncertainty of the context preceding a target in terms of the *entropy* of the

prediction-probability vector. This metric reflects the extent to which a context is neutral or constraining in nature. Low entropy values are typical of constraining contexts as only a few words have a very high probability of occurring next (e.g. *The elephant drank using its long ____*). In contrast, neutral contexts are represented by high entropy values as many words are equally likely to follow (e.g. *The man looked for the ____*). The second metric quantified the predictability of a target word given the preceding context as *surprisal*, the negative natural log of the language model's probability of that word. It is important to stress that in the computation of a context's entropy, no information about the succeeding target word is provided to the language model. Entropy reflects the state of the language model's predictions *before* the target word is encountered, while surprisal reflects the state *when the target word is encountered*. Thus, while an effect of surprisal on eye-movement behaviour might be ascribed to prediction or integration processes, an effect of entropy *must* be ascribed to prediction (Pickering & Gambi, 2018). Though the concepts of entropy and surprisal are related, they can be distinguished in natural language (see Table 1 for examples), and we monitored our statistical models closely for excessive collinearity (see Results).

Entropy	Surprisal	Context	Predictions	Target
High	High	All that the	<i>world government media</i>	brain
Low	High	Binge drinking may not necessarily kill or even damage	<i>your the a</i>	brain
High	Low	... But it can block key receptors in the	<i>brain hippocampus body</i>	brain
Low	Low	... But it can block key receptors in the brain and trigger production of a steroid that interferes with	<i>the brain normal</i>	brain

Table 1. Examples from passages in the Provo corpus with varying entropy and surprisal metrics as derived from the language model's predictions.

If readers predict specific upcoming words, then we would expect to observe a *prediction error cost* when these predictions are inaccurate (Staub, 2015). Thus, the impact of

target surprisal on fixation behaviour should vary as a function of entropy; surprising targets should be particularly detrimental when entropy is low. To examine whether the semantic relationship between the specific prediction and the target modulates any prediction error cost, we computed a final metric for each word in the corpus referred to as *target-prediction similarity*. Target-prediction similarity was operationalised as the cosine similarity between embeddings of the top prediction in the language model and the target. It is important to remember that these are *contextual embeddings* (i.e. representations of a word in a particular context).

Results

Eye-tracking data were analysed using (generalized) linear mixed-effects models (lme4 package by Bates et al., 2020). In the analysis, reading time measures were log-transformed, as indicated by a Box-Cox test (from the R package car; Fox & Weisberg, 2019). For visualisation purposes, these were back-transformed by exponentiation. In addition to the fixed factors of interest, models included word frequency estimates from SUBTLEX (Van Heuven et al., 2014), word length in letters, and the number of the word in a sentence as fixed factors as well as random intercepts for participants, items and text passages. The significance of the fixed effects was determined using the Anova function in the car package (type III model comparisons, Fox & Weisberg, 2019). The models reported below all had a variance inflation factor below 2, suggesting that they do not suffer from excessive collinearity (Johnston et al., 2018). Full results of all models are available on the OSF site for this project [link provided to reviewers].

Validity of the Language Model Predictions

Before analysing the influence of the language model predictions on eye movements during natural reading, we examined the validity of these metrics against cloze data for each word reported in the Provo corpus (Luke & Christianson, 2018). The distribution of target probability derived from the cloze task and the language model are compared in Figure 2a. This figure plots target probability using logit scores as these are more normally distributed and allow superior comparison of the two distributions in the low range. The probability values from the cloze task reflect an average across participants, while the values from the language model reflect the probability of that target occurring relative to every other word

in the language model's vocabulary. It is immediately apparent that the probability distributions of the language model and the cloze task are very similar in the upper range but completely diverge in the lower range of the distribution. Whereas the language model provides a continuous distribution of probabilities, the cloze task estimates 32% of the entire dataset to have a target probability of zero. This analysis suggests that the cloze task dramatically underestimates low probability values when compared to the language model.

Figure 2. Panel A shows the target probability distribution in logits derived from the language model and cloze task (where logit scores of -inf correspond to probabilities of zero and +inf to probabilities of 1). Panel B is a frequency distribution of next-word prediction accuracy for the 470 human participants who contributed cloze data (red) and the language model (blue). Panel C shows the accuracy of the language model in predicting the next word on top k predictions, for k from 1 to 10 (error bars indicate the standard error of the mean).

We then compared the performance of the language model in predicting the next word (given a preceding context) to human performance in doing the same task. Human performance was based on 470 participants each completing the word-by-word cloze task for 5 passages (an average of 227 words each; Luke & Christianson, 2018). On average, participants had an accuracy of 22% (standard deviation 6%; maximum accuracy 56%), whereas the top prediction of the language model was accurate 31% of the time (Figure 2b). This figure shows that the language model was more accurate in its predictions than most but not all human participants. Figure 2c shows that when we increase the prediction window from the top prediction to the top 10 predictions, the accuracy of the language model increases to 61% accuracy.

These assessments of GPT-2 performance advance understanding in several ways. They reveal that the cloze task is insensitive to probabilities at the lower end of the distribution (see also Smith & Levy, 2013), suggesting that cloze data are not appropriate for investigating prediction in natural reading contexts, in which low probability sequences are frequently encountered. Conversely, these data reveal that the language model not only produces a graded distribution of probability values but also is superior in its predictions to almost all individual participants. Finally, the fact that the language model was able to predict the next word 31% of the time in a natural text corpus, and that this figure rose to over 60% when the model's top 10 predictions were considered, suggests that natural language may be more predictable than has previously been asserted (Jackendoff, 2007).

Eye Movements Analysis

Having established the language model as a valid source of predictions, we then sought to determine how internal model predictions are related to fine-grained eye-movement behaviour. Thus, we examined the influence of the language model's state (a) *before* participants reach a target (entropy) and (b) *when* participants reach a target (surprisal). We also analysed whether there is evidence that eye-movement behaviour is influenced by the semantic relationship between the predicted and observed target (target-prediction similarity). We investigated the effects of these metrics as appropriate on skipping, reading times, and regression probability. Eye-tracking measures of reading times included first fixation duration (length of the first fixation on a word), gaze duration (sum of all fixation durations in the first pass), and total reading time (sum of all fixation durations).

Skipping behaviour. It is well known that not all words are fixated directly, and that the probability of skipping a word is modulated by factors including its length, frequency, and cloze probability (see Reichle & Drieghe, 2013 for discussion). Words are more likely to be skipped when they are short, high frequency, and predictable from the preceding context. One influential model of eye-movement behaviour suggests that the decision to skip an upcoming word may be based on analysis of that word in the parafovea (Reichle et al., 2003), hence the impact of these target properties on skips.

Our aim was to determine whether entropy influences the probability of skipping an upcoming word over and above any information that might be accessed about that word parafoveally. We therefore tested the influence of entropy on whether an upcoming target is skipped but included in the model factors pertaining to the target including length, frequency, and surprisal. We specifically tested for an interaction between entropy and surprisal to determine whether any effect of entropy is modulated by whether the prediction derived from preceding context is accurate. We used the following model formula: *SkippingProbability* ~ *Entropy*Surprisal* + *Freq* + *Length* + *WordNumberInSentence* + (*1|ParticipantID*) + (*1|WordID*) + (*1|TextID*). Descriptive statistics indicate that 113,690 words were skipped out of the 249,984 observations in the Provo Corpus (Luke & Christianson, 2018).

Results revealed that the entropy of the preceding context influenced the probability that a target would be skipped ($\chi^2 = 75.95, OR = 0.93, SE = 0.01, Z = -8.95, p < 0.001$). Readers are more likely to skip the upcoming word if the preceding context has lower prediction-probability entropy (more constraining contexts) than higher prediction-probability entropy (more neutral contexts). However, this analysis also revealed a significant interaction between entropy and surprisal on skipping probability ($\chi^2 = 4.25, OR = 0.99, SE = 0.01, Z = -2.06, p = 0.04$; see Figure 3). This result indicates that target surprisal has an impact on skipping probability but only when readers are more uncertain of the next word (higher entropy). In lower uncertainty contexts (lower entropy), the reader acts on the strength of their prediction without regard to whether that prediction is accurate.¹ Critically, none of these effects cannot be ascribed to frequency, length, or position in the sentence as these variables are included in the model.

¹ Further analysis revealed that 43% of the low entropy / high surprisal cases were skipped (where these categories were determined by equal frequency binning). These are cases in which there is a strong but inaccurate prediction. This result suggests that a substantial proportion of skips may be based on incorrect predictions.

Figure 3. Model estimates of the interaction between entropy and surprisal, expressed in information units (bits), on human skipping probability. Surprisal was split into high and low values for visualisation purposes only with equal frequency binning (to create equally distributed categories).

Fixation durations. For those targets that were not skipped (136,294 of the original 249,984 observations), we investigated how reading times were influenced by uncertainty of the preceding context (entropy) and target surprisal. We used the following model formula: $DV \sim Surprisal * Entropy + Freq + Length + WordNumberInSentence + (1|ParticipantID) + (1|WordID) + (1|TextID)$.

Results on first fixations (duration of the first fixation on a word) revealed an interaction between surprisal and entropy ($\beta < 0.01, SE < 0.01, t = -3.86, p < 0.001$). High surprisal targets lengthened first fixations most strongly when there was high certainty that the target should be something else (low entropy; see Figure 4a). This result is crucial because it reveals a *prediction error cost* in natural reading, a processing penalty when a strong prediction is inaccurate.

Figure 4. Model estimates of the effect of surprisal and entropy of the preceding context, expressed in information units (bits), on human reading times. Entropy was split into high and low values for visualisation purposes only with equal frequency binning. First fixations show a significant interaction between surprisal and entropy (Panel A) with surprisal impacting fixations more strongly when entropy is low. The effects of surprisal and entropy are additive on later measures of gaze duration (Panel B) and total reading time (Panel C).

Results on later measures revealed significant effects of entropy ($\beta = 0.01, SE < 0.01, t = 5.17, p < 0.001$) and surprisal ($\beta = 0.02, SE < 0.01, t = 9.00, p < 0.001$) on gaze duration (sum of fixations in the first pass). Likewise, there were significant effects of entropy ($\beta = 0.01, SE < 0.01, t = 6.16, p < 0.001$) and surprisal ($\beta = 0.05, SE < 0.01, t = 17.82, p < 0.001$) on total reading time (sum of all fixations; see Figure 4). In both cases, fixations were longer when there was greater uncertainty in the preceding context, and when targets were more surprising. There was no interaction between these factors on gaze duration ($\beta < 0.01, SE < 0.01, t = 1.72, p = 0.08$) or total reading time ($\beta < 0.01, SE < 0.01, t = 0.86, p = 0.39$).

Regressions. We examined the influence of entropy and surprisal on the probability of regressing back to the target using the following model formula: *RegressionProbability* ~ *Surprisal*Entropy* + *Freq* + *Length* + *WordNumberInSentence* + (*1|ParticipantID*) + (*1|WordID*) + (*1|TextID*). Results revealed that readers were more likely to regress back to words that occur in higher entropy contexts ($OR = 1.08, SE = 0.01, z = 5.83, p < 0.001$) and that are more surprising ($OR = 1.25, SE = 0.02, z = 15.09, p < 0.001$; see Figure 5). These data suggest that the ramping up of total reading times in low entropy and high surprisal contexts (Figure 4c) is likely due to regressions back to targets. There was no interaction between entropy and surprisal on regression probability ($OR = 0.99, SE = 0.01, z = -1.04, p = 0.29$).

Figure 5. Model estimates of the effect of entropy (panel A) and surprisal (panel B) on human regression probability. Entropy and surprisal are expressed in information units (bits).

These findings advance understanding by showing unambiguously that predictions *prior* to encountering a target drive aspects of fixation behaviour at very short timescales. They also provide new insight into the nature of those predictions. Specifically, they indicate that in cases in which strong predictions are made (low entropy), these predictions are sufficiently specific to generate a processing penalty on first fixation duration when they are incorrect. To our knowledge this is the first demonstration of a prediction error cost on eye-movement behaviour during natural reading. One question is why this prediction error cost does not also arise in the gaze duration and total reading time data. Our hypothesis is that the prediction error cost in low entropy / high surprisal situations is mitigated in later measures through *integration* processes if the target is semantically similar to the predicted word (see also Frisson et al., 2017). To illustrate, while the sentences below would both be expected to generate prediction error on first fixation durations, this prediction error would be mitigated in sentence (a) in later measures due to the semantic similarity between the prediction (denoted in square brackets) and the observed target. Sentences like (b) do not occur in our dataset because the Provo corpus (Luke & Christianson, 2018) uses natural language.

(a) The boy was hungry and picked up the shiny red [apple] cherry.

(b) The boy was hungry and picked up the shiny red [apple] button.

Target-prediction similarity. To test this hypothesis, we investigated whether the prediction error cost observed on first fixation durations (Figure 4a) is modulated by the semantic similarity between the prediction and the target (*target-prediction similarity*). To do so, we selected a subset of cases in which the language model had a strong prediction (low entropy) that was ultimately incorrect (high surprisal). To select these cases, we converted the continuous measures of entropy and surprisal to categorical variables with an equal frequency binning (this creates equally distributed categories; arules R package; Hahsler et al., 2005), then selected targets with low entropy and high surprisal values, resulting in a subset of 34,020 observations representing 84 participants x 405 tokens.

We then analysed whether target-prediction similarity modulates the impact of prediction error on eye movements. In other words, when readers have strong predictions given the preceding context, are they less surprised to see an unexpected target word if it is semantically related to the prediction than when it is unrelated? The effect of target-prediction similarity was analysed with the following formula: $DV \sim Similarity + Freq + Length + WordNumberInSentence + (1|ParticipantID) + (1|WordID) + (1|TextID)$. Results showed a main effect of target-prediction similarity on gaze duration ($\beta = -0.02, SE = 0.01, t = -2.94, p = 0.003$) and total reading time ($\beta = -0.03, SE = 0.01, t = -5.13, p < 0.001$). These results suggest that prediction error is reduced in later reading time measures when there is greater semantic similarity between predictions and targets (see Figure 4). There was no statistically significant effect of target-prediction similarity on first fixation duration ($\beta = -0.01, SE < 0.01, t = 1.76, p = 0.08$) or regressions ($OR = 0.96, SE = 0.04, z = -1.22, p = 0.22$).

Figure 6. Model estimates of the effect of target-prediction similarity on reading times for 405 tokens (34,020 observations) in which the language model made a strong prediction that was incorrect (low entropy / high surprisal). Target-prediction similarity had a significant effect on human gaze durations and total reading times for these items.

The fact that target-prediction similarity effects are observed on later measures (gaze duration and total reading time) supports our hypothesis that integration processes mitigate the downstream impact of prediction error. However, our analyses also suggest that these integration processes are not sufficiently rapid to mitigate the impact of prediction error on first fixation durations.

Discussion

The field of natural language processing and its applications have been revolutionised by the development of language models that are sensitive to context (Devlin et al., 2019; Hayashi et al., 2019). These models do not claim to solve natural language processing problems in the same way as people do, and they have rarely been used to improve understanding of human language processing. However, use of GPT-2 (Radford et al., 2019) in this instance has allowed us to make a series of important advances in our understanding of the mechanisms that underpin skilled reading.

These advances arise because GPT-2 (Radford et al., 2019) allows us to quantify the prediction distribution *prior* to encountering a target (including for the substantial proportion of words with a cloze probability of zero; see Figure 2a), and to separate this from the degree of surprise *when* a target is reached. Our analyses of GPT-2 predictions revealed a high degree of accuracy: the model's top prediction was accurate for 31% of words in the Provo

corpus (Luke & Christianson, 2018), and this figure rose to 61% when its top-10 predictions were considered (Figures 2b and 2c). This high degree of accuracy counters long-standing and influential claims that natural language is too unconstrained for routine prediction to be viable (Jackendoff, 2007). Moreover, the prediction performance achieved by the model is based on relatively superficial language input: the occurrence of word strings and their co-occurrence with other word strings. One might argue that humans' broader experience of language and their grounded understanding of meaning would permit even better prediction performance. In short, we believe that an argument against routine use of prediction that invokes the unconstrained nature of language is untenable.

There has been increasing interest in relating language model surprisal to reading times (e.g. Aurnhammer & Frank, 2019; Wilcox et al., 2020). However, the distinction that we make between entropy and surprisal in this work is important because it allows us to disentangle prediction from integration processes. Entropy is a measure of what the reader expects to see, and its influence on reading *must* be ascribed to prediction. In contrast, surprisal reflects a reader's reaction to the word that they actually encounter, and its influence on reading may be ascribed to prediction and/or integration (Pickering & Gambi, 2018). Investigating how both of these metrics map onto events in the eye-movement record has allowed us to make a series of observations that yield new insight into how prediction and integration processes work together across the time-course of reading. These observations ultimately lead to a new perspective on how prediction is used to facilitate rapid, skilled reading.

Our results showed that the degree of uncertainty prior to encountering a target (entropy) influences the earliest moments of word recognition, before a target is even fixated (Figure 3). Entropy influenced the probability of skipping the next word; however, this effect was qualified by an interaction with target surprisal. Higher surprisal targets were less likely to be skipped than lower surprisal targets, but only in higher entropy (higher uncertainty) contexts. These results challenge decades-old thinking about the computations involved in the planning of saccades in reading. Previous research has converged on a view in which analysis of information in the parafovea determines whether a subsequent target will be skipped (Reichle et al., 2003; see Staub, 2015 for discussion). Our findings suggest instead that the strength of predictions determines how much information readers sample from the parafovea in planning the next saccade. If readers' predictions are strong then they act on

the basis of those predictions; this strategy is likely to support reading speed although a substantial proportion of words in our analysis appear to have been skipped based on inaccurate predictions. Our findings suggest that when readers' predictions are not as strong, they sample more information from the parafovea, and hence are less likely to skip higher surprisal targets.

These observations imply that readers may be using predictions dynamically to guide how the eyes interrogate upcoming text. Our findings suggest that it is possible to make accurate next-word predictions in a substantial proportion of cases in natural language (Figure 2). Thus, it may be that the routine use of prediction allows readers to adopt a relatively superficial reading strategy, in which deep analysis is reserved for higher uncertainty contexts. If so, this would go some way to explaining the remarkable speed of skilled reading. Our results pertaining to first fixation durations provide further insight into the nature of those predictions. In low uncertainty situations in which readers' predictions are accurate (low entropy / low surprisal), readers enjoy a first fixation benefit (Figure 4a). However, in low uncertainty situations in which predictions are inaccurate (low entropy / high surprisal), a prediction error cost is observed (Figure 4a). The observation of this processing cost is important because it provides some of the first evidence that the predictions being made are for specific words (rather than more general semantic or syntactic categories; Staub, 2015).² Our novel analysis of target-prediction similarity further suggests that this prediction error may be resolved through integration processes operating later in the time-course of recognition. Specifically, we observe significant reduction in gaze duration and total reading times when the incorrect prediction is semantically similar to the target; this is likely to be the usual state of affairs in natural language. However, this integration process is not fast enough to mitigate the impact of prediction error on first fixations (Figure 6).

The prediction error cost observed arises in situations in which it is possible to make strong next-word predictions. However, there will also be instances in which predictions are weak (high entropy). We have already seen that in these cases, readers' decisions about whether to skip the next word appear to involve greater sampling of information about the

² We believe that the insensitivity of the cloze task undermined previous attempts to demonstrate prediction error cost on eye fixations (Frisson et al., 2017; see https://osf.io/ypeb2/?view_only=6080e254f06f4faa99c30716de838849, Appendix S.1 for discussion).

identity of that word in the parafovea (as compared to low entropy cases). The sampling of information in the parafovea may help us to understand why the effect of surprisal appears to be attenuated in first fixation durations in higher entropy contexts (Figure 4a). If readers are sampling enough information from the parafovea to generate an influence of surprisal on the decision to skip the target, then perhaps that sampling is also sufficient to mitigate the effect of surprisal if the eyes land on the target. Further research is necessary to develop this account and to determine why the effect of surprisal in high entropy cases returns in the gaze duration measure.

Our findings give pause to consider how predictive processes may contribute to human mastery of the reading skill. Unlike spoken language, reading is a cultural invention and a learned skill. Yet, literate adults consume linguistic information via reading at speeds that eclipse their engagement with spoken language (Brysbaert, 2019). Our findings suggest that the language used in text is relatively predictable based on the text statistics available to GPT-2, and that readers routinely use prediction to guide their processing of upcoming words. It may be that literate adults' reading speeds are achieved because deep analysis of individual words is required only infrequently; for example, when uncertainty is unusually high or when predictions are inaccurate. Superficial, prediction-guided reading may be sufficient for understanding in most cases. We use the word *superficial* because the online predictive process that we envisage does not involve deep analysis of the word being encountered. However, we recognize that a readers' ability to engage in this type of predictive processing most likely requires extensive language knowledge and understanding. If the ability to deploy accurate predictions rapidly is a hallmark of skilled reading, then it would be important to understand how the statistical knowledge that permits accurate prediction is acquired, and how this knowledge contributes to the development of reading fluency.

In summary, our work demonstrates that the internal states of a modern language model predict fine-grained aspects of eye-movement behaviour: the decision to land on a word, how long to spend on it, and whether to regress back to it. These models do not strive to mimic human language processing mechanisms. However, they do offer a compelling description of next-word prediction based on statistical information pertaining to the occurrence of words in text. By studying how predictability metrics derived from these

models relate to human eye-movement behaviours, we have been able to develop a new perspective on skilled reading that involves coordination of rapid prediction based on statistical information and slower integration based on deeper analysis. It will be important to operationalise and test this account through next-generation computational models of reading, and we believe that language models like GPT-2 may provide a foundation for some of that computational work. Likewise, we believe that understanding how internal model states map onto online reading difficulty opens new opportunities for advances in natural language applications such as readability and summarisation. Research at the interface of cognition, neuroscience, and natural language processing is beginning to emerge (e.g. Schrimpf et al., 2021) and offers significant opportunity for novel theoretical insight and application.

Acknowledgements

We thank Steven Luke for sharing individual cloze data from the Provo corpus. We are grateful to the Economic and Social Research Council for funding support (ES/P001874/1)

References

- Aurnhammer, C., & Frank, S. L. (2019). Evaluating information-theoretic measures of word prediction in naturalistic sentence reading. *Neuropsychologia*, 134, 107198. <https://doi.org/10.1016/J.NEUROPSYCHOLOGIA.2019.107198>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/doi:10.18637/jss.v067.i01>
- Brothers, T., Swaab, T. Y., & Traxler, M. J. (2017). Goals and strategies influence lexical prediction during sentence comprehension. *Journal of Memory and Language*, 93, 203–216. <https://doi.org/10.1016/j.jml.2016.10.002>
- Brysbaert, M. (2019). How many words do we read per minute? A review and meta-analysis of reading rate. *Journal of Memory and Language*, 109, 104047. <https://doi.org/10.1016/j.jml.2019.104047>
- Cevoli, B., Watkins, C., Gao, Y., & Rastle, K. (2022). Shades of meaning: Natural language models offer insights and challenges to psychological understanding of lexical ambiguity. *PsyArXiv*. <https://doi.org/10.31234/OSF.IO/Z8RMP>
- Clifton, C., & Staub, A. (2012). Syntactic influences on eye movements during reading. In *The Oxford Handbook of Eye Movements*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199539789.013.0049>
- Dambacher, M., Dimigen, O., Braun, M., Wille, K., Jacobs, A. M., & Kliegl, R. (2012). Stimulus onset asynchrony and the timeline of word recognition: Event-related potentials during sentence reading. *Neuropsychologia*, 50(8), 1852–1870. <https://doi.org/10.1016/J.NEUROPSYCHOLOGIA.2012.04.011>
- DeLong, K. A., Urbach, T. P., & Kutas, M. (2005). Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience*, 8(8), 1117–1121. <https://doi.org/10.1038/nn1504>
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *NAACL HLT 2019 - 2019*

- Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference, 1*, 4171–4186.
- Federmeier, K. D., & Kutas, M. (1999). A Rose by Any Other Name: Long-Term Memory Structure and Sentence Processing. *Journal of Memory and Language*, 41(4), 469–495. <https://doi.org/10.1006/jmla.1999.2660>
- Fox, J., & Weisberg, S. (2019). *Applied Regression 3E*. Sage Publications. <https://socialsciences.mcmaster.ca/jfox/Books/Companion/>
- Frisson, S., Harvey, D. R., & Staub, A. (2017). No prediction error cost in reading: Evidence from eye movements. *Journal of Memory and Language*, 95, 200–214. <https://doi.org/10.1016/j.jml.2017.04.007>
- Hahsler, M., Grün, B., & Hornik, K. (2005). Arules - A computational environment for mining association rules and frequent item sets. *Journal of Statistical Software*, 14. <https://doi.org/10.18637/JSS.V014.I15>
- Hale, J. (2001). *A Probabilistic Earley Parser as a Psycholinguistic Model*. <https://www.aclweb.org/anthology/N01-1021>
- Hayashi, H., Oda, Y., Birch, A., Konstas, I., Finch, A., Luong, M.-T., Neubig, G., & Sudoh, K. (2019). *Findings of the Third Workshop on Neural Generation and Translation*. 1–14. <https://doi.org/10.18653/v1/d19-5601>
- Huettig, & Guerra, E. (2015). Testing the limits of prediction in language processing: Prediction occurs but far from always. *21st Annual Conference on Architectures and Mechanisms for Language Processing (AMLaP 2015)*. https://pure.mpg.de/pubman/faces/ViewItemOverviewPage.jsp?itemId=item_2351842
- Huettig, & Mani, N. (2016). Is prediction necessary to understand language? Probably not. *Language, Cognition and Neuroscience*, 31(1), 19–31. <https://doi.org/10.1080/23273798.2015.1072223>
- Jackendoff, R. (2007). Foundations of Language. In *Foundations of Language*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198270126.001.0001>
- Johnston, R., Jones, K., & Manley, D. (2018). Confounding and collinearity in regression

- analysis: a cautionary tale and an alternative procedure, illustrated by studies of British voting behaviour. *Quality & Quantity*, 52(4), 1957–1976. <https://doi.org/10.1007/S11135-017-0584-6>
- Kuperberg, G. R., & Jaeger, T. F. (2016). What do we mean by prediction in language comprehension? *Language, Cognition and Neuroscience*, 31(1), 32–59. <https://doi.org/10.1080/23273798.2015.1102299>
- Kutas, M., & Hillyard, S. A. (1984). Brain potentials during reading reflect word expectancy and semantic association. *Nature*, 307(5947), 161–163. <https://doi.org/10.1038/307161a0>
- Landauer, T. K., & Dumais, S. T. (1997). A Solution to Plato’s problem: the latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211–240. <https://doi.org/10.1037/0033-295X.104.2.211>
- Luke, S. G., & Christianson, K. (2018). The Provo Corpus: A large eye-tracking corpus with predictability norms. *Behavior Research Methods*, 50(2), 826–833. <https://doi.org/10.3758/s13428-017-0908-4>
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *ArXiv*. <http://arxiv.org/abs/1301.3781>
- Nieuwland, M. S., Politzer-Ahles, S., Heyselaar, E., Segaert, K., Darley, E., Kazanina, N., Von Grebmer Zu Wolfsturn, S., Bartolozzi, F., Kogan, V., Ito, A., Mézière, D., Barr, D. J., Rousselet, G. A., Ferguson, H. J., Busch-Moreno, S., Fu, X., Tuomainen, J., Kulakova, E., Husband, E. M., ... Huettig, F. (2018). Large-scale replication study reveals a limit on probabilistic prediction in language comprehension. *ELife*, 7. <https://doi.org/10.7554/eLife.33468>
- Ong, J. K. Y., & Kliegl, R. (2008). Conditional co-occurrence probability acts like frequency in predicting fixation durations. *Journal of Eye Movement Research*, 2(1). <https://doi.org/10.16910/jemr.2.1.3>
- Pickering, M. J., & Gambi, C. (2018). Predicting while comprehending language: A theory and review. *Psychological Bulletin*, 144(10), 1002–1044. <https://doi.org/10.1037/BUL0000158>
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). *Language models*

are *unsupervised multitask learners*. OpenAI Blog.
<https://github.com/codelucas/newspaper>

Rayner, K., Ashby, J., Pollatsek, A., & Reichle, E. D. (2004). The effects of frequency and predictability on eye fixations in reading: Implications for the E-Z reader model. *Journal of Experimental Psychology: Human Perception and Performance*, 30(4), 720–732. <https://doi.org/10.1037/0096-1523.30.4.720>

Rayner, K., Slattery, T. J., Drieghe, D., & Liversedge, S. P. (2011). Eye Movements and Word Skipping During Reading: Effects of Word Length and Predictability. *Journal of Experimental Psychology: Human Perception and Performance*, 37(2), 514–528. <https://doi.org/10.1037/a0020990>

Reichle, E. D., & Drieghe, D. (2013). Using E-Z Reader to Examine Word Skipping During Reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(4), 1311–1320. <https://doi.org/10.1037/A0030910>

Reichle, E. D., Rayner, K., & Pollatsek, A. (2003). The E-Z reader model of eye-movement control in reading: Comparisons to other models. In *Behavioral and Brain Sciences* (Vol. 26, Issue 4, pp. 445–476). <https://doi.org/10.1017/S0140525X03000104>

Schotter, E. R., Tran, R., & Rayner, K. (2014). Don't believe what you read (Only Once): Comprehension is supported by regressions during reading. *Psychological Science*, 25(6), 1218–1226. <https://doi.org/10.1177/0956797614531148>

Schrimpf, M., Blank, I. A., Tuckute, G., Kauf, C., Hosseini, E. A., Kanwisher, N., Tenenbaum, J. B., & Fedorenko, E. (2021). The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45), e2105646118. <https://doi.org/10.1073/pnas.2105646118>

Smith, N. J., & Levy, R. (2011). Cloze but no cigar: The complex relationship between cloze, corpus, and subjective probabilities in language processing. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 33, Issue 33).

Smith, N. J., & Levy, R. (2013). The effect of word predictability on reading time is logarithmic. *Cognition*, 128(3), 302–319. <https://doi.org/10.1016/j.cognition.2013.02.013>

- Staub, A. (2015). The Effect of Lexical Predictability on Eye Movements in Reading: Critical Review and Theoretical Interpretation. *Language and Linguistics Compass*, 9(8), 311–327. <https://doi.org/10.1111/lnc3.12151>
- Staub, A., Grant, M., Astheimer, L., & Cohen, A. (2015). The influence of cloze probability and item constraint on cloze task response time. *Journal of Memory and Language*, 82, 1–17. <https://doi.org/10.1016/j.jml.2015.02.004>
- Taylor, W. L. (1953). "Cloze Procedure": A New Tool for Measuring Readability. *Journalism Quarterly*, 30(4), 415–433. <https://doi.org/10.1177/107769905303000401>
- Van Berkum, J. J. A., Brown, C. M., Zwitserlood, P., Kooijman, V., & Hagoort, P. (2005). Anticipating upcoming words in discourse: Evidence from ERPs and reading times. *Journal of Experimental Psychology: Learning Memory and Cognition*, 31(3), 443–467. <https://doi.org/10.1037/0278-7393.31.3.443>
- Van Heuven, W. J. B., Mandera, P., Keuleers, E., & Brysbaert, M. (2014). SUBTLEX-UK: A new and improved word frequency database for British English. *Quarterly Journal of Experimental Psychology*, 67(6), 1176–1190. <https://doi.org/10.1080/17470218.2013.850521>
- Van Petten, C., & Luka, B. J. (2012). Prediction during language comprehension: Benefits, costs, and ERP components. In *International Journal of Psychophysiology* (Vol. 83, Issue 2, pp. 176–190). Int J Psychophysiol. <https://doi.org/10.1016/j.ijpsycho.2011.09.015>
- Wang, L., Kuperberg, G., & Jensen, O. (2018). Specific lexico-semantic predictions are associated with unique spatial and temporal patterns of neural activity. *ELife*, 7. <https://doi.org/10.7554/eLife.39061>
- Wicha, N. Y. Y., Moreno, E. M., & Kutas, M. (2004). Anticipating words and their gender: An event-related brain potential study of semantic integration, gender expectancy, and gender agreement in Spanish sentence reading. *Journal of Cognitive Neuroscience*, 16(7), 1272–1288. <https://doi.org/10.1162/0898929041920487>
- Wilcox, E. G., Gauthier, J., Hu, J., Qian, P., & Levy, R. (2020). *On the Predictive Power of Neural Language Models for Human Real-Time Comprehension Behavior*. <http://arxiv.org/abs/2006.01912>

Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., ... Rush, A. (2020). Transformers: State-of-the-Art Natural Language Processing. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38–45. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>