

Relative Importance in Sentence Processing

Nora Hollenstein

Center for Language Technology
University of Copenhagen
nora.hollenstein@hum.ku.dk

Lisa Beinborn

Computational Linguistics & Text Mining Lab
Vrije Universiteit Amsterdam
l.beinborn@vu.nl

Abstract

Determining the relative importance of the elements in a sentence is a key factor for effortless natural language understanding. For human language processing, we can approximate patterns of relative importance by measuring reading fixations using eye-tracking technology. In neural language models, gradient-based saliency methods indicate the relative importance of a token for the target objective. In this work, we compare patterns of relative importance in English language processing by humans and models and analyze the underlying linguistic patterns. We find that human processing patterns in English correlate strongly with saliency-based importance in language models and not with attention-based importance. Our results indicate that saliency could be a cognitively more plausible metric for interpreting neural language models. The code is available on github: https://github.com/beinborn/relative_importance.

1 Introduction

When children learn to read, they first focus on each word individually and gradually learn to anticipate frequent patterns (Blythe and Joseph, 2011). More experienced readers are able to completely skip words that are predictable from the context and to focus on the more *relevant* words of a sentence (Schroeder et al., 2015). Psycholinguistic studies aim at unraveling the characteristics that determine the relevance of a word and find that lexical factors such as word class, word frequency, and word complexity play an important role, but that the effects vary depending on the sentential context (Rayner and Duffy, 1986).

In natural language processing, the relative importance of words is usually interpreted with respect to a specific task. Emotional adjectives are most relevant in sentiment detection (Socher et al.,

2013), relative frequency of a term is an indicator for information extraction (Wu et al., 2008), the relative position of a token can be used to approximate novelty for summarisation (Chopra et al., 2016), and function words play an important role in stylistic analyses such as plagiarism detection (Stamatatos, 2011). Neural language models are trained to be a good basis for any of these tasks and are thus expected to represent a more general notion of relative importance (Devlin et al., 2019).

Relative importance of the input in neural networks can be modulated by the so-called “attention” mechanism (Bahdanau et al., 2014). Analyses of image processing models indicate that attention weights reflect cognitively plausible patterns of visual saliency (Xu et al., 2015; Coco and Keller, 2012). Recent research in language processing finds that attention weights are not a good proxy for relative importance because different attention distributions can lead to the same predictions (Jain and Wallace, 2019). Gradient-based methods such as saliency scores seem to better approximate the relative importance of input words for neural processing models (Bastings and Filippova, 2020).

In this work, we compare patterns of relative importance in human and computational English language processing. We approximate relative importance for humans as the relative fixation duration in eye-tracking data collected in naturalistic language understanding scenarios. In related work, Sood et al. (2020a) measure the correlation between attention in neural networks trained for a document-level question-answering task and find that the attention in a transformer language model deviates strongly from human fixation patterns. In this work, we instead approximate relative importance in computational models using gradient-based saliency and find that it correlates much better with human patterns.

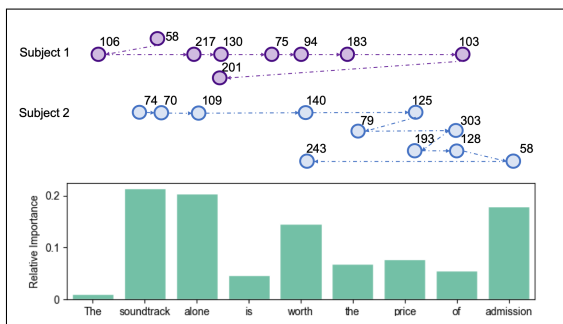


Figure 1: Example fixations for two subjects in the ZuCo dataset for the sentence “The soundtrack alone is worth the price of admission”. The numbers indicate the fixation duration and the circles represent the approximate horizontal position of the fixation (positions are simplified for better visualization). The plot at the bottom indicates the relative importance of each token averaged over all subjects.

2 Determining Relative Importance

The concept of relative importance of a token for sentence processing encompasses several related psycholinguistic phenomena such as relevance for understanding the sentence, difficulty and novelty of a token within the context, semantic and syntactic surprisal, or domain-specificity of a token. We take a data-driven perspective and approximate the relative importance of a token by the processing effort that can be attributed to it compared to the other tokens in the sentence.

2.1 In Human Language Processing

The sentence processing effort can be approximated indirectly using a range of metrics such as response times in reading comprehension experiments (Su and Davison, 2019), processing duration in self-paced reading (Linzen and Jaeger, 2016), and voltage changes in electroencephalography recordings (Frank et al., 2015). In this work, we approximate relative importance using eye movement recordings during reading because they provide online measurements in a comfortable experimental setup which is more similar to a normal, uncontrolled reading experience. Eye-tracking technology can measure with high accuracy how long a reader fixates each word. The fixation duration and the relative importance of a token for the reader are strongly correlated with reading comprehension (Rayner, 1977; Malmaud et al., 2020).

Language models that look ahead and take both the left and right context into account are often considered cognitively less plausible because humans process language incrementally from left to

right (Merks and Frank, 2020). However, in human reading, we frequently find regressions: humans fixate relevant parts of the left context again while already knowing what comes next (Rayner, 1998). In Figure 1, subject 1 first reads the entire sentence and then jumps back to the token “alone”. Subject 2 performs several regressions to better understand the second half of the sentence. The fixation duration is a cumulative measure that sums over these repeated fixations. Absolute fixation duration can vary strongly between subjects due to differences in reading speed but the relative fixation duration provides a good approximation for the relative importance of a token as it abstracts from individual differences. We average the relative fixation duration over all subjects to obtain a more robust signal (visualized in the plot at the bottom of Figure 1).

2.2 In Computational Language Processing

In computational language models, the interpretation of a token depends on the tokens in its context but not all tokens are equally important. To account for varying importance, so-called attention weights regulate the information flow in neural networks (Bahdanau et al., 2014). These weights are optimized with respect to a target objective and higher attention for an input token has been interpreted as higher importance with respect to the output (Vig, 2019). Recent research indicates that complementary attention distributions can lead to the same model prediction (Jain and Wallace, 2019; Wiegrefe and Pinter, 2019) and that the removal of input tokens with large attention weights often does not lead to a change in the model’s prediction (Serrano and Smith, 2019). In transformer models, the attention weights often approximate an almost uniform distribution in higher model layers (Abnar and Zuidema, 2020). Bastings and Filippova (2020) argue that saliency methods are more suitable for assigning importance weights to input tokens.

Saliency methods calculate the gradient of the output corresponding to the correct prediction with respect to an input element to identify those parts of the input that have the biggest influence on the prediction (Lipton, 2018). Saliency maps were first developed for image processing models to highlight the areas of the image that are discriminative with respect to the tested output class (Simonyan et al., 2014). Li et al. (2016) adapt this method to calculate the relative change of the output probabilities with respect to individual input tokens in

text classification tasks and Ding et al. (2019) calculate saliency maps for interpreting the alignment process in machine translation models.

In general-purpose language models such as BERT (Devlin et al., 2019), the objective function tries to predict a token based on its context. A saliency vector for a masked token thus indicates the importance of each of the tokens in the context of correctly predicting the masked token (Madsen, 2019).

We iterate over each token vector \mathbf{x}_i in our input sequence x_1, x_2, \dots, x_n . Let \mathbf{X}_i be the input matrix with x_i being masked. The saliency s_{ij} for input token x_j for the prediction of the correct token t_i is then calculated as the Euclidean norm of the gradient of the logit for x_i .

$$s_{ij} = \|\nabla_{\mathbf{x}_j} f_{t_i}(\mathbf{X}_i)\|_2 \quad (1)$$

The saliency vector \mathbf{s}_i indicates the relevance of each token for the correct prediction of the masked token t_i .¹ The saliency scores are normalized by dividing by the maximum. We determine the relative importance of a token by summing over the saliency scores for each token. For comparison, we also approximate importance using attention values from the last layer of each model as Sood et al. (2020a).

2.3 Patterns of Relative Importance

Relative importance in human processing and in computational models is sensitive to linguistic properties. Rayner (1998) provides a detailed overview of token-level features that have been found to correlate with fixation duration such as length, frequency, and word class. On the contextual level, lexical and syntactic disambiguation processes cause regressions and thus lead to longer fixation duration (Just and Carpenter, 1980; Lowder et al., 2018). Computational models are also highly susceptible to frequency effects and surprisal metrics calculated using language models can predict the human processing effort (Frank et al., 2013).

The inductive bias of language processing models can be improved using the eye-tracking signal (Barrett et al., 2018; Klerke and Plank, 2019) and the modification leads to more “human-like” output in generative tasks (Takmaz et al., 2020; Sood et al., 2020b). This indicates that patterns of relative importance in computational representations

¹Our implementation adapts code from <https://pypi.org/project/textualheatmap/>. An alternative would be to multiply saliency and input (Alammar, 2020).

	Dataset	BERT	Distil	ALBERT	Rand
Saliency	GECO	.54	.51	.48	.00
	ZuCo	.68	.64	.62	.00
Attention	GECO	.18	.06	.26	.00
	ZuCo	.11	.03	.37	.00

Table 1: Spearman correlation between relative fixation duration by humans and attention and saliency in the language models. Correlation values are averaged over all sentences. *Rand* is a permutation baseline.

differ from human processing patterns. Previous work focused on identifying links between the eye-tracking signal and attention (Sood et al., 2020a). To our knowledge, this is the first attempt to correlate fixation duration with saliency metrics.

The eye-tracking signal represents human reading processes aimed at language understanding. In previous work, we have shown that contextualized language models can predict eye patterns associated with human reading (Hollenstein et al., 2021), which indicates that computational models and humans encode similar linguistic patterns. It remains an open debate to which extent language models are able to approximate language understanding (Bender and Koller, 2020). We are convinced that language needs to be cooperatively grounded in the real world (Beinborn et al., 2018). Purely text-based language models clearly miss important aspects of language understanding but they can approximate human performance in an impressive range of processing tasks. We aim to gain a deeper understanding of the similarities and differences between human and computational language processing to better evaluate the capabilities of language models.

3 Methodology

We extract relative importance values for tokens from eye-tracking corpora and language models as described in section 2 and calculate the Spearman correlation for each sentence.² We first average the correlation over all sentences to analyze whether the importance patterns of humans and models are comparable and then conduct token-level analyses.

3.1 Eye-tracking Corpora

We extract the relative fixation duration from two eye-tracking corpora and average it over all readers for each sentence. Both corpora record natural reading and the text passages were followed by

²Kendall’s τ and KL divergence yield similar results.

multiple-choice questions to test the readers’ comprehension.

GECO contains eye-tracking data from 14 native English speakers reading the entire novel *The Mysterious Affair at Styles* by Agatha Christie (Cop et al., 2017). The text was presented on the screen in paragraphs.

ZuCo contains eye-tracking data of 30 native English speakers reading full sentences from movie reviews and Wikipedia articles (Hollenstein et al., 2018, 2020).³

3.2 Language Models

We compare three state-of-the-art language models trained for English: BERT, ALBERT, and DistilBERT.⁴ BERT was the first widely successful transformer-based language model and remains highly influential (Devlin et al., 2019). ALBERT and DistilBERT are variants of BERT that require less training time due to a considerable reduction of the training parameters while maintaining similar performance on benchmark datasets (Lan et al., 2019; Sanh et al., 2019).⁵ We analyze if the lighter architectures have an influence on the patterns of relative importance that the models learn.

4 Results

The results in Table 1 show that relative fixation duration by humans strongly correlates with the saliency values of the models. In contrast, attention-based importance does not seem to be able to capture the human importance pattern. A random permutation baseline that shuffles the importance assigned by the language model yields no correlation (0.0) in all conditions.⁶ As the standard deviations of the correlation across sentences are quite high (ZuCo: ~ 0.22 , GECO: ~ 0.39), the small differences between models can be neglected (although they are consistent across corpora). For the subsequent analyses, we focus only on the BERT model

³We combine ZuCo 1.0 (T1, T2) and ZuCo 2.0. (T1).

⁴We use the *Huggingface* transformers implementation (Wolf et al., 2020) and the models `bert-based-uncased`, `albert-base-v2`, and `distilbert-base-uncased`.

⁵Reduction is achieved by parameter sharing across layers (ALBERT) and by distillation which approximates the output distribution of the original BERT model using a smaller network (DistilBERT). See model references for details.

⁶We repeat the permutation 100 times and average the correlation over all iterations.

		Length		Frequency	
		Sent	Tok	Sent	Tok
GECO	Human	.69	.31	-.36	-.25
	BERT	.65	.27	-.48	-.28
ZuCo	Human	.75	.47	-.52	-.36
	BERT	.72	.36	-.65	-.40

Table 2: Spearman correlation between relative importance and word length and frequency. For the *Sent* condition, correlation is calculated per sentence and averaged. For *Tok*, importance is normalized by sentence length and correlation is calculated over all tokens.

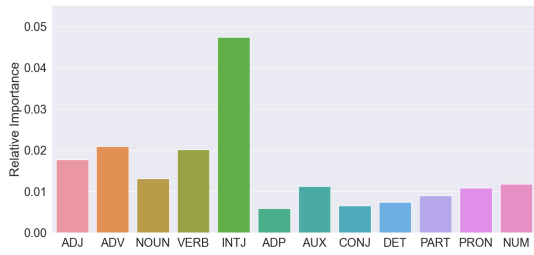
which yields the best results. The differences between the corpora might be related to the number of sentences and the differences in average sentence length (ZuCo: 924, 19.5, GECO: 4,926, 12.7).

Length and Frequency In eye-tracking data, word length correlates with fixation duration because it takes longer to read all characters. The correlation for frequency is inverse because high-frequency words (e.g. “the”, “has”) are often skipped in processing as they carry (almost) no meaning (Rayner, 1998). For English, word frequency and word length are both closely related to word complexity (Beinborn et al., 2014). Language models do not directly encode word length but they are sensitive to word frequency.


Our results in Table 2 show that both token length and frequency are strongly correlated with relative importance on the sentence level. Interestingly, the correlation decreases when it is calculated directly over all tokens indicating that the token-level relation between length and importance is more complex than the correlation might suggest.

Word Class Figure 2 shows the average relative importance of all tokens belonging to the same word class (normalized by sentence length). We see that both humans and BERT clearly assign higher importance to content words (left) than to function words (right). Interjections such as “Oh” in figure 3 receive the highest relevance which is understandable because they interrupt the reading flow. When we look at individual sentences, we note that the differences in importance are more pronounced in the model saliency while human fixation duration yields a smoother distribution over the tokens.

Novelty We extract the language model representations for each sentence separately whereas the readers processed the sentences consecutively. If tokens are mentioned repeatedly such as “Sherlock



(a) Human Fixation



(b) Model saliency

Figure 2: Relative importance of tokens with respect to word class. Relative importance is measured as relative fixation duration for humans in GECO (left) and as relative gradient-based saliency in the BERT model (right).




Figure 3: Relative importance values for an example sentence from the GECO corpus for the BERT model and the human values.

Holmes” which also occurred in the sentence preceding the example in Figure 3), processing ease increases for the reader, and not for the model. Some language models are able to process multiple sentences, but establishing semantic links across sentences remains a challenge.

5 Conclusion

We find that human sentence processing patterns in English correlate strongly with saliency-based importance in language models and not with attention-based importance. Our results indicate that saliency could be a cognitively more plausible metric for interpreting neural language models. In future work, it would be interesting to test the robustness of the approach with different variants for calculating saliency (Bastings and Filippova, 2020; Ding and Koehn, 2021). As we conducted our analyses only for English data, it is not yet clear whether our results generalize across languages. We will address this in future work using eye-tracking data from non-English readers (Makowski et al., 2018; Laurinavichyute et al., 2019) and comparing mono- and multilingual models (Beinborn and Choenni, 2020). We want to extend the token-level analyses to syntactic phenomena and cross-sentence effects. For example, it would be interesting to see how a language model encodes relative importance for sentences that are syntactically correct but not se-

mantically meaningful (Gulordava et al., 2018).

Previous work has shown that the inductive bias of recurrent neural networks can be modified to obtain cognitively more plausible model decisions (Bhatt et al., 2020; Shen et al., 2019). In principle, our approach can also be applied to left-to-right models such as GPT-2 (Radford et al., 2019). In this case, the tokens at the beginning of the sentence would be assigned disproportionately high importance as the following tokens cannot contribute to the prediction of preceding tokens in incremental processing. It might thus be more useful to only use the first fixation duration of the gaze signal for analyzing importance in left-to-right models. However, we think that the regressions by the readers provide valuable information about sentence processing.

6 Ethical Considerations

Data from human participants were leveraged from freely available datasets (Hollenstein et al., 2018, 2020; Cop et al., 2017). The datasets provide anonymized records in compliance with ethical board approvals and do not contain any information that can be linked to the participants.

Acknowledgements

Lisa Beinborn’s research was partially funded by the Dutch National Science Organisation (NWO) through the project CLARIAH-PLUS (CP-W6-19-005). We thank the anonymous reviewers for their constructive feedback.

References

Samira Abnar and Willem Zuidema. 2020. [Quantifying attention flow in transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4190–4197, Online. Association for Computational Linguistics.

- Jay Alamar. 2020. [Interfaces for explaining transformer language models](#).
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Maria Barrett, Joachim Bingel, Nora Hollenstein, Marek Rei, and Anders Søgaard. 2018. [Sequence classification with human attention](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 302–312, Brussels, Belgium. Association for Computational Linguistics.
- Jasmijn Bastings and Katja Filippova. 2020. [The elephant in the interpretability room: Why use attention as explanation when we have saliency methods?](#) In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 149–155, Online. Association for Computational Linguistics.
- Lisa Beinborn, Teresa Botschen, and Iryna Gurevych. 2018. [Multimodal grounding for language processing](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2325–2339, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Lisa Beinborn and Rochelle Choenni. 2020. [Semantic drift in multilingual representations](#). *Computational Linguistics*, 46(3):571–603.
- Lisa Beinborn, Torsten Zesch, and Iryna Gurevych. 2014. [Predicting the difficulty of language proficiency tests](#). *Transactions of the Association for Computational Linguistics*, 2:517–530.
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Gantavya Bhatt, Hritik Bansal, Rishubh Singh, and Sumeet Agarwal. 2020. [How much complexity does an RNN architecture need to learn syntax-sensitive dependencies?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 244–254, Online. Association for Computational Linguistics.
- Hazel Blythe and Holly Joseph. 2011. *Children’s Eye Movements during Reading*.
- Sumit Chopra, Michael Auli, and Alexander M. Rush. 2016. [Abstractive sentence summarization with attentive recurrent neural networks](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 93–98, San Diego, California. Association for Computational Linguistics.
- Moreno I. Coco and Frank Keller. 2012. [Scan patterns predict sentence production in the cross-modal processing of visual scenes](#). *Cognitive Science*, 36(7):1204–1223.
- Uschi Cop, Nicolas Dirix, Denis Drieghe, and Wouter Duyck. 2017. Presenting GECO: An eyetracking corpus of monolingual and bilingual sentence reading. *Behavior Research Methods*, 49(2):602–615.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Shuoyang Ding and Philipp Koehn. 2021. [Evaluating saliency methods for neural language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5034–5052, Online. Association for Computational Linguistics.
- Shuoyang Ding, Hainan Xu, and Philipp Koehn. 2019. [Saliency-driven word alignment interpretation for neural machine translation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 1–12, Florence, Italy. Association for Computational Linguistics.
- Stefan L. Frank, Leun J. Otten, Giulia Galli, and Gabriella Vigliocco. 2013. Word surprisal predicts n400 amplitude during reading. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 878–883. ACL.
- Stefan L. Frank, Leun J. Otten, Giulia Galli, and Gabriella Vigliocco. 2015. The ERP response to the amount of information conveyed by words in sentences. *Brain and Language*, 140:1–11.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. [Colorless green recurrent networks dream hierarchically](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1195–1205, New Orleans, Louisiana. Association for Computational Linguistics.
- Nora Hollenstein, Federico Pirovano, Ce Zhang, Lena Jäger, and Lisa Beinborn. 2021. [Multilingual language models predict human reading behavior](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 106–123, Online. Association for Computational Linguistics.

- Nora Hollenstein, Jonathan Rotsztein, Marius Troendle, Andreas Pedroni, Ce Zhang, and Nicolas Langer. 2018. ZuCo, a simultaneous EEG and eye-tracking resource for natural sentence reading. *Scientific Data*.
- Nora Hollenstein, Marius Troendle, Ce Zhang, and Nicolas Langer. 2020. [ZuCo 2.0: A dataset of physiological recordings during natural reading and annotation](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 138–146, Marseille, France. European Language Resources Association.
- Sarthak Jain and Byron C. Wallace. 2019. [Attention is not Explanation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3543–3556, Minneapolis, Minnesota. Association for Computational Linguistics.
- Marcel A Just and Patricia A Carpenter. 1980. A theory of reading: From eye fixations to comprehension. *Psychological review*, 87(4):329.
- Sigrid Klerke and Barbara Plank. 2019. [At a glance: The impact of gaze aggregation views on syntactic tagging](#). In *Proceedings of the Beyond Vision and LAnguage: inTEgrating Real-world kNowledge (LANTERN)*, pages 51–61, Hong Kong, China. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. [ALBERT: A lite BERT for self-supervised learning of language representations](#). *CoRR*, abs/1909.11942.
- AK Laurinavichyute, Irina A Sekerina, SV Alexeeva, and KA Bagdasaryan. 2019. Russian Sentence Corpus: Benchmark measures of eye movements in reading in Cyrillic. *Behavior research methods*, 51(3):1161–1178.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016. [Understanding neural networks through representation erasure](#). *CoRR*, abs/1612.08220.
- Tal Linzen and T. Florian Jaeger. 2016. [Uncertainty and expectation in sentence processing: Evidence from subcategorization distributions](#). *Cognitive Science*, 40(6):1382–1411.
- Zachary C. Lipton. 2018. [The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery](#). *Queue*, 16(3):31–57.
- Matthew W Lowder, Wonil Choi, Fernanda Ferreira, and John M Henderson. 2018. Lexical predictability during natural reading: Effects of surprisal and entropy reduction. *Cognitive science*, 42:1166–1183.
- Andreas Madsen. 2019. [Visualizing memorization in rnns](#). *Distill*. <https://distill.pub/2019/memorization-in-rnns>.
- Silvia Makowski, Lena A Jäger, Ahmed Abdelwahab, Niels Landwehr, and Tobias Scheffer. 2018. A discriminative model for identifying readers and assessing text comprehension from eye movements. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 209–225. Springer.
- Jonathan Malmaud, Roger Levy, and Yevgeni Berzak. 2020. Bridging information-seeking human gaze and machine reading comprehension. In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 142–152.
- Danny Merx and Stefan L Frank. 2020. Comparing transformers and RNNs on predicting human sentence processing data. *arXiv preprint arXiv:2005.09471*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Keith Rayner. 1977. Visual attention in reading: Eye movements reflect cognitive processes. *Memory & Cognition*, 5(4):443–448.
- Keith Rayner. 1998. Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, 124(3):372.
- Keith Rayner and Susan A Duffy. 1986. Lexical complexity and fixation times in reading: Effects of word frequency, verb complexity, and lexical ambiguity. *Memory & Cognition*, 14(3):191–201.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Sascha Schroeder, Jukka Hyönä, and Simon Liv-ersedge. 2015. [Developmental eye-tracking research in reading: Introduction to the special issue](#). *Journal of Cognitive Psychology*, 27:500–510.
- Sofia Serrano and Noah A. Smith. 2019. [Is attention interpretable?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2931–2951, Florence, Italy. Association for Computational Linguistics.
- Yikang Shen, Shawn Tan, Alessandro Sordani, and Aaron Courville. 2019. [Ordered neurons: Integrating tree structures into recurrent neural networks](#). In *International Conference on Learning Representations*.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. [Deep inside convolutional networks: Visualising image classification models and saliency maps](#). In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings*.

- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642.
- Ekta Sood, Simon Tannert, Diego Frassinelli, Andreas Bulling, and Ngoc Thang Vu. 2020a. [Interpreting attention models with human visual attention in machine reading comprehension](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 12–25, Online. Association for Computational Linguistics.
- Ekta Sood, Simon Tannert, Philipp Mueller, and Andreas Bulling. 2020b. [Improving natural language processing tasks with human gaze-guided neural attention](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Efstathios Stamatatos. 2011. Plagiarism detection using stopword n-grams. *Journal of the American Society for Information Science and Technology*, 62(12):2512–2527.
- Shiyang Su and Mark L Davison. 2019. Improving the predictive validity of reading comprehension using response times of correct item responses. *Applied Measurement in Education*, 32(2):166–182.
- Ece Takmaz, Sandro Pezzelle, Lisa Beinborn, and Raquel Fernández. 2020. [Generating Image Descriptions via Sequential Cross-Modal Alignment Guided by Human Gaze](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4664–4677, Online. Association for Computational Linguistics.
- Jesse Vig. 2019. [A multiscale visualization of attention in the transformer model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 37–42, Florence, Italy. Association for Computational Linguistics.
- Sarah Wiegrefe and Yuval Pinter. 2019. [Attention is not not explanation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 11–20, Hong Kong, China. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame,
- Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Ho Chung Wu, Robert Wing Pong Luk, Kam Fai Wong, and Kui Lam Kwok. 2008. [Interpreting tf-idf term weights as making relevance decisions](#). *ACM Trans. Inf. Syst.*, 26(3).
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. [Show, attend and tell: Neural image caption generation with visual attention](#). In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2048–2057, Lille, France. PMLR.

A Additional Results




Figure 4: Relative importance of tokens with respect to word class in the GECO dataset. Relative importance is measured as relative fixation duration for humans (top) and as relative gradient-based saliency in the BERT model (bottom). This is the same figure as Figure 2 in the paper but it includes the number of instances per word class on top of the respective bar.




Figure 5: Relative importance values with respect to word length from human readers and from the BERT model for the GECO corpus.




Figure 6: Relative importance values with respect to word frequency from human readers and from the BERT model for the GECO corpus.