

A Comparative Study on Textual Saliency of Styles from Eye Tracking, Annotations, and Language Models

Karin de Langis
University of Minnesota
dento019@umn.edu

Dongyeop Kang
University of Minnesota
dongyeop@umn.edu

Abstract

There is growing interest in incorporating eye-tracking data and other implicit measures of human language processing into natural language processing (NLP) pipelines. The data from human language processing contain unique insight into human linguistic understanding that could be exploited by language models. However, many unanswered questions remain about the nature of this data and how it can best be utilized in downstream NLP tasks. In this paper, we present eyeStyliency, an eye-tracking dataset for human processing of stylistic text (e.g., politeness). We develop a variety of methods to derive style saliency scores over text using the collected eye dataset. We further investigate how this saliency data compares to both human annotation methods and model-based interpretability metrics. We find that while eye-tracking data is unique, it also intersects with both human annotations and model-based importance scores, providing a possible bridge between human- and machine-based perspectives. We propose utilizing this type of data to evaluate the cognitive plausibility of models that interpret style. Our eye-tracking data and processing code are publicly available.¹

1 Introduction

Human perception and understanding of text is critical in NLP. Typically, this understanding is leveraged in the form of ground-truth human annotations in supervised learning pipelines, or in the form of human evaluations of generated text. However, human language understanding is complex; multiple cognitive processes work together to enable reading, many of which occur automatically and unconsciously (DeVito, 1970).

Because of the complexity, disciplines concerned with understanding and modeling how humans read – e.g., psycholinguistics and cognitive science – heavily utilize *implicit* measures of the

¹<https://github.com/minnesotanlp/eyeStyliency>

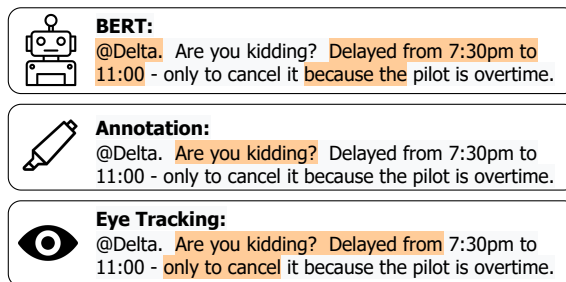


Figure 1: Salient words for *impoliteness* from three different perspectives. We find that eye tracking data contains some overlap between machine and human-annotated saliency.

human reading experience that capture signals from these automatic processes in real time. Examples of implicit measures include event-related potential, reaction times, and eye movements. In contrast, *explicit* measures include surveys and other methods that directly ask people to report their perceptions and experiences. We posit that traditional NLP pipelines, which have widely used explicit measures of human understanding, can also benefit from implicit measures. In this paper, we focus specifically on *the use of eye movements as an implicit measure* of textual saliency.

Recent research in NLP has demonstrated the feasibility of incorporating various types of eye movement data into NLP models in order to improve performance on a number of tasks (see Table 2 for an overview). However, this is still an underexplored area: best practices remain unclear, and it's not obvious whether there are tasks that are unsuitable for eye movement data, or how eye movement data should be balanced with traditional annotation data. In this work, we address two main research questions: RQ1: Does eye-tracking-based saliency meaningfully differ from simply gathering word-level human annotations, or from model-based word importance measures? RQ2: How can we measure eye movements specific to a high-level textual feature like style, and which eye tracking metrics and data processing methods are best suited to capturing textual saliency?

To address these questions, we conduct an eye tracking case study in which participants read texts the HummingBird dataset (Hayati et al., 2021). We choose this dataset because it contains lexical-level human annotations indicating which words contribute to the text’s style and because its domain (textual styles) has not to our knowledge been widely explored for eye tracking applications – although prior work investigates eye tracking and sentiment analysis, it does not extend to other linguistic styles such as politeness.

We collect style-specific eye movements through a carefully designed experiment (see Section 3 for details), and we use these eye movements to derive saliency scores over the text. We compare this eye-based saliency to human annotations as well as two large language model (LLM)-derived importance scores: integrated gradient scores from a BERT model fine-tuned on style datasets (Hayati et al., 2021), and word-surprisal scores from GPT-2 (Radford et al., 2019) (see Figure 1 for an example). Our findings indicate that eye-tracking-based saliency highlights some unique areas of the text, but it also intersects with both saliency from model-based metrics and saliency from human annotations, making a bridge of sorts between the human- and machine-based perspectives. We discuss some implications of these findings for NLP research.

Specifically, our contributions are:

- An experimental paradigm for obtaining eye tracking-based signals for specific features of text (in our case, textual style).
- A first-of-its-kind eye movement dataset on style saliency, collected from 20 participants and consisting of both control readings and style-focused readings for polite, impolite, positive, and negative textual styles.
- An illustration of the distinction between this dataset’s **explicit** human annotations and **implicit** human eye data through a unique comparison between salient text obtained via annotation and via eye tracking.

2 Related Work

Eye tracking has been a staple of psycholinguistic investigations of reading for decades (Rayner, 1978; Just and Carpenter, 1980). Eye movement data is compelling because it provides realtime information about how people process language in a natural, ecologically valid setting (i.e., there is no

	NLP Area	H	M	learning from eye data
Ours	Textual Style	✓	✓	✗
Kuribayashi et al. (2021)	Perplexity	✗	✓	✗
Malmaud et al. (2020)	QA	✗	✗	Joint learning
Bolotova et al. (2020)	QA	✗	✓	✗
Sood et al. (2020b)	QA	✗	✓	✗
Sood et al. (2020a)	Paraphrasing	✗	✗	Joint learning
Hollenstein et al. (2019)	Sentiment Clf., NER	✗	✗	Joint learning
Barrett et al. (2018)	PoS tagging	✗	✗	HMM
Tokunaga et al. (2017)	NER	✗	✗	✗
Klerke et al. (2015)	Summarization	✓	✗	✗
Green (2014)	Parsing	✗	✗	✗

Table 1: A summary of prior work applying eye tracking methods to NLP. The **H** column indicates whether traditional human annotations are considered in relation to the eye tracking data, and the **M** indicates whether model attention is considered. Most prior research has focused on either (a) comparing and contrasting eye movements with various models’ attention mechanisms, or (b) using eye movements for multi-task learning, where NLP task performance can be improved by a model that jointly learns to predict eye movements in addition to the relevant NLP task. To our knowledge, there have not been three-way comparisons between attention mechanisms from eye tracking, large language models, and manual human annotations.

explicit experimental task, such as question answering, for participants to complete) (Kaiser, 2013). Eye data provides insight into cognitive processes through the eye-mind assumption, which posits that (1) our eyes fixate on whatever our brains are currently processing, and (2) as cognitive effort to process an item increases, the amount of time that the eyes fixate on that item also increases (Just and Carpenter, 1980). Analysis of eye data under this framework has led to important insights into many unconscious phenomena in human language comprehension, e.g. the mechanisms involved in ambiguity resolution during reading (Traxler and Frazier, 2008).

Eye Tracking in NLP. Due to the eye-mind assumption, eye-tracking data is particularly well-suited to inferring patterns of reader attention, or saliency, over text. This saliency information has so far shown promising results when integrated into NLP models for question answering (e.g. Malkin et al. (2022); Sood et al. (2020a); Malmaud et al. (2020)). However, this is still a developing research area: there is limited available data, and there is little consensus regarding how to effectively collect

data and incorporate it into NLP pipelines. To our knowledge there is no previous research that investigates saliency for style via eye tracking, nor any previous research that compares saliency from eye tracking to human annotations (Table 1 compares our work with the prior work).

Outside of textual saliency, eye-tracking data has been leveraged for a variety of NLP tasks. Mishra et al. (2013) quantify the difficulty of sentences in machine translation tasks using eye movement data; Mishra et al. (2016) determine whether a reader understands sarcasm in text, and Sjøgaard (2016) evaluate the quality of word embeddings and text generations, respectively. Other work uses existing datasets, sometimes augmenting the data with a learned gaze predictor model, and uses this eye movement data as an additional signal when training models for various NLP tasks, including named entity recognition (Hollenstein et al., 2019; Tokunaga et al., 2017), paraphrasing (Sood et al., 2020b), part-of-speech tagging (Barrett et al., 2018), and sentiment analysis (see also Mathias et al. (2020) for a review).

Saliency in Linguistic Styles. People apply styles to language in order to express attitudes, reflect interpersonal intentions or goals, or convey social standings of the speaker or listener. (Note that while many sociolinguistics theories distinguish between textual style and textual attributes, in this work, we follow the common convention in recent NLP papers of broadly using ‘style’ to encompass both of these ideas (Jin et al., 2022).) The meaning expressed by textual styles can be significant; in fact, there is strong evidence that effective communication requires an understanding of both style and literal semantic meaning (Hovy, 1987). Although BERT (Devlin et al., 2018) based fine-tuned models show strong performance on style classification, there are notable differences between how BERT perceives style at the lexical level and how humans perceive it, and that using data about these differences during training improves model performance (Hayati et al., 2023).

3 eyeStyliency: A Dataset of Eye Movement for Textual Saliency

We describe the data collection procedure for eyeStyliency dataset from 20 participants and methods for computing saliency scores over text.

3.1 Data Setups

Our dataset consists of items from the Hummingbird dataset (Hayati et al., 2021) in the following stylistic categories: *polite*, *impolite*, *positive sentiment*, and *negative sentiment*.² We chose this subset because of the small correlation between categories (other categories, e.g. *anger*, *disgust*, and *negative sentiment* are all highly correlated).

In this study, we limit participants’ total time commitment to one hour. To achieve this, the dataset size is 90 items across the four style categories. (The average word count per item in the dataset is 21.6 overall; for the impolite, polite, negative, and positive styles average word count is 21.3, 22.8, 21.4, and 20.8, respectively.) Most participants finished the experiment in 40-60 minutes, depending on both the individual’s reading speed and the time needed to calibrate the individual to the eye tracker.

3.2 Eye-Tracking Measures

Monocular eye movement data is collected with an EyeLink 1000 Plus³ at a rate of 1000Hz. We look at the following eye-tracking metrics:

- First Fixation Duration (FFD): The duration of the first fixation in an interest area.
 - First Run Dwell Time (FRD): The time interval beginning with the first fixation in the interest area and ending when the eye exits an interest area (whether to the right or left).
 - Go Past Time (GP): The time interval beginning with the first fixation in an interest area and ending when the eye exits the interest area to the *left* (i.e., to reread).
 - Dwell Time (DT): The total fixation duration for all fixations in an interest area. Also known as gaze duration.
 - Reread Time (RR): The total fixation duration for all fixations in an interest area after the area has already been entered and exited once.
 - Pupil Size (PS): The average pupil size over all fixations in an interest area.
- (Note that First Run Dwell Time + Reread Time = Dwell Time.)

These measures can broadly be categorized into early measures (first fixation duration, pupil size) that reflect more low-level reading processes and

²Politeness and sentiment datasets in Hummingbird are originally sourced from Danescu-Niculescu-Mizil et al. (2013) and Socher et al. (2013).

³Made by SR Research, Ontario, Canada; <https://www.sr-research.com/eyelink-1000-plus/>

	Applications	N	FFD	FC	FRD	DT	RR	RC	PL
eyeStyliency (Ours)	Textual Style	20	✓	✗	✓	✓	✓	✗	✓
Kuribayashi et al. (2021)	Language model perplexity	✗	✗	✗	✗	✓	✗	✗	✗
Malmaud et al. (2020)	Question Answering	269	✗	✗	✗	✓	✗	✗	✗
Bolotova et al. (2020)	Question Answering	20	✗	✓	✗	✓	✓	✗	✗
Sood et al. (2020b)	Paraphrasing	✗	✗	✓	✗	✗	✗	✗	✗
Sood et al. (2020a)	Question Answering	23	✗	✗	✗	✓	✗	✗	✗
Hollenstein et al. (2019)	NER, Sentiment/Relation Classification	✗	✓	✓	✓	✓	✓	✓	✓
Barrett et al. (2018)	PoS tagging	✗	✓	✗	✗	✓	✓	✓	✗
Tokunaga et al. (2017)	Named entity recognition (NER)	✗	✗	✗	✗	✓	✗	✗	✗
Mishra et al. (2016)	Sarcasm detection	7	✗	✗	✗	✓	✗	✗	✗
Klerke et al. (2015)	NLG evaluation	24	✗	✓	✗	✓	✓	✓	✓
Green (2014)	Phrase-structure parsing	40	✗	✗	✗	✗	✓	✗	✗

Table 2: A comparison of prior works with respect to the eye tracking metrics studied, data processing techniques, and number of participants whose eye tracking data is collected. FFD = first fixation duration, FC = fixation count, RC = regression count, RR = reread time, PL = pupil size, N = number of participants if new eye data collected.

late measures (go past time, dwell time, reread time) that reflect higher-level processing and meaning integration (Conklin et al., 2021). Previous eye tracking applications for NLP have commonly used dwell time, but a variety of measures have been examined (see Table 2). In this study, we aim to compare a wide variety of measures in order to estimate which may be best-suited to capturing textual saliency. Note that to avoid redundancy, we chose to omit fixation counts from our analysis after finding high correlations between this measure and dwell time (pearson’s $r = 0.93$, $p < 0.01$). We also chose to omit regression counts from our analysis after finding that regression counts were extremely sparse – specifically, 1.8% of the dataset had a non-zero regression count.

3.3 Experimental Procedure

The experiment follows a between-subjects, blocked design. The key part of our experiment is the technique to isolate eye movements that are specifically relevant to the text’s style. In order to do this, we inform participants at the beginning of each block that the block will contain only stimuli that share a style (polite, impolite, positive, or negative) and source (Twitter, IMDb, or Stack Exchange/Wikipedia forums) – but in fact, we will occasionally present an *incongruent* style in the block (e.g., present an impolite Tweet during the polite Tweet block). We expect that incongruency to cause readers to pay more attention to style-specific aspects of the text, as they are unexpected. We are interested in comparing the eye movements

of participants who read a stimulus in the congruent condition with those of participants who read that stimulus in the incongruent condition. Note that the experiment has a between-subjects design, i.e. the same participant does not see the same text in both conditions. The congruent reading of the text provides a control. Figure 2 shows a concrete example of these two conditions, while Figure 3 shows a visualization of these contrasted eye movements.

Figure 4 shows a procedure of our experiments. The experimental procedure is as follows (more details in Appendix A). Participants complete nine blocks. At the beginning of block, the participant is informed of the style and source, and asked to pay attention to the style of the following texts. Each block contains 10 items, eight of which are *congruent* with the target style. The remaining two items are *incongruent* with the target style. Incongruent items are counterbalanced across participants. Blocks are presented in a random order, and items within the blocks are pseudorandomized to ensure adequate spacing between congruent and incongruent trials (Egner, 2007) (there is also a *context-free* text as an added control). Participants are asked True/False comprehension questions pseudorandomly after 30% of the items in order to maintain motivation to read carefully. After the experiment concludes, participants complete the Perceived Awareness of Research Hypothesis Scale (PARH) (Rubin, 2016) to evaluate whether demand characteristics (Nichols and Maner, 2008) of the experiment may have influenced reading behavior. The study procedure was approved by the

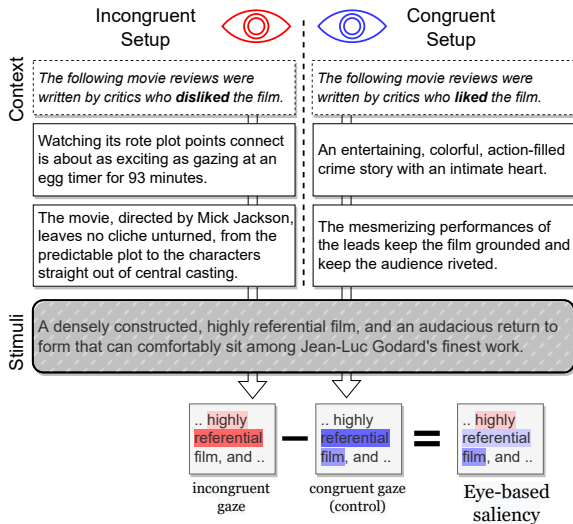


Figure 2: Illustrative example of congruent vs incongruent presentation of the same stimulus. We rely on expectation effects to induce participants to attend to the unexpected style (in this case, positive sentiment); in other words, we assume that the surprise regarding the style will result in longer gaze durations for words that contribute to the perception of that style — in this case, words relating to positive sentiment.

institutional review board (IRB).

Participants We collect data from 20 participants (12 male, 7 female, 1 non-binary; median age 23 years) recruited from the University community and word-of-mouth. An additional 6 participants were recruited but unable to complete the study due to problems with eye calibration. Participants were compensated with a \$15 Amazon gift card.

Apparatus Monocular eye movement data is collected with an EyeLink 1000 Pro, using the desktop mount, at a rate of 1000Hz. Participants use a chin-rest while reading in order to stabilize the head. We use the Experiment Builder software to present stimuli to participants in a 16pt serif font with 1.5 line spacing, on our display monitor with a 508mm display area and a 1680x1050 resolution. Participants are seated with their eyes 50-60cm away from the display monitor.

Study Design Rationale Based on the well-documented phenomenon of expectancy effects in cognition (see Schwarz et al. (2016) for further discussion), we assume that the *incongruent* texts that subvert the stylistic expectation will lead to participants reacting with surprise and increased processing difficulty in response to parts of the text associated with the unexpected style.

Alternative designs that explicitly ask participants to classify an item’s style were strongly considered, but were rejected for two reasons: first, we are interested in observing a relatively natural read-




Figure 3: Exemplary eye-tracking data showing saliency for *polite* style, with comparison to human word-level style importance highlighting. The eye-tracking data is visualized as a heat map showing gaze data from the incongruent style condition, with the gaze data from the congruent style (control) condition subtracted.




Figure 4: Experimental procedure.

ing process and introducing a classification task runs counter to that goal; second, the style classification task could increase the saliency of not only the target style but also its opposing style, as both can be relevant to the decision (e.g., the presence of an impolite word is relevant to the decision of whether a statement is polite). We also considered designs in which congruency is established via explicit text labels rather than implicit expectations, but decided to instead choose an experimental paradigm that adheres as closely as possible to an ecologically valid reading task.

3.4 Pre-processing Eye Tracking Data

Eye data was delineated into fixations and saccades using the DataViewer software with EyeLink’s standard algorithm and default velocity and acceleration thresholds. We further cleaned the data by removing trials with significant track loss (i.e. trials with track loss in over 50% of the text area); 1.5% of trials were removed due to track loss. An outlier analysis showed that 0.5% of fixations were outliers and were removed in our analysis.

3.5 Calculating Saliency Scores

We divide the text into interest areas (IAs) and calculate saliency scores for each IA. We do not segment the IAs such that each IA contains a single word, because in a single fixation people can read a span of about 21 surrounding characters (Rayner, 1978), meaning that many short words are not fixated on, leading to difficulties with our desired

analyses. Instead, we use the natural language processing toolkit (NLTK)’s stopwords list (Bird et al., 2009) to define each IA such that stopwords share an IA with the closest non-stopword. Specifically, each stopword is combined with the closest non-stopword, with non-stopwords to the right being preferred in the case of a tie. We also ensure that no IA contains a line break.

We utilize two techniques for calculating each eye tracking-based metric for each IA_i . Note that these techniques are applied across all eye tracking measures $x \in \{DT, FRD, GP, DT, RR, PS\}$ as defined in Section 3.2.

- **z-score:** For each participant p_k , denote the eye tracking measurement in IA_i as x_{ki} . We calculate the participant-specific z-score of eye tracking measurement from IA_i as $z_k(IA_i) = \frac{x_{ki} - \mu_k}{\sigma_k}$, where μ_k and σ_k are the participant-specific arithmetic mean and standard deviation, respectively. Then, the saliency score for IA_i is given by $\frac{\sum_{k=0}^n z_k(IA_i)}{n}$.
- **raw:** We aggregate the raw values of the eye tracking measurements from each IA. The saliency score for IA_i is given by $\frac{\sum_{k=0}^n x_{ki}}{n}$.

4 Experimental Results

4.1 Comparison with Other Saliency Metrics

We investigate how eye tracking metrics compare with other existing measures for lexical-level significance – namely, human annotations, integrated gradient scores, and large language model surprisal scores (see Figure 5 for a visualization of these scores):

- **Surprisal scores:** For the text in the i^{th} interest area, denoted IA_i , the surprisal is $P(IA_i | IA_0, IA_1, \dots, IA_{i-1})$. We obtain this probability estimate from the pre-trained GPT-2 model (Radford et al., 2019).⁴ In the event that an IA includes multiple tokens, we sum the surprisal of those tokens.
- **Model gradient scores:** The integrated gradient method (Sundararajan et al., 2017) is often used to obtain scores over the input tokens to a deep neural network, where a token’s score reflects how much that token influenced the network’s final output. We obtain these scores with the Captum codebase (Kokhlikyan et al., 2020), using the fine-tuned BERT model

⁴We include word-surprisal scores from GPT-2 as they have previously been found to correlate with human reading times (Wilcox et al., 2020).

from Hayati et al. (2021). For IA_i , the integrated gradient score is the average of the individual tokens within IA_i .

- **Human annotations:** Human annotations come from the Hummingbird dataset (Hayati et al., 2021). Three annotators per item were asked to highlight words that contribute to the text’s style. We averaged these binary highlighting scores over each annotator to arrive at a saliency score for each interest area.

Throughout the comparison, we answer the following two questions: How much do the salient IAs derived from each measure overlap and how much does each measure agree on the saliency strength of each IA?

To find the overlap between salient interest areas derived from different measures, we compute a binary saliency map over the dataset for each measure. We then compute the pairwise Jaccard similarity coefficient for each possible pairing of salient text sets (Fig 6), where the Jaccard similarity coefficient is their intersection over union. We use the median saliency score as the threshold that determines whether the IA is labeled “salient” so that each measure results in the same number of salient words, allowing a more straightforward comparison between measures.

We find that the intersection over union of salient interest areas from eye tracking methods and both integrated gradient scores and human annotations falls between 0.26 and 0.31. Critically, the three-way intersection over union between salient text from integrated gradients, human annotations, and eye tracking metrics falls between 0.05 and 0.06, indicating that each metric captures a relatively unique set of text within the dataset (see Fig 7).

We also investigate what types of words are selected as salient by each method by performing part-of-speech (POS) tagging on the salient interest areas for each measure, finding that while distributions of parts of speech are similar, humans select proportionally more adjectives while eye tracking metrics select proportionally more verbs and adverbs (Figure 9). This discrepancy may be explained by human annotators focusing more on single words with high stand-alone style (oftentimes these are adjectives such as *happy*, *gracious*), while people’s eyes attend to the context surrounding that word (oftentimes this context includes verbs and adverbs). For example, in the polite phrase “Thank you for removing...,” human annotators highlight

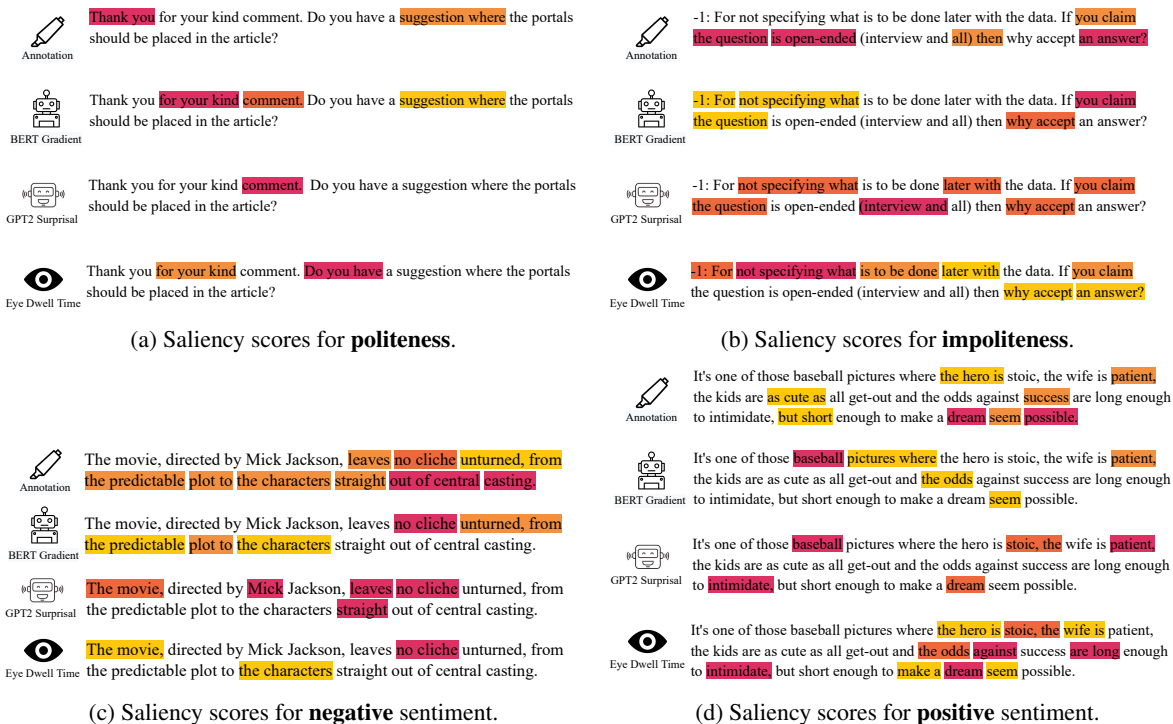


Figure 5: A comparison of saliency scores from various methods: manual human annotations, language model introspection, and eye tracking. Darker highlights indicate stronger saliency scores.

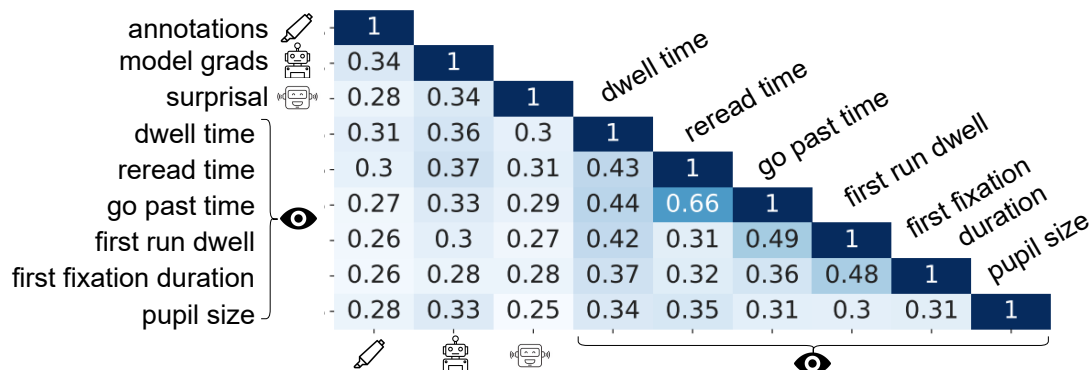


Figure 6: Confusion matrix of the Jaccard similarity score for salient text derived from each metric. (See Appendix for the correlation coefficient for saliency scores derived from each metric.)

only “thank you” whereas eye gaze also focuses on the gerund verb “removing.”

To measure agreement between different measures with respect to saliency strength, we compute a saliency score for each IA in the dataset derived from each measure. We then compute the pairwise Pearson’s r correlation coefficient, finding most coefficients are near 0 (see Appendix). In other words, while there is some agreement across human-, machine-, and eye-based methods with respect to which IAs are above median saliency, there is little correlation with respect to the saliency scores themselves.

4.2 Qualitative Results

For a qualitative visualization of saliency over the politeness style, see Figure 8. In general, human annotations have a tendency to focus on segments of text with clear style markers. For instance, phrases such as “please” are consistently highlighted by human annotators. Our eye tracking data indicates that these phrases do not reliably draw the reader’s gaze during the realtime reading process. We notice that the eyes often focus on the object of the politeness marker rather than the politeness marker itself: For instance, the polite text “Thank you for your kind comment,” human annotators highlight only “thank you” whereas gaze data focuses on “your kind comment.”

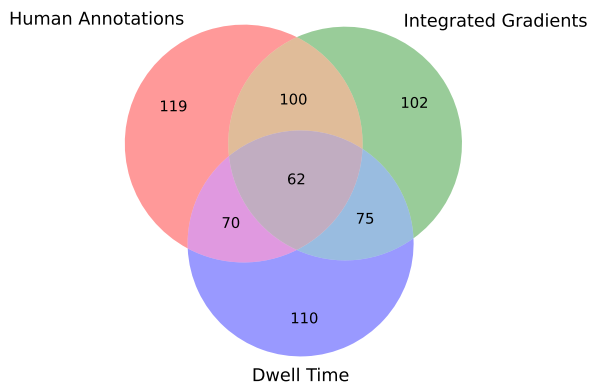


Figure 7: Venn diagram illustrating the intersection of sets of salient interest areas derived from Dwell Time (blue), integrated gradients (green), and human annotations (red).

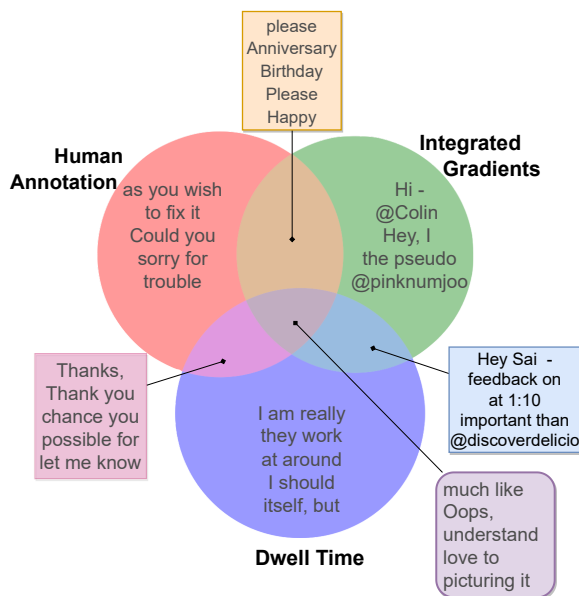


Figure 8: Venn diagram showing interest areas salient to the **polite** style. For each section of the Venn diagram, the interest areas with the top five highest saliency scores are shown.

We also observe that eye data, and in particular dwell time, shows high attention to certain nouns – i.e., names, usernames, and movie titles. This cannot be explained by word frequency effects, as participants in the control condition did not spend as long attending to these nouns.

4.3 “Eye-in-the-loop” few-shot learning

We utilize “eye-in-the-loop” few-shot learning in order to roughly probe the cognitive plausibility of GPT-3 (Brown et al., 2020). Our prompts present a classification task and include zero to four examples from our dataset, including an “important words” section that contains the salient text as defined by each eye-tracking measure, human annotations, and integrated gradient scores (see Section 3.5 for details). As a baseline, we omit the “important words.” We expect that if GPT-3 has a particularly strong cognitive understanding of style

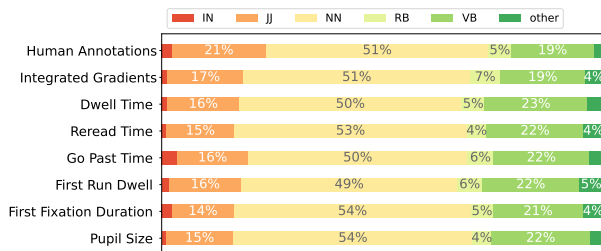


Figure 9: Top 5 most common parts of speech for each measure’s salient IA set. IN: prepositions and subordinating conjunctions, JJ: adjectives, NN: nouns, RB: adverbs, VB: verbs.

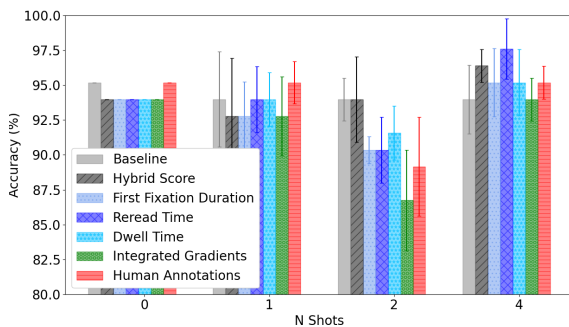


Figure 10: Few-shot learning classification experiment accuracy scores, averaged over 5 rounds with randomly selected demonstrations. Error bars indicate 95% confidence interval.

processing, “important words” from eye movement data may improve its task performance (in these experiments, we use the text-davinci-002 model). Results are relatively inconsistent across each of the four shots, but in most cases, it seems that including salient words has little effect on the model accuracy on the style classification task. A subset of the results are shown in Figure 10; see Appendix for full results and prompt details.

5 Key Findings and Discussion

Here we discuss the relationship between our results and our research questions:

RQ1: Does eye tracking data for saliency meaningfully differ from simply gathering word-level human annotations, or from model-based word importance measures? Our data show a substantial difference between eye-tracking-based saliency, model-based saliency, and human annotations. It is perhaps unintuitive that reading behavior would differ from self-reports after reading, but this is consistent with findings in psycholinguistics that establish strong distinctions between explicit measures (i.e., human annotations) and implicit measures (i.e., eye tracking) of human language processing. Interestingly, there is some intersection between eye tracking-based saliency and model-

based saliency that is **not** shared with human annotators. This suggests that some automatic aspects of human language processing, accessible through eye tracking but not necessarily survey methods, may be shared with large language models.

RQ2: How can we measure eye movements specific to a high-level textual feature like style, and which eye tracking metrics and data processing methods are best suited to capturing textual saliency? The results from our experiment indicate that our experimental paradigm exploiting congruency effects may be effective in finding eye movements specific to certain text features. In a linear mixed effect model analyzing the data, we find significant effects of the congruency condition on dwell time and pupil size (see Appendix A.2). This suggests that the congruency effect does impact reading patterns – whether this impact is directly linked to the textual style is difficult to definitively answer, but given the overlap between eye-tracking-based style saliency and other style saliency measures, it seems reasonable to believe that the experimental manipulation resulted in an implicit measure of style perception. Experiments based on congruency effects may be a promising route for capturing eye movements related to other high-level textual features such as sarcasm and metaphor. We find that dwell time appears to be the strongest eye-tracking metrics for capturing textual saliency, as it has both the highest overlap with human- and machine-based saliency and most strongly responded to the experimental manipulation. Using the same criteria, we also find that using participant-level z-scores to represent the eye movement data yields the best results.

6 Limitations

In this exploratory study, our dataset and sample size are both small, limiting the possibilities for a more thorough evaluation of the data e.g. by fine-tuning a language model. We also note that by design, our experiment presents incongruent items rarely, and consequently we have considerably more congruent datapoints than incongruent datapoints – an inherent limitation of the proposed experimental paradigm. In light of our results, which suggest that eye-tracking data contains useful and unique information, we plan to develop methods for collecting this kind of real-time human reading data at scale – i.e., without the constraints of costly in-person eye tracking – in future work.

Finally, eye tracking analysis in general is limited by the eye-mind assumption, which holds that the eye fixates on what the mind is currently processing. While there is strong evidence supporting the eye-mind assumption during reading, there is a notable exception: retrieval processes (i.e. accessing memory) are not reflected in eye movements (Anderson et al., 2004).

Acknowledgements

We would like to thank Jeffrey Bye, Andrew Elfenbein, Charles Fletcher, Shirley Hayati, Brooke Lea, Andreas Schramm, and Mariya Toneva for their valuable feedback and insightful suggestions regarding the experimental procedure and data analysis. We would also like to thank Miguel Miguelez Diaz, Risako Owan, Faziela Khan, and Josh Spitzer-Resnick for testing and critiquing the initial experimental pipeline. This research received funding from the Sony Research Innovation Award.

References

- John R Anderson, Dan Bothell, and Scott Douglass. 2004. Eye movements do not reflect retrieval processes: Limits of the eye-mind hypothesis. *Psychological Science*, 15(4):225–231.
- Maria Barrett, Joachim Bingel, Nora Hollenstein, Marek Rei, and Anders Sjøgaard. 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.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. O’Reilly Media, Inc.
- Valeria Bolotova, Vladislav Blinov, Yukun Zheng, W Bruce Croft, Falk Scholer, and Mark Sanderson. 2020. Do people and neural nets pay attention to the same words: Studying eye-tracking data for non-factoid QA evaluation. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 85–94.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, and Amanda et al. Askell. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901.
- Henry Conklin, Bailin Wang, Kenny Smith, and Ivan Titov. 2021. [Meta-learning to compositionally generalize](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*

- and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 3322–3335, Online. Association for Computational Linguistics.
- Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. A computational approach to politeness with application to social factors. In *51st Annual Meeting of the Association for Computational Linguistics*, pages 250–259. ACL.
- Joseph A DeVito. 1970. *The Psychology of Speech and Language: An Introduction to Psycholinguistics*. Random House, New York.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Tobias Egner. 2007. Congruency sequence effects and cognitive control. *Cognitive, Affective, & Behavioral Neuroscience*, 7(4):380–390.
- Matthew J Green. 2014. An eye-tracking evaluation of some parser complexity metrics. In *Proceedings of the 3rd Workshop on Predicting and Improving Text Readability for Target Reader Populations (PITR)*, pages 38–46.
- Shirley Hayati, Dongyeop Kang, and Lyle Ungar. 2021. Does BERT learn as humans perceive? Understanding linguistic styles through lexica. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6323–6331.
- Shirley Anugrah Hayati, Kyumin Park, Dheeraj Rajagopal, Lyle Ungar, and Dongyeop Kang. 2023. Stylex: Explaining style using human lexical annotations. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2835–2848.
- Nora Hollenstein, Maria Barrett, Marius Troendle, Francesco Bigioli, Nicolas Langer, and Ce Zhang. 2019. Advancing NLP with cognitive language processing signals. *arXiv preprint arXiv:1904.02682*.
- Eduard Hovy. 1987. Generating natural language under pragmatic constraints. *Journal of Pragmatics*, 11(6):689–719.
- Di Jin, Zhijing Jin, Zhiting Hu, Olga Vechtomova, and Rada Mihalcea. 2022. [Deep learning for text style transfer: A survey](#). *Computational Linguistics*, 48(1):155–205.
- Marcel A Just and Patricia A Carpenter. 1980. A theory of reading: From eye fixations to comprehension. *Psychological Review*, 87(4):329.
- Elsi Kaiser. 2013. Experimental paradigms in psycholinguistics. *Research methods in linguistics*, pages 135–168.
- Sigrid Klerke, Héctor Martínez Alonso, and Anders Sjøgaard. 2015. Looking hard: Eye tracking for detecting grammaticality of automatically compressed sentences. In *Proceedings of the 20th Nordic Conference of Computational Linguistics (NODALIDA 2015)*, pages 97–105.
- Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds, Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, et al. 2020. Captum: A unified and generic model interpretability library for pytorch. *arXiv preprint arXiv:2009.07896*.
- Tatsuki Kuribayashi, Yohei Oseki, Takumi Ito, Ryo Yoshida, Masayuki Asahara, and Kentaro Inui. 2021. [Lower perplexity is not always human-like](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5203–5217, Online. Association for Computational Linguistics.
- Nikolay Malkin, Zhen Wang, and Nebojsa Jojic. 2022. Boosting coherence of language models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*. (To appear).
- 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.
- Sandeep Mathias, Diptesh Kanojia, Abhijit Mishra, and Pushpak Bhattacharya. 2020. [A survey on using gaze behaviour for natural language processing](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 4907–4913. International Joint Conferences on Artificial Intelligence Organization. Survey track.
- Abhijit Mishra, Pushpak Bhattacharyya, and Michael Carl. 2013. Automatically predicting sentence translation difficulty. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 346–351.
- Abhijit Mishra, Diptesh Kanojia, Seema Nagar, Kuntal Dey, and Pushpak Bhattacharyya. 2016. Harnessing cognitive features for sarcasm detection. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1095–1104.
- Austin Lee Nichols and Jon K Maner. 2008. The good-subject effect: Investigating participant demand characteristics. *The Journal of general psychology*, 135(2):151–166.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Keith Rayner. 1978. Eye movements in reading and information processing. *Psychological Bulletin*, 85(3):618.

- Mark Rubin. 2016. The perceived awareness of the research hypothesis scale: Assessing the influence of demand characteristics. *Figshare*, 10:m9.
- Katharina A Schwarz, Roland Pfister, and Christian Büchel. 2016. Rethinking explicit expectations: connecting placebos, social cognition, and contextual perception. *Trends in cognitive sciences*, 20(6):469–480.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y 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.
- Anders Søgaard. 2016. Evaluating word embeddings with fmri and eye-tracking. In *Proceedings of the 1st workshop on evaluating vector-space representations for NLP*, pages 116–121.
- 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.
- 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*, volume 33, pages 6327–6341.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International conference on machine learning*, pages 3319–3328. PMLR.
- Takenobu Tokunaga, Hitoshi Nishikawa, and Tomoya Iwakura. 2017. An eye-tracking study of named entity annotation. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 758–764.
- Matthew J Traxler and Lyn Frazier. 2008. The role of pragmatic principles in resolving attachment ambiguities: Evidence from eye movements. *Memory & cognition*, 36(2):314–328.
- Ethan Gotlieb Wilcox, Jon Gauthier, Jennifer Hu, Peng Qian, and Roger Levy. 2020. On the predictive power of neural language models for human real-time comprehension behavior. In *Proceedings of the 42nd Annual Meeting of the Cognitive Science Society*.
- Alain F Zuur, Elena N Ieno, and Chris S Elphick. 2010. A protocol for data exploration to avoid common statistical problems. *Methods in ecology and evolution*, 1(1):3–14.

A Appendix

A.1 Experimental Materials

The following materials were presented to participants during the experiment. Informed consent was obtained from each participant before the experiment began. Instructions were displayed as shown in Figure 11.

The practice items, which participants completed after reading the instructions and before beginning the experiment, were as follows:

```
Text: What does this have to do with programming
? Are you trying to solve this problem
with a program?
Question: None

Text: this is source code... what is the
question? Do you really think that throwing
code at us will solve your problem?!
Question: Do you agree or disagree with the
following statement: The writer of the post
seems upset.
```

See also Figure 11 for screenshots of the display shown to participants at various points in the experiment.

A.2 Mixed Effect Modeling

We fit linear mixed effect models to predict our eye tracking measures, using the R packages lme4 and lmetest. Our fixed effects are the number of characters in the interest area, the HAL frequency of the interest area, whether the previous interest area was viewed, and whether the interest area is in the congruent or incongruent condition. Our random effect is the participant ID. All variables are normalized prior to analysis.

```
model = lmer(EYE_TRACKING_MEASURE ~ 1 +
congruent + previous_viewed+ LENGTH +
HAL_FREQ + (1 | RECORDING_SESSION_LABEL))
```

The Dwell Time and Pupil Size eye tracking measure showed significance for the the fixed congruency effect. The other eye tracking measures – First Run Dwell Time, First Fixation Duration, Reread Time, and Go Past Time – result in a singular fit, likely because they are considerably more sparse (i.e., many interest areas have a null values for these metrics).

	t value	$Pr(> t)$	Sig.	VIF
(Intercept)	-19.114	< 0.001	***	
frequency	-18.238	< 0.001	***	2.53
length	31.858	< 0.001	***	2.53
congruency	2.449	< 0.05	*	1.00
previous IA	26.662	<0.001	*	1.00

Table 3: Fixed Effects: predicting dwell time

	t value	$Pr(> t)$	Sig.	VIF
(Intercept)	-4.098	< 0.001	***	
frequency	1.865	0.06	.	2.28
length	3.056	< 0.01	**	2.27
congruency	-8.382	< 0.001	***	1.00
previous IA	9.915	<0.001	***	1.00

Table 4: Fixed Effects: predicting pupil size

We tested variables for collinearity using the variance inflation factor (VIF) (Zuur et al., 2010) (none exceeded the recommended threshold of 3).

A.3 Additional Saliency Comparisons

A.3.1 Saliency Scores

Figure 12 shows the Pearson’s r value for saliency score over interest areas derived from each method. We also include more example items from the dataset with associated saliency scores in Figure 5b.

A.4 Few-Shot Learning Experiment Details and Results

The full few-shot learning results can be found in Table 5. The experiment was conducted with the OpenAI API⁵ completion endpoint and the following parameters: the text-davinci-002 model, a temperature of 0, and a top_p of 1.

We generated in-context learning prompts over our dataset by including *important words* as follows:

```
Decide whether the following text is Polite or
Impolite.
Text: Thank you for your kind comment. Do you have a
suggestion where the portals should be placed?
Important words: thank you, suggestion
Polite or Impolite:
```

⁵<https://openai.com/api>, accessed in accordance with OpenAI’s terms of use

We are studying how people convey their emotions and attitudes in text. In this experiment, you will read short texts extracted from Twitter, online forums, and movie reviews.

The texts have been organized into eight different sections. At the beginning of each section, you will see a description of (a) what kind of author wrote the texts in the section, and (b) what kind of emotions or attitudes the authors are trying to convey.

Sometimes you will be asked a question about the text after you finish reading.

Press any key to continue.

(a) Experiment instructions screen.

Important

The following tweets were written by people who were feeling bad in some way. The writers were feeling dissatisfied, discouraged, or otherwise negative; they tried to express that feeling through their tweets.

Press any key to continue.

(b) One of the “context” screens shown at the beginning of each block. This information makes participants aware of what type of text to expect in the following screens.

@Ginoceno Great prediction bro! Nice optimism #FAKEFAN

(c) One of the screens displaying an item from the dataset.

Do you agree or disagree with the following statement?

The writer of the post seems upset.

Press the right key to agree.
Press the left key to disagree.

(d) One of the comprehension question screens.

Figure 11: Screenshots from the experiment program.

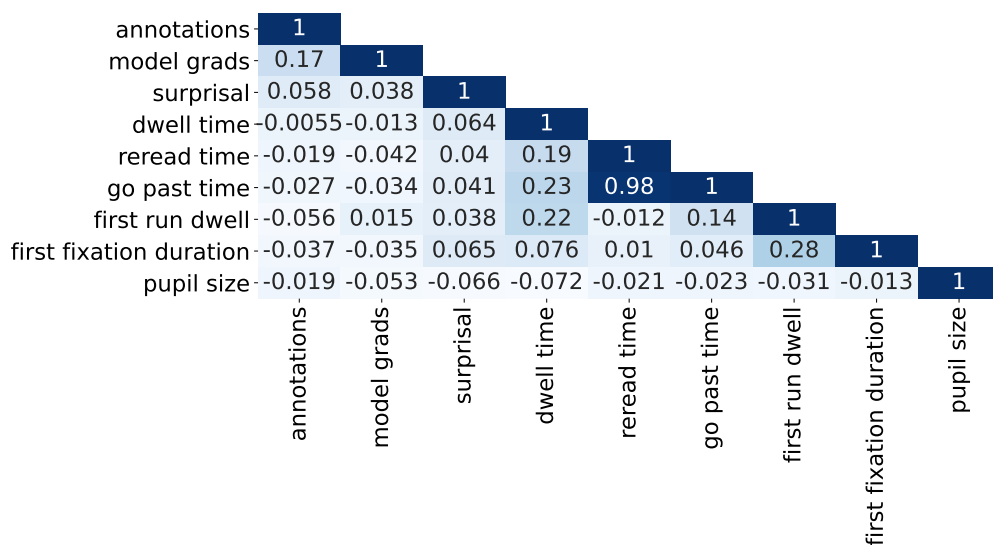


Figure 12: Correlations (Pearson’s r) between the saliency scores derived from each method.

Metric for Saliency	Data aggregation (eye-tracking only)	Experimental Conditions	0-shot	1-shot	2-shot	4-shot
Baseline	NA	NA	95.18	93.98 (2.46)	90.36 (0.96)	95.18 (0.96)
Human Annotations	NA	NA	93.98	91.57 (2.89)	90.36 (3.27)	93.98 (1.80)
Integrated Gradients	NA	NA	93.98	93.98 (1.93)	92.77 (2.46)	96.39 (0.96)
GPT2 Surprisal	NA	NA	93.98	92.77 (0.96)	92.77 (0.96)	97.59 (1.18)
Dwell Time	z score	All	93.98	92.77 (1.80)	93.98 (0.96)	96.39 (2.89)
Dwell Time	z score	Incongruent - Congruent	93.98	93.98 (1.93)	91.57 (1.93)	95.18 (2.36)
Dwell Time	LME	All	93.98	92.77 (0.96)	92.77 (0.96)	97.59 (1.18)
Dwell Time	LME	Incongruent - Congruent	93.98	92.77 (1.18)	91.57 (1.93)	95.18 (2.36)
Dwell Time	raw	All	93.98	93.98 (1.18)	95.18 (1.93)	96.39 (1.80)
Dwell Time	raw	Incongruent - Congruent	93.98	92.77 (1.80)	89.16 (2.89)	95.18 (2.46)
Reread Time	z score	All	93.98	92.77 (0.96)	92.77 (0.96)	97.59 (1.18)
Reread Time	z score	Incongruent - Congruent	93.98	93.98 (2.36)	90.36 (2.36)	97.59 (2.16)
Reread Time	raw	All	93.98	92.77 (1.18)	91.57 (2.46)	93.98 (2.81)
Reread Time	raw	Incongruent - Congruent	92.77	92.77 (2.46)	86.75 (2.81)	96.39 (1.80)
Go Past Time	z score	All	93.98	92.77 (0.96)	92.77 (0.96)	97.59 (1.18)
Go Past Time	z score	Incongruent - Congruent	93.98	91.57 (2.89)	87.95 (3.86)	92.77 (2.46)
Go Past Time	raw	All	92.77	92.77 (0.96)	91.57 (2.46)	96.39 (1.18)
Go Past Time	raw	Incongruent - Congruent	93.98	92.77 (3.27)	90.36 (3.20)	93.98 (2.46)
First Run Dwell Time	z score	All	93.98	92.77 (0.96)	92.77 (0.96)	97.59 (1.18)
First Run Dwell Time	z score	Incongruent - Congruent	93.98	92.77 (2.46)	92.77 (1.18)	92.77 (2.46)
First Run Dwell Time	raw	All	93.98	92.77 (1.18)	92.77 (2.46)	96.39 (2.36)
First Run Dwell Time	raw	Incongruent - Congruent	93.98	92.77 (1.80)	89.16 (3.54)	93.98 (2.46)
First Run Dwell Time	LME	All	93.98	92.77 (1.93)	92.77 (2.36)	95.18 (2.46)
First Run Dwell Time	LME	Incongruent - Congruent	93.98	92.77 (2.46)	89.16 (3.54)	93.98 (2.46)
First Fixation Duration	z score	All	93.98	92.77 (0.96)	92.77 (0.96)	97.59 (1.18)
First Fixation Duration	z score	Incongruent - Congruent	93.98	92.77 (2.46)	90.36 (0.96)	95.18 (2.46)
First Fixation Duration	raw	All	93.98	93.98 (2.64)	89.16 (1.80)	96.39 (0.96)
First Fixation Duration	raw	Incongruent - Congruent	93.98	92.77 (2.81)	90.36 (3.27)	95.18 (1.93)
Pupil Size	z score	All	93.98	92.77 (0.96)	92.77 (0.96)	97.59 (1.18)
Pupil Size	z score	Incongruent - Congruent	92.77	93.98 (2.46)	86.75 (4.31)	93.98 (1.80)
Pupil Size	raw	All	93.98	92.77 (1.18)	92.77 (1.80)	95.18 (2.81)
Pupil Size	raw	Incongruent - Congruent	93.98	91.57 (1.93)	86.75 (4.03)	96.39 (2.46)
Pupil Size	LME	All	93.98	91.57 (2.46)	95.18 (2.36)	92.77 (1.18)
Pupil Size	LME	Incongruent - Congruent	93.98	92.77 (2.16)	86.75 (4.03)	96.39 (2.46)
Hybrid (Human + Dwell Time)	z score	All	95.18	93.98 (1.18)	93.98 (2.46)	96.39 (1.52)
Hybrid (Human + Dwell Time)	z score	Incongruent - Congruent	93.98	92.77 (4.15)	93.98 (3.05)	96.39 (1.18)

Table 5: Accuracy results on few-shot learning experiments over dataset. For 1-, 2-, and 4-shot learning, five different randomly selected prompts were chosen and the average accuracy is reported (the 95% confidence interval is reported in parentheses after the accuracy score).