







Does BERT Learn as Humans Perceive? Understanding Linguistic Styles through Lexica

Shirley Anugrah Hayati  Dongyeop Kang  Lyle Ungar 
 University of Pennsylvania  Georgia Institute of Technology  University of Minnesota
shirley@gatech.edu dongyeop@umn.edu ungar@cis.upenn.edu

Abstract

People convey their intention and attitude through linguistic styles of the text that they write. In this study, we investigate lexicon usages across styles throughout two lenses: *human perception* and *machine word importance*, since words differ in the strength of the stylistic cues that they provide. To collect labels of human perception, we curate a new dataset, HUMMINGBIRD, on top of benchmarking style datasets. We have crowd workers highlight the *representative words* in the text that makes them think the text has the following styles: politeness, sentiment, offensiveness, and five emotion types. We then compare these human word labels with word importance derived from a popular fine-tuned style classifier like BERT. Our results show that the BERT often finds content words not relevant to the target style as important words used in style prediction, but humans do not perceive the same way even though for some styles (e.g., positive sentiment and joy) human- and machine-identified words share significant overlap for some styles.¹

1 Introduction

To express their interpersonal goal and attitude, people often use different styles in their communication. The style of a text can be as important as its literal meaning for effective communication (Hovy, 1987). NLP researchers have built many models to identify different styles in text, including politeness (Danescu-Niculescu-Mizil et al., 2013), emotion (Alm et al., 2005; Mohammad et al., 2018), and sentiment (Socher et al., 2013). Recently, transformer-based (Vaswani et al., 2017) pretrained language models, such as BERT (Devlin et al., 2019), have achieved impressive performance on many NLP tasks, including stylistic studies. However, explaining what these deep learning models learn remains

* research conducted at the University of Pennsylvania

¹Our dataset and code are available at <https://github.com/sweetpeach/hummingbird>



Figure 1: Both humans and BERT models label the sentence (a) as “polite”, whereas in sentence (b), the humans label it as “anger” but BERT does not. Pink highlight: high human perception score. Blue: BERT’s important words. Purple: the word is seen as a strong cue by both human and BERT. The darker color means that the score for human perception or machine word importance is higher. Best seen in color.

a challenge. Thus, there is a growing effort to understand how these models behave (Rogers et al., 2021; Rajagopal et al., 2021).

In this work, we attempt to understand style variation through the contrasting words identified by humans and BERT as determining a style. Given the subjective nature of styles, we are interested in capturing human’s inherent perception of stylistic cues in the text and compare this with the BERT’s “perception”. Specifically, we investigate the extent to which BERT’s word importance, as estimated using Shapley value-based attribution scores (Mudrakarta et al., 2018), aligns with human perception in stylistic text classification.

When humans identify styles in a text, specific words play an important role in recognizing the style, such as hedges for identifying politeness (Danescu-Niculescu-Mizil et al., 2013). We call such words *stylistic cues*. For example, in Figure 1(a), humans perceive the words “understand,” “like,” and “accept” as strong stylistic cues for politeness. But does the BERT model learn the same

words as indicative? It turns out that although the model learns that the word “*accept*” is an important feature for classifying the text as polite, it disagrees with humans for “*understand*” and “*like*” by identifying these words as signals for impoliteness. This leads to a concern that lexical explanation from BERT could be unreliable and motivates us to investigate more deeply into the lexical cues used by humans and BERT. Since styles overlap significantly (Kang and Hovy, 2021), we cover multiple styles: politeness, sentiment, offensiveness, anger, disgust, fear, joy, and sadness.

Our contributions are as follows:

- This is the first comparative study to examine stylistic lexical cues from human perception and BERT. To characterize their discrepancy, we developed a dataset, called HUMMINGBIRD, where crowd-workers relabeled benchmarking datasets for style classification tasks.
- We found that human and BERT cues are quite different; BERT pays more attention to *content words*, and word-level human labels provide more accurate multi-style correlations than sentence-level machine predictions.
- Our work differs from previous works which have generated stylistic lexica from manually-curated seed words or thesauri (Davidson et al., 2017; Mohammad and Turney, 2010); Instead, in our work, the full text is given to annotators, providing more context to the selection of the cue words.

2 Collection of Human and BERT’s Importance Scores on Stylistic Words

While there are many datasets with stylistic labels, to the best of our knowledge, there is no available dataset of stylistic texts with *human labels on the individual words* that drive the human perception. Therefore, on top of existing benchmark style datasets, we we develop HUMMINGBIRD, a new dataset with human-identified stylistic words in those stylistic sentences .

Dataset We use the following datasets for our style classification tasks and as a starting point for collecting human perception scores on lexical level: StanfordPoliteness (Danescu-Niculescu-Mizil et al., 2013), Sentiment Treebank (Socher et al., 2013), Davidson et al. (2017)’s tweet dataset for offensiveness, and SemEval 2018’s dataset for emotion classification (Mohammad et al., 2018).

Style	Label Distribution	Interannotator Agreement	F1
Politeness	22.8%(+) / 41.2%(-)	62.8	69.4
Sentiment	24.6%(+) / 53.6%(-)	71.1	96.5
Offensiveness	33.6%	75.7	98.0
Anger	35.0%	73.5	82.0
Disgust	41.6%	71.2	80.7
Fear	16.4 %	76.1	84.6
Joy	22.6%	82.7	86.5
Sadness	26.4%	72.4	78.2

Table 1: Dataset statistics: % of stylistic texts labeled by annotators in 500 texts. (+) refers to polite or positive, (-) refers to impolite or negative. Inter-annotator agreement: percent agreement scores for two or more people. F1 score is the performance of BERT models on the existing test sets.

Human Perception Scores To collect human perception scores, we first pick 500 stylistically-diverse texts from the four style datasets by the following method. First, we fine-tune BERT on the training sets of the exiting datasets using the original train/dev/test splits. The models’ performance is shown in Table 1. We then run each model on every development set. For example, we run a sentiment classifier on our emotion dataset. From this, we obtain the probability score from the model for predicting each style.

To encourage that the chosen texts exhibit diverse styles, we sort them based on their probability scores and compute the standard deviation of these scores across the eight styles, following Kang and Hovy (2021). We then select the 50 most polite texts, 50 most impolite texts, 50 positive texts, 50 negative texts, 100 offensive texts, and 200 emotional texts (40 from each emotion style), resulting in total 500 texts from the four different style datasets.

We hired 622 workers to annotate them with human perception on Prolific² from November to December 2020. We required the workers to be in the United States and payed them an average of \$9.6/hour. Each worker was asked what styles they perceive each of the texts to exhibit. If they think the text has certain styles, workers then highlight the words in the text which they believe make them think the text had those styles (Pink highlights in Figure 1).

Three workers label the same pair of sentence and style, and we take majority voting for the style labels.³ Crowd-workers obtained an average per-

²<https://www.prolific.co/>

³See Appendix for original dataset details.

Politeness			Positive Sentiment			Joy		
👤↑🤖↑	👤↑	🤖↑	👤↑🤖↑	👤↑	🤖↑	👤↑🤖↑	👤↑	🤖↑
lovely	hilarious	disappointed	delightful	deep	shocking	excited	moved	movies*
delightful	thank	scenes*	lovely	thanks	scare	love	share	managing
loving	moved	suffers	smart	fun	move	entertaining	performances	referring
smart	good	hi [#]	solid	deftly	absolutely	great	congrats	documentary
trouble	clear	optimism	excited	best	wow [#]	perfect	smile	baseball*

Table 2: Top 5 words where humans and BERT agree or disagree. 👤↑🤖↑: both human and BERT agree. 👤↑: high human perception score but low BERT’s importance score. 🤖↑: high BERT’s importance score but low human perception score. **BERT-only agreement includes more content words (*) or interjections (#) than human-only agreement.**

centage agreement of 73.2% on majority labeling, which is a substantial agreement, for text-level as shown in Table 1 and an average percentage agreement of 27.7% for word-level agreement.

Then, for a word w_i in a text $t = w_1..w_N$, the human perception score is defined as:

$$H(w_i) = \frac{\sum_{j=1}^{\#\text{annotators}} h_j(w_i)}{\#\text{annotators}} \quad (1)$$

where $h_j \in \{-1, 0, 1\}$ is the score given by the j^{th} annotator. Each annotator’s label will contribute a score of either 1 for a word that is perceived as a positive cue, -1 for a negative cue, and otherwise 0 (neutral or no emotion).

BERT’s Importance Scores To obtain the word importance (attribution) scores from BERT, we first trained BERT-based models, yielding with F1 scores in Table 1. We then use the popular technique of layered integrated gradients (Mudrakarta et al., 2018) provided by Captum (Kokhlikyan et al., 2020). This technique is a variant of integrated gradients, an interpretability algorithm that attributes an importance score to each input feature by approximating the integral of the gradients of the model’s output with respect to the inputs along a straight line from given baselines to the inputs (Sundararajan et al., 2017).

Since BERT could tokenize a word w into several word pieces, the importance of a word is an average of the scores of the word pieces x that make up the word. For an input of word pieces x , if we have a function $F : \mathbb{R}^n \rightarrow [0, 1]$ as a neural network, and an input $x = (x_1, \dots, x_n) \in \mathbb{R}^n$, an attribution of the prediction at input x relative to a baseline input x' is a vector $A_F = (x, x') = (a_1, \dots, a_n) \in \mathbb{R}^n$ where a_i is the attribution of x_i to the prediction $F(x)$. We use the default setting of Captum for the baseline input x' which is zero

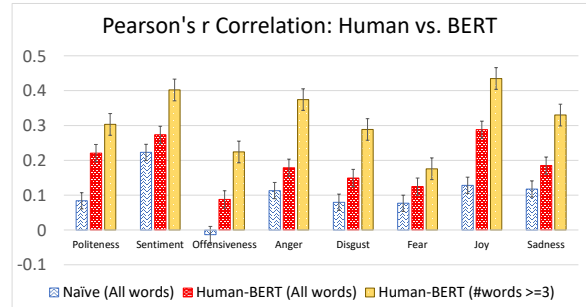


Figure 2: Pearson’s r between human and BERT for the eight styles ($p < 0.001$).

scalar. Finally, we obtain $[-1,1]$ attribution score for each token like the blue highlights in Figure 1.

3 Human-BERT Agreement through Lexical Analysis

We study how similar human perception and BERT’s word importance are, within each style (intra-style) and across styles (multi-styles).

Intra-stylistic Analyses We measure the correlation between human perception of stylistic words and BERT’s word importance, by computing the Pearson’s r for them across all words in the vocabulary, as shown in Figure 2. Naïve refers to our baseline which is that we simply count word frequencies in the stylistic text. For example, if the style is positive sentiment, for a word w , we computed how many times w appears for sentences labeled as “positive”. We calculated the Pearson’s r between this word count and the sentences’ styles across all sentences. This Pearson’s r score is the baseline score of the word importance for word w .

We find that BERT’s word importances correlate more highly with human judgements than this baseline; neither BERT nor humans rely purely on co-occurrence frequencies. Some styles are easier to identify by both human and BERT, such as joy and sentiment with Pearson $r=0.288$ and 0.273 . The yellow bar suggests that human-BERT agree-

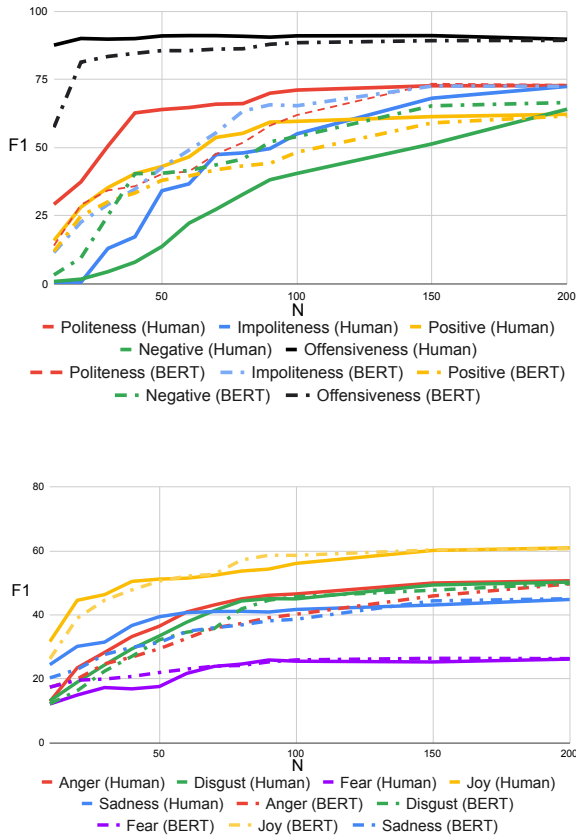


Figure 3: Simple classification using top-N human and BERT features for all the eight styles. Best seen in color.

ment is higher when the word appears more often, especially for offensiveness (0.088 vs. 0.224).

We now look into which words BERT and humans agree and disagree on. Table 2 shows such words selected based on the difference of the word ranks of the human perception score and those from BERT’s word importance. To include only highly stylistic words, words are selected only if their scores are greater than a threshold of 0.3. When humans and BERT agree (👤↑🤖↑), they attend to words that are clearly associated with the styles (e.g joy, positive) and are general (“lovely”, “delightful”, “excited”).

In contrast, BERT often finds words that suggest contexts in which the sentiment is likely to occur. For example, top-5 words from BERT-only agreement (🤖↑) contain more content words such as “scenes” for politeness and “movies” and “baseball” for joy than those from human-only agreement (👤↑). In particular, we see that for politeness and positive sentiment, BERT pays more attention to interjections (e.g., “hi”, “wow”) than humans. For offensiveness and fear in Table 4 in the Ap-

pendix, humans perceive hashtags as important cues but BERT does not. Interestingly, humans perceive a seemingly positive word, “charming,” as offensive while BERT does not, perhaps missing sarcasm. These content words or words irrelevant to the target style are mostly learned due to the biased training dataset, leading to inaccurate prediction by the machine.

Then, we evaluate the impact of important words perceived by human and BERT in the existing test set using a simple occurrence-based classification method. From the ranked word list by their human perception score and BERT’s word importance scores, we label a text as having the target style, if at least one word in the test sentence exists in the top-N word list. For this study, we only select words that appear three times or more in the dataset.

In Figure 3, human’s word list outperforms BERT’s for most styles, even with this small size of annotations compared to the large size of original datasets used for training the BERT model. Interestingly, for some negative styles (e.g., impoliteness, negative sentiment, fear), BERT’s word list performs better. We observe that words from offensive dataset (mostly swear words) are more consistently labeled as impolite and negative by human annotators. However, these words are not often seen in the original politeness and sentiment datasets. It explains why features from BERT models which are trained on the original, large datasets get higher F1 score. As for fear, we found that content words, such as “facebook” and “theatre”, appear in the test data. Here we see that BERT relies on content words (topic-related words) to help predict the style, which is fragile to out-of-domain samples.

Multi-stylistic Analyses As we extend our analyses on multi-style correlation from a lexical viewpoint, we found that humans and machines give similar correlations among the styles. For instance, joy, positive sentiment, and politeness are all positively correlated, as are anger, disgust, and offensiveness (Figure 4). However, the multi-style correlation strength is greater for human perceptions than for machine importance.

The weaker correlation across styles for machines is confirmed in Figure 5, which presents a lower-dimensional visualization for the stylistic representation of each word. Stylistic words are more clustered in human perception, while




Figure 4: Pearson’s r word correlation matrix across styles. The upper triangle (blue and red) represents human perception scores, while the lower triangle (green and brown) represents machine word importances.

for BERT, the separation between highly stylistic words and non-stylistic words is less clear. Figure 5 also shows the geometric closeness across the style clusters, giving extra information beyond the pairwise correlations in Figure 4. In human scores, styles cluster into two extremes: politeness, positive sentiment, and joy to the left, and anger, negative sentiment, offensiveness, and impoliteness to the right, with disgust, fear, and sadness, between them. This leads to more accurate style correlation analysis than machine-based analysis on the text level (Kang and Hovy, 2021).

4 Conclusion

We showed that BERT’s word importances for style prediction, as calculated using integrated gradients, correspond only very loosely with the word importances given by human annotators. These differences likely result from several factors: 1) Word-importances computed for words which appear rarely in the text tend to be noisy. 2) BERT, as a contextual pretrained model, take more context into account for deciding the style of the text while human intuitively choose the most obvious “stylistic” words to judge the style of the text. 3) Styles are subjective matter, so human annotators may have different perception toward the style of a sentence.

Future Directions This work also provides a public dataset as a first step for researchers to further investigate these issues. We plan to scale up our data collection in their size and style types including other higher-level of styles such as sarcasm and humor. We also explore a possibility of informing BERT to pay more attention on human-annotated lexica.

Limitations We acknowledge that while the inter-annotator agreement for the sentence-level


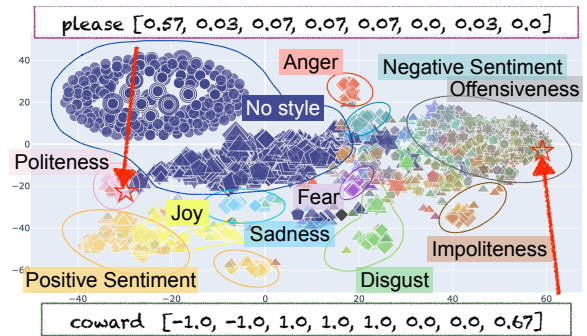


Figure 5: T-SNE (Van der Maaten and Hinton, 2008) for human (top) and machine (bottom). Each word is represented as a vector of its perception score for the styles in this order: politeness, sentiment, offensiveness, anger, disgust, fear, joy, and sadness.

style is quite high, there is a huge variation for the word-level agreement. As a caveat, the annotators could be unreliable. We do find that annotators label different words as being important than those that drive BERT predictions. Note that we do not claim that BERT is “wrong” and humans are “always reliable”; only that they are different. BERT’s important words can help the model predict correctly, but they are perceived as stylistic features as humans do. Studying this difference is our major goal of this paper. We believe that if a word is perceived as “stylistic” by the majority of people, this word can be regarded as an important cue for the model. Learning this variability of human perception on styles could be an interesting future work using HUMMINGBIRD.

5 Ethical Considerations

A full analysis of style, such as politeness or expression of anger, depends upon the context of the utterance: who is saying it to whom in what situation. Such analysis is beyond the scope of this work, which looks only at how the style of the utterance is perceived without context by a small number of crowd workers. Methods such as we have used here should be extended to look at the

more subtle contextual interpretations of style and, eventually, at the ways in which perceived styles may differ from intended styles.

Many people have (correctly) drawn attention to the role that (mis)perceptions of style can foster gender or racial discrimination (Kang and Hovy, 2021). Closer attention to the words which drive style perception is an important first step towards addressing such problems.

Commercial platforms such as Crystal, Grammarly, and Textio offer "style checkers". Such software would benefit from analyses that extend the work presented here, in that they could compare the words that human editors suggest indicate a given style to the words that NLP methods select as most important for recognizing different styles. Such comparisons, particularly when contextualized, should allow construction of better software to help writers control the effect their writing has on the people reading it.

Acknowledgments

We would like to thank Garrick Sherman for helping with the server setup during data collection and the anonymous reviewers for their thoughtful comments.

References

- Cecilia Ovesdotter Alm, Dan Roth, and Richard Sproat. 2005. Emotions from text: machine learning for text-based emotion prediction. In *Proceedings of human language technology conference and conference on empirical methods in natural language processing*, pages 579–586.
- 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 *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 250–259, Sofia, Bulgaria. Association for Computational Linguistics.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 11.
- 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.
- Eduard Hovy. 1987. Generating natural language under pragmatic constraints. *Journal of Pragmatics*, 11(6):689–719.
- Dongyeop Kang and Eduard Hovy. 2021. Style is not a single variable: Case studies for cross-stylistic language understanding. Association for Computational Linguistics.
- 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*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. SemEval-2018 task 1: Affect in tweets. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 1–17, New Orleans, Louisiana. Association for Computational Linguistics.
- Saif Mohammad and Peter Turney. 2010. Emotions evoked by common words and phrases: Using Mechanical Turk to create an emotion lexicon. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 26–34, Los Angeles, CA. Association for Computational Linguistics.
- Pramod Kaushik Mudrakarta, Ankur Taly, Mukund Sundararajan, and Kedar Dhamdhere. 2018. Did the model understand the question? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1896–1906, Melbourne, Australia. Association for Computational Linguistics.
- Dheeraj Rajagopal, Vidhisha Balachandran, Eduard Hovy, and Yulia Tsvetkov. 2021. Selfexplain: A self-explaining architecture for neural text classifiers. *arXiv preprint arXiv:2103.12279*.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2021. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- 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, Seattle, Washington, USA. Association for Computational Linguistics.

- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, pages 3319–3328. PMLR.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(11).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008. NIPS.

A Existing Datasets for Style Classification

We use existing style datasets from StanfordPoliteness (Danescu-Niculescu-Mizil et al., 2013) for politeness, Sentiment TreeBank (Socher et al., 2013) for sentiment, (Davidson et al., 2017)’s dataset for offensiveness, and SemEval 2018 Task 1: Affect in Tweets for emotion classification (Mohammad et al., 2018). We convert non-binary labels or scores to binary labels to standardize the multi-style analysis, resulting in eight styles. Table 3 shows their dataset sizes and train/dev/test splits.

Styles	Train	Dev	Test
Politeness	9,859	530	567
Sentiment	236,077	1,045	2,126
Offensiveness	22,277	1,251	1,255
Anger	6,839	887	3,260
Disgust	6,839	887	3,260
Fear	6,839	887	3,260
Joy	6,839	887	3,260
Sadness	6,839	887	3,260

Table 3: Dataset Statistics

StanfordPoliteness is collected from StackExchange and Wikipedia requests. The labels are continuous values of $[-2, 2]$ so we convert it to binary labels of “polite” and “impolite” by converting all values greater than 0 as polite and the rest are impolite. Sentiment TreeBank dataset consists of movie review texts, and we only use the coarse label of “positive” and “negative” labels for training. Davidson et al. (2017) collected their data from Twitter, and we only consider “offensive” and “none” labels. SemEval 2018 dataset is collected from tweets and it has total 11 emotions for the same 10.9k instances: anger, anticipation, disgust, fear, joy, love, optimism, pessimism, sadness, surprise, and trust. We select anger, disgust, fear, joy, and sadness since these emotions have the highest F1-score compared to the rest. Each emotion has two labels: “anger” or “not anger”, “disgust” or “not disgust”, and so on.

B Training Configuration

We use the lower-cased BERT-base model with 12 hidden layers, 12 attention heads, hidden size 768, for training our style classifiers on GeForce GTX TITAN X GPU. Drop-out rate is 0.1, learning rate is 2×10^{-5} , and the optimizer is AdamW (Loshchilov and Hutter, 2017). Vocabulary size is 30,522 and max position embeddings is 512. Training ran for

Instructions

Thank you for participating in this study! Please read the instructions carefully.

(WARNING: This study may contain offensive words which may be upsetting)

You will be given 20 pairs of a sentence and a style from this list: politeness (polite or impolite), senti (positive or negative), emotions (anger, disgust, fear, happiness, sadness), and offensiveness.

1. Choose the most appropriate style for the given sentence.
2. Choose the word(s) which make you think the sentence has that style.

Example 1:

This cake is **delicious**

What is the sentiment of this text?: **Positive**

Word that makes it **positive**: **delicious**

Example 2:

Their **encouragement** **helps** her finish the task

What is the sentiment of this text?: **Positive**

Words that make it **positive**: **encouragement, help**

Example 3:

The movie is **boring**

What is the sentiment of this text?: **Negative**

Word that makes it **negative**: **boring**

3. Don't reload the page.

4. There will be a practice section and a short survey before you begin. You need to pass this practice section to do this task. **You need to fill in your prolific ID to get paid.**

Figure 6: Instruction page for crowd workers.

Input Text:

This cake is **delicious**

What is the **sentiment** of the input text?

Positive

Neutral

Negative

If you select "Positive" or "Negative", please click the word(s) in the input text that makes you think the sentence is "Positive" or "Negative" until the word color becomes blue.

Figure 7: Annotation page for crowd workers.

Demographic Survey
* - required

1. Age* (in years) <input type="radio"/> < 18 <input type="radio"/> 25 - 34 <input type="radio"/> 55 - 74 <input type="radio"/> 18 - 24 <input type="radio"/> 35 - 44 <input type="radio"/> > 74 <input type="radio"/> 45 - 54	3. Ethnicity* <input type="radio"/> Asian <input type="radio"/> Pacific Islander <input type="radio"/> Caucasian <input type="radio"/> Middle Eastern <input type="radio"/> African <input type="radio"/> Others <input type="radio"/> Hispanic/Latino <input type="radio"/> Native American	4. Last Education Degree* <input type="radio"/> No Degree <input type="radio"/> High School <input type="radio"/> Vocational Training/Associate Degree <input type="radio"/> Bachelor Degree <input type="radio"/> Master's Degree <input type="radio"/> Doctoral Degree <input type="radio"/> Professional Degree (Law, Medical etc..)
--	--	---

Figure 8: Demographic survey for crowd workers.

3 epochs, and each epoch took around 4 minutes.

C Annotation Interface

For each text-style pair (total: 500 texts \times 8 styles = 4,000 pairs), we ask three different annotators to select the style label for text and highlight 463

Politeness			Positive Sentiment			Joy		
👤↑👤↑	👤↑	👤↑	👤↑👤↑	👤↑	👤↑	👤↑👤↑	👤↑	👤↑
lovely	hilarious	disappointed	delightful	deep	shocking	excited	moved	movies
delightful	thank	scenes	lovely	thanks	scare	love	share	managing
loving	moved	suffers	smart	fun	move	entertaining	performances	referring
smart	good	hi	solid	deftly	absolutely	great	congrats	documentary
trouble	clear	optimism	excited	best	wow	perfect	smile	baseball
happy	share	weather	hilarious	high	optimism	loving	example	audience
charming	friend	sounds	great	pretty	news	amazing	morning	scenes
compelling	rest	genre	how's	friday	journey	happy	how's	fan
serious	chance	grief	#comfort	smile	nice	delightful	among	dream
Offensiveness			Anger			Fear		
👤↑👤↑	👤↑	👤↑	👤↑👤↑	👤↑	👤↑	👤↑👤↑	👤↑	👤↑
negro	charming	suffers	awful	ignore	negro	despair	rage	shocking
pussy	haunting	used	insult	works	goat	loss	tired	haunting
filth	murderous	ma	bitches	throw	exam	horrific	#heartbreaking	childish
bitch	wee	those	nasty	dumb	closed	creepy	sucks	insult
shit	#porn#android...	got	#horrific	mental	baby	smashed	mental	journey
smashed	holy	made	bitch	bust	trump	#horrific	spectacle	hate
fat	worst	bad	fuck	broke	sex	feeling	#sad	doors
nasty	#dreadful	men	smashed	rage	gotta	nervous	imma	delayed
ass	cheating	ho	fucking	scare		dreadful	war	movies
dreadful	tf	sex	discussing	murderous		murderous	depression	nightmare
Disgust			Sadness					
👤↑👤↑	👤↑	👤↑	👤↑👤↑	👤↑	👤↑			
insult	shocking	trouble	crying	sucks	smashed			
fuck	#terrified	business	#depressing	#horrific	closed			
terrible	failing	referring	miss	absolutely	often			
sucks	#saddened	trump	#disappointment	jail	ways			
damn	#nervous	correct	sad	flight	decline			
utterly	horror	negro	shocking	#hatred	creepy			
#horrific	mad	ruins	despair	nasty	insult			
hate	cold	church	crash	ignore	clumsy			
pussy	broke	#terrifying	#sad	negro	white			
bitches	bitter	exam	ruins	gives	emotional			

Table 4: Top 10 words where humans and BERT agree and disagree for all the eight styles. We only select words that appear ≥ 2 . 👤↑👤↑: both human and BERT agree. 👤↑: high human perception score but low word importance score. 👤↑: high word importance score but low human perception score.

the words which make them think the text has that style with instructions shown in Figure 6. To guarantee that the workers are serious with this task, we provide a screening practice session which resembles the exact task but with a text that is very obvious to be annotated as in Figure 7. The real task interface is also the same as Figure 7. Figure 8 displays an interface where we also ask the worker’s demographic profile.

D Important Words Perceived by Humans and the Machine

Table 4 shows top twenty words where humans and BERT agree and disagree for all styles.