



HAL
open science

On the Evaluation of the Plausibility and Faithfulness of Sentiment Analysis Explanations

Julia El Zini, Mohamad Mansour, Basel Mousi, Mariette Awad

► **To cite this version:**

Julia El Zini, Mohamad Mansour, Basel Mousi, Mariette Awad. On the Evaluation of the Plausibility and Faithfulness of Sentiment Analysis Explanations. 18th IFIP International Conference on Artificial Intelligence Applications and Innovations (AIAI), Jun 2022, Hersonissos, Greece. pp.338-349, 10.1007/978-3-031-08337-2_28 . hal-04668678

HAL Id: hal-04668678

<https://inria.hal.science/hal-04668678v1>

Submitted on 7 Aug 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



This document is the original author manuscript of a paper submitted to an IFIP conference proceedings or other IFIP publication by Springer Nature. As such, there may be some differences in the official published version of the paper. Such differences, if any, are usually due to reformatting during preparation for publication or minor corrections made by the author(s) during final proofreading of the publication manuscript.

On the Evaluation of the Plausibility and Faithfulness of Sentiment Analysis Explanations

Julia El Zini, Mohamad Mansour, Basel Mousi, and Mariette Awad

Electrical and Computer Engineering Department
American University of Beirut, Beirut, Lebanon
{jwe04,mgm35,bam20}@mail.aub.edu, mariette.awad@aub.edu.lb

Abstract. With the pervasive use of Sentiment Analysis (SA) models in financial and social settings, performance is no longer the sole concern for reliable and accountable deployment. SA models are expected to explain their behavior and highlight textual evidence of their predictions. Recently, Explainable AI (ExAI) is enabling the “third AI wave” by providing explanations for the highly non-linear black-box deep AI models. Nonetheless, current ExAI methods, especially in the NLP field, are conducted on various datasets by employing different metrics to evaluate several aspects. The lack of a common evaluation framework is hindering the progress tracking of such methods and their wider adoption.

In this work, inspired by offline information retrieval, we propose different metrics and techniques to evaluate the explainability of SA models from two angles. First, we evaluate the strength of the extracted “rationales” in *faithfully* explaining the predicted outcome. Second, we measure the agreement between ExAI methods and human judgment on a homegrown dataset¹ to reflect on the rationales *plausibility*. Our conducted experiments comprise four dimensions: (1) the underlying architectures of SA models, (2) the approach followed by the ExAI method, (3) the reasoning difficulty, and (4) the homogeneity of the ground-truth rationales.

We empirically demonstrate that *anchors* explanations are more aligned with the human judgment and can be more confident in extracting supporting rationales. As can be foreseen, the reasoning complexity of sentiment is shown to thwart ExAI methods from extracting supporting evidence. Moreover, a remarkable discrepancy is discerned between the results of different explainability methods on the various architectures suggesting the need for consolidation to observe enhanced performance. Predominantly, transformers are shown to exhibit better explainability than convolutional and recurrent architectures. Our work paves the way towards designing more interpretable NLP models and enabling a common evaluation ground for their relative strengths and robustness.

1 Introduction

Sentiment Analysis (SA) is instrumental to the financial services industry [28, 29] as it develops techniques to interpret customer feedback, monitor product reputations, understand the customers’ needs, and conduct market research. Harnessing the power of Deep Learning (DL) in understanding general contexts, the performance of SA models

¹ Dataset and code available at <https://gitlab.com/awadailab/exai-nlp-eval>

is considerably boosted [12, 11, 23]. However, the non-linearity and the black-box nature of such models hinder the interpretation of the predictions [20, 32]. Besides providing guarantees on reliability, generalization, robustness, and fairness, the interpretability of the SA models can be of service to behavioral marketing and personalized advertisement.

Recently, **Explainable Artificial Intelligence (ExAI)** algorithms are breathing a new flexibility in general AI applications by developing methods to explain model’s prediction [5, 41, 30]. Numerical data frameworks and computer vision applications have witnessed an explosive growth of ExAI nurtured by the ease of expression of features as interpretable components [26, 22, 10]. However, only a few ExAI methods are applied to textual classifiers, embeddings, and language models [7]. In the SA framework, researchers integrated data augmentation techniques to improve the interpretability of SA models [6], studied attention mechanisms in SA through an explainability lens [3] and applied ExAI on aspect-based SA models [37]. To date, ExAI methods on Natural Language Processing (NLP) tasks are not evaluated on standardized benchmarking datasets through common metrics which hinders the progress and adoption of such methods in the NLP field. Evaluating explainability methods is two-fold. First, it helps assess the extent to which a deep model can be made explainable. Second, it provides a common ground to measure the contrast between explanations produced by diverse ExAI approaches.

In this work, we inspect two human aspects of explainability methods: (1) faithfulness to the model being explained and (2) plausibility from a human lens. For this purpose, we select eight state-of-the-art SA models with underlying architectures of recurrent, convolutional, and attention layers. We generate explanations of the predictions of these models on three ExAI methods that can be applied in NLP; mainly LIME [30], *anchors* [31] and SHAP [18]. The generated explanations are then evaluated through two procedures. First, *faithfulness*² is evaluated by examining the degradation in the model’s performance when only extracted rationales are fed to the model. Second, the *plausibility*¹ of extracted rationales is evaluated via comparison to the human judgment of what a sufficient explanation is. This experiment entails a homegrown dataset of manually labeled explanations on SA data aggregated through conjunction and disjunction means. The comparison is achieved on six proposed metrics, inspired by information retrieval, to evaluate the precision and fallout of exAI methods on the SA models. Hence, our evaluation is carried out over four different dimensions: (1) SA model, (2) ExAI method, (3) reasoning complexity, and (4) human judgment homogeneity.

The contributions of this work are: (1) a dataset for SA explainability labeled on different dimensions (2) the first faithfulness and plausibility evaluation inspired from information retrieval (3) a thorough four-dimensional ExAI evaluation on SA models.

Our empirical analysis allows us to draw conclusions on the faithfulness of LIME rationales and the plausibility of the anchors model which is found to be more confident in extracting supporting evidence. Moreover, we highlight the consistency of different attention architectures in deriving relatively more plausible explanations.

Next, we provide a general background on the sentiment analysis and ExAI models used in this work in Section 2. Then, we present our evaluation dataset and framework

² refers to the metric hereafter

in Sections 3 and 4. We report our comparative analysis in Section 5 before concluding with final remarks in Section 6.

2 Background

Little has been done on the evaluation of ExAI in NLP settings. Recently, a framework to evaluate rationalized explanations is introduced in ERASER [9]. ERASER provides benchmarking data for 7 NLP tasks and suggests sufficiency and comprehensiveness as evaluation metrics. While ERASER considers a wider range of NLP tasks; it is narrow in terms of the deep architectures and angles that it considers. In contrast, we study the explainability of SA from four different perspectives. Prior to ERASER, the work of [1] evaluates the explainability of SVMs compared to CNNs to find that the latter models yield more interpretable decisions. Other attempts only consider the attention mechanism and the debate concerning its inherent interpretability [34, 21, 14]. Additionally, the concept of explanation faithfulness has been introduced before [13, 30] with no explicit evaluation of textual classifiers.

Next, we provide the background on sentiment analysis and explainability methods.

2.1 Sentiment Analysis Models

SA models currently exploit deep architectures, word embeddings, transfer learning, and attention mechanisms [35]. In this work, we experiment with eight state-of-the-art SA models with different architectures that vary between convolution, recurrent, and attention networks. The models are chosen with a diversity of word embeddings and some of the models leverage transfer learning techniques during training.

First, *CNN-MC*, a CNN for sentiment classification [15], consists of a simple CNN trained on top of static *word2vec* word embeddings. To experiment with different embedding models, the Universal Sentence Encoder (USE)[4], is used to train additional convolution layers. USE obtains sentence embeddings through a deep averaging network. The recurrent architecture is tested on the Byte-multiplicative LSTMs (*bmLSTM*), used in [27] to generate reviews and discover sentiments.

The attention architecture [39] is studied in five transformer models. We first study the parent transformer, *BERT* [8] which pre-trains deep bidirectional representations by jointly conditioning on the left as well as the right context extracted by each layer on unlabeled data. Then, we study an optimized version of BERT, *RoBERTa* [17], which is trained longer and on prolonged sequences. We further study a different optimization of BERT, *ALBERT* [16], that factorizes the embedding matrix into smaller matrices, separating thus the size of vocabulary from the size of the hidden layers. Moreover, we consider *DistilBERT* [33], a distilled version of BERT trained on very large batches while leveraging the computation of the gradients through dynamic masking. Finally, a non-BERT transformer, *XLNET* [40], introduces the auto-regressive formulation to transformers and considers all permutations of the factorization order of the context while learning of bidirectional contexts.

A comparison of the performance of the studied models on the IMDB movie review dataset is reported under *accuracy* in Table 3.

2.2 Explainability Models

Explainability entails various realizations such as visualization, attention interpretation and alignment [41, 5, 38, 25] as surveyed in [42]. We focus on *rationale* extraction that supports a particular decision by highlighting causal input segments. Hence, we consider three explainability methods: Local Interpretable Model-agnostic Explanations (LIME) [30], high-precision model-agnostic explanations (anchors) [31] and SHapley Additive exPlanations (SHAP) [18]. All these explainability models are black-box algorithms with a profound theoretical ground and a sound practical implementation.

LIME [30] approximates *any* model locally in a model-agnostic fashion. In NLP settings, LIME presents the interpretation in terms of bag-of-words regardless of what the NLP model originally accepts. This is achieved by approximating the classifier with a locally-more-explainable, potentially linear, model which is then trained on perturbed inputs in a local neighborhood.

Later, [31] define “anchors” as if-then rules to generate model agnostic explanations. These rules are a set of predicates, A , that are defined on interpretable representations in such a way that $A(x)$ returns 1 if all its feature predicates are satisfied on x . Consequently, a rule is an anchor if $A(x) = 1$ and A is a sufficient condition for the prediction $f(x)$. Then, random words, having the same Part-of-Speech (POS) TAG as “absent” tokens, are sampled and the “absent” tokens are then replaced by the random words to apply input perturbations and study their effect on the model prediction.

Inspired by game theoretical concepts, [18] developed a unified approach to interpret model predictions. Analogous to the Shapley value computation, SHAP values are proposed as a unified measure that different explainability models, such as LIME [30], Deep LIFT [36] and layer-wise relevance propagation [2], approximate. SHAP is theoretically proven to satisfy *local accuracy*, *missingness* and *consistency* as defined in [18].

3 Evaluation Dataset

We anticipate the support of a group of ten data scientists who contributed to providing ground truth rationales. The labelers are between 21 and 26 years old balanced across gender. We consider the Rotten Tomatoes dataset [24] with binary labels and we complement it with two additional labels: *reasoning difficulty*, and *extracted rationales*³.

First, the *reasoning difficulty* reflects the complexity of reasoning about the sentiment and takes three values: 1, 2, and 3 with 3 being the most difficult. A special tag 4 for *reasoning difficulty* is used when a sentence requires additional context understanding. For example, the following review is labeled as 4: “*not even the hanson brothers can save it*”. Explaining this sentence would require a higher-level knowledge of what the reviewer meant by *hanson brothers* in a different context. It is noteworthy that the *reasoning difficulty* does not reflect the readability nor the semantic complexity. It rather shows the difficulty of extracting evidence of a particular sentiment.

Second, the *extracted rationales* are a list of particular words in the sentence that contributed to the sentiment prediction. In curating the dataset, we alleviate human bias and subjectivity by asking more than one labeler to extract the rationales of each sentence.

³ Instructions to labelers are provided in the supplementary material

The aggregation of these explanations is done using the conjunction and disjunction operations. An additional correction step was applied to the aggregated explanations.

Table 1 shows samples from our dataset and Table 2 represents the dataset description aggregated by reasoning difficulty. Specifically, we have 1973 sentences with 20.9 words per sentence on average. Rationales consist of 3.1 words on average when merging (union) across labelers. When considering rationales that different labelers agree on (intersection), the sentences have 0.9 words on average. 934 sentences yield an empty set of rationale intersections. The latter statistic is important to show the degree of agreement between the labelers. A high number of empty intersections reflects diversity in the human thinking in the group of labelers.

Sentence	Label	Reason- ing Diff.	Rationales (union)	Rationales (intersection)
beautiful to watch and holds a certain charm	1	1	beautiful charm	beautiful charm
it turns out to be smarter and more diabolical than you could have guessed at the beginning	1	1	smarter more diabolical	smarter diabolical
none of this is very original , and it isn't particularly funny.	0	1	none original is n't funny	none is n't funny
the film tunes into a grief that could lead a man across centuries.	1	2	grief	grief
the film is about the relationships rather than about the outcome . and it sees those relationships, including that between the son and his wife, and the wife and the father, and between the two brothers, with incredible subtlety and acumen.	1	3	incredible subtlety acumen	incredible subtlety
hard as this may be to believe, here on earth, a surprisingly similar teen drama, was a better film.	0	3	better	better
not even the hanson brothers can save it	0	4	-	-
more than simply a portrait of early extreme sports , this peek into the 1970s skateboard revolution is a skateboard film as social anthropology...	1	4	-	-

Table 1: Dataset sample. Rationales are individual words separated by a space.

4 Methodology

In what follows, we describe the design of our evaluation of faithfulness and plausibility and the formulation of our information retrieval-based metrics for *plausibility* evaluation.

4.1 Experimental Design

Once the explanations are derived, two evaluation methods are adopted.

Difficulty	1	2	3	4	Total
Number of Sentences	1535	208	148	82	1973
Words per sentence	20.1	24.2	23.7	23.2	20.9
Words per explanation (union)	3.4	2.9	2.6	0.0	3.1
Words per explanation (intersection)	1.07	0.7	0.45	0.0	0.9
Number of Empty intersections	629	121	102	82	934

Table 2: Dataset description

Faithfulness To evaluate the significance of the words in E to the sentiment classification task objectively, we feed E as an input to \mathcal{M} and we compute the prediction accuracy. This experiment does not require a comparison to ground-truth data; it rather studies the validity of the explanations through prediction. In this experiment, we investigate the effect of the ExAI model and its confidence on the explanation accuracy.

Plausibility After getting the explanations E and their corresponding contributions W , a comparison with the ground truth labels L allows us to compute our proposed metrics as in Section 4.2. Figure 1 summarizes the experimental design.

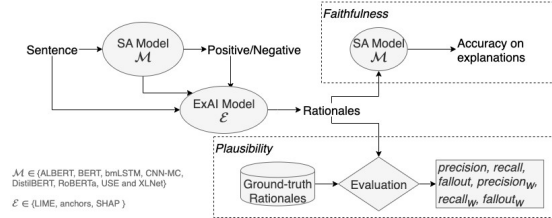


Fig. 1: Experimental Design

4.2 Plausibility Evaluation Metrics

In what follows we assume the following: an explanation model $\mathcal{E} \in \{ \text{SHAP, anchors, LIME} \}$ is outputting a set of words $E = \{e_1, e_2, \dots, e_{N_1}\}$ with their corresponding contributions $W = \{w_1, w_2, \dots, w_{N_1}\}$ when explaining a machine learning model \mathcal{M} , where $\mathcal{M} \in \{ \text{ALBERT, BERT, bmLSTM, CNN-MC, DistilBERT, RoBERTa, USE, XLNet} \}$, on an input sentence $S = \{s_1, s_2, \dots, s_n\}$. To evaluate *plausibility*, we compare E to a sequence of ground truth words provided by human labelers $L = \{l_1, l_2, \dots, l_{N_2}\}$. Consequently, we propose three metrics inspired from information retrieval: *precision*, *recall* and *fallout*. The *precision* computes the fraction of the words retrieved by \mathcal{E} that are relevant from a human perspective. The *recall* computes the fraction of words that are relevant from a human perspective and that are retrieved by \mathcal{E} .

The *fallout* computes the fraction of non-relevant explanations that are retrieved from all the non-relevant words.

$$precision = |L \cap E|/|E| \quad (1)$$

$$recall = |L \cap E|/|L| \quad (2)$$

$$fallout = |(S - L) \cap E|/|S - L| \quad (3)$$

where $|\cdot|$, \cap , and $-$ are the set cardinality, intersection, and set difference respectively.

A plausible rationale extraction is reflected in explanations that match human judgment, hence high precision and recall and low fallout rates. A high precision suggests that it is unlikely that the model will provide the word s_i in the explanation if s_i is not provided in the annotations. A high recall reflects that it is unlikely for the model to miss a word if it is in the annotations. We integrate the contribution scores W to generate the weighted version of our metrics as follows:

$$precision_w = |L \cap E|_W / |E|_W \quad (4)$$

$$recall_w = |L \cap E|_W / |L| \quad (5)$$

$$fallout_w = |(S - L) \cap E|_W / |S - L| \quad (6)$$

where $|\cdot|_W$ is the weighted set cardinality computed as $|s_1, s_2, s_M|_W = w_1 + w_2 + w_M$, with $w_i \in W$ is the weight assigned to the word s_i by the explainability method.

5 Results and Discussion

5.1 Back-end Setup

Before generating the explanations, the SA models were re-trained on the same dataset, IMDB movie reviews [19], according to their official code repositories. The obtained models are then explained by LIME, anchors, and SHAP on our provided dataset. All the experiments are run on Nvidia K80 GPU with 12 GB RAM and 0.82GHz memory clock rate. Following their official repositories, the considered back-end of the SA models varies between *keras*, *torch* and *tensorflow*. Repositories implemented on Python 2.x were translated to Python 3.x and models were retrained accordingly.

5.2 Evaluation of faithfulness

Table 3 shows the initial model validation accuracy and its reduction when only the explanations are fed to the model \mathcal{M} with different confidences ϵ . The parameters ϵ are tailored to the confidences that each explainer \mathcal{E} produces. Anchors models produce explanations with confidence ≈ 1 and thus no thresholds are considered. One can see that the explanations provided by LIME consistently outperform those provided by *anchors* and SHAP on all SA models. This consistency is also maintained with different thresholds per model. Predominantly, *anchors* produces the least *faithful* explanations yielding significant degradation of SA accuracy. Furthermore, the most interpretable models are predominantly transformers as bmLSTM, RoBERTa and ALBERT scored

the highest according to LIME, anchors, and SHAP respectively. The consistency is not maintained on different SA models \mathcal{M} . For instance, bmLSTM is more explainable than BERT when LIME is used to explain but less explainable when SHAP is used.

More strikingly, LIME generally enhances the SA accuracy. This can suggest that the sentiment is highly concentrated in the rationales derived by LIME. SHAP, on the other, hand reduces this accuracy and the anchors model further significantly deteriorates it. This can be explained by two factors. First, the models are all trained on full sentences; discontinuous chunks might mislead the model. Second, in some cases, the ExAI models fail to correctly label the negation in their provided explanations which might lead to misclassification. These results suggest that LIME’s explanations are more absolute, hence useful, and can serve as standalone representations of the original sentences.

Model	Accuracy	LIME			anchors		SHAP		
		$\epsilon = 0.1$	$\epsilon = 0.2$	$\epsilon = 0.3$	$\epsilon = 0.1$	$\epsilon = 0.2$	$\epsilon = 0.3$	$\epsilon = 0.5$	
ALBERT	89.55	93.15	<u>93.20</u>	93.16	68.30	88.54	88.34	88.54	88.54
BERT	89.86	92.60	<u>92.65</u>	92.44	67.28	87.63	87.53	87.20	86.56
bmLSTM	82.61	93.93	93.93	93.83	77.38	76.84	76.28	75.87	75.22
CNN-MC	71.36	88.35	88.89	<u>88.90</u>	55.40	69.08	68.52	68.02	68.12
DistilBERT	88.77	<u>91.05</u>	90.84	90.70	63.91	88.40	88.10	88.35	87.97
RoBERTa	90.26	88.34	<u>88.39</u>	<u>88.39</u>	82.32	88.25	88.24	88.21	88.19
USE	83.05	92.74	92.78	<u>92.87</u>	80.43	75.45	74.65	74.48	74.48
XLNet	90.11	92.61	<u>92.65</u>	92.60	71.50	87.83	86.87	86.87	85.85

Table 3: Accuracy of \mathcal{M} on explanations of \mathcal{E} . Highest accuracies per explanation model are highlighted in bold. Highest accuracies per SA model are underlined.

5.3 Plausibility Evaluation on Ground Truth Data

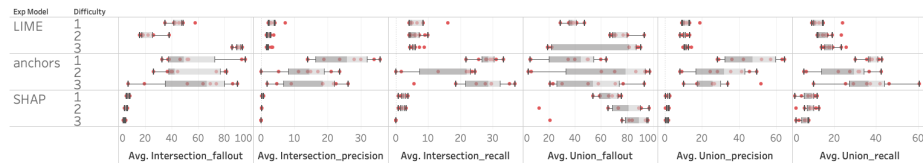


Fig. 2: Explainability scores achieved by ExAIs (for different SA models)

Figure 2 shows the distribution *precision*, *recall* and *fallout* computed on the SA models explained by LIME, anchors and SHAP. Regardless of the reasoning difficulty, the anchors model consistently achieves higher scores on all models \mathcal{M} with the union

and intersection operations. This observation can be explained by the fact that anchors' explanations are in harmony with human rationales, especially since anchors are defined as "if-then" rules. However, the anchors model exhibits higher variability and a consistent improvement is seen with the union operation giving the explanation a higher chance to be matched in the ground truth dataset. Finally, the reasoning difficulty level has an effect on the plausibility of explanations but the effect is not significant when aggregating across different SA and ExAI models.

As in the previous experiment, the explanation consistency is not maintained by \mathcal{E} . For instance, the most explainable model according to anchors is XLNet followed by RoBERTa (as in the previous experiment), where CNN-MC and RoBERTa are the most explainable according to LIME and SHAP respectively. However, similar to the previous findings, the majority of transformer models disclose higher explainability scores compared to convolutional and recurrent ones. These results are consistent with the weighted metrics which further accentuate the outperformance of anchors' explanations over those of LIME and SHAP. Besides exhibiting better explainability, the anchors model is demonstrated to be more confident in generating its explanations.

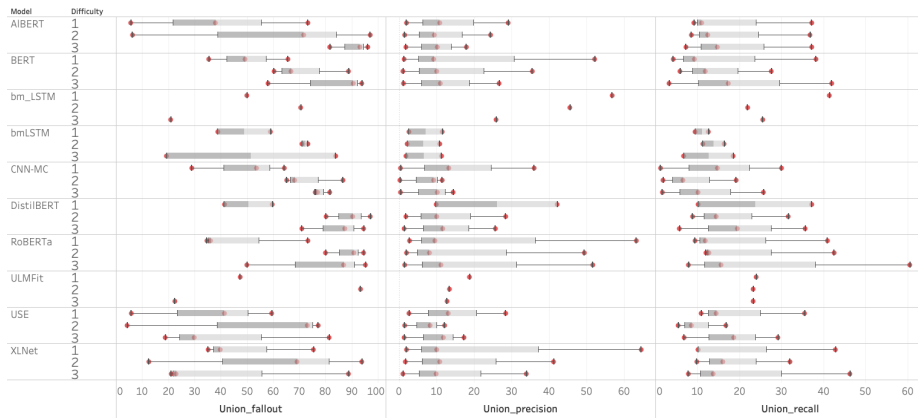


Fig. 3: (Unweighted) explainability scores of SA models when varying the ExAI model

Figure 3 shows the explainability scores achieved by the different SA models while varying \mathcal{E} . Once again, transformers such as RoBERTa, BERT, and AIBERT achieve high agreements with ground-truth rationales. These results point to the crucial role of attention in matching human reasoning. Additionally, reasoning difficulty thwarts explainability, especially in non-transformer architectures such as CNN-MC and USE.

5.4 Qualitative Evaluation

Inherently, instances with reasoning difficulty 4 are missing ground-truth explanations. We conduct a qualitative assessment of their explanations. For the sentence: “*watching queen of the damned is like reading a research paper, with special effects tossed in*”,

LIME produces *tossed, paper, in, like, is, effects* and *damned* as explanations with confidences < 0.3 on RoBERTa. Clearly, the provided explanation is of low confidence and requires context knowledge such as boredom from reading a research paper. For sentences with a difficulty of 4, SHAP also provides an explanation with negligible weights (order of $10^{-3} - 10^{-7}$).

5.5 Take-Away Messages

All dimensions have been considered, our evaluation allows us to conclude with the following speculations. LIME produces rationales with high “sentiment concentration” leading to more faithful explanations and *useful* interpretations from a human perspective. The anchors model extracts supporting evidence that is more aligned with human judgment with higher confidence. On the level of the SA architecture, transformers can be, generally, deemed more faithful in deriving explanations. To a significant extent, attention can emulate human reasoning at the back of sentiment extraction. Nonetheless, no single tailor-made deep architecture is highly interpretable when changing the explainability lenses which motivates the ensembling of ExAI results. Finally, high confidence in the evidence extraction can be linked to reasoning simplicity rather than complexity.

6 Conclusion

In this work, we set the grounds for a rigorous evaluation framework for ExAI methods on SA models. We develop metrics to evaluate the explainability of SA models for their faithfulness and plausibility on our annotated rationales. We study the explainability of the SA methods from different angles: reasoning difficulty level, degree of agreement in the annotated data, ExAI model and its confidence level, and the underlying architecture of the SA model. Our empirical study derives important insights concerning the explainability of NLP models, SA specifically. Mainly the plausibility of anchors’ explanations and faithfulness of LIME is demonstrated in our set of experiments. More importantly, we highlight the potential that attention mechanisms have in matching human reasoning.

This work paves the way for a sound explainability evaluation framework that would engender user trust in NLP models and support behavioral marketing and market analysis systems. An immediate step in this line of work would be the extension to multi-target and multi-aspect sentiment analysis tasks and general NLP models.

7 Acknowledgment

This work was supported by the University Research Board (URB) and the Maroun Semaan Faculty of Engineering and Architecture (MSFEA) at the American University of Beirut.

References

1. Arras, L., Horn, F., Montavon, G., Müller, K., Samek, W.: “what is relevant in a text document?”: An interpretable machine learning approach. CoRR **abs/1612.07843** (2016)

2. Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.R., Samek, W.: On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one* **10**(7), e0130140 (2015)
3. Bodria, F., Panisson, A., Perotti, A., Piaggese, S.: Explainability methods for natural language processing: Applications to sentiment analysis (discussion paper) (2020)
4. Cer, D.e.a.: Universal sentence encoder for English. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing. pp. 169–174. Association for Computational Linguistics, Brussels, Belgium (Nov 2018)
5. Chattopadhyay, A., Sarkar, A., Howlader, P., Balasubramanian, V.N.: Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. 2018 IEEE Winter Conference on Applications of Computer Vision (WACV) (Mar 2018)
6. Chen, H., Ji, Y.: Improving the explainability of neural sentiment classifiers via data augmentation. arXiv preprint arXiv:1909.04225 (2019)
7. Danilevsky, M., Qian, K., Aharonov, R., Katsis, Y., Kawas, B., Sen, P.: A survey of the state of explainable ai for natural language processing (2020)
8. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
9. DeYoung, J., Jain, S., Rajani, N.F., Lehman, E., Xiong, C., Socher, R., Wallace, B.C.: Eraser: A benchmark to evaluate rationalized nlp models. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 4443–4458 (2020)
10. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., Pedreschi, D.: A survey of methods for explaining black box models. *ACM computing surveys* **51**(5), 1–42 (2018)
11. Habimana, O., Li, Y., Li, R., Gu, X., Yu, G.: Sentiment analysis using deep learning approaches: an overview. *Science China Information Sciences* **63**(1), 1–36 (2020)
12. Hassan, A., Mahmood, A.: Deep learning approach for sentiment analysis of short texts. In: 2017 3rd international conference on control, automation and robotics (ICCAR). pp. 705–710. IEEE (2017)
13. Jacovi, A., Goldberg, Y.: Towards faithfully interpretable nlp systems: How should we define and evaluate faithfulness? In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 4198–4205 (2020)
14. Jain, S., Wallace, B.C.: 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). pp. 3543–3556 (2019)
15. Kim, Y.: Convolutional neural networks for sentence classification. arXiv preprint arXiv:1408.5882 (2014)
16. Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., Soricut, R.: Albert: A lite bert for self-supervised learning of language representations (2020)
17. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach (2019)
18. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: Advances in neural information processing systems. pp. 4765–4774 (2017)
19. Maas, A.L., Daly, R.E., Pham, P.T., Huang, D., Ng, A.Y., Potts, C.: Learning word vectors for sentiment analysis. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies. pp. 142–150. Portland, Oregon, USA (June 2011)
20. Mishra, P.: Explainability for Non-Linear Models, pp. 93–127. Apress, Berkeley, CA (2022). https://doi.org/10.1007/978-1-4842-7158-2_4
21. Mullenbach, J., Wiegrefe, S., Duke, J., Sun, J., Eisenstein, J.: Explainable prediction of medical codes from clinical text. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers). pp. 1101–1111 (2018)

22. Nguyen, A., Yosinski, J., Clune, J.: Understanding neural networks via feature visualization: A survey. In: *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, pp. 55–76. Springer (2019)
23. ONAN, A.: Sentiment analysis on massive open online course evaluations: a text mining and deep learning approach. *Computer Applications in Engineering Education* **29**(3), 572–589 (2021)
24. Pang, B., Lee, L.: Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In: *Proceedings of the ACL* (2005)
25. Patro, B.N., Lunayach, M., Patel, S., Namboodiri, V.P.: U-cam: Visual explanation using uncertainty based class activation maps. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 7444–7453 (2019)
26. Qin, Z., Yu, F., Liu, C., Chen, X.: How convolutional neural network see the world - a survey of convolutional neural network visualization methods (2018)
27. Radford, A., Jozefowicz, R., Sutskever, I.: Learning to generate reviews and discovering sentiment. arXiv preprint arXiv:1704.01444 (2017)
28. Rambocas, M., Gama, J., et al.: Marketing research: The role of sentiment analysis. Tech. rep., Universidade do Porto, Faculdade de Economia do Porto (2013)
29. Rambocas, M., Pacheco, B.G.: Online sentiment analysis in marketing research: a review. *Journal of Research in Interactive Marketing* (2018)
30. Ribeiro, M.T., Singh, S., Guestrin, C.: "why should i trust you?": Explaining the predictions of any classifier (2016)
31. Ribeiro, M.T., Singh, S., Guestrin, C.: Anchors: High-precision model-agnostic explanations. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 32 (2018)
32. Samek, W., Müller, K.R.: Towards explainable artificial intelligence. In: *Explainable AI: interpreting, explaining and visualizing deep learning*, pp. 5–22. Springer (2019)
33. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter (2020)
34. Serrano, S., Smith, N.A.: Is attention interpretable? In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. pp. 2931–2951 (2019)
35. Shi, Y., Zhu, L., Li, W., Guo, K., Zheng, Y.: Survey on classic and latest textual sentiment analysis articles and techniques. *International Journal of Information Technology & Decision Making* **18**(04), 1243–1287 (2019)
36. Shrikumar, A., Greenside, P., Kundaje, A.: Learning important features through propagating activation differences (2019)
37. Silveira, T.D.S., Uszkoreit, H., Ai, R.: Using aspect-based analysis for explainable sentiment predictions. In: *CCF International Conference on Natural Language Processing and Chinese Computing*. pp. 617–627. Springer (2019)
38. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. arXiv preprint arXiv:1703.01365 (2017)
39. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: *Advances in neural information processing systems*. pp. 5998–6008 (2017)
40. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R., Le, Q.V.: Xlnet: Generalized autoregressive pretraining for language understanding (2020)
41. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. *CoRR abs/1311.2901* (2013)
42. Zini, J.E., Awad, M.: On the explainability of natural language processing deep models. *ACM Computing Surveys* (2022)