

Concept-Based Explanations to Test for False Causal Relationships Learned by Abusive Language Classifiers

Isar Nejadgholi, Svetlana Kiritchenko, Kathleen C. Fraser, and Esma Balkir

National Research Council Canada

Ottawa, Canada

{Isar.Nejadgholi, Svetlana.Kiritchenko, Kathleen.Fraser, Esma.Balkir}@nrc-cnrc.gc.ca

Abstract

Classifiers tend to learn a false causal relationship between an over-represented concept and a label, which can result in over-reliance on the concept and compromised classification accuracy. It is imperative to have methods in place that can compare different models and identify over-reliances on specific concepts. We consider three well-known abusive language classifiers trained on large English datasets and focus on the concept of *negative emotions*, which is an important signal but should not be learned as a sufficient feature for the label of abuse. Motivated by the definition of *global sufficiency*, we first examine the unwanted dependencies learned by the classifiers by assessing their accuracy on a challenge set across all decision thresholds. Further, recognizing that a challenge set might not always be available, we introduce concept-based explanation metrics to assess the influence of the concept on the labels. These explanations allow us to compare classifiers regarding the degree of false global sufficiency they have learned between a concept and a label.

Content Warning: This paper presents examples that may be offensive or upsetting.

1 Introduction

In various natural language classification tasks, particularly in abusive language detection, certain concepts are known to be strong signals for the label of interest. These concepts are often over-represented in the respective class of the training set, making them susceptible to being learned as potential causes for the label. Consequently, the classifier over-relies on these concepts and ignores the broader context, leading to reduced generalizability (Yin and Zubiaga, 2021). Hence, to ensure models are robust and reliable, it is crucial to develop methods that can detect these over-reliances in various natural language classification tasks.

In the context of abusive language detection, we consider the concept of *negative emotions*. The

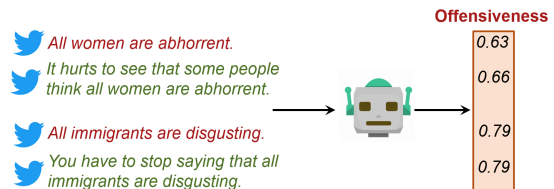


Figure 1: Probability of offensiveness generated by the TweetEval Classifier (Barbieri et al., 2020). The classifier has learned a false global sufficiency between negative emotions and the label of offense. It over-relies on this concept and ignores the broader context.

presence of an expression associated with *negative emotion* is an important signal for detecting abusive language and has been used in feature-based systems before (Chiril et al., 2022; Fortuna and Nunes, 2018). Crucially, in some examples, negative emotion words might be the cause for the abusive label, i.e., the sentence might not be abusive if the negative emotion word is replaced with other words (e.g., *I know these people. They are disgusting*). However, at the global level, the relationship between negative emotion and abusive language is a strong correlation, not causation, as it is neither globally necessary nor globally sufficient for the label of abuse.¹ Negative emotions are not globally necessary for the label of abuse because there are abusive sentences that do not contain any negative emotion words (e.g., offensive jokes, stereotyping and microaggressions). Also, words evoking negative emotions are not globally sufficient for a sentence to be abusive when interpreted in a broader context (e.g., *We should admit that in our society, they are oppressed*). But, an end-to-end model might learn that negative emotion in a sentence is globally sufficient for that sentence to be abusive. Such a classifier will struggle in classifying non-abusive sentences that contain negative emotion

¹Phenomenon P is globally sufficient for the Phenomenon Q, if whenever P happens, Q happens too. P is globally necessary for Q if whenever Q happens, P happens, too (Zaeem and Komeili, 2021).

words leading to a lack of generalizability. An example of such a case is shown in Figure 1. Specifically, classifiers’ over-reliance on the negative emotion signal can inadvertently discriminate against marginalized groups since their communications (e.g., discussing their experiences of discrimination and marginalization) can contain negative emotion words and, therefore can be wrongly considered abusive.

We explore a scenario where a user, aware of the importance of negative emotions for their use case, wants to evaluate and compare a set of trained models. Their goal is to identify and eliminate those models that are more prone to generating inaccurate results due to an overemphasis on negative emotions as primary indicators. For that, we use concept-based explanations to test if a model has learned a false global causal relationship between a user-identified concept and a label where the true relationship is a correlation. Note that global causal relationships explain the model’s output across an entire dataset, as opposed to local causal explanations, which concern the dependency of an individual prediction on a specific input feature.

Concept-based explanations are a class of explainability methods that provide global explanations at the level of human-understandable concepts (Yeh et al., 2022). While local explanations help the users understand the model’s reasoning for an individual decision with respect to the input features, global explanations are critical in comparing the processes learned by models and selecting the one that best suits the needs of a use case (Balkir et al., 2022; Burkart and Huber, 2021). Global explanations might be obtained at the level of input features through aggregating local explanations (Lundberg et al., 2020). Alternatively, global-by-design methods (e.g., probing classifiers (Conneau et al., 2018)) can be used to gain insights at higher levels of abstractions, such as linguistic properties or human-defined concepts.²

Similar to most feature importance explainability methods (e.g, Ribeiro et al. (2016); Lundberg and Lee (2017a)), concept-based explanations are originally designed to measure the *importance* of a concept. The intuitive meaning of *importance* usually refers to correlation, and it can be interpreted differently based on two notions of causality: necessity and sufficiency (Galhotra et al., 2021). Local

explainability methods usually focus on features that are of high local necessity or high local sufficiency for the label (Watson et al., 2021; Balkir et al., 2022; Joshi et al., 2022), thus considered important by human users. However, at the global level, all features must be interpreted in a larger context for accurate decision-making. We aim to determine if concept-based explanations can be utilized to evaluate whether a trained binary classifier for abusive language detection has learned a false global sufficiency relationship between the label and the concept of negative emotion. Our code and data are available at <https://github.com/IsaNejad/Global-Sufficiency/tree/main>. Our main contributions are:

- We formalize the issue of over-reliance on a concept as falsely learned global sufficiency. For the task of an abusive language classifier, we consider concepts related to *negative emotion* as being important but not globally sufficient for the label of abuse. We discuss how learning these concepts as globally sufficient results in compromised classification accuracies.
- Based on our formalization of false global sufficiency, as a baseline method, we measure the over-reliance of models on a human-defined concept using an unseen challenge set that contains the concept in both classes. Recognizing that various classifiers may have a distinct range of optimal decision thresholds, we assess the over-reliance on a concept across all possible decision thresholds and show that one of the classifiers over-relies on emotion-related concepts significantly more than the other two classifiers.
- Taking the challenge set approach as a baseline for comparison, we propose novel concept-based explanation metrics, demonstrating that similar conclusions about the degree of false global sufficiency can be drawn using these metrics. Building on previous work, we modify the TCAV procedure to measure not only the feature’s importance but also the extent of its impact on the label. We conclude that a concept-based method is preferable as it eliminates the need for manual data curation.

2 Concept-Based Explanations

Concept-based explanations evaluate the model’s decision-making mechanism at the level of a

²Here, we use the term “feature” to refer to the latent representations of a semantic concept learned by a classifier.

human-defined concept expected to be important for the task (Koh et al., 2020). Specifically, we use the Testing Concept Activation Vectors (TCAV) method to measure the influence of a human-defined concept on the model’s predictions (Kim et al., 2018). The idea of TCAV is based on the observation that human-understandable concepts can be encoded as meaningful and insightful information in the linear vector space of trained neural networks (Mikolov et al., 2013). A Concept Activation Vector (CAV), which represents the concept in the embedding space, is a vector normal to a hyperplane that separates concept and non-concept examples. Such a hyperplane is obtained by training a linear binary classifier to separate the representations of concept and non-concept examples in the embedding space.

Although TCAV can be applied to all neural network classifiers, for simplicity we limit our experiments to binary RoBERTa-based abusive language classifiers. We choose the RoBERTa-based models for their superior performance in processing social media data compared to other base language models (Liu et al., 2019). The concept, C , is defined by N_C concept examples. Also, N_R random examples are used to define non-concept examples. The RoBERTa representations for all these examples are calculated using f_{emb} , which maps an input text to its [CLS] token representation. Then, P number of CAVs, v_C^p , are generated, each through training a linear classifier that separates a sub-sample (with size N_c) of concept examples from a sub-sample of random examples (with size N_r) in the RoBERTa embedding space. The *conceptual sensitivity* of a label to the CAV, v_C^p , at input x can be computed as the directional derivative $S_{C,p}(x)$:

$$\begin{aligned} S_{C,p}(x) &= \lim_{\epsilon \rightarrow 0} \frac{h(f_{emb}(x) + \epsilon v_C^p) - h(f_{emb}(x))}{\epsilon} \\ &= \nabla h(f_{emb}(x)) \cdot v_C^p \end{aligned} \quad (1)$$

where h maps the RoBERTa representation to the logit value of the class of interest.

In this work, we use two metrics to specify the influence of the concept on the model’s prediction. First, we calculate $TCAV_{dir}$, the fraction of inputs in a set of input examples X , for which the directional derivative $S_{C,p}(x)$ is positive, i.e.:

$$TCAV_{dir}^{C,p} = \frac{|x \in X : S_{C,p}(x) > 0|}{|X|} \quad (2)$$

$TCAV_{dir}$ indicates the fraction of input examples for which the prediction scores of the model

increase if the input representation is infinitesimally moved towards the concept representation. This metric has been widely used to identify if the label has learned the concept as an important signal for the label (Yeh et al., 2020).

Besides the widely used metric of $TCAV_{dir}$ (referred to as $TCAV$ score in previous work), we introduce a new metric, $TCAV_{mag}$, which considers the size of the directional derivatives, and measures the magnitude of the influence of the concept on the label for the positive directional derivatives:

$$TCAV_{mag}^{C,p} = \frac{\sum_{x \in X, S_{C,p}(x) > 0} S_{C,p}(x)}{|X|} \quad (3)$$

We demonstrate in our results that $TCAV_{mag}$ can be an indicator of the over-reliance of the label on the concept. When calculated for all CAVs, Equations 2 and 3 generate two distributions of scores with size P for the concept C . Using a t-test, these distributions are compared with the distributions of $TCAV_{dir}$ and $TCAV_{mag}$ calculated for random examples to check for statistical significance (Kim et al., 2018).

3 False Global Sufficiency

Phenomenon P is considered globally sufficient for phenomenon Q ($P \Rightarrow Q$) if, whenever P occurs, Q also occurs (Zaeem and Komeili, 2021). In other words, global sufficiency refers to the extent to which a concept can explain the model’s output across all instances in a held-out dataset, as opposed to the more studied topic of local sufficiency, which concerns the stability of an individual prediction for a given feature in perturbed contexts (Balkir et al., 2022).

In a real-world setting, it is very unlikely that any single concept is truly sufficient for the label at a global level. In a binary classifier, a concept C is falsely learned as sufficient for the positive label if all inputs containing C are classified as positive by the classifier, regardless of context. This undesired dependency of the label on the concept suggests that the model has failed to learn how the concept interacts with context to influence the label. While this issue is closely related to spurious correlation, we use the term *false global sufficiency* because spurious correlation typically implies that the feature is irrelevant to the label, and a correlation is learned due to a confounding factor. In contrast, we consider the cases where the feature is relevant and important but not globally sufficient.

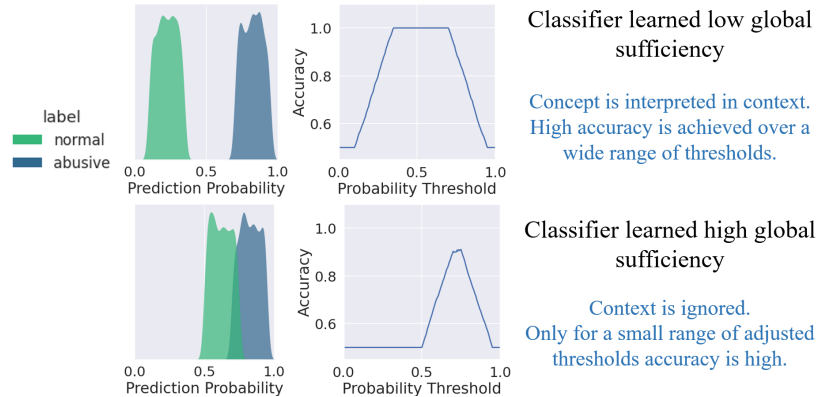


Figure 2: Illustration of the potential distribution of probabilities generated by a trained binary classifier for a challenge set that represents an important concept, along with accuracy versus threshold curves.

To make this clearer, consider the case of abusive language detection and the concept of negative emotions; if the mere presence of negative emotions in a sentence always guarantees the prediction of the positive label (abuse), then the model has learned a false sufficiency relation between the concept and the label. It over-relies on this feature and ignores the context.

To quantify falsely learned global sufficiency, we consider two scenarios: 1) where a balanced challenge set is available, which contains C in all of its examples (both classes), and 2) where no challenge set is available. For the first scenario we use the traditional approach of assessing accuracy of the classifier on a held-out test set. This approach provides a baseline in our evaluations. For the second scenario, we propose concept-based explanation metrics and compare them with the baselines obtained with the challenge sets.

3.1 Quantifying the Falsely Learned Global Sufficiency with a Challenge Set

Based on our definition of global sufficiency, one way to assess a model’s over-reliance on a concept is to evaluate its performance on a held-out challenge set, \mathbb{F} , containing both positive and negative examples of the concept of interest (Yin and Zubiaga, 2021). For simplicity, we assume that this challenge set consists of equal numbers of positive and negative examples. If a model learns a high global sufficiency between the concept C and the label of abuse, all examples in both positive and negative classes of a challenge set \mathbb{F} will be labeled as abusive. However, if the model interprets the concept in context, only the positive examples of \mathbb{F} will receive the abusive label. This indicates that in

cases where the decision threshold of the classifier clearly separates the probability distributions of the two classes, the model has learned a low global sufficiency between the concept and the label.

However, when comparing different classifiers, it is important to note that a reliable classifier should perform well (high precision and high recall) over a broad range of decision thresholds. This is because different applications may require different thresholds depending on the desired trade-off between precision and recall. For example, a classifier used to moderate social media content may need to prioritize precision over recall, which could mean using a high threshold to avoid false positives. On the other hand, a classifier used to detect all instances of abusive language may need to prioritize recall over precision, which would mean using a lower threshold to catch as many instances of abuse as possible, even if it means tolerating more false positives. Therefore, a classifier that is reliable over a wide range of decision thresholds can be more effective in different use cases, making it more practical and adaptable.

Figure 2 demonstrates two hypothetical cases for the distribution of probabilities that the classifiers might generate for the challenge set \mathbb{F} . A classifier that learned low global sufficiency between C and the positive label generates easily separable distributions of probabilities for the positive and negative examples of \mathbb{F} . In other words, for a large range of decision thresholds, the two classes of \mathbb{F} are separable, and high accuracy is achieved. Conversely, the classifier that has learned high global sufficiency between C and the positive label assigns a similar distribution of probabilities to both negative and positive examples. The two classes

of \mathbb{F} are hardly separable, and for a wide range of thresholds, the accuracy is low. Note that in order for this classifier to be accurate, it requires a careful adjustment of the decision threshold with a labeled dataset. However, this process can be very costly.

Based on this discussion, we argue that *AUC_Challenge*, the area under the curve of accuracy vs threshold, is a quantitative indicator of the separability of two classes of \mathbb{F} for all decision thresholds. According to our definition above, global sufficiency is negatively correlated with the separability of these classes. Therefore, *False_Suff*, described in Equation 4, is a quantitative metric that can be used to compare the degree of sufficiency learned by the classifiers based on \mathbb{F} :

$$False_Suff = 1 - AUC_Challenge \quad (4)$$

3.2 Quantifying the Falsely Learned Global Sufficiency with Concept-Based Explanations

The practical application of the method detailed in Section 3.1 can be limited due to the necessity of creating a custom challenge set. In this section, we use concept-based explanation to measure the falsely learned global sufficiency in a scenario where a challenge set is not available, but a lexicon representing the concept of interest exists. Following the approach of Nejadgholi et al. (2022a), we employ short templates and the concept lexicon to generate unlabeled concept examples. Then, we utilize the method described in Section 2 to compute two metrics: $TCAV_{dir}$ and $TCAV_{mag}$. If the $TCAV_{dir}$ value for the concept significantly deviates from that of random concepts, it indicates that the classifier has learned an association between the label and the concept. A significant difference in $TCAV_{mag}$ compared to random concepts suggests a strong influence of the concept on the label, potentially causing the classifier to disregard the context when the concept is present. While the absolute values of these metrics might not be definitive, we show that they can be used to compare various classifiers in terms of the degree of global sufficiency they have learned for a concept.

4 Sufficiency of the Concept of Describing Protected Groups with Negative Emotion

In this section, we evaluate the metrics introduced in Section 3 in explaining the extent of the falsely learned sufficiency between a human-defined concept and the positive label of the classifiers. We




Figure 3: Probability distributions generated by the classifiers for a challenge set (F2 and F21 of HateCheck).

specifically consider the concept of *describing a protected group with negative emotion words* and refer to it as *DesNegEm* for brevity. We chose this concept because it is tightly related to hate speech and is expected to be important for more general definitions of harmful language, such as toxic, abusive or offensive. Still, it is not a sufficient concept for these labels and has to be interpreted in the broader context (as shown by examples in Table 1).

We consider three RoBERTa-based binary classifiers, publically available and trained with large English datasets. The models are trained for general definitions of abusive language, toxicity or offensive language. We refer to these classifiers by their training datasets: Jigsaw, Civil Comments (or Civil for brevity) and TweetEval. These models are described in detail in Appendix A

Quantifying Sufficiency with a Challenge Set:

To calculate the metric described in Section 3.1, we first use the HateCheck (Röttger et al., 2021) test cases to build a challenge set for the concept of *DesNegEm*. For that, we use the F2 and F21 functionalities of HateCheck, i.e., the hateful and non-hateful examples that include this concept (Table 1). Figure 3 shows the distribution of probabilities that the three classifiers generate for this challenge set. We observe that, for a large range of decision thresholds, all three classifiers label the majority of the examples of both classes of the challenge set with a positive label. In other words, all three classifiers have learned a high sufficiency between *DesNegEm* and the label of abuse. However, the extent of the learned sufficiency is different among the classifiers. The TweetEval classifier makes the least differentiation between the two classes and generates similar distributions of probabilities for negative and positive examples with the *DesNegEm* concept. Because of this overlap between probability distributions of positive and negative classes, the accuracy of this classifier is low over all ranges of thresholds, as shown in Fig-

Functionality	Example Test Case	Gold Label	n
F2: Description using very negative attributes	“[IDENTITY] are disgusting.” “All [IDENTITY] like them are vile.”	hateful	140
F21: Denouncement of hate that make direct reference to it	“You have to stop calling [IDENTITY] disgusting.”	non-hate	140

Table 1: Used test cases from HateCheck (Röttger et al., 2021).




Figure 4: Accuracy vs threshold curve for the challenge set (F2 and F21 of HateCheck).

ure 4. The false sufficiency learned by the Jigsaw and the Civil Comments classifiers is less extreme, and Jigsaw makes the most differentiation between the two classes.

This observation can be quantified with the *False_Suff* metric (Equation 4) using the area under the curves in Figure 4. We obtain *False_Suff* of 0.41, 0.40, and 0.50 for Civil, Jigsaw and TweetEval, respectively. This metric shows a higher falsely learned sufficiency score for TweetEval than the two other classifiers, as expected. Based on these observations, we expect TCAV metrics to show lower scores for Civil and Jigsaw than the TweetEval classifier.

Global Sufficiency with Concept-Based Explanations: Here, we use the results obtained with the challenge set as a baseline to evaluate the TCAV-based metrics. Concept examples are generated using the template ‘<protected_group> are <emotion_word>.’, where <protected_group> is one of the protected groups *women, trans people, gay people, black people, disabled people, Muslims* and *immigrants* as identified by Röttger et al. (2021). For <emotion_word>, we use the *disgust* and *anger* categories of the NRC Emotion Intensity Lexicon (NRC-EIL) (Mohammad, 2018). We use the NLTK package³ to filter out words other than adjectives, past tense verbs and past participles, and also remove the words with emotion intensity lower than

³<https://www.nltk.org/>

0.5. After these steps, we are left with 368 concept words. We calculate the $TCAV_{dir}$ and $TCAV_{mag}$ scores for the concept of *DesNegEm* and compare those to the metrics calculated for random concepts with t-test for statistical significance. For random concepts, the concept examples are random tweets collected with stop words. In our implementation of the TCAV procedure, $N_R = 1000$, $N_c = 50$, $N_r = 200$ and $N_C = 386$ (number of filtered lexicon words). For input examples, X , we use 2000 tweets collected with stop words.

As presented in Table 2, for the Civil classifier, $TCAV_{dir}$ is not significantly different from the random concept, indicating that the concept information might not always be encoded as a coherent concept in the embedding space of this classifier. However, $TCAV_{mag}$ is significantly higher than random, indicating that when the information is encoded well, the presence of this concept has a significant influence on the label of abuse. The other two classifiers have learned a strong association between the concept and the label, i.e., when the concept is added to a neutral context, the likelihood of the positive label increases. However, only in the case of the TweetEval classifier, $TCAV_{mag}$ is significantly different from the random concepts, indicating a strong influence of the concept on the label, which might override the context. Therefore, for TweetEval the distribution of generated probabilities is mostly determined by the concept, not the context (similar distributions are obtained for the positive and negative examples of the challenge set). The other two classifiers consider the context to some extent and generate relatively different distributions of probabilities for the two classes.

Discussion: For all classifiers, the presence of the concept *describing a protected group with negative emotion words* is a strong signal for the label of abuse. All classifiers struggle in considering the broader contexts in sentences such as ‘*It is not acceptable to say <protected_group> are disgusting.*’ Among the three classifiers, TweetEval has learned a higher degree of sufficiency, leading to its worse performance on a challenge set containing this con-

classifier	$TCAV_{dir}$		$TCAV_{mag}$	
	DesNegEm	random	DesNegEm	random
Civil	0.67(0.05)	0.5(0.4)	0.05(0.03)	0.00(0.01)
Jigsaw	1(0)	0.7(0.4)	0.04(0.01)	0.05(0.05)
TweetEval	1(0)	0.7(0.4)	0.15(0.03)	0.01(0.01)

Table 2: Mean and standard deviation of the TCAV score for explaining the sufficiency of *Describing Protected Groups with Negative Emotion* (DesNegEm) for the three classifiers. All scores statistically significantly different from random concepts are in boldface.

cept. The TCAV metrics can be used to compare the classifiers regarding the false sufficiency relationships they have learned. These metrics provide similar insights to what is learned from assessing global sufficiency with a challenge set.

5 Global Sufficiency of Fine-Grained Negative Emotions Concepts

In the previous section, we considered the concept of *describing protected groups with negative emotions*, which is tightly related to hate speech, and thus prone to be mistakenly learned as sufficient for the label of abuse. In this section, we test our proposed method for a less obvious case by disentangling the concept of emotions and hate speech. We focus on the concept of *describing a (non-protected) group of people with negative emotions*, which differs from the previous section in 1) removing the protected groups and replacing them with unprotected groups and 2) breaking down the emotion concept to more fine-grained levels.

For fine-grained emotion concepts, we first develop a compact challenge set, examples of which are presented in Table 3. Since we consider non-protected groups in this challenge set, the examples are labeled as abusive/non-abusive as opposed to hateful/non-hateful in HateCheck (shown in Table 1). We assess the sufficiency of these concepts with the challenge set first and then compare the results to those of the proposed concept-based explanation metrics. Our goal is to investigate if the findings for the broad concept of describing protected groups with negative emotions can also be replicated at a more nuanced level of emotional granularity. We analyze the models for fine-grained categories of negative emotions, identified by Mohammad (2018), namely *disgust*, *anger*, *sadness*, and *fear*. Similar pre-processing steps to what was described in Section 4 were performed to filter the lexicon in each category of emotions.

For the challenge set, we write five abusive and five non-abusive example templates for each

emotion. Then we generate 40 abusive and 40 non-abusive examples by replacing <group> with one of the terms *Canadians*, *Chinese people*, *doctors*, *teachers*, *school children*, *football players*, *my neighbours*, and *men* to represent non-protected groups.⁴ Full list of examples of this challenge set is available in our GitHub repository mentioned in Section 1. Equivalently, for the TCAV procedure for concept templates, we use ‘*They are <emotion_word>*’, instead of ‘*<protected_group> are <emotion word>*’, which we used in Section 4.

5.1 Results

We first compare the three classifiers in handling negative emotions by investigating the results they produce for the challenge set. The $False_{suff}$ scores in Table 5 show that TweetEval has learned the highest sufficiency between these concepts and the label of abuse and therefore achieves the lowest separability between the positive and negative classes of the challenge set. To further clarify this we show the accuracy vs threshold curve for the *disgust* category of the challenge set in Figure 5. We observe that TweetEval only reaches high accuracies for a small range of thresholds, i.e, it generates a similar distribution of probabilities for the positive and negative classes that contain the emotion of *disgust*. On the other hand, Jigsaw has learned the least global sufficiency and reaches high accuracy over a wide range of thresholds.

Then we turn to the TCAV scores shown in Table 4. First, $TCAV_{dir}$ shows that the Civil Comments classifier is not significantly sensitive to negative emotions, i.e., the feature of negative emotions is not fully learned as a coherent feature by this classifier. TweetEval, on the other hand, shows significant $TCAV_{dir}$ and $TCAV_{mag}$ scores, indicating that this classifier is not only sensitive to these concepts but the influence of the concept on

⁴Though nationality may be considered a protected characteristic in some contexts, we include “Canadian” and “Chinese” here since nationality was not included in HateCheck and therefore not covered in the previous section.

Emotion	Example	Gold Label
disgust	"<group> are known for being trash ."	abusive
	"It's reasonable that <group> are feeling vindictive ."	non-abusive
anger	"All <group> are abhorrent ."	abusive
	"<group> were infuriated by the injustice. "	non-abusive
sadness	"<group> are known for being pathetic ."	abusive
	"<group> were crushed by the injustice."	non-abusive
fear	"<group> are just so scary ."	abusive
	"<group> are threatened by these policies."	non-abusive

Table 3: Challenge set created for emotion categories. In each sentence, we use a word from the corresponding category of the NRC Emotion Intensity Lexicon (NRC-EIL) (Mohammad, 2018), shown in bolded font.

classifier	$TCAV_{dir}$				$TCAV_{mag}$			
	disgust	anger	sadness	fear	disgust	anger	sadness	fear
Civil	0.48(0.37)	0.26(0.27)	0.31(0.31)	0.19(0.27)	0.05(0.06)	0.02(0.03)	0.02(0.03)	0.01(0.03)
Jigsaw	0.98(0.09)	0.93(0.2)	0.91(0.2)	0.95(0.18)	0.08(0.03)	0.05(0.03)	0.03(0.02)	0.05(0.03)
TweetEval	1(0)	1(0)	1(0)	1(0)	0.20(0.04)	0.17(0.04)	0.11(0.03)	0.13(0.03)

Table 4: Mean and standard deviation of concept-based metrics for four negative emotion concepts. Scores that are significantly different from random concepts are in boldface.




Figure 5: Accuracy vs threshold for the disgust category of the challenge set.

the label is also significantly high. Jigsaw is the classifier that has learned the dependency between negative emotions and the label of abuse and therefore is sensitive to it (as indicated by $TCAV_{dir}$), but the magnitude of the influence of concept on the label is not significantly high, and the concept is interpreted in the larger context. Interestingly, the magnitude of the influence of *disgust* and *anger* is higher than *fear* and *sadness* for all classifiers, stating a higher association of *disgust* and *anger* with abusive language. These results are in line with conclusions drawn from assessing global sufficiency with a challenge set.

6 Related Works

Most of the explainability works in NLP focus on feature importance methods to measure the importance of an input feature for the prediction at the

local level (Bahdanau et al., 2015; Sundararajan et al., 2017; Ribeiro et al., 2016; Lundberg and Lee, 2017b). However, recent works highlight that models should be assessed beyond feature importance criteria and that the reasoning behind the model’s decisions should be investigated through explainability methods. Some examples of such explainability methods include counterfactual reasoning (Wu et al., 2021; Kaushik et al., 2021; Ribeiro et al., 2020; Ross et al., 2020) or necessity and sufficiency metrics (Balkır et al., 2022; Joshi et al., 2022). Also, there is a need to compare various classifiers at the global level. Although local explanations can be aggregated to generate global explanations, they are usually obtained through costly interventions and are not practical to be applied on a large scale. For global explanations, a popular approach is to train probing classifiers (Conneau et al., 2018). However, probes only identify whether a classifier has learned a feature but stay silent about whether the feature is used in predictions (Belinkov, 2022; Tenney et al., 2019; Rogers et al., 2020). Amnesic probing is an extension of probing classifiers that identifies whether removing a feature influences the model’s predictions, which relates to the notion of the global necessity of a human-understandable concept for a prediction (Ravfogel et al., 2020; Elazar et al., 2021). Our work, on the other hand, focuses on the global sufficiency of concepts. While probing classifiers are applied to linguistic properties such as POS

classifier	<i>False_Suff</i>			
	disgust	anger	sadness	fear
Civil	0.13	0.35	0.19	0.25
Jigsaw	0.08	0.28	0.14	0.22
TweetEval	0.36	0.36	0.35	0.35

Table 5: The global sufficiency of emotion categories learned by classifiers with respect to the challenge set described in Table 3.

tagging, which are necessary for accurate language processing, we focus on human-defined semantic concepts that are known to be important for the label and test if they have been falsely learned as a sufficient cause for the label.

Concept-based explanations have been introduced in computer vision and are mostly used to explain image classification models (Graziani et al., 2018; Ghorbani et al., 2019; Yeh et al., 2020). In NLP, concept-based explanations were used to measure the sensitivity of an abusive language classifier to the emerging concept of *COVID-related anti-Asian hate speech* (Nejadgholi et al., 2022b), to assess the fairness of abusive language classifiers in using the concept of sentiment (Nejadgholi et al., 2022a), and to explain a text classifier with reference to the concepts identified through topic modelling (Yeh et al., 2020). To the best of our knowledge, our work is the first that uses concept-based explanations to assess the sufficiency of human-defined concepts in text classification.

7 Conclusion

Concept-based explanations can assess the influence of a concept on a model’s predictions. We used two metrics based on the TCAV method: the TCAV *direction* score identifies whether the classifier has learned an association between a concept and a label, and the TCAV *magnitude* score measures the extent of the influence of the concept on the label. We showed that the best-performing abusive language classifiers learned that negative emotion is associated with abuse (positive direction) but did not over-rely on this concept (low magnitude); that is, they did not overestimate the global sufficiency of that concept.

Our method can potentially be used for other NLP classification tasks. This approach is suitable for tasks where certain concepts are closely related to the label, but not enough to make a definitive determination. For example, in sentiment analysis, the price of products may have a strong connection to negative sentiment, but is insufficient to

determine it. Further research should explore how concept-based explanations can help identify cases where certain concepts are relied upon too heavily in abusive language detection or other NLP classification tasks.

8 Limitations

Our work has limitations. First, we use the TCAV framework, which assumes that concepts are encoded in the linear space of semantic representations. However, recent works show that in some cases, linear discriminants are not enough to define the semantic representations of concepts in the embedding spaces (Koh et al., 2020). Future work should consider nonlinear discriminants to accurately represent concepts in the hidden layers of NLP neural networks.

In this study, we used simple challenge sets to obtain a baseline for assessing the effectiveness of concept-based explanations in measuring false global sufficiency. Future work should focus on curating challenge sets by annotating user-generated data for the label and the concepts, in order to achieve a stronger baseline.

Our work is limited to pre-defined concepts and requires human input to define the concepts with examples. However, defining concepts in TCAV is less restrictive than pre-defining features in other explainability methods, in that concepts are abstract ideas that can be defined without requiring in-depth knowledge of the model’s inner workings or the specific features it is using. This allows for a more flexible approach where users can test the model regarding their concept of interest.

Our method can only be applied to concepts that are known to be important for the classifier and are prone to being over-represented in training sets. It’s important to check this condition independently before using our metrics. In cases where this condition does not hold true, the metrics we use in our work may be interpreted differently and may not be reliable indicators of global sufficiency. Also, we only considered two variations of emotion-related concepts. Other variations such as *expression of negative emotions by the writer of the post* should be investigated in future work.

Further, our metrics are limited to cases where different classifiers are being compared since the most important information is in the relative value of the metrics. Our metrics should not be used as absolute scores for testing a classifier.

Testing a classifier for false causal relationships is most valuable for detecting the potential flaws of the models. If our metrics do not reveal a false relationship between the concept and the label, that should not be interpreted as an indicator of a flawless model.

Ethical Statement

As with most AI technology, this approach can be used adversely to exploit the system’s vulnerabilities and produce toxic texts that would be undetectable by the studied classifier. Specifically, for methods that require access to the model’s inner layers, care should be taken so that only trusted parties could gain such access. The obtained knowledge should only be used for model transparency purposes, and the security concerns should be adequately addressed.

Regarding environmental concerns, contemporary NLP systems based on pre-trained large language models, such as RoBERTa, require significant computational resources to train and fine-tune. Larger training datasets, used for fine-tuning, usually result in better classification performance but also an even higher computational cost. To lower the cost of this study and its negative impact on the environment, we chose to use existing, publicly available classification models.

References

- Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the 3rd International Conference on Learning Representations*.
- Esma Balkir, Svetlana Kiritchenko, Isar Nejadgholi, and Kathleen Fraser. 2022. [Challenges in applying explainability methods to improve the fairness of NLP models](#). In *Proceedings of the 2nd Workshop on Trustworthy Natural Language Processing (TrustNLP 2022)*, pages 80–92, Seattle, U.S.A. Association for Computational Linguistics.
- Esma Balkır, Isar Nejadgholi, Kathleen Fraser, and Svetlana Kiritchenko. 2022. [Necessity and sufficiency for explaining text classifiers: A case study in hate speech detection](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2672–2686, Seattle, United States. Association for Computational Linguistics.
- Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. [TweetEval: Unified benchmark and comparative evaluation for tweet classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, Online. Association for Computational Linguistics.
- Yonatan Belinkov. 2022. Probing classifiers: Promises, shortcomings, and advances. *Computational Linguistics*, 48(1):207–219.
- Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2019. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion Proceedings of The 2019 World Wide Web Conference*, pages 491–500.
- Nadia Burkart and Marco F Huber. 2021. A survey on the explainability of supervised machine learning. *Journal of Artificial Intelligence Research*, 70:245–317.
- Patricia Chiril, Endang Wahyu Pamungkas, Farah Benamara, Véronique Moriceau, and Viviana Patti. 2022. Emotionally informed hate speech detection: a multi-target perspective. *Cognitive Computation*, 14(1):322–352.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. What you can cram into a single vector: Probing sentence embeddings for linguistic properties. *arXiv preprint arXiv:1805.01070*.
- Yanai Elazar, Shauli Ravfogel, Alon Jacovi, and Yoav Goldberg. 2021. Amnesic probing: Behavioral explanation with amnesic counterfactuals. *Transactions of the Association for Computational Linguistics*, 9:160–175.
- Paula Fortuna and Sérgio Nunes. 2018. A survey on automatic detection of hate speech in text. *ACM Computing Surveys (CSUR)*, 51(4):1–30.
- Sainyam Galhotra, Romila Pradhan, and Babak Salimi. 2021. Explaining black-box algorithms using probabilistic contrastive counterfactuals. In *Proceedings of the 2021 International Conference on Management of Data*, pages 577–590.
- Amirata Ghorbani, James Wexler, James Y Zou, and Been Kim. 2019. Towards automatic concept-based explanations. *Advances in Neural Information Processing Systems*, 32.
- Mara Graziani, Vincent Andrearczyk, and Henning Müller. 2018. Regression concept vectors for bidirectional explanations in histopathology. In *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*, pages 124–132. Springer.
- Laura Hanu. 2020. [How well can we detoxify comments online?](#) Unitary, accessed on 15 June, 2022.
- Laura Hanu and Unitary team. 2020. Detoxify. Github. <https://github.com/unitaryai/detoxify>.

- Nitish Joshi, Xiang Pan, and He He. 2022. Are all spurious features in natural language alike? an analysis through a causal lens. *arXiv preprint arXiv:2210.14011*.
- Divyansh Kaushik, Amrith Setlur, Eduard H Hovy, and Zachary Chase Lipton. 2021. Explaining the efficacy of counterfactually augmented data. In *Proceedings of the International Conference on Learning Representations*.
- Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, et al. 2018. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *Proceedings of the International Conference on Machine Learning*, pages 2668–2677. PMLR.
- Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. 2020. Concept bottleneck models. In *Proceedings of the International Conference on Machine Learning*, pages 5338–5348. PMLR.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Scott M Lundberg, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. 2020. From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1):56–67.
- Scott M Lundberg and Su-In Lee. 2017a. [A unified approach to interpreting model predictions](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4765–4774. Curran Associates, Inc.
- Scott M Lundberg and Su-In Lee. 2017b. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Saif M. Mohammad. 2018. Word affect intensities. In *Proceedings of the 11th Edition of the Language Resources and Evaluation Conference (LREC-2018)*, Miyazaki, Japan.
- Isar Nejadgholi, Esmā Balkır, Kathleen C Fraser, and Svetlana Kiritchenko. 2022a. Towards procedural fairness: Uncovering biases in how a toxic language classifier uses sentiment information. In *Proceedings of the Workshop on Analyzing and Interpreting Neural Networks for NLP (BlackboxNLP)*.
- Isar Nejadgholi, Kathleen Fraser, and Svetlana Kiritchenko. 2022b. [Improving generalizability in implicitly abusive language detection with concept activation vectors](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5517–5529, Dublin, Ireland. Association for Computational Linguistics.
- Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. Null it out: Guarding protected attributes by iterative nullspace projection. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7237–7256.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. “Why should i trust you?” Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with CheckList](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Alexis Ross, Ana Marasović, and Matthew E Peters. 2020. Explaining NLP models via minimal contrastive editing (MiCE). *arXiv preprint arXiv:2012.13985*.
- Paul Röttger, Bertie Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, and Janet Pierrehumbert. 2021. [HateCheck: Functional tests for hate speech detection models](#). 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 41–58, Online. Association for Computational Linguistics.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *Proceedings of the International Conference on Machine Learning*, pages 3319–3328.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R Bowman, Dipanjan Das, et al. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. *arXiv preprint arXiv:1905.06316*.
- David S Watson, Limor Gultchin, Ankur Taly, and Luciano Floridi. 2021. Local explanations via necessity

and sufficiency: Unifying theory and practice. In *Uncertainty in Artificial Intelligence*, pages 1382–1392. PMLR.

Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel Weld. 2021. [Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models](#). 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 6707–6723, Online. Association for Computational Linguistics.

Ellery Wulczyn, Nithum Thain, and Lucas Dixon. 2017. [Ex machina: Personal attacks seen at scale](#). In *Proceedings of the 26th International Conference on World Wide Web, WWW '17*, pages 1391–1399, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.

Chih-Kuan Yeh, Been Kim, Sercan Arik, Chun-Liang Li, Tomas Pfister, and Pradeep Ravikumar. 2020. On completeness-aware concept-based explanations in deep neural networks. *Advances in Neural Information Processing Systems*, 33:20554–20565.

Chih-Kuan Yeh, Been Kim, and Pradeep Ravikumar. 2022. Human-centered concept explanations for neural networks. In P. Hitzler and M. K. Sarker, editors, *Neuro-Symbolic Artificial Intelligence: The State of the Art*, volume 342, page 2. IOS Press.

Wenjie Yin and Arkaitz Zubiaga. 2021. Towards generalisable hate speech detection: a review on obstacles and solutions. *PeerJ Computer Science*, 7:e598.

Mohammad Nokhbeh Zaeem and Majid Komeili. 2021. Cause and effect: Concept-based explanation of neural networks. In *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 2730–2736. IEEE.

Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. [Predicting the type and target of offensive posts in social media](#). 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 1415–1420, Minneapolis, Minnesota. Association for Computational Linguistics.

A Models

We include the following publicly available abusive language classification models in this study:

- **Jigsaw**⁵: a RoBERTa-based binary toxicity classifier fine-tuned on the combination of two datasets created by Jigsaw and used in

⁵https://huggingface.co/SkolkovoInstitute/roberta_toxicity_classifier/tree/main

Kaggle competitions on toxicity prediction in 2018–2020. The first dataset, Wikipedia Toxic Comments (Wulczyn et al., 2017), includes 160K comments from Wikipedia talk pages. The second dataset, Civil Comments (Borkan et al., 2019), comprises over 1.8M online comments from news websites. Both datasets are annotated for toxicity (and its subtypes) by crowd-sourcing. The model creators report the AUC of 0.98 and F1-score of 0.76 on the Wikipedia Toxic Comments test set. The model is released under CC BY-NC-SA 4.0.

- **Civil Comments**⁶ (Hanu and Unitary team, 2020): a multi-class RoBERTa-based model fine-tuned on the Civil Comments dataset to predict toxicity and six toxicity subtypes (severe toxicity, obscene, threat, insult, identity attack, and sexual explicit). A part of the dataset is annotated for identity groups targeted in toxic comments. The prediction model is trained to optimize the outcome fairness for the groups in addition to the overall accuracy. This is achieved through the loss function that combines the weighted loss functions for two tasks, toxicity prediction and identity prediction (Hanu, 2020).
- **TweetEval**⁷ (Barbieri et al., 2020): a RoBERTa-based binary classifier to detect offensive language, released as part of the TweetEval evaluation benchmark. The model was trained on 58M tweets and then fine-tuned on the Offensive Language Identification Dataset (OLID) (Zampieri et al., 2019). The OLID training set comprises about 12K tweets. The model achieved the macro-averaged F1-score of 77.1 on the OLID test set.

⁶<https://huggingface.co/unitary/unbiased-toxic-roberta>

⁷<https://huggingface.co/cardiffnlp/twitter-roberta-base-offensive>