

The Language of Trauma: Modeling Traumatic Event Descriptions Across Domains with Explainable AI

Miriam Schirmer¹, Tobias Leemann¹, Gjergji Kasneci¹, Jürgen Pfeffer¹, David Jurgens²

¹Technical University of Munich

²University of Michigan

miriam.schirmer@tum.de, tobias.leemann@tum.de, gjergji.kasneci@tum.de,

juergen.pfeffer@tum.de, jurgens@umich.edu

Abstract

Psychological trauma can manifest following various distressing events and is captured in diverse online contexts. However, studies traditionally focus on a single aspect of trauma, often neglecting the transferability of findings across different scenarios. We address this gap by training language models of progressing complexity on trauma-related datasets, including genocide-related court data, a Reddit dataset on post-traumatic stress disorder (PTSD), counseling conversations, and Incel forum posts. Our results show that the fine-tuned RoBERTa model excels in predicting traumatic events across domains, slightly outperforming large language models like GPT-4. Additionally, SLALOM-feature scores and conceptual explanations effectively differentiate and cluster trauma-related language, highlighting different trauma aspects and identifying sexual abuse and experiences related to death as a common traumatic event across all datasets. This transferability is crucial as it allows for the development of tools to enhance trauma detection and intervention in diverse populations and settings.

1 Introduction

Post-Traumatic Stress Disorder (PTSD) is a significant mental health condition that can develop after experiencing a traumatic event. For an event to potentially lead to PTSD, it must involve actual or threatened death, serious injury, or a threat to one's physical integrity, causing intense fear, helplessness, or horror (Friedman et al., 2007; Gold, 2017). Although about 70% of Americans will encounter such traumatic events in their lifetime, only about 5-7% develop PTSD, highlighting that PTSD is relatively rare despite high trauma exposure. This figure could be even higher, as many cases may go undiagnosed (Bonn-Miller et al., 2022; Atwoli et al., 2015).

This discrepancy suggests that various factors, including psychological resilience, the nature of

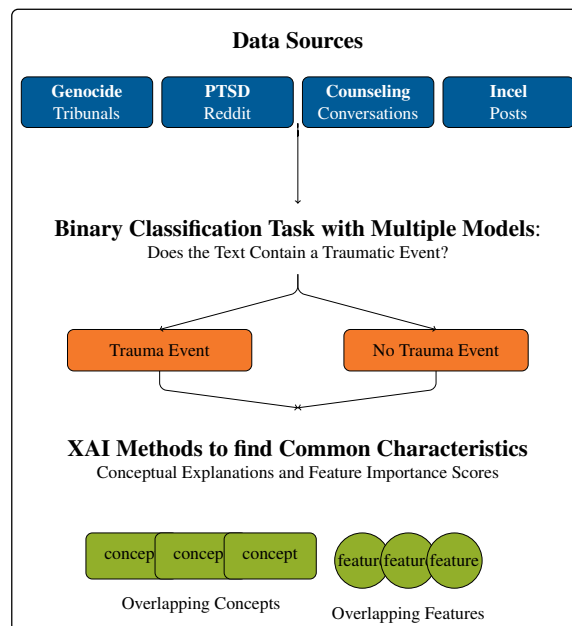


Figure 1: We (1) create a cross-domain trauma dataset, (2) classify traumatic events with models of different complexity, and (3) use XAI methods to identify overlapping characteristics of traumatic events.

the trauma, and access to mental health support, influence the development of PTSD. Definitions of trauma and responses to it can vary widely across cultures and social contexts, affecting the prevalence and expression of PTSD.

To investigate the interplay of these factors, we are proposing a Natural Language Processing (NLP) approach to identify traumatic events across different domains. Understanding the cross-domain mechanisms of trauma is crucial for developing comprehensive support systems and interventions that are adaptable to various contexts. To address this, we are following up on these research questions:

RQ1: Given the diverse forms of trauma, what are the most effective methods for modeling and predicting its manifestations?

RQ2: How transferable is the detection of multi-faceted traumatic events across domains?

RQ3: What are the cross-cutting mechanisms related to trauma that can be identified across different types and contexts of traumatic events?

Our work advances trauma detection by applying NLP and Explainable AI (XAI) methods to offer detailed insights not yet explored in the literature. We contribute by:

- (1) Identifying key trauma concepts from psychological literature and replicating them using NLP methods.
- (2) Modeling traumatic event detection with various language models and creating a dataset that includes genocide court transcripts, PTSD-related Reddit posts, counseling conversations, and “Involuntary Celibates” (Incel) forum posts.
- (3) Developing a three-stage XAI framework that approximates Shapley values, assesses feature importance, and identifies task-relevant concepts, providing a comprehensive understanding of trauma at both the instance and dataset levels.
- (4) Automating trauma detection to enhance online psychological support by displaying hot-line information and resources in forums where trauma is frequently discussed.¹

2 Traumatic Events & Language

2.1 Definition & Scope

Psychological trauma, as defined by the American Psychological Association (APA), encompasses experiences of "exposure to actual or threatened death, serious injury, or sexual violence," whether directly encountered or witnessed. This includes instances where individuals "learn that the traumatic event(s) occurred to a close family member or close friend" (American Psychiatric Association, 2013).

While psychological trauma and PTSD are frequently discussed in the context of childhood abuse and the military, trauma can manifest in a variety of situations (Van der Kolk, 2003; Yehuda, 1998). It can arise in interpersonal violence, like domestic abuse and sexual assault, and accidents or natural disasters. Trauma can also result from medical

issues, bereavement and loss, emotional and psychological abuse, and its manifestation can vary depending on cultural beliefs and values (Smelser et al., 2004).

2.2 Trauma Contexts & Categorization

Within the psychological literature, key events have been identified that are typical for specific trauma contexts. In armed conflict and mass atrocities, exposure to severe violence and death is prevalent. This often includes the death of close family members, forced displacement, and sexual abuse (Powell et al., 2003). For instance, Dyregrov et al. (2000) found that most child survivors of the Rwandan genocide had witnessed severe injuries and deaths, with more than half witnessing massacres.

In domestic trauma, the most common forms are physical abuse (e.g., intimate partner violence), emotional abuse, and neglect (McCloskey and Walker, 2000). Emotional abuse is particularly hard to detect due to its subtle nature, including consistent belittling, criticizing, or bullying (Dye, 2020; Idsoe et al., 2021). Sexual violence, whether in war or domestic contexts, is an especially devastating form of trauma (Kiser et al., 1991). This includes childhood sexual abuse, rape, and exploitation.

The range of traumatic events makes conceptualizations of trauma complex. Researchers have categorized trauma in line with diagnostic manuals like the Diagnostic and Statistical Manual of Mental Disorders (DSM) into types such as assaultive violence (e.g., military combat, rape, threats with weapons), other injuries or shocking events (e.g., serious car accidents and life-threatening illnesses) (Breslau et al., 2004). Identifying these events is crucial, as most subsequent issues are linked to the initial trauma due to the development of trauma-specific fears in PTSD (Terr, 2003).

2.3 NLP for Trauma Detection

Given the variety and subjective nature of traumatic experiences, detecting them in text is complex. Despite these challenges, recent research has shown that NLP methods can improve the detection of psychological disorders and aid in treatment adaptation (Ahmed et al., 2022; De Choudhury and De, 2014; Le Glaz et al., 2021; Malgaroli et al., 2023; Zhang et al., 2022).

NLP and Mental Health. Major areas in this field include promoting better health and early disorder identification for intervention (Calvo et al., 2017; Swaminathan et al., 2023). For example,

¹We make all data and code available at <https://github.com/MiriamSchirmer/trauma-language>

Levis et al. (2021) associated linguistic markers from psychotherapist notes with treatment duration. Analyzing mental health chat conversations, Hornstein et al. (2024) found that words indicating younger age and female gender were associated with a higher chance of re-contacting.

Recently, the use of Large Language Models (LLMs) has led to the development of specific models for mental health applications (Xu et al., 2024; Yang et al., 2024). While LLMs effectively detect mental health issues and provide eHealth services, their clinical use poses risks, such as the lack of expert-annotated multilingual datasets, interpretability challenges, and issues regarding data privacy and over-reliance (Guo et al., 2024).

Specifically for social media data, there has been research on using sentiment analysis and semantic structures to detect anxiety (Low et al., 2020) or depression (Tejaswini et al., 2024) on Reddit posts. In suicide prevention on social media, Sawhney et al. (2020) developed a superior model for suicidal risk screening that identifies emotional and temporal cues, outperforming competitive methods (c.f., Ji (2022) on suicidal risk detection).

Trauma Detection. In trauma research, progress is being made in analyzing patient narratives (He et al., 2017) and identifying cases of post-traumatic stress disorder (PTSD) through speech (Marmar et al., 2019). Miranda et al. (2024) developed an NLP workflow using a pre-trained transformer-based model to analyze clinical notes of PTSD patients, revealing consistent reductions in trauma criteria post-psychotherapy. Disruptions in lexical characteristics and emotional valence have been found to contribute to identifying PTSD (Quillivic et al., 2024). Using Twitter data, Ul Alam and Kapadia (2020) investigated whether posts can complete clinical PTSD assessments, achieving promising accuracy in PTSD classification and intensity estimation validated with veteran Twitter users (cf. Coppersmith et al. (2014); Reece et al. (2017)).

2.4 Trauma Event Detection in this Study

Previous work has identified language markers of PTSD, such as overuse of first-person singular pronouns, increased use of words related to depression, anxiety, and death, and more negative emotions (Low et al., 2020). However, these markers are not specific to trauma and can also be associated with other psychological disorders, complicating accurate identification. Additionally, the transferability of detection methods is often lacking (Coppersmith

et al., 2014; Quillivic et al., 2024).

Trauma detection in NLP is distinct in that it involves identifying a specific traumatic event that precedes a PTSD diagnosis, unlike the detection of depression or anxiety, which do not require a concrete event in their definitions. This study focuses on detecting such events in online resources, avoiding symptom or diagnosis analysis. Drawing conclusions about mental health from public text data alone is impossible without additional psychological information. We aim to identify instances meeting the APA’s definition of trauma, minimizing subjectivity by closely following their criteria.

3 Data & Labeling

3.1 Data Sources

Our final dataset is built from four datasets, each offering unique perspectives on traumatic experiences (Table 1) to identify common characteristics of trauma that extend beyond specific events, such as those related to war: The Genocide Court Transcripts (GTC; Schirmer et al., 2023a) dataset comprises text from genocide tribunals, providing insights into severe human rights violations and the profound trauma experienced by victims and witnesses. This encompasses 90 cases across the International Criminal Tribunal for Rwanda, the International Criminal Tribunal for the former Yugoslavia, and the Extraordinary Chambers in the Courts of Cambodia. The Reddit PTSD Dataset includes posts from the PTSD subreddit of the Reddit Mental Health Dataset (Low et al., 2020), where individuals discuss their experiences with post-traumatic stress disorder, sharing personal stories and support. The Mental Health Counseling Conversations Dataset (Amod, 2024) features questions and answers sourced from online counseling and therapy platforms. The questions cover a wide range of mental health topics, and qualified psychologists provide the answers.

The Incel Posts Dataset (Matter et al., 2024) contains posts from Incel community forums and reflects extreme misogynistic viewpoints. This dataset serves as a control in our study: Though not explicitly trauma-related, it includes posts on depression, bullying, and violence directed towards women. The violent and aggressive language in this dataset helps quantify our models’ ability to distinguish explicit trauma from related emotional distress.

Dataset	Description	Size & Balance	AA
Genocide Transcript Corpus (GTC)	Witness statements from 90 different cases across three different genocide tribunals.	15,845 samples (trauma: 13.54%)	n/a
PTSD Subreddit (PTSD)	Post-Traumatic Stress Disorder (PTSD) subset of the Reddit Mental Health Dataset.	1,200 samples (trauma: 47.19%)	(1) $\alpha = .63$ (2) $F1 = .77$
Counseling Dataset	Queries submitted by users seeking advice, with answers provided by professionals.	1,200 samples (trauma: 8.16%)	(1) $\alpha = .69$ (2) $F1 = .95$
Incel Dataset	Posts from the Incel online forum <i>incels.is</i> .	300 samples (trauma: 2.67%)	(1) $\alpha = .43$ (2) $F1 = .78$

Table 1: Dataset Overview. Note: Annotator agreement (AA) was calculated (1) among crowd workers (Krippendorff’s α) and (2) for the crowd worker majority vote vs. the expert vote (Binary F1).

3.2 The Trauma Event Dataset TRACE

We present the final trauma event dataset TRACE (Trauma Event Recognition Across Contextual Environments). To that end, all source datasets were pre-processed to ensure comparability for the detection task, including the removal of URLs and standardization of formatting. Due to their varied origins, the samples from each dataset differ in size, with instances ranging from single-word sentences to more elaborate descriptions of events and personal thoughts across all datasets. For compatibility with the BERT-architecture, we split instances exceeding the 512-token limit into smaller segments. Our approach treats each segment as independent, with trauma classification based solely on its content. While some segments from the same text may appear in both training and test sets, we expect label leakage to be minimal since the model must rely on the segment’s content for accurate prediction. 7-20% (depending on the dataset) of segments were split overall.

Our study aims to demonstrate cross-domain transferability on realistic data, making it crucial to use datasets with their expected class distribution, even if they differ in context and trauma event rates. We matched the size of all datasets to the Counseling Dataset, which had the fewest samples and the most significant class imbalance. Despite these constraints, the Counseling Dataset remains highly valuable for its unique perspective on online mental health conversations, particularly in seeking expert advice.

Annotation Process. The GTC already contains a binary trauma variable that psychologists have annotated according to the APA definition of trauma. For the PTSD and Counseling datasets, 1,200 instances each were annotated by crowdworkers. We used the Portable Text Annotation Tool (Potato;

Pei et al., 2022) to set up an annotation interface for crowdworkers using Prolific as a recruitment platform for annotators. Each instance was labeled by three annotators, and all annotators received an hourly reimbursement of approximately 12 US\$. The crowdworkers were provided detailed instructions, the APA definition of a traumatic event, and three examples. Both the Prolific pre-screening and the instructions contained a trigger warning, ensuring that participants were free to pause or stop the study at any time (Appendix B, Figure 6). Annotators were based in either the US or the UK and fulfilled English language requirements.

We conducted a pilot study comparing single-choice and span annotation setups, where participants highlighted traumatic events in the text. The final annotation task used the span setup to ensure accurate detection (Appendix B, Figure 7). Annotations were quality-checked, resulting in the removal of two annotator entries that featured an unlikely number of trauma labels. This did not affect the final sample count (e.g., 1,200). For the Incel dataset, we only labeled 300 instances since it serves as a control test set. To ensure quality, two researchers with psychology degrees annotated a subset of 200 instances from each dataset and resolved disagreements through discussion (Cohen’s $\kappa = .82$).

Annotator Agreement. To assess annotator consistency, we report Krippendorff’s α for agreement among crowdworkers and provide Binary F1 scores to measure agreement between the crowdworker majority vote and the expert vote, with the latter serving as the ‘true’ reference (Table 1). Both agreements were best for the Counseling Dataset. All agreement scores indicate at least moderate agreement (Krippendorff, 2018). Despite the variability, our primary focus is on the accuracy of labels from majority voting. The moderate F1 scores

indicate that majority votes are reliable labels, supporting the robustness of our annotation process. Given the subjective nature of interpreting trauma-related constructs, some disagreement is expected, similar to lower agreement seen in tasks like hate speech detection (Li et al., 2024). This level of agreement, while not perfect, provides a solid foundation for the study.

4 Methods

4.1 Models and Hyperparameters

In this work, we implement five sequence classification models for natural language inputs. The suitability of these models for trauma detection in different contexts is defined by criteria such as complexity, interpretability, hyperparameter optimization, and scalability. To help in understanding the trade-offs and strengths of each approach, we provide an overview of the models considered in Table 2. The hyperparameters given are optimized with a hyperparameter optimization framework (see Appendix A.2).

BoW-Naive-Bayes Model. The simplest model is obtained by fitting a Naive-Bayes model on the word counts in both classes. Let $\mathbf{t} = [t_1, t_2, \dots, t_N]$ be an input sequence. We model the log-odds by combining two key components.

First, we calculate the prior odds, which is the log of the initial ratio of the probabilities of the two categories. Second, we add the word-specific weights, which are summed over all elements in the input sequence. Each weight represents the log of the ratio of the probabilities of that element occurring in each category. We obtain the weight of a term by counting its occurrences in documents from both classes and applying Laplace smoothing with a specified hyperparameter α . The main advantage of this linear model is its interpretability due to the individual weights of each token that are explicitly computed.

N-Gram Logistic Regression Model. We compute n -grams for the datasets and fit a logistic regression model on the TF-representation of the n -Grams, where n is $[1, 2, 3]$.

TF-IDF Fully-Connected Model. Furthermore, we compute TF-IDF vectors for the samples and train a fully connected neural network using this representation as an input. We use either one or two hidden layers, with the number of hidden layers and their width as a hyperparameter.

BERT-based Models. We train the popular

encoder-only transformer models BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019). We experiment with both pretrained and non-pretrained versions of these models. We find that the pretrained models yield superior performance, which is why we restrict our analysis to these models for the main paper. We use the learning rate, number of layers, and number of heads as hyperparameters.

Black-box API models (GPT-3.5/GPT-4). We use a prompt template to access publicly available foundation model APIs for GPT-3.5 and GPT-4 (Achiam et al., 2023). We rephrase the classification tasks as a sequence completion tasks by using a prompt template, which instructs the model to either output “0” or “1”, and apply basic prompt engineering, including a task definition, the trauma definition, and labeling instructions (see Appendix A.2). We use the top token log-probabilities returned by the API to compute class log-odds, which can be used to compute calibration measures and ROC curves.

4.2 Explainable AI Methods

We apply XAI approaches to gather insights into how trauma is described and recognized across different domains. Feature-based explanations allow us to gain insights into the importance of individual input features, i.e., tokens. We chose model-agnostic approaches that treat the predictive model as a black-box function and can be applied to any model (SHAP values) and model-specific, mechanistic approaches that are only applicable to specific models but can more faithfully describe the output of certain model classes. Additionally, concept-based explanations allow us to move beyond individual feature attributions to a higher level of abstraction and help us identify interpretable concepts that are crucial for trauma detection without requiring extensive supervision. These methods collectively enhance our ability to interpret model predictions and validate their reliability.

SHAP Explanations. Shapley values originate from game theory and have been proposed to compute the contribution of individual features to the output of a non-linear function. They are a form of feature attribution explanation that assigns each input token a numerical score. The score corresponds to the average contribution to the output obtained when this feature is added. We compute SHAP values using an efficient sampling-based algorithm with the implementation of Lundberg and Lee (2017).

Model	Complexity	Interpretability	Hyperparameters	Scalability	Prediction
BoW-Naive-Bayes	Low	High	binary, smoothing param. α	High	After training
N-Gram Logistic Regression	Low	Medium	TF-IDF, n-grams	High	After training
TF-IDF Fully-Connected NN	Medium	Medium	Hidden layers, layer width	Medium	After training
BERT-based Models	High	Low	Learning rate, layers, heads	Low	One-shot or after fine-tuning
Black-box API (GPT-3.5/4)	High	Low	Prompt template, API settings	Low	One-shot or after fine-tuning

Table 2: Model Categorization According to General Suitability Criteria

SLALOM Explanations. Leemann et al. (2024) have shown that single attribution scores cannot fully describe the inner workings of modern transformer language models. The authors propose SLALOM, a model to assess the role of input tokens along two dimensions: The *token value* score describes the effect each token has on its own, while the *token importance* describes how much weight is placed on each token when tokens are concatenated to sequences. While SLALOM can be used to approximate any model’s behavior in principle, it is particularly suited for transformer models, like the BERT and RoBERTa models used in this work.

Concept-based Explanations. Concept-based explanations have been proposed as an alternative to feature-wise explanations. They do not reason over individual input features (tokens, pixels, etc.) but instead use a higher level of abstraction (Kim et al., 2018; Koh et al., 2020). However, it is difficult to discover meaningful concepts from the data without supervision (Leemann et al., 2023). In case no concept annotations are present in the data, they identify clusters in a model’s latent space that best describe a model’s decision. In this work, we turn to Completeness-Aware Concept-Based Explanations (Yeh et al., 2019), which are one of the few conceptual explanation techniques that are applicable to textual inputs and do not require supervision in terms of the data. The concepts are represented as a set of salient examples, i.e., sample snippets that most strongly exhibit the discovered concept.

We focus on the RoBERTa architectures for concept-based text classification, which proved reliable across all datasets. We use the logit outputs of this model to obtain SHAP and SLALOM explanations and use the latent representation before the classification head as the latent space where the concept vectors are identified. Details on expla-

nation approaches and their hyperparameters are provided in Appendix A.1.

5 Model Performance Results

Classification Performance We fit all the models to the respective datasets after performing hyperparameter optimization (cf. Appendix A.2) and report their performance metrics in Table 3. The evaluation across GTC, PTSD, and Counseling datasets shows clear trends. Transformer-based models, especially fine-tuned BERT and RoBERTa, significantly outperform traditional models and feedforward neural networks. The Naive-Bayes-BoW and N-Gram Logistic Regression models show moderate performance but lag behind due to their simpler architectures. The feedforward model performs reasonably well but is outclassed by transformer models. Fine-tuned BERT and RoBERTa exhibit substantial improvements in all metrics, with RoBERTa achieving the highest F1 scores in the GTC dataset ($F1 = .74$) and the PTSD dataset ($F1 = .71$), highlighting its effective language comprehension capabilities. To control for dataset size effects, we ran an additional experiment using 1,000 randomly selected GTC samples in the training set to match the size of other datasets. The performance remained consistent, indicating that our findings on smaller datasets likely extend to larger ones (Appendix C, Table 7).

OpenAI’s GPT-4 also performs particularly well on the PTSD and Counseling datasets and even outperforms BERT in the F1 metric on Counseling, showcasing its strong generalization abilities despite not being further fine-tuned and relying on a single prompt for these tasks. Interestingly, all models perform reasonably well, which may be attributed to the specific task of trauma event detection. However, the Counseling dataset proved more challenging due to its very imbalanced class distri-

Dataset	GTC		PTSD		Counseling		
	LM	F1 (bin.)	AU-ROC	F1 (bin.)	AU-ROC	F1 (bin.)	AU-ROC
NaiveBayes-BoW		0.53 ± 0.09	0.82 ± 0.09	0.56 ± 0.04	0.70 ± 0.02	0.17 ± 0.01	0.70 ± 0.02
NGramLogisticRegression		0.51 ± 0.10	0.83 ± 0.09	0.58 ± 0.02	0.70 ± 0.02	0.15 ± 0.05	0.79 ± 0.01
FeedForwardModel		0.52 ± 0.10	0.84 ± 0.09	0.52 ± 0.05	0.74 ± 0.01	0.03 ± 0.03	0.78 ± 0.01
BERT (finetuned)		0.71 ± 0.01	0.96 ± 0.00	0.66 ± 0.02	0.80 ± 0.01	0.35 ± 0.05	0.91 ± 0.01
RoBERTa (finetuned)		0.74 ± 0.01	0.97 ± 0.00	0.71 ± 0.01	0.83 ± 0.01	0.18 ± 0.09	0.88 ± 0.02
OpenAI GPT-4		0.64	0.94	0.69	0.82	0.36	0.85

Table 3: Classification performance of the language models used in this work. We report Binary F1-Scores, and Area under the Receiver-Operator Curve (“AU-ROC”). We report standard errors over cross-validation with 5 runs for all models but the Black-box API models, where computation costs are prohibitive.

bution and the presence of very few trauma event samples. This is reflected in the GPT-4 F1 score of .36, which was the highest for this dataset but still indicates the difficulty of the task. RoBERTa achieves strong performance metrics overall, highlighting the impact of architectural improvements and extensive training on larger datasets, though it does not outperform BERT on the Counseling dataset.

Cross-Domain Performance Figure 2 presents the cross-domain results of RoBERTa models fine-tuned on one dataset and evaluated on other datasets, using the AUC-ROC metric (cf., Appendix A, Table 5). Models trained on the GTC dataset showed the highest generalizability, performing well across all test sets. Those trained on the PTSD dataset excelled on their own test set and performed strongly on others. Models trained on the Counseling dataset achieved top performance on their own set but did less well on others. The model trained on all combined datasets showed robust and consistent performance across all test sets, maintaining high accuracy and reliability. Despite differences in trauma types across datasets, significant overlaps contribute to strong cross-testing results. For example, both the GTC and PTSD datasets include trauma related to death, acute stress reactions, and physical violence, aiding models’ cross-dataset performance. However, the GTC dataset’s unique military component may cause some performance differences. Overall, high cross-domain performance suggests that shared trauma themes enable effective generalization across different contexts.

The results show that the RoBERTa model fine-tuned on the PTSD dataset has the best generalizability across different datasets, with models trained on the full data also performing well. Given the diversity of traumatic events across datasets,

this result suggests the trauma features in the PTSD dataset are broadly applicable for learning a general event type rather than causing models to pick up on only keywords. Counseling-trained models perform well on their own dataset but do not generalize as effectively. Performance on the Incel dataset indicates all models effectively differentiate trauma-related vocabulary from control data.




Figure 2: Cross-domain performance (AUC-ROC) when a RoBERTa model is trained on one dataset and tested on other datasets.

SHAP Explanations To understand how the models attribute feature importance to the trauma label, we calculated SHAP values for some samples from all datasets, focusing on comparing RoBERTa and GPT-4 due to their high performances and the interesting differences in how these language models classify trauma. While most classifications aligned (see Figure 8 in Appendix A), we found that, in several instances, GPT-4 provided more non-trauma attributions for certain features compared to RoBERTa.

Figure 3 shows a counseling dataset example where RoBERTa and GPT-4 disagree. RoBERTa assigns high relevance to words like *yells*, *abuse*, and *depressed*, while GPT-4 does not, possibly due to the forum user’s uncertainty about defining abuse. This discrepancy may stem from GPT-4’s closer adherence to the APA definition of trauma, with less variation and personal bias than human anno-

tators, who may classify events based on their own experiences and interpretations.

These findings, though based on exemplary instances, highlight the challenge of detecting mental abuse. RoBERTa may rely more on specific keywords related to abuse, whereas GPT-4 seems to consider contextual nuances. Human annotators might interpret such incidents as traumatic based on subjective judgment and empathy, while GPT-4, adhering strictly to the APA definition of trauma, did not classify these incidents as trauma.

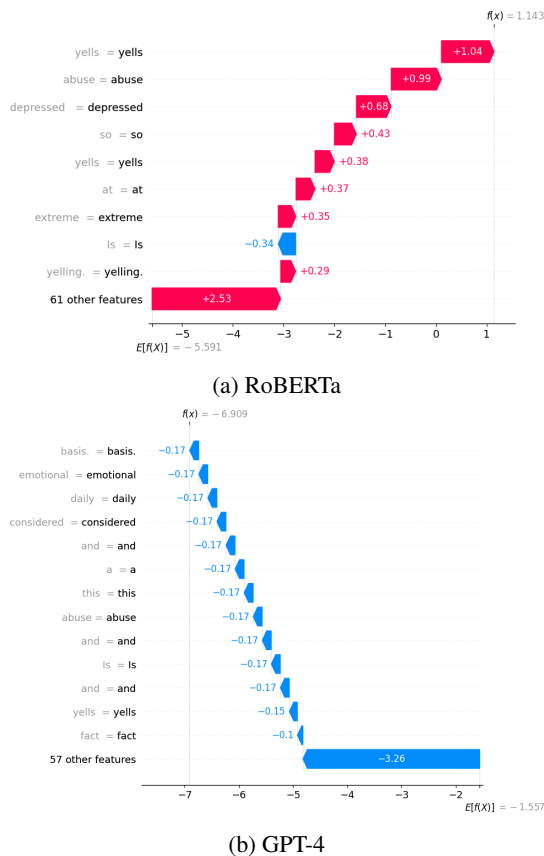


Figure 3: SHAP value, for instance, from the **Counseling Dataset**: “My dad doesn’t like the fact that I’m a boy. He yells at me daily because of it and he tells me I’m extreme and over dramatic. I get so depressed because of my dad’s yelling. He keeps asking me why I can’t just be happy the way I am and yells at me on a daily basis. Is this considered emotional abuse?”

6 Characteristics of Trauma Across Domains

Feature Characteristics with SLALOM The SLALOM feature importance scores from all datasets focus on the highest value features for trauma classification. Features like *dream* and *shattered*, in the top right corner, contribute most to

the trauma classification. For clarity, overlapping features were excluded (blue dots remain in the figure) (Figure 4).

Notable feature variability includes war-related vocabulary (e.g., *bombardment*, *bullets*) likely from genocide-related data, and more generalizable words (e.g., *dreams*, *accident*, *dead*) applicable across domains. Amplifying words like *intense*, *suddenly*, and *gloomy* also appear, fitting traumatic contexts without specific events.

Groups of thematically related words are evident: *dead* and *assassinated* represent death, *wounded*, *choking*, and *slapped* indicate physical injury and violence, and *dreams*, *shattered*, and *replay* are associated with trauma’s psychological impact.




Figure 4: SLALOM feature importance scores based on the full dataset and the RoBERTa model.

Conceptual Explanations For each dataset, we assessed conceptual explanations to detect context-specific trauma concepts. We select the concepts that have the highest number of traumatic instances in the neighborhood closely associated with the corresponding concept (Figure 5).

In the genocide dataset, concepts related to killings, death, and severe injuries were prominent, reflecting the extreme nature of the content. In contrast, the PTSD and counseling datasets, which address more everyday trauma, contained more references to domestic violence and abuse. The smaller size of the counseling dataset made it challenging to identify unique concepts without overlap.

Across all contexts, death and sexual violence were prevalent. In the genocide dataset, these were

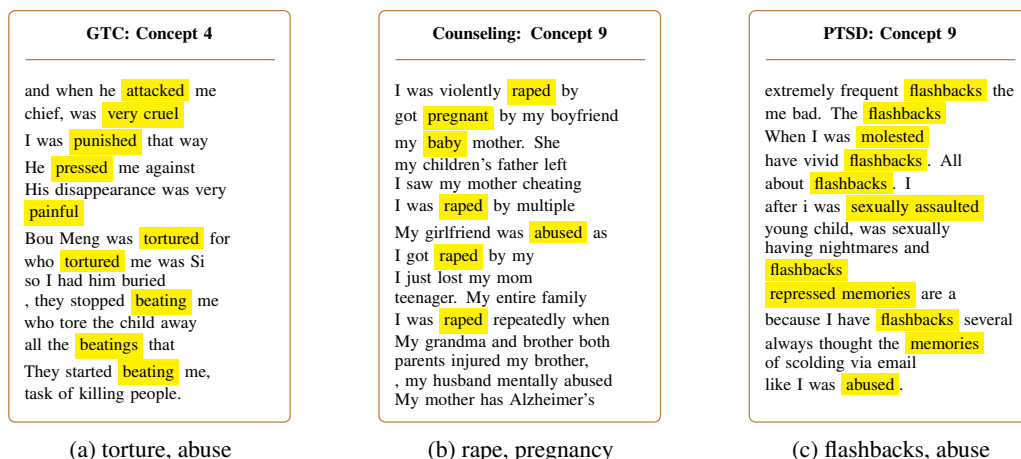


Figure 5: Trauma-related concepts found in the three datasets (most salient examples, RoBERTa Model). For more examples, see Appendix A.

depicted through killings and executions, whereas in other datasets, they were associated with grief, loss, and suicide. Sexual violence, particularly rape, consistently appeared as a common source of PTSD, which is consistent with the psychological literature (Atwoli et al., 2015).

7 Conclusion

Traumatic events shape millions of lives. Computational tools to recognize these events can help third parties provide support. However, their diversity makes classification challenging. This paper introduces a new dataset for recognizing traumatic events and analyzes (i) NLP models' performance, (ii) their generalizability across domains, and (iii) whether they learn general trauma features using XAI techniques. We show that transformer-based models, particularly fine-tuned BERT and RoBERTa, outperform traditional models and neural networks, although simpler models achieve very close scores. Models trained on diverse datasets, like the combined dataset model, perform well across various contexts, emphasizing the importance of exposure to varied data. Despite unique aspects in datasets like GTC's military component, shared trauma themes enable strong cross-domain generalizability. Notably, large language models like GPT-4 do not necessarily improve performance on this task without fine-tuning and domain-specific adjustments.

Our analysis shows that while certain features of trauma are context-specific, there are also universal elements across different experiences. However, certain types of traumatic events—notably mental abuse—are particularly challenging to classify

due to their less defined nature and greater variability, highlighting the need for clear definitions and enhanced model performance.

Limitations

The different contexts of the datasets and label imbalance, especially in the Counseling dataset, affect the cross-testing results and overall model performance in trauma detection. Label imbalance is particularly challenging because models may become biased towards the more frequent non-trauma events, leading to poorer performance in detecting the less common trauma events. It is normal to have a smaller number of trauma event samples, making it harder for models to learn and accurately identify these underrepresented cases. However, given that the primary goal of this study is to demonstrate cross-domain transferability on realistic data, it is essential to use datasets with an expected and realistic class distribution.

Technical limitations include the summative nature of the explanations, which only provide high-level insights into the different natures of trauma across domains. Additionally, sampling-based explanations such as SLALOM and SHAP are only approximations of the true model behavior, and their fidelity can be increased with more samples, though this incurs higher computational costs.

Another limitation is that people discuss traumatic events differently depending on the context, which might limit the comparability of the datasets used in this study. Conversations with mental health professionals often use clinical terms, focusing on symptoms, triggers, and coping mechanisms (Tong et al., 2019), while online forums blend in-

formal and semi-formal language where anonymity allows for candid sharing, but responses may vary in depth and understanding (Lahnala et al., 2021; Stana et al., 2017). This contrasts with court testimonies, which require precise, factual language focused on specific events and details for legal documentation (Ciorciari and Heindel, 2011; Schirmer et al., 2023b).

We chose the span annotation method, where annotators select the text indicating a traumatic event because pilot experiments showed it improved performance by focusing attention on specific events rather than a simple "yes" or "no" decision. Although this was a design choice and not a central research question, analyzing these spans could offer insights into annotation quality and inform future training. Investigating the detection of specific traumatic event spans rather than general segments is a promising direction for future research.

Finally, our analysis partially relies on social media data. This type of data provides vast, real-time insights into public mental health trends but can be noisy and less reliable. It would be important for future studies to replicate our results with clinical data to ensure the findings' robustness and applicability in medical settings.

Ethics Statement

Our data processing procedures did not involve any handling of private information. No user names were obtained at any point of the data collection process. The human annotators were informed of and aware of the potentially violent content before the annotation process, with the ability to decline annotation at any time. The same is true for crowdworkers, who were presented several trigger warnings throughout the process. Both expert annotators were given the chance to discuss any distressing material encountered during annotation. As discussions on the potential trauma or adverse effects experienced by annotators while dealing with distressing material become more prevalent (Kennedy et al., 2022), we have proactively provided annotators with a recommended written guide designed to aid in identifying changes in cognition and minimizing emotional risks associated with the annotation process.

Acknowledgments

This work was supported by the National Science Foundation under Grant No. IIS-2143529.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Arfan Ahmed, Sarah Aziz, Carla T Toro, Mahmood Alzubaidi, Sara Irshaidat, Hashem Abu Serhan, Alaa A Abd-Alrazaq, and Mowafa Househ. 2022. *Machine learning models to detect anxiety and depression through social media: A scoping review. Computer Methods and Programs in Biomedicine Update*, 2:100066.
- American Psychiatric Association. 2013. *Diagnostic and Statistical Manual of Mental Disorders, 5th Edition*. American Psychiatric Publishing.
- Amod. 2024. [mental_health_counseling_conversations \(revision 9015341\)](#).
- Lukoye Atwoli, Dan J Stein, Karestan C Koenen, and Katie A McLaughlin. 2015. Epidemiology of posttraumatic stress disorder: prevalence, correlates and consequences. *Current opinion in psychiatry*, 28(4):307–311.
- Marcel O Bonn-Miller, Megan Brunstetter, Alex Simonian, Mallory J Loflin, Ryan Vandrey, Kimberly A Babson, and Hal Wortzel. 2022. The long-term, prospective, therapeutic impact of cannabis on post-traumatic stress disorder. *Cannabis and cannabinoid research*, 7(2):214–223.
- Naomi Breslau, EL Peterson, LM Poisson, LR Schultz, and VC Lucia. 2004. Estimating post-traumatic stress disorder in the community: lifetime perspective and the impact of typical traumatic events. *Psychological medicine*, 34(5):889–898.
- Rafael A Calvo, David N Milne, M Sazzad Hussain, and Helen Christensen. 2017. Natural language processing in mental health applications using non-clinical texts. *Natural Language Engineering*, 23(5):649–685.
- John D Ciorciari and Anne Heindel. 2011. Trauma in the courtroom. *Cambodia's hidden scars: Trauma psychology in the wake of the Khmer Rouge. Phnom Penh: Documentation Center of Cambodia (DC-Cam)*.
- Glen Coppersmith, Craig Harman, and Mark Dredze. 2014. Measuring post traumatic stress disorder in twitter. In *Proceedings of the international AAAI conference on web and social media*, volume 8, pages 579–582.
- Munmun De Choudhury and Sushovan De. 2014. Mental health discourse on reddit: Self-disclosure, social support, and anonymity. In *Proceedings of the international AAAI conference on web and social media*, volume 8, pages 71–80.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of NAACL-HLT 2019*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Heather L Dye. 2020. Is emotional abuse as harmful as physical and/or sexual abuse? *Journal of Child & Adolescent Trauma*, 13(4):399–407.
- Atle Dyregrov, Leila Gupta, Rolf Gjestad, and Eugenie Mukanoheli. 2000. Trauma exposure and psychological reactions to genocide among rwandan children. *Journal of traumatic stress*, 13:3–21.
- Matthew J Friedman, Terence M Keane, and Patricia A Resick. 2007. *Handbook of PTSD: Science and practice*. Guilford press.
- Steven N Gold. 2017. *APA handbook of trauma psychology: foundations in knowledge, Vol. 1*. American Psychological Association.
- Zhijun Guo, Alvina Lai, Johan Hilge Thygesen, Joseph Farrington, Thomas Keen, and Kezhi Li. 2024. Large language model for mental health: A systematic review. *arXiv preprint arXiv:2403.15401*.
- Qiwei He, Bernard P Veldkamp, Cees AW Glas, and Theo de Vries. 2017. [Automated assessment of patients’ self-narratives for posttraumatic stress disorder screening using natural language processing and text mining](#). *Assessment*, 24(2):157–172.
- S. Hornstein, J. Scharfenberger, U. Lueken, et al. 2024. [Predicting recurrent chat contact in a psychological intervention for the youth using natural language processing](#). *npj Digital Medicine*, 7:132.
- Thormod Idsoe, Tracy Vaillancourt, Atle Dyregrov, Kristine Amlund Hagen, Terje Ogden, and Ane Nærde. 2021. Bullying victimization and trauma. *Frontiers in psychiatry*, 11:480353.
- Shaoxiong Ji. 2022. Towards intention understanding in suicidal risk assessment with natural language processing. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4028–4038.
- Brendan Kennedy, Mohammad Atari, Aida Mostafazadeh Davani, Leigh Yeh, Ali Omrani, Yehsong Kim, Kris Coombs, Shreya Havaldar, Gwenyth Portillo-Wightman, Elaine Gonzalez, et al. 2022. Introducing the Gab Hate Corpus: defining and applying hate-based rhetoric to social media posts at scale. *Language Resources and Evaluation*, pages 1–30.
- 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 *International Conference on Machine Learning*, pages 2668–2677. PMLR.
- Laurel J Kiser, Jerry Heston, Pamela A Millsap, and David B Pruitt. 1991. Physical and sexual abuse in childhood: Relationship with post-traumatic stress disorder. *Journal of the American Academy of Child & Adolescent Psychiatry*, 30(5):776–783.
- Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. 2020. Concept bottleneck models. In *International Conference on Machine Learning*, pages 5338–5348. PMLR.
- Klaus Krippendorff. 2018. *Content analysis: An introduction to its methodology*. Sage publications.
- Allison Lahnala, Yuntian Zhao, Charles Welch, Jonathan K Kummerfeld, Lawrence An, Kenneth Resnicow, Rada Mihalcea, and Verónica Pérez-Rosas. 2021. Exploring self-identified counseling expertise in online support forums. *arXiv preprint arXiv:2106.12976*.
- Aziliz Le Glaz, Yannis Haralambous, Deok-Hee Kim-Dufor, Philippe Lenca, Romain Billot, Taylor C Ryan, Jonathan Marsh, Jordan Devylder, Michel Walter, Sofian Berrouiguet, et al. 2021. [Machine learning and natural language processing in mental health: systematic review](#). *Journal of Medical Internet Research*, 23(5):e15708.
- Tobias Leemann, Alina Fastowski, Felix Pfeiffer, and Gjergji Kasneci. 2024. [Attention mechanisms don’t learn additive models: Rethinking feature importance for transformers](#). *arXiv preprint arXiv:2405.13536*.
- Tobias Leemann, Michael Kirchof, Yao Rong, Enkelejda Kasneci, and Gjergji Kasneci. 2023. [When are post-hoc conceptual explanations identifiable?](#) In *Uncertainty in Artificial Intelligence*, pages 1207–1218. PMLR.
- Maxwell Levis, Christine Leonard Westgate, Jiang Gui, Bradley V Watts, and Brian Shiner. 2021. Natural language processing of clinical mental health notes may add predictive value to existing suicide risk models. *Psychological medicine*, 51(8):1382–1391.
- Lingyao Li, Lizhou Fan, Shubham Atreja, and Libby Hemphill. 2024. “hot” chatgpt: The promise of chatgpt in detecting and discriminating hateful, offensive, and toxic comments on social media. *ACM Transactions on the Web*, 18(2):1–36.
- 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*.
- Daniel M Low, Laurie Rumker, John Torous, Guillermo Cecchi, Satrajit S Ghosh, and Tanya Talkar. 2020. Natural language processing reveals vulnerable mental health support groups and heightened health anxiety on reddit during covid-19: Observational study. *Journal of medical Internet research*, 22(10):e22635.

- Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- Matteo Malgaroli, Thomas D Hull, James M Zech, and Tim Althoff. 2023. Natural language processing for mental health interventions: a systematic review and research framework. *Translational Psychiatry*, 13(1):309.
- Charles R Marmar, Adam D Brown, Meng Qian, Eugene Laska, Carole Siegel, Meng Li, Duna Abu-Amara, Andreas Tsiartas, Colleen Richey, Jennifer Smith, et al. 2019. [Speech-based markers for post-traumatic stress disorder in us veterans](#). *Depression and Anxiety*, 36(7):607–616.
- Daniel Matter, Miriam Schirmer, Nir Grinberg, and Jürgen Pfeffer. 2024. Investigating the increase of violent speech in incel communities with human-guided gpt-4 prompt iteration. *Frontiers in Social Psychology*, 2:1383152.
- Laura Ann McCloskey and Marla Walker. 2000. Post-traumatic stress in children exposed to family violence and single-event trauma. *Journal of the American Academy of Child & Adolescent Psychiatry*, 39(1):108–115.
- Oshin Miranda, Sophie Marie Kiehl, Xiguang Qi, M Daniel Brannock, Thomas Kosten, Neal David Ryan, Levent Kirisci, Yanshan Wang, and LiRong Wang. 2024. Enhancing post-traumatic stress disorder patient assessment: leveraging natural language processing for research of domain criteria identification using electronic medical records. *BMC medical informatics and decision making*, 24(1):1–14.
- Jiaxin Pei, Aparna Ananthasubramaniam, Xingyao Wang, Naitian Zhou, Apostolos Dedeloudis, Jackson Sargent, and David Jurgens. 2022. Potato: The portable text annotation tool. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*.
- Steve Powell, Rita Rosner, Willi Butollo, Richard G Tedeschi, and Lawrence G Calhoun. 2003. Posttraumatic growth after war: A study with former refugees and displaced people in sarajevo. *Journal of clinical psychology*, 59(1):71–83.
- Robin Quillivic, Frédérique Gayraud, Yann Auxéméry, Laurent Vanni, Denis Peschanski, Francis Eustache, Jacques Dayan, and Salma Mesmoudi. 2024. Interdisciplinary approach to identify language markers for post-traumatic stress disorder using machine learning and deep learning. *Scientific reports*, 14(1):12468.
- Andrew G Reece, Andrew J Reagan, Katharina LM Lix, Peter Sheridan Dodds, Christopher M Danforth, and Ellen J Langer. 2017. Forecasting the onset and course of mental illness with twitter data. *Scientific reports*, 7(1):13006.
- Ramit Sawhney, Harshit Joshi, Saumya Gandhi, and Rajiv Shah. 2020. A time-aware transformer based model for suicide ideation detection on social media. In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 7685–7697.
- Miriam Schirmer, Isaac Misael Olguín Nolasco, Edoardo Mosca, Shanshan Xu, and Jürgen Pfeffer. 2023a. [Uncovering trauma in genocide tribunals: An nlp approach using the genocide transcript corpus](#). In *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, pages 257–266.
- Miriam Schirmer, Jürgen Pfeffer, and Sven Hilbert. 2023b. [Talking about torture: A novel approach to the mixed methods analysis of genocide-related witness statements in the khmer rouge tribunal](#). *Journal of Mixed Methods Research*, page 15586898231218463.
- Neil J Smelser et al. 2004. Psychological trauma and cultural trauma. *Cultural trauma and collective identity*, 4:31–59.
- Alexandru Stana, Mark A Flynn, and Eugenie Almeida. 2017. Battling the stigma: Combat veterans’ use of social support in an online ptsd forum. *International Journal of Men’s Health*, 16(1).
- Akshay Swaminathan, Iván López, Rafael Antonio Garcia Mar, Tyler Heist, Tom McClintock, Kaitlin Caoili, Madeline Grace, Matthew Rubashkin, Michael N Boggs, Jonathan H Chen, et al. 2023. Natural language processing system for rapid detection and intervention of mental health crisis chat messages. *NPJ Digital Medicine*, 6(1):213.
- Vankayala Tejaswini, Korra Sathya Babu, and Bibhuddatta Sahoo. 2024. Depression detection from social media text analysis using natural language processing techniques and hybrid deep learning model. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 23(1):1–20.
- Lenore C Terr. 2003. Childhood traumas: An outline and overview. *Focus*, 1(3):322–334.
- Janet Tong, Katrina Simpson, Mario Alvarez-Jimenez, and Sarah Bendall. 2019. Talking about trauma in therapy: Perspectives from young people with post-traumatic stress symptoms and first episode psychosis. *Early intervention in psychiatry*, 13(5):1236–1244.
- Mohammad Arif Ul Alam and Dhawal Kapadia. 2020. [Laxary: A trustworthy explainable twitter analysis model for post-traumatic stress disorder assessment](#). In *2020 IEEE International Conference on Smart Computing (SMARTCOMP)*, pages 308–313.
- Bessel A Van der Kolk. 2003. *Psychological trauma*. American Psychiatric Pub.

Xuhai Xu, Bingsheng Yao, Yuanzhe Dong, Saadia Gabriel, Hong Yu, James Hendler, Marzyeh Ghassemi, Anind K Dey, and Dakuo Wang. 2024. Mental-llm: Leveraging large language models for mental health prediction via online text data. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 8(1):1–32.

Kailai Yang, Tianlin Zhang, Ziyang Kuang, Qianqian Xie, Jimin Huang, and Sophia Ananiadou. 2024. Mental-lama: interpretable mental health analysis on social media with large language models. In *Proceedings of the ACM on Web Conference 2024*, pages 4489–4500.

Chih-Kuan Yeh, Been Kim, Sercan O Arik, Chun-Liang Li, Tomas Pfister, and Pradeep Ravikumar. 2019. On completeness-aware concept-based explanations in deep neural networks. In *Advances in Neural Information Processing Systems*, volume 32.

Rachel Yehuda. 1998. *Psychological trauma*. American Psychiatric Pub.

Tianlin Zhang, Annika M Schoene, Shaoxiong Ji, and Sophia Ananiadou. 2022. [Natural language processing applied to mental illness detection: a narrative review](#). *NPJ Digital Medicine*, 5(1):46.

A Implementation Details

A.1 Implementation Details: Explanation Methods

In this section, we give more details on how we computed the explanations shown in this paper.

SHAP Values. To obtain SHAP values, we use the official shap² package. We use the TextExplainer class.

SLALOM. We use the SGD algorithm proposed in [Leemann et al. \(2024\)](#) to estimate the SLALOM model on 100k background samples of length 2. We use all the tokens that appear in the samples from the datasets used and fit one global SLALOM model.

Conceptual Explanations. We use the completeness-aware loss proposed by [Yeh et al. \(2019\)](#) with snippets of length of 5 token as snippets for the algorithm. We trained with the concept discovery module to discover $K = 10$ concepts using the Adam optimizer at an initial learning rate of 1×10^{-3} , decaying to 5×10^{-4} and 1×10^{-4} in subsequent epochs. Training lasted three epochs with a batch size of 12. The model weights used were obtained from the best-performing model. We identified the 25 closest activations per concept.

²<https://github.com/shap/shap>

Evaluation on a separate test set involved dot products between latent representations and concept vectors, selecting the top activations.

A.2 Implementation Details: Models

We use the optuna³ framework for hyperparameter optimization with 50 steps for each model/dataset. We then train the models using different seeds and on five random data splits using the discovered hyperparameters. Through the optimization, we obtain the parameters given in Table 4.

Prompt Template. We use the following prompt template to prompt the GPT models as the system prompt.

"You are tasked with detecting trauma in text segments of transcripts of genocide tribunals. Specifically, detect instances that meet the APA's definition of trauma. Psychological trauma, as defined by the APA, includes experiences of exposure to actual or threatened death, serious injury, or sexual violence, either directly encountered or witnessed. It also includes instances where individuals learn that the traumatic event(s) occurred to a close family member or friend. Label the text with '1' if there are indicators of trauma based on this definition and '0' if there are no indicators of trauma. Note that trauma is rare and occurs in less than 20% of the cases. Only answer with either '0' or '1'."

The samples are then passed as a user prompt.

B Annotation Details

Participants were prescreened using Prolific based on self-reported English-language proficiency. We did not collect demographic data from the annotators as such data was not central to the questions our study is focused on, and Prolific does not normally include this metadata. Figure 6 and 7 in the Appendix further illustrate both annotator instructions and the layout of the span annotation task.

C Metrics

For completeness, we additionally report accuracy, recall, and precision for the trained models in Table 6. To further study the effect of dataset size, we conducted an additional experiment based on the Genocide Transcript Data that compares the performance using only 1,000 (randomly selected) samples in the training dataset. Table 7 displays the ROC-AUC scores and shows that values do not differ significantly even with a smaller dataset size,

³<https://optuna.org/>

Model	Parameters		
	GTC	PTSD	Counseling
NaiveBayes-BoW	multiplicities: true alpha: 1.01	multiplicities: true alpha: 5.97	multiplicities: false alpha: 1.01
NGramLogisticRegression	n_gram_range: [1, 2] C: 0.92 penalty: l2	n_gram_range: [2, 3] C: 0.0 penalty: none	n_gram_range: [1, 2] C: 9.36 penalty: l2
FeedForwardModel	hidden_dim1: 50 hidden_dim2: 80 lr: 5.72e-05	hidden_dim1: 50 hidden_dim2: none lr: 1.79e-04	hidden_dim1: 200 hidden_dim2: 50 lr: 5.72e-05
BERT (finetuned)	n_layers: 5 lr: 2.32e-05	n_layers: 12 lr: 1.10e-05	n_layers: 6 lr: 1.41e-05
RoBERTa (finetuned)	n_layers: 12 lr: 2.04e-06	n_layers: 7 lr: 6.43e-06	n_layers: 4 lr: 9.54e-05
OpenAI	target_model: gpt-4-turbo	target_model: gpt-4-turbo	target_model: gpt-4-turbo

Table 4: Automatically selected hyperparameters for the different datasets

Dataset	Test Dataset			
	GTC	PTSD	Counsel.	Incels
GTC	0.967 ± 0.000	0.734 ± 0.005	0.812 ± 0.020	0.847 ± 0.003
PTSD	0.885 ± 0.010	0.830 ± 0.006	0.872 ± 0.014	0.894 ± 0.010
Counsel.	0.740 ± 0.017	0.738 ± 0.018	0.881 ± 0.016	0.725 ± 0.027
All	0.966 ± 0.001	0.833 ± 0.013	0.922 ± 0.012	0.878 ± 0.005

Table 5: Cross-Testing models trained on one dataset on other datasets. Model: RoBERTa finetuned with AU-ROC metric

suggesting that our findings obtained on smaller dataset sizes may be valid for larger dataset sizes as well.

Model	Accuracy	Precision	Recall
GTC			
NaiveBayesBOWmodel	0.84 ± 0.03	0.44 ± 0.08	0.69 ± 0.12
NGramLogisticRegression	0.88 ± 0.02	0.60 ± 0.12	0.44 ± 0.09
FeedForwardModel	0.88 ± 0.02	0.60 ± 0.12	0.46 ± 0.09
BERTmodel	0.88 ± 0.03	0.58 ± 0.12	0.46 ± 0.10
RoBERTamodel	0.91 ± 0.00	0.70 ± 0.02	0.59 ± 0.05
BERTPretrainedmodel	0.92 ± 0.00	0.74 ± 0.03	0.70 ± 0.04
RoBERTaPretrainedmodel	0.93 ± 0.00	0.75 ± 0.03	0.74 ± 0.04
OpenAI GPT-4	0.91	0.68	0.61
PTSD			
NaiveBayesBOWmodel	0.69 ± 0.01	0.63 ± 0.02	0.52 ± 0.06
NGramLogisticRegression	0.68 ± 0.01	0.62 ± 0.03	0.54 ± 0.03
FeedForwardModel	0.70 ± 0.01	0.71 ± 0.04	0.42 ± 0.05
BERTPretrainedmodel	0.72 ± 0.01	0.64 ± 0.01	0.69 ± 0.06
RoBERTaPretrainedmodel	0.75 ± 0.01	0.66 ± 0.02	0.78 ± 0.04
OpenAI GPT-4	0.69	0.58	0.84
Counseling			
NaiveBayesBOWmodel	0.26 ± 0.01	0.09 ± 0.01	0.99 ± 0.01
NGramLogisticRegression	0.92 ± 0.01	0.55 ± 0.17	0.09 ± 0.03
FeedForwardModel	0.92 ± 0.01	0.10 ± 0.10	0.02 ± 0.02
BERTPretrainedmodel	0.93 ± 0.01	0.54 ± 0.04	0.27 ± 0.05
RoBERTaPretrainedmodel	0.91 ± 0.01	0.36 ± 0.19	0.20 ± 0.12
OpenAI GPT-4	0.91	0.42	0.31

Table 6: Additional model performance metrics. We see that the non-pretrained versions of BERT/RoBERTa do not perform on par with the pretrained ones on GTC. Therefore, we consider only the pretrained versions for the rest of the paper.

LM	GTC-1000		GTC-All	
	F1 (bin.)	AU-ROC	F1 (bin.)	AU-ROC
FeedForwardModel	0.38 ± 0.01	0.86 ± 0.00	0.52 ± 0.10	0.84 ± 0.09
BERTPretrainedmodel	0.61 ± 0.03	0.93 ± 0.00	0.71 ± 0.01	0.96 ± 0.00
RoBERTaPretrainedmodel	0.66 ± 0.03	0.95 ± 0.00	0.74 ± 0.01	0.97 ± 0.00

Table 7: Additional experiments with a smaller GTC (Genocide Transcript Corpus) sample size to control for dataset size effects.

Overview

Instructions:

This project aims to detect and understand how individuals express trauma in different contexts in various online platforms, such as online forums, blogs, and social media posts.

You will be presented with short text snippets from different online sources and asked to highlight if there is a TRAUMATIC EVENT being described in the text. If you think a text contains a traumatic event, click on "trauma" and then highlight the event in the text. You can also mark a text span with the "uncertain" option. If there is NO TRAUMATIC EVENT, MOVE FORWARD without selecting anything. For the best experience, please ensure your browser window is maximized to full screen. You can use the left and right arrow keys to navigate between snippets.

Trigger Warning: Please be aware that the texts you will review may contain references to traumatic events, which could be distressing. If you feel uncomfortable at any point, you are free to take a break or discontinue your participation in the study.

Note: "trauma" is generally LESS FREQUENT than "no trauma"

General references to depression or anxiety DO NOT COUNT AS TRAUMA when there is no traumatic event. The event also needs to be related to the person who has written the text.

Examples of Traumatic Events:

"I witnessed my mother have a heart attack when I was a child and she had to be hospitalized for a very long time."

"I got into a really bad accident about 9 years ago, and after nearly causing another accident a few days later after swerving when I saw an oncoming car I gave up driving. For about 2 years I could hardly get into a car without panicking."

"But I didn't even realize how fucked certain "minor" things were. Like being forced to wake up at 6 am when I got a cold or how I wasn't allowed to even sleep during "reflection time" which was basically a 3 hour time out style punishment for breaking rules or trash talking the facility, or refusing to eat the expired food."

Move forward

Figure 6: Instructions for Annotators. Note: We selected these examples because they were the most frequently mislabeled in the pilot, making them particularly relevant. Additionally, we kept the instruction page concise to avoid overwhelming the annotators, as excessive detail could deter them or lead to less careful reading.


How do you self-soothe? Tw: childhood sexual abuse & neglect Long story short, the sexual abuse happened around 4-5 years of age and the emotional neglect went on for over a decade after. I have amnesia from the sexual trauma, but am currently dealing with crippling night terrors that are confusing me as to what I thought happened. I don't know if they're flashbacks, I know they're dreams, but they're traumatic in themselves because of the disgusting images I'm left with when I wake up. It's horrifying and overwhelming. I recently made huge headway in therapy and have finally learned to stop dissociating/going numb, something I've done for my entire life. Now that I can feel, I'm drowned in emotion, sadness, anger, hurt, but that's good because I'm finally really feeling everything. It's the not being able to calm myself down and help myself feel safe again that's the problem. I'm going to see a psychiatrist to see if I can get meds for the chronic nightmares (it's been an entire year now), but in the meantime I don't know how to calm myself when I'm emotionally overwhelmed. No one taught me how to properly soothe myself and all I've ever known is resorting to numbness. So here's my question, friends: how do you self soothe when you're overwhelmed and can't stop sobbing? How do you take care of yourself and help yourself feel safe? Thank you for reading~

Does this text contain a traumatic event? Click on "trauma" and highlight the traumatic event or on "uncertain" if you are uncertain. Skip if there is no such event. Traumatic events are defined as "exposure to actual or threatened death, serious injury, sexual violence, or severe mental abuse," whether directly encountered or witnessed. It also counts as a traumatic event if something like that has happened to a close family member or friend.


trauma
 uncertain

Move backward Move forward


Figure 7: Interface of the Span Annotation Task.



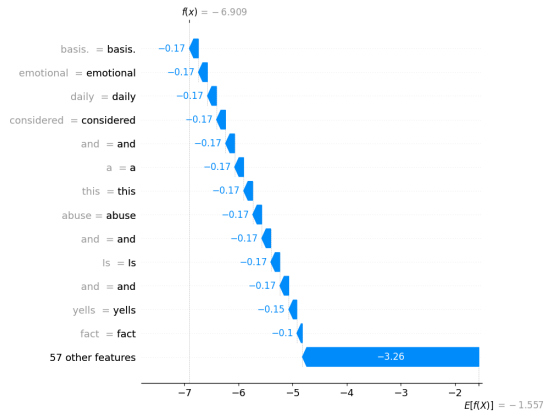
(a) RoBERTa - Genocide Transcript Corpus




(b) GPT-4 - Genocide Transcript Corpus




(c) RoBERTa - Counseling Dataset (Instance 1)




(d) GPT-4 - Counseling Dataset (Instance 1)




(e) RoBERTa - Counseling Dataset (Instance 2)



(f) GPT-4 - Counseling Dataset (Instance 2)



(g) RoBERTa - PTSD Dataset



(h) GPT-4 - PTSD Dataset

Figure 8: SHAP Values for various instances from different datasets. See Table 8 for the full text of each instance.

Dataset	Instance
Genocide Transcript Corpus	I can feel that the person committed any wrongdoing would be burned alive, and I would also see that one day if I committed any wrongdoing I would experience the same fate.
Counseling Dataset (Instance 1)	My dad doesn't like the fact that I'm a boy. He yells at me daily because of it and he tells me I'm extreme and over dramatic. I get so depressed because of my dad's yelling. He keeps asking me why I can't just be happy the way I am and yells at me on a daily basis. Is this considered emotional abuse?
Counseling Dataset (Instance 2)	I was raped by multiple men, and now I can't stand the sight of myself. I wear lingerie to get my self excited enough to have sex with my wife.
PTSD Dataset	It's nearly been 4 years (trigger warning) It's almost been 4 years since he died. I can't look at hospitals without the memories coming back. Seeing him half dead. His body was all sorts of fucked up. I can't deal with this any longer. I'm going to go insane. Every day it gets worse.

Table 8: Instances from various datasets used for SHAP value analysis (see Figure 8).

<p style="text-align: center;">Concept 2</p> <p>some cows were wounded. or diarrhoea. supply themselves with food. so itchy everywhere. bullet in the head; those various gunshots. once I was wounded. , one hit me in three bursts of gunfire. themselves from the bullets. to work I kept crying and killed animals for wedding the gunshots and we were I even saw dead bodies I cry every night.</p>	<p style="text-align: center;">Concept 4</p> <p>and when he attacked me chief, was very cruel I was punished that way He pressed me against His disappearance was very painful Bou Meng was tortured for who tortured me was Si and then my eyes would olded me for being blinded so I had him buried , they stopped beating me who tore the child away all the beatings that They started beating me, task of killing people.</p>	<p style="text-align: center;">Concept 6</p> <p>long you fell unconscious. falling into the ground. another place two bodies. them were lying on the weapon in my mouth. I started freezing, and up the dead bodies and us drowned in the river a lot of dead and people who perished there. children die of hunger, then we all got stuck people and throwing their bodies her sister burned to death who died of hunger,</p>	<p style="text-align: center;">Concept 7</p> <p>December he called me and my husband called cadres sometimes he called me to his phone call because he or called me and then was called Lucia, Ruk know what happened to Ph man called Rukara went since I was busy looking school was called Hasan Ve really is still in my one called me up and my wife asked them what tell you exactly when this It arrived in Kosovo Pol</p>
--	---	---	---

(a) firearms, physical injuries (b) torture, abuse (c) death, motionless bodies (d) communication, asking questions (non-traumatic)

(a) Concepts in the **GTC Dataset** (Examples, RoBERTa Model): (a)-(c) Trauma-Related, (d) Non-Trauma Related

<p style="text-align: center;">Concept 0</p> <p>in the dream with the In this dream. awake in the dream. I survived veteran suicide. Or in a dream rel toxic relationships with men. .500 military sexual assaults shot and killed himself in ago I had a dream only dated for a few after having a dream that harmed while in a was a dream. This shot herself in the head , on a medical discharge</p>	<p style="text-align: center;">Concept 5</p> <p>. Our car clipped the his car multiple times. my skull. It was raped at least three times hit another car, causing 22 year old man (were later found and her a lot of death. SA and death. of sexual abuse, two murder-suicides each tried to apprehend the two /SuicideWatch and after service. Fires are experiences with SSRIs</p>	<p style="text-align: center;">Concept 6</p> <p>my parents were kids, my childhood over to my my mom divorced my violent I was a child. had an abusive ex and my mother died 2 years memories of my childhood. abusive and cheating gay ex first serious boyfriend repeatedly realize TW: childhood sexual my trauma from childhood. my now ex of 4 My parents were divorced when my parents raised me and throughout my childhood by my</p>	<p style="text-align: center;">Concept 9</p> <p>extremely frequent flashbacks the me bad. The flashbacks When I was molested have vivid flashbacks. All aphobia in the last about flashbacks. I phobia related to the after i was sexually assaulted young child, was sexually having nightmares and flashbacks repressed memories are a because I have flashbacks several always thought the memories of scolding via email like I was abused.</p>
--	---	--	---

(e) dreams, military, suicide (f) accidents, death, sexual abuse (g) childhood abuse, family trauma (h) flashbacks, memories, sexual abuse, phobia

(b) Concepts Found in the **PTSD Dataset** (Examples, RoBERTa Model): (a)-(d) Trauma-Related

<p style="text-align: center;">Concept 4</p> <p>My boyfriend lost his dad my mom. My dad shot and killed my rapist His mom, my grandma My dad cheated on my I lost my mother recently his wife. My uncle , my mother and father I was kidnapped at fourteen . My father cheated on My ex-wife married niece whom my sister abandoned My daughter was overly tired my mom. Years later misses mom and dad in</p>	<p style="text-align: center;">Concept 5</p> <p>My boyfriend lost his dad my mom. My dad shot and killed my rapist his wife. My uncle His mom, my grandma My dad cheated on my I lost my mother recently , my mother and father . My father cheated on My ex-wife married I was kidnapped at fourteen Post Traumatic Stress Disorder Post Traumatic Stress Disorder niece whom my sister abandoned posttraumatic stress disorder</p>	<p style="text-align: center;">Concept 6</p> <p>My boyfriend lost his dad my mom. My dad abusive father and his wife his wife. My uncle His mom, my grandma My dad cheated on my I lost my mother recently , my mother and father I was kidnapped at fourteen My sister never defended me , my doctor gave me . He bought me a all of my family left broken apart after she got rehabilitation program and got kicked</p>	<p style="text-align: center;">Concept 9</p> <p>I was violently raped by got pregnant by my boyfriend my baby mother. She my children's father left I saw my mother cheating I was raped by multiple My girlfriend was abused as I got raped by my I just lost my mom teenager. My entire family I was raped repeatedly when My grandma and brother both parents injured my brother, , my husband mentally abused My mother has Alzheimer's</p>
--	--	---	---

(i) family members: cheating, loss (j) PTSD (overlap Concept 4) (k) clinical context, dependencies (overlap Concept 4) (l) rape, abuse, pregnancy

(c) Concepts Found in the **Counseling Dataset** (Examples, RoBERTa Model): (a)-(d) Trauma-Related

Figure 9: Examples of Concepts Discovered on Various Datasets for the RoBERTa Model: (a) GTC dataset, (b) PTSD dataset, (c) Counseling dataset.