

If there’s a Trigger Warning, then where’s the Trigger? Investigating Trigger Warnings at the Passage Level

Matti Wiegmann¹ Jennifer Rakete² Magdalena Wolska¹ Benno Stein¹ Martin Potthast³

¹Bauhaus-Universität Weimar ²Leipzig University ³University of Kassel, hessian.AI, and ScaDS.AI

Abstract

Trigger warnings are labels that preface documents with sensitive content if this content could be perceived as harmful by certain groups of readers. Since warnings about a document intuitively need to be shown before reading it, authors usually assign trigger warnings at the document level. What parts of their writing prompted them to assign a warning, however, remains unclear. We investigate for the first time the feasibility of identifying the triggering passages of a document, both manually and computationally. We create a dataset of 4,135 English passages, each annotated with one of eight common trigger warnings. In a large-scale evaluation, we then systematically evaluate the effectiveness of fine-tuned and few-shot classifiers, and their generalizability. We find that trigger annotation belongs to the group of subjective annotation tasks in NLP, and that automatic trigger classification remains challenging but feasible.

Warning. This paper shows example text relating to *Death*.

1 Introduction

Online content is considered harmful if it has a negative emotional, psychological, or physical effect on people (Kirk et al., 2022). This applies regardless of whether the effect is intentional (as with hate speech) or unintentional. Since the publication of harmful content is in some cases justified or even important (as with news), affirmative action seems warranted. Several (online) communities have therefore developed a new type of affirmative action borrowed from trauma therapy: trigger warnings. Trigger warnings are freeform labels prefaced to a document by its author to indicate that it contains potentially harmful content. Commonly used warnings range from harmful concepts that can affect any individual, such as *Aggression*, *War*, or *Death*, to those that only affect certain groups of people, such as *Discrimination* and its

sub-concepts *Misogyny*, *Racism*, or *Homophobia*. Given the availability of large-scale resources of author-labeled documents with triggering content, the task of automatic trigger warning assignment has recently been picked up (Section 2).

Trigger warnings usually do not specify where the triggering content occurs in a document. This poses both theoretical and practical questions for the manual and automatic annotation or assignment of trigger warnings to text documents. For example, a concept such as *Death* may not be addressed throughout an entire document, but only occasionally. From an application perspective, document-level warnings prevent readers from reading the non-triggering parts of a document, which may undermine harm mitigation (Bridgland et al., 2023). From a linguistic perspective, triggers are at the pragmatic level of discourse and depends on context: Figure 1 shows two examples addressing *Death*, their context influences annotator perception. Author-supplied ground truth, too, may hence comprise label noise, warranting closer inspection.

In this paper, we lay the foundation to investigate if and how trigger warnings can be reliably annotated and assigned at the passage level: (i) We conduct a large annotation study to understand the challenges of trigger annotation, resulting in a dataset of 4,135 English five-sentence passages from the Webis Trigger Warning Corpus 2022 (Wiegmann et al., 2023), annotated with eight triggers (Section 4). (ii) We then systematically evaluate state-of-the-art classifiers in assigning various trigger warnings and analyze their behavior regarding training data availability, label subjectivity, and generalization to unseen concepts (Section 5).¹ We find that trigger warning annotation belongs to the subjective annotation tasks in NLP, and that classifying triggering passages requires a careful choice of the right model per warning (Section 6).

¹All code and data are shared for reproducibility: github.com/MattiWe/passages-level-trigger-warnings

2 Related Work

Trigger warnings originate from trauma therapy (Knox, 2017) and have been studied in both clinical and educational contexts. A recent meta-study by Bridgland et al. (2023) provides an overview of the state of the art in these areas. The first computational trigger warning assignment approach (Stratta et al., 2020) investigated an interaction design in a user study using a browser plugin (DeText) on generic websites, which was limited to a dictionary-based assignment of the *Sexual Assault* warning. Subsequently, Wolska et al. (2022) examined the binary document classification of fan fiction documents using the trigger *Graphic Violence*. A taxonomy for multimedia triggers is the Narrative Experiences Online (NEON) taxonomy proposed by Charles et al. (2022), which contains 90 labels based on 136 guides from the web. Wiegmann et al. (2023) later presented a 36-label taxonomy of trigger warnings, based on academic guidelines and a large-scale annotated corpus of fan fiction documents. Since the former does not provide annotated works, we draw on the latter’s taxonomy and data. All these works consider the assignment of trigger warnings at the document level, while our focus is on the assignment of warnings at the passage level.

The classification of triggers also relates to that of other harmful content, such as toxic comments on Wikipedia articles (Wulczyn et al., 2017; Adams et al., 2017), verbal violence in YouTube and Reddit comments (Mollas et al., 2020), and the taxonomy of harmful online content (Banko et al., 2020), which overlaps with the above-mentioned taxonomies in the verbal categories. These works differ mainly in the delineation of harmfulness: who is affected (everyone, groups, or individuals), whether the harm is intentional or collateral, how the content is conveyed (in online speech such as microblogs or chats, or in long-form texts, such as comments, blogs, or narratives). The distinctive feature of trigger warnings is that they focus primarily on collateral harms to individuals, regardless of the form of transmission. Kirk et al. (2022) present a more comprehensive discussion of harmful (online) texts and their various aspects; we have adopted their recommendations for annotation.

As our research shows, another common feature of triggers and other harmful content is label subjectivity and annotator disagreement. In Sandri et al.’s 2023 taxonomy of causes of disagree-

The moans and cries of the damned assaulted one’s ears. The stench of rotting and burning flesh assaulted one’s nose. *The disfigurement of each hapless undead slave, some missing limbs, covered in blood and ooze, some naked, some with their skin missing, and more assaulted one’s eyes.* Dark Magic, claws, and weapons assaulted one’s body. And finally, the horror of it all assaulted one’s mind.

Annotation: positive instance

Zach would move onto whoever he met at the next frat party. His friends warned him beforehand that Jack is kinda weird. *Jack was in Corbyn’s anatomy class and a cadaver has gone missing recently.* Of course, it was pure coincidence despite Corbyn insisting otherwise. It was just an unfortunate mishap, and it wouldn’t deter Zach from his goal tonight.

Annotation: negative instance

Figure 1: Two example passages from fan fiction stories annotated for the trigger warning *Death*. The upper example was unanimously annotated as ‘positive’, the lower as ‘negative’. The central sentence in italics was retrieved via keywords; preceding and following sentences serve as context. Figure 3 shows more examples.

ment in detecting offensive language use, ‘Sloppy Annotation’ (which we mitigate by monitoring annotators and reducing task complexity to binary annotations) and ‘Missing Information’ (which we mitigate by using passages instead of sentences) are mentioned as possible causes, but also the hard-to-avoid ‘Subjectivity’ and ‘Ambiguity’. (Röttger et al., 2022) note that there are several valid beliefs about labeling harmful content and that the “descriptive” annotation paradigm used in academic work (including ours) attempts to capture all of these beliefs, leading to disagreement among annotators. They suggest isolating a particular belief using a “prescriptive” annotation paradigm when desirable for subsequent application. Another approach is learning with disagreement, as studied in the shared task by that name (Uma et al., 2021) as well as that on ‘sexism recognition’ (Plaza et al., 2023), captures multiple beliefs and models them explicitly. Similarly, Davani et al. (2022) study multitask learning over multiple annotators’ votes to classify hate speech and emotions without losing effectiveness.

(a)					(b)					(c)				
Warning	Passages		Keyword		Warning	Number of Positive Votes				Time	α	Warning	Train	
	num	len	src	clean		0	1	2	3				ID	OOD
Violence	1,041	92	29	28	Violence	198(53%)	98(26%)	60(16%)	21(6%)	39	0.41	Violence	1,001	675
Death	544	83	122	50	Death	82(31%)	60(22%)	37(14%)	88(33%)	30	0.36	Death	504	178
War	827	95	38	25	War	86(34%)	82(32%)	53(21%)	34(13%)	18	0.22	War	787	202
Abduction	511	83	20	18	Abduction	97(31%)	78(25%)	97(31%)	41(13%)	29	0.25	Abduction	471	252
Racism	267	90	43	37	Racism	310(56%)	120(22%)	71(13%)	43(8%)	29	0.52	Racism	227	68
Homophobia	313	79	38	27	Homophobia	678(65%)	175(17%)	119(11%)	69(7%)	27	0.24	Homophobia	273	129
Misogyny	377	84	62	55	Misogyny	238(47%)	141(28%)	94(18%)	38(7%)	31	0.25	Misogyny	337	152
Ableism	255	83	47	31	Ableism	545(66%)	180(22%)	83(10%)	19(2%)	24	0.25	Ableism	215	162
Total	4,135	88	399	271	Total	2,234(54%)	935(23%)	613(15%)	353(9%)	33	0.35	Total	3,815	1,818

Table 1: Descriptive dataset statistics: **(a)** Number and length (words) of the passages and the number of phrases in the source and cleaned keyword lists. **(b)** Number and ratio of examples with n positive votes, average annotation time per passage in seconds, and Krippendorff’s α inter-annotator agreement. **(c)** Number of passages in the training and test splits for the in-distribution (ID) and out-of-distribution (OOD) settings.

3 Task Design

Our approach to trigger annotation is the result of several small test runs and pilot studies that provided us with the insights to make design decisions for our annotation task. From these studies, we derived the following three key constraints for the task of trigger annotation at the passage level:

Trigger Diversity The trigger warnings used in practice are based on the personal opinions and experiences of many authors and thus very diverse. The WTWC-22 organizes them into seven categories containing 36 warnings, encompassing hundreds of thousands of variants. Annotating a given passage with all 36 warnings, let alone all variants, has proven infeasible at scale under reasonable budget constraints. We therefore select eight common warnings for annotation, four each from the two most frequently assigned warning categories *Aggression* and *Discrimination* of the WTWC-22 taxonomy. From *Aggression* we use all four warnings *Death*, *Violence*, *Abduction* and *War*. From *Discrimination* we use the four most frequently assigned warnings *Misogyny*, *Racism*, *Homophobia* and *Ableism*. The former relate primarily to physical harm, the latter to psychological harm.

Trigger Sparsity Harmful passages are often sparsely distributed over long documents,² leading to class imbalance between positive and negative cases. We therefore resort to the commonly used approach of dictionary-based retrieval to obtain a sufficient number of positive candidate passages as detailed below.³ This entails selection bias, since

²E.g., many fan fiction works are as long as books.

³Retrieval-guided annotation have previously been employed to create NLP resources.

very subtle cases of triggering content may not contain any explicit mention of one of the dictionary’s phrases, and since some important phrases might be missing from our dictionaries, both leading to false negatives. We analyze this possibility with a targeted out-of-distribution experiment.

Trigger Severity Authors and readers are not equally sensitive to a potentially triggering concept (see Section 4.4). A reader who enjoys horror may find fewer references to *Death* in fiction triggering than readers who do not. For instance, the positive example in Figure 1 might be considered enticing rather than harmful. Therefore, when annotators are asked to make a descriptive binary decision, they often disagree. This is not uncommon in work on harmful content and, according to Röttger et al. (2022), even desirable in an initial study like ours. However, this presents difficulties when assessing annotations and annotators, as inter-annotator agreement measures become less reliable due to different reader sensitivity to a particular trigger. This prevents the use of crowdsourcing platforms for annotation, as we cannot (i) quantitatively evaluate annotators, (ii) train them reliably ahead of time, and (iii) provide appropriate support during annotation. We therefore recruited local annotators and provided them with personal support.

4 Dataset Construction

We constructed a dataset of 4,135 passages with five consecutive sentences, each annotated with three binary human labels for one of eight selected trigger warnings. All passages originate from the Webis Trigger Warning Corpus 2022 (WTWC-22) (Wiegmann et al., 2023), which compiles fan fiction documents with document-level trigger warn-

Persona	You have suffered through <i>death</i> in the past and want to be warned before reading text that contains <i>death</i> .
Definition	<i>Death includes graphic or implied descriptions of dead people or creatures, actions leading to death, experiences of dying, and descriptions of grief.</i> A Warning is required if the text evokes a negative experience through the actions or speech of characters, graphic descriptions, or through the narrative. A Warning is required if a vulnerable reader is harmed by reading the text.
Demo.	Consider the following examples: Text: <i>He laid the cadaver back down, uselessly recalling every single thing he learned in first aid in an ...</i>
Instruction	Assign ‘Warning: yes’ if the following story piece needs a trigger warning for death, else assign ‘Warning: no’.
Passage	Text: <i>The stench of rotting and burning flesh assaulted one’s nose. The disfigurement of each ...</i>

Figure 2: **Warning Death.** Structure of both the annotation instructions and the prompts used in the few-shot experiments. Shown here is an example passage for the Death warning. The italics sections vary by instance.

ings. Figure 1 shows example passages and Table 1 overviews our dataset.

4.1 Retrieving Triggering Passage Candidates

We collected the passages for annotation using a keyword-based retrieval approach. We first constructed a keyword list for each of the eight considered warnings, then retrieved matching documents from WTWC-22 using a BM25-based initial retrieval⁴ where the keywords served as query terms. From each document, we extracted the first sentence with a keyword match and added the two pre- and succeeding sentences as context. Adding context is often necessary understand the central sentence. We did not use the original (paragraph) segmentation of the documents, since they greatly varied in length, particularly since dialog turns are often written as single paragraph. The resulting five-sentence passages were de-duplicated across warnings and annotated in order of the BM25 score of the originating document.

Keyword List Construction. We build a list of keywords and phrases for each of the eight trigger warnings by prompting gpt-3.5-turbo-0301 using the following prompt with slight morpho-syntactic adaptations for each warning:

Provide a list of <warning> language that must be avoided in a story for people that are triggered by <warning>. Return the list in JSON format.

We cleaned this initial list manually to remove redundant (i.e. lexically similar) phrases, phrases that better match a different warning, or ambiguous phrases that produce many false positives, reducing the initial keywords by ca. 32%. Table 1a shows the number of words and phrases on each warning’s list, before and after cleaning. We split each of the

cleaned lists into two sets of keywords (in the middle according to the order returned by the model). The set identity is used to distinguish between in and out-of-distribution examples (cf. Section 5.1). Table 3 (A.1) illustrates the list construction process for the *Death* and *Misogyny* warnings.

4.2 Passage Annotation

Figure 1 shows two passages, one for each annotation decision. The collected passages were annotated as either positive (requires a warning) or negative (does not require a warning) by three different annotators each. We decided on the final annotation based on the majority of the votes.

Instead of annotating a fixed number of passages for each warning, the first annotator of each warning rated continuously, in order of the BM25 score of the source document, until 50 passages were marked positive in each set (100 per warning). All passages marked either positive or negative by the first annotator were then also rated by the other annotators. This step is necessary because the ratio of positive-to-negative passages varies between the warnings and is very low (1:10) for some warnings like *Death* and *War* in set 2, so having an equal number of passages for each warning would result either in a very high number of annotations or a low number of positive examples for some classes.

Annotation Task Design The annotation task was designed as a binary classification decision.⁵ Annotators were presented with one passage at a time, an example of which is shown in Figure 2. The markable consisted of five parts: (i) a description of the ‘Persona’ which the annotators should adapt for their decision, (ii) the ‘Definition’ of the warning which we created manually to have a similar length and scope, (iii) two ‘Demonstrations’, one positive and one negative, selected by the au-

⁴We retrieved the documents via elasticsearch 7.17.4 with the default BM25 scoring on the chapter text. We applied html-removal and stemming after the default pre-processing.

⁵We used Label Studio 1.6.0rc5 as the annotation system.

Warning	Pairwise κ			Pairwise Overlap			Mean κ			Positive Rate		
	1-2	1-3	2-3	1-2	1-3	2-3	1	2	3	1	2	3
Violence	0.38	0.33	0.53	0.82	0.80	0.82	0.36	0.46	0.43	0.09	0.25	0.25
Death	0.34	0.29	0.44	0.78	0.74	0.78	0.32	0.39	0.36	0.19	0.25	0.28
War	0.18	0.08	0.38	0.78	0.77	0.81	0.13	0.28	0.23	0.12	0.19	0.18
Abduction	0.27	0.31	0.26	0.77	0.67	0.65	0.29	0.26	0.29	0.20	0.20	0.47
Racism	0.66	0.43	0.46	0.83	0.71	0.73	0.55	0.56	0.44	0.44	0.49	0.56
Homophobia	0.23	0.15	0.35	0.60	0.61	0.67	0.19	0.29	0.25	0.34	0.56	0.37
Misogyny	0.39	0.14	0.25	0.74	0.69	0.73	0.26	0.32	0.20	0.31	0.30	0.14
Ableism	0.13	0.31	0.35	0.57	0.66	0.71	0.22	0.24	0.33	0.49	0.33	0.31
Total	0.36	0.28	0.41	0.76	0.73	0.76	0.32	0.38	0.34	0.21	0.28	0.29

Table 2: Additional measures of annotation quality. Shown are the pairwise Cohen’s κ between the annotators (label wise, since an annotator always rated one label completely), the fraction of overlapping annotations between two annotators, the mean Cohen’s κ agreement of each annotator with the other two, and each annotator’s rate of positive annotations. Note that the annotators differ across warnings.

thors, (iv) the ‘Instructions’ explained the binary classification decision, and (v) the ‘Passage’ to be annotated.

Annotator Instruction and Monitoring In total, we recruited seven permanent annotators of different genders and backgrounds. The passages were assigned based on the warning, i.e. annotator 1 rated all *Misogyny* passages, annotator 2 all *Racism* passages, and so on. The annotators were allowed to freely choose the warning they wanted to annotate, however, they were also asked to choose a warning that they were familiar with, if possible.

Following Kirk et al. (2022), we tried to reduce the stress on the annotators in two ways: First, we set no deadlines and we encouraged annotators to work in small batches over a longer period (6 weeks). Second, we arranged for weekly personal meetings with the annotators to discuss unclear cases, to monitor their well-being regarding the task, and to refine the shared understanding of the task through open discussion. The annotators were allowed to re-iterate and modify their annotations at any time.

4.3 Evaluation

We evaluate (i) if the keyword-based passage retrieval is effective, (ii) if the dataset is large and balanced enough for machine learning, and (iii) if the task design facilitates high-quality annotations.

Passage Retrieval The annotation results in Table 1b and Table 2 show that 1,901 (46%) of the retrieved passages received at least one positive vote, so a sensitive reader might desire a warning. The highest positive rate (PR) is 69% for *Homophobia* and *Racism*, the lowest is 34% for

War. Considering that most passages in any given document are negative, we consider our keyword-based passage retrieval effective in recalling good annotation candidates. However, some ambiguous keywords like ‘hit’ (*Violence*) or ‘occupy’ (*War*) retrieve many off-topic examples, lowering precision. Although these are often easy to annotate, better filtering would reduce annotation costs.

Size Table 1b shows that negative instances are more frequent than positives. *Discrimination* warnings have a higher positive ratio than *Aggression* warnings. While the complete dataset is nearly balanced (45% PR) under a 1-vote threshold, the data skew towards negatives (23% PR) under a 2-vote threshold. The balance varies between keywords, which is a problem for standard splits because the test or validation sets may have few positive (Set 2 *War*) or negative (Set 1 *Death*) instances. Instead, we opted for cross-validation.

Causes for the imbalance are, first, some keywords retrieve more severe passages. For example, *Racism* Set 1 contains many strong slurs and has a high positive rate (99%), while *Death* Set 2 (31% positive rate) contains many fantasy concepts that are also used in harmless contexts. Second, annotators differ in sensitivity to some warnings (see Table 2), like Rater 1 on *Violence* (9% PR compared to 25%) or Rater 3 on *Abduction* (47% to 20%).

Annotation Quality Table 1b and Table 2 show a chance-corrected inter-annotator agreement (0.22–0.52, mean 0.35 Krippendorff’s α). This indicates a “fair agreement”, which is consistent with similar tasks (0.34–0.58 Cohen’s κ over two annotators for binary offensive language by Pitenis et al. (2020); 0.20 Fleiss’ κ over 20 annotators for binary hateful language by Röttger et al. (2022)).

We expected a reasonable degree of disagreement due to the subjective nature of trigger warnings. Text appears harmful to individuals based on their personal lived experience, which, naturally, varies between annotators. We still evaluated our annotators quantitatively by measuring the pairwise overlap (0.74–0.77 mean across all warnings) and the pairwise agreement (0.31–0.44 Cohen’s κ). One annotator systematically disagreed with the other two (*Racism*, Annotator 3 with mean κ of -0.13), so that the annotations were repeated by another.

4.4 Subjectivity and Annotator Beliefs

Table 1b shows the passage count by number of positive votes. Out of all passages, 55 % were unanimously negative, but only 9 % unanimously positive. The negatives are partially explained by off-topic passages retrieved by ambiguous keywords, although this does not account for all negatives (see Figure 5). The variance across positive votes is explained by varying annotator sensitivity and a different belief of when a warning is required. Figure 3 shows typical positive passages. Unanimous positive passages are often severe cases with heavy use of slurs and very graphic descriptions. Non-unanimous examples feature unfamiliar or unrelated settings (fantasy, science fiction) or concepts (‘conversion therapy’), background or implied mentions, or a mix between warnings which confuse the annotators. Contributing to this are keywords like ‘undead’, ‘reanimation’, or ‘necromancy’ often contained in fantastical scenes (see Table 3). Removing those keywords would reduce annotation load and annotator uncertainty but would also reduce recall, remove positive passages, and further limit the scope of the concepts for model training.

The vote aggregation threshold impacts the characteristics of a classifier. A low threshold (1 vote) produces recall-oriented classifiers for warning assignment, while a high threshold (2–3 votes) produces precision-oriented classifiers suitable for moderation tools. We assume that resolving annotator disagreement will require either a form of personalization with a per-warning known sensitivity, or a prescriptive transformation of the annotation task, where the triggers are intensionally defined.

5 Passage Classification

We structure our experiments around four relevant design decisions and analyze their effect on the performance across models and warnings.

Class Modeling We comparatively evaluate binary, multiclass, and multilabel modeling via fine-tuned classification. We formulated the annotation task as binary classification: given a passage and a warning, decide if the warning should be assigned. Hence, our reference baseline and our evaluation setting is binary classification. However, binary classification is expensive to scale since each new warning needs a new classifier, and new training data, and each classifier does not use most of the annotated data. Under multiclass modeling, only one classifier is needed so the training data is better utilized, although such a classifier can only predict one warning per passage. Multilabel modeling combines these advantages of binary and multiclass, but is more difficult to train and our (binary) data might not be sufficient.

Fine-tuning vs. Few-shot Learning We comparatively evaluate six few-shot prompted generative LLMs which can scale to a high-dimensional warning taxonomy like NEON. Trigger warnings are potentially open-class and might require personalization, which is difficult to scale to with model fine-tuning.

Vote Aggregation We comparatively evaluate majority (2 votes count as positive) and minority voting (1 vote counts as positive) across all other experimental configurations. Our dataset evaluation suggested that examples containing clear and intense harm often receive multiple positive votes (24% with 2–3), while examples with one vote (23%) often contain only mild or implicit harm. The minimum number of votes required for an example to count as positive for classification will thus influence the sensitivity of the classifier.

Keyword Distribution We evaluate all experiments across in-distribution (ID) and out-of-distribution (OOD) samples of the dataset. Using keyword-based filtering to pre-select passages may limit how well the classifiers generalize to unseen examples that match the respective warnings but not the concepts captured by the keywords. We simulate this situation by splitting the keywords (and the passages retrieved using them) into two non-overlapping sets (cf. 4.1) and sampling datasets once from both sets combined and once separately.

5.1 Experiment Datasets

We compiled a total of 16 initial, unbalanced datasets, one for each of the eight warnings with

<p>That’s not unusual. There’s a lot of turnover. <i>She could be a wop or a spic, or she could have some nigger blood in her</i></p> <p>(1) <i>like Obama, and she’s old.</i> Coke-bottle-bottom eyeglasses, shriveled as a prune, and almost as small as one. 0 out of 10.</p> <p>Racism (3 Votes) – Heavy Slurs</p>
<p>Marcy’s last moments had not been pleasant ones. She’d looked down at the fiery sword protruding from her chest and knew this wasn’t something she’d recover from, that this was it. <i>As she apologized to anne and collapsed, she thought that would be the end of it, but dying was a lot slower than she anticipated.</i></p> <p>(2) <i>As her vision blurred and her limbs grew numb, she watched as Sasha rushed over to her, screaming something she couldn’t hear. She tried to reach out to her and say something, anything, but all she managed was wheezing and gurgling as blood poured from her mouth.</i></p> <p>Death (3 Votes) – Graphic Descriptions</p>
<p>One night, however, changed everything. Bella had a spinal stroke, paralyzing her from the waist down. <i>Bella’s husband tried to be there for her, but caring for a wife who was an invalid proved to be too much for him.</i></p> <p>(3) <i>He divorced his wife, unable to cope with her long recovery. Despite being abandoned by her husband, Bella was determined to not let her disability define her.</i></p> <p>Ableism (3 Votes) – Clear Description</p>
<p>For Blighttown humanity was a foreign concept: disease fueled mutation stole the human forms of most who live there. Their bodies stretched in gruesome ways, minds lost in the process. <i>Others lost theirs from once desperate cannibalism.</i></p> <p>(4) <i>A conscience stricken self preservation morphed into an insatiable greed for flesh. Fear had turned to panic and then to hopelessness as years ago disease washed over the rickety mining town.</i></p> <p>Death (2 Votes) – Unreliable Setting</p>
<p>Two days. Two goddamn days that burned every good thing in his life to ash. <i>Two days wherein he’d killed another member of his family (they were all dead so long ago, why did i let myself have new family to kill?), physically and emotionally devastated his life’s only love, and saw final death come at last to his greatest (until what i did to Buffy) sin.</i></p> <p>(5) <i>And they said it took centuries to shake the world to its very core. The evil inside him had managed to rock the foundation of his world off its axis in two short days.</i></p> <p>Violence (1 Vote) – Concept Mixing with Death</p>
<p>They were disgusting and meaningless, and they were nothing but torture to the poor people who were forced to go through with them. And what had Terry done that one awful morning? <i>Was what Mickey lived through anything but a conversion</i></p> <p>(6) <i>therapy? ‘She’s gonna fuck the faggot out of you.’</i></p> <p>Maybe Ian wasn’t just fighting for himself and the others.</p> <p>Homophobia (1 Vote) – Unfamiliar or Unreliable Concept</p>
<p>She remembered a tall, strong man, with a handsome face framed by black hair. The man she looked at now didn’t match that description in the least. <i>He was confined to a wheelchair loaded with machinery.</i></p> <p>(7) <i>Frail, thin face surrounded by a mane of unkempt grey hair. His eyes however had not changed.</i></p> <p>Ableism (1 Vote) – Background Mentions of Concept</p>

Figure 3: **Trigger Warning**. Selected example passages with assigned warning and number of positive votes. The center sentence that was retrieved via keywords is in italics.

majority and minority vote aggregation respectively. To get comparable results, we sampled balanced test sets with 20 positive and 20 negative instances for each of the 16 unbalanced datasets, which is the limit imposed by the smallest set (Misogyny Set 2). Since 40 instances are too few to get stable results, we created six random folds for a 5-fold Monte Carlo cross-validation with a sixth fold for parameter tuning.

For the in-distribution experiments, the test data are randomly drawn from all examples (i.e. they share passages retrieved by both keyword sets), while all other instances remain for training. In the out-of-distribution experiments, the test data are randomly drawn from all passages retrieved by keyword set 2, all other set 2 instances are discarded and all set 1 instances remain for training. Table 1c shows the number of examples in the training and

test datasets.

5.2 Models and Training

We evaluate five fine-tuning-based classification models Binary, Binary+ (which includes negative instances from other classes), Multiclass, Multilabel (without any all-negative instances) and Multilabel+ (with all-negative instances) and six generative LLMs GPT 3.5, GPT 4, Mistral 7B, Mixtral 7x8B, Llama 7B, and Llama 13B. We implemented the models using Huggingface’s *transformers* library (4.35.0). We prompted GPT via OpenAI’s API.

Binary We trained eight binary classification models for each configuration, one for each warning, on all positive and negative warning-passage pairs of the respective warning. We also trained a

version Binary+ where each classifier was trained on the positive instances of the respective warning and the negative instances of all other warnings. We hypothesize that this expansion might improve the classification in the out-of-distribution configurations.

Multiclass We trained one multiclass classifier by combining the eight training datasets: We assigned each positive instance a class label for its warning (class 0–7) and all negative instances to class 8. We test the models on the binary datasets. We count a prediction as positive if it predicts the class corresponding to the dataset’s warning and all other classes as negative (class 8). For example, when predicting multiclass for *Death*, then all predicted classes except *Death* count as negative predictions. This is a limitation imposed by annotating the dataset in a binary fashion.

Multilabel We trained one multilabel classifier by combining the eight training datasets and converting the positive binary class labels to respective one-hot label vectors and discarding all negative labels. This corresponds to a one-class paradigm which can be easily expanded for new warnings. We also trained a version Multilabel+ which includes all negative examples with a zero vector as label (see limitations below). We test the models, similarly to multiclass, by only considering the models prediction for a class if the instance was annotated for that class, i.e. we ignore all predictions for other classes.

Since the training passages are only annotated in a binary way for the respective warning and there is no overlap between passages across warnings, converting the examples to multilabel class vectors introduces errors: If a passage would be positive for a warning it was not annotated for, the class vector becomes incomplete and might confuse the classifier. Although the multilabel class modeling is the most practically convenient, it is also not very promising to reach a high performance due to these issues. Adding negative examples for the multilabel+ classifier will likely exacerbate this problem (cf. Figure 3(4)).

However, since the passages are short, they are usually only positive for multiple warnings from within the same category (either *Aggression* or *Discrimination*). An example selected for *Death* may be positive for *War* or *Violence*, but is likely not positive for *Racism*. That means we can mit-

igate adding many false positives by splitting the multilabel+ classifier into two: one multilabel classifier for all *Aggression* warnings where only negatives from *Discrimination* are added, and vice versa.

GPT 3.5 and GPT 4 We prompted the models ‘gpt-3.5-turbo-0125’ and ‘gpt-4-0125-preview’ using the OpenAI API as described in Section B.3. We only prompted the models for the respective test splits across all warnings and folds.

Mistral 7B and Mixtral 7x8B We implemented both Mistral models via Huggingface’s *transformers* library, using the ‘mistralai/Mistral-7B-Instruct-v0.2’ and ‘mistralai/Mixtral-8x7B-Instruct-v0.1’ checkpoints respectively. We encoded the prompt using Mistral’s chat template (without system message) and prefixed the prompt with: *You are a classification model that only answers with ‘yes’ or ‘no’*. Mistral 7B was tested on one A100 40GB. Mixtral 7x8B was loaded with 8-bit quantization and tested on three A100 40GB.

Llama 7B and 13B We implemented both Llama models via Huggingface’s *transformers* library, using the ‘meta-llama/Llama-2-7b-chat-hf’ and ‘meta-llama/Llama-2-13b-chat-hf’ checkpoints respectively. We used the prompt directly and without the chat template, because the models produces difficult-to-parse output when using the chat template as we did with Mistral. Llama 7B was tested on two A100 40GB with a batch size of 12. Llama 12B was loaded using the A100 GPUs and batch sizes of 8.

5.3 Evaluation

We test all models on the balanced, binary test datasets for ease of comparison and report accuracy and positive rate. For multiclass and multilabel models we only count the predictions relating to the respective warning (cf ??). We report the mean performance across all five folds for each model and macro-averages across models and warnings.

6 Results and Discussion

Figure 4 and Table 5 (A.2) show that models typically score between ca. 0.6–0.7 accuracy on average across all warnings, depending on the model and data configuration, while the top models per warning score between ca. 0.7–0.8. The highest-scoring models are Multiclass on in-distribution data (ID) with minority voting (MinV) (0.82),

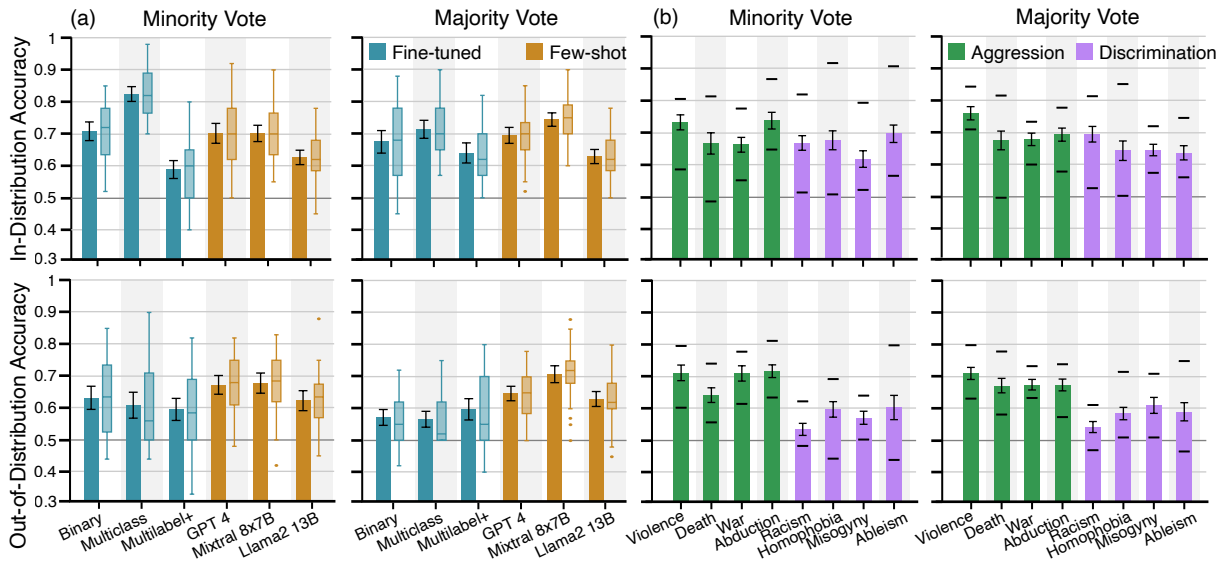


Figure 4: Overview of the experimental results across distribution (in vs. out-of) and vote aggregation (minority vs. majority). **(a)** The bar charts show the mean accuracy for the models across all warnings with 95% t-estimated confidence interval. The box plots show the quartiles and outliers of all folds and warnings. **(b)** The bar charts show the mean accuracy for the warnings across all models with 95% t-estimated confidence interval. The upper and lower ticks show the best and worst model’s performance. Table 5 (Appendix A) shows the full results.

Mixtral on ID (0.74), and out-of-distribution (OD) (0.71) majority voting (MajV), and GPT 4 on OD/MinV. The lowest mean scores are barely above random and occur for some worst-case configurations. The best per-warning score is 0.91 (Multiclass on *Homophobia*). Multiclass is the most effective model in 10 of 32 cases across warnings and configurations, followed by Mixtral with 9. The scores vary between warnings, where *Violence* has the highest (0.70–0.72 mean accuracy) and *Racism* the lowest scores (0.52–0.69). *Aggression* warnings score higher than *Discrimination*.

Class Modeling Model efficiency varies significantly between modeling strategies: Multiclass outperforms Binary (0.11 pp. for ID/MinV;) and Multilabel+ (0.23 pp. for ID/MinV; 0.07 pp. ID/MajV) and Binary outperforms Multilabel (0.09 for OOD/MinV). All other differences are not significant. Adding more negative examples to Binary shows no significant difference and adding negative examples to Multilabel improves the measured performance within the 95% CI by 0.01–0.06. Although Multiclass is the most effective fine-tuned model, it works only if the warnings are separable and rarely overlap. Since this is often not the case (cf. Figure 3 (5)), the binary model is a more realistic alternative.

Fine-tuning vs. Few-shot Learning The results show that all fine-tuned models perform equal or

worse than the best few-shot models (Mixtral) in the OOD configuration. In the ID configuration, the fine-tuned models score comparative or higher (Multilabel for ID/MinV). The fine-tuned models are also more efficient in time and energy and should be preferred whenever possible.

Another notable difference is that all few-shot models have a higher average positive rate (PR) than the fine-tuned model (9–49 pp.), so the generative LLMs can be considered more sensitive. The reason for this is not clear, although explicit harm avoidance during training is a possible explanation.

Vote Aggregation Increasing the aggregation threshold from one to two significantly increases the accuracy for Multilabel+ and decreases the accuracy for Multiclass, for *Ableism*, and for *Abduction* (all ID). In addition, the inter-quartile range of the folds (box plots in Figure 4) is smaller for MajV. This is likely explained by the reduced number of positive examples, which reduced the variance of the examples in the folds since there are much fewer positive examples to draw from.

Increasing the threshold also reduces the PR (3–28 pp.) across all fine-tuned models and increases the PR (6–10 pp.) across all few-shot models. The increase is likely explained by the overall increases of instances with 1+ votes increases with higher thresholds, since these get samples as the negative class and, since the the prompt did not change, this

<p>It's easier for him to carry what he needs to avoid suspicions. Kunimi is required to pay for his medical school tuition tomorrow. Which is a cadaver. Kunimi doesn't understand what's the point of finding dead bodies anymore. His classmates practically dug all the bodies from all the cemeteries, but it's still not enough.</p> <p>(1)</p> <p>Death (0 Votes) – Not Perceived as Harmful</p>
<p>Many, many things did happen to me. the hardest part is simply in starting, because once you do, the words come flying out. I asked myself for a long time whether i would be able to return to a civilian life. How could i get up in the morning if</p> <p>(2) i didn't have a purpose, if i wasn't useful? I'd been a soldier, and a doctor, and a bullet through my shoulder ripped me of both titles.</p> <p>War (0 Votes) – Implicit or Background Mentions</p>
<p>The military bases also held the space force academy, a supposedly nonpartisan association that was clearly held and run by the allies of the United States. If Luna decided to put sanctions on Chinese trade, that also included chinese traders. Directly, that would affect thousands of people, but with sanctions came racism. The workers that lived on Luna, Mars,</p> <p>(3) Deimos, Phobos, and the other new colonies would be ostracized, meaning millions would be left in the cold. Luna would be cutting off all shipments to and from China, which meant China wouldn't be trading with any colonies in the solar system.</p> <p>Racism (0 Votes) – Implicit or Background Mentions</p>
<p>'I remember my last holder complaining about that too...' 'It sucks and it hurts and there's no point! and i'd go to the doctor and get that pill thing to stop it, but my dad won't let me, and... and speaking of my dad, you heard what he was saying earlier...' Fluff nodded again. Sitting in Alix's pocket all day, there was no way she couldn't have heard it - the usual why can't you be more ladylike? You're not a little child anymore, you need to stop being so immature! Can't you be more like your friends?</p> <p>(4)</p> <p>Misogyny (1 Vote) – Edge Cases with Disagreement</p>
<p>Usually, they lasted longer than that, a month, at most. Whenever she started to trust them, they would disappear. Some found murdered, others never seen again. Ramsay liked to play with her, giving and taking anything she found comfort in.</p> <p>(5) There was Theon, but she found no comfort in him, only pity.</p> <p>Violence (1 Vote) – Edge Cases with Disagreement</p>

Figure 5: **Trigger Warning**. Selected example passages that were always misclassified by all models.

increases the few-shot models' positive rates.

Keyword Distribution Overall, all models are less effective in the OOD configuration. This is expected behavior for fine-tuned models, since they can rely less on learned lexical cues, but not necessarily for few-shot models. A possible explanation is that we split the keywords in the order of generation, so keywords that are strongly associated with a trigger are in set 1. The OOD test set then contains passages that are less strongly associated with the trigger, thus reducing the few-shot scores.

In the OOD configurations, all fine-tuned models are significantly worse by ca. 0.1–0.2 pp. on average with an increased spread of per-warning scores while the few-shot models score as on ID (except for a reduction on *Racism*). Fine-tuned models are still competitive for most *Aggression* but ineffective for *Discrimination* warnings with barely above random accuracy and very low PRs. This is likely due to systematic differences between the keywords (cf. 3). While the *Aggression* keywords are somewhat consistent between sets, *Discrimination* keywords often start with a series of slurs (Set 1) and become more conceptual later (Set 2), leading to a topical gap between OOD training and test. This was the

goal of the configuration and shows that fine-tuned models generalize poorly in this case, highlighting the importance of diverse training data.

6.1 Causes of Misclassification

Table 5 shows selected passages that are always misclassified. Cases of systematic misclassification contain passages that (i) are on-topic but are never rated as triggering by the annotators, (ii) mention the topic implicitly or in the background (like historic discrimination, descriptions of the setting) but are never rated as triggering by the annotators, and (iii) with edge cases which are rated positively by only one annotator. Misclassified instances seem to align with annotator disagreement. Across all few-shot models, many (56%/75%) instances with 1 or 3 votes are always classified correctly, but only 15% with 1 vote. Mixtral classifies 70–91% of passages with 1/2/3-votes correctly, but only 30% of 1-vote instances. The fine-tuned models, too, more often misclassify examples with disagreement but less extreme: Binary (ID) is correct for 66% of 3-vote and only 44% on 2-vote instances.

7 Conclusion

In this paper, we seek to identify the exact text passages that prompt authors to preface their works with trigger warnings. We model the task as binary classification and create a dataset of 4,135 English passages, each annotated by three human votes across 8 trigger warnings. We investigate how different beliefs about triggers contribute to the assignment of warnings by quantitatively and qualitatively analyzing annotator disagreement and the behavior of 11 classifiers across different design decisions: Vote Aggregation, Keyword Distribution, Class Modeling, and Training vs. Prompting.

Our keyword-based passage retrieval identifies many positive instances, from severe and graphic passages with high agreement to mild and implicit ones where the annotators’ opinions differ based on their personal sensitivities and beliefs about the need for trigger warnings. We find that where annotators disagree, classification errors occur more frequently. Furthermore, we can experimentally show that, first, diverse training data is required for models to generalize well to unseen concepts and rare triggers, second, few-shot models like Mixtral are competitive (albeit computationally expensive), especially for unseen triggers, and, third, fine-tuned models are often still the best models for individual warnings and in certain configurations. It is advantageous to select a model specifically for the targeted warning or configuration, depending on the goal, like deleting content vs. adding warnings, or on the targeted sensitivity.

In conclusion, we question whether authors can determine a trigger warning equally well for every reader. Future work may include corresponding investigations. Personalizing trigger warning assignment seems a fruitful computational direction.

Limitations

We consider four limitations. First, we only consider eight warnings from two categories in this study, all of them from more frequent cases. The WTWC-22 corpus contains 36 warnings from seven categories, while other taxonomies list up to 90. Consequentially, our strategy to generate keywords and collect passages to annotate might not work for rare or very individual warnings. Besides, the annotation instructions and prompts may also not generalize to some of the warnings, and model performance may differ.

Second, we recruited annotators from a population of computer science students which is a shared

demographic and hence a source of bias. All annotations were done by individuals who are not (clinically) affected by any trigger but who have different sensitivities to the different triggering concepts. The latter limitation holds for almost all trigger warnings which are usually assigned by the authors of the content. However, more reliable annotations would require that annotators are actually triggered by the concepts described by the annotation categories. We know of no practical and ethically feasible way to create such annotations.

Third, we cannot make any claims about which parts of the annotation instructions and the prompts (cf. Figure 2) influence the performance of the annotators or the models beyond our ablation study.

Fourth, we did not quantitatively evaluate if our annotation instructions are the most optimal or if a different annotation paradigm would reduce the subjectivity of the task without limiting the captured beliefs about trigger warnings.

Impact Statement

As with any work on harmful content in the age of generative AI, our data, code, and insights may be used in bad faith to generate harmful content or for adversarial engineering to trick detection algorithms. The artifacts we use (WTWC-22) are for academic use only, as are the artifacts we created: our data and code.

It is safe to assume that due to the inherently subjective nature of the phenomenon of triggering, future work will focus on personalized decision-making for harmful content classification. However, personalization here is a highly sensitive issue since it requires deep knowledge about the individuals and their vulnerabilities.

We followed best practices when annotating harmful content for this work. The student annotators were not paid, besides teaching credit of which annotations were only part. Annotators were not pressured to complete annotations, were given a possibility of opting out of categories that they found particularly disturbing, and it was made clear that grading was independent of the completion of the annotation work.

References

- C. J. Adams, Jeffrey Sorensen, Julia Elliott, Lucas Dixon, Mark McDonald, Nithum, and Will Cukierski. 2017. [Toxic comment classification challenge](#).

- Michele Banko, Brendon MacKeen, and Laurie Ray. 2020. [A unified taxonomy of harmful content](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms, WOAHA 2020, Online, November 20, 2020*, pages 125–137. Association for Computational Linguistics.
- Victoria M. E. Bridgland, Payton J. Jones, and Benjamin W. Bellet. 2023. [A meta-analysis of the efficacy of trigger warnings, content warnings, and content notes](#). *Clinical Psychological Science*, 0(0):21677026231186625.
- Ashleigh Charles, Laurie Hare-Duke, Hannah Nudds, Donna Franklin, Joy Llewellyn-Beardsley, Stefan Rennick-Egglestone, Onni Gust, Fiona Ng, Elizabeth Evans, Emily Knox, et al. 2022. Typology of content warnings and trigger warnings: Systematic review. *PLoS one*, 17(5):e0266722.
- Aida Mostafazadeh Davani, Mark Díaz, and Vinodkumar Prabhakaran. 2022. Dealing with disagreements: Looking beyond the majority vote in subjective annotations. *Trans. Assoc. Comput. Linguistics*, 10:92–110.
- Hannah Kirk, Abeba Birhane, Bertie Vidgen, and Leon Derczynski. 2022. Handling and presenting harmful text in NLP research. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 497–510, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Emily Knox. 2017. *Trigger Warnings: History, Theory, Context*. Rowman & Littlefield.
- Ioannis Mollas, Zoe Chrysopoulou, Stamatis Karlos, and Grigorios Tsoumakas. 2020. ETHOS: an online hate speech detection dataset. *CoRR*, abs/2006.08328.
- Zesis Pitenis, Marcos Zampieri, and Tharindu Ranasinghe. 2020. [Offensive language identification in Greek](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 5113–5119, Marseille, France. European Language Resources Association.
- Laura Plaza, Jorge Carrillo-de-Albornoz, Roser Morante, Enrique Amigó, Julio Gonzalo, Damiano Spina, and Paolo Rosso. 2023. Overview of EXIST 2023 - learning with disagreement for sexism identification and characterization. In *CLEF*, volume 14163 of *Lecture Notes in Computer Science*, pages 316–342. Springer.
- Paul Röttger, Bertie Vidgen, Dirk Hovy, and Janet B. Pierrehumbert. 2022. [Two contrasting data annotation paradigms for subjective NLP tasks](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 175–190. Association for Computational Linguistics.
- Marta Sandri, Elisa Leonardelli, Sara Tonelli, and Elisabetta Jezeq. 2023. Why don't you do it right? analysing annotators' disagreement in subjective tasks. In *EACL*, pages 2420–2433. Association for Computational Linguistics.
- Manuka Stratta, Julia Park, and Cooper deNicola. 2020. [Automated content warnings for sensitive posts](#). In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems, CHI 2020, Honolulu, HI, USA, April 25-30, 2020*, pages 1–8. ACM.
- Alexandra Uma, Tommaso Fornaciari, Anca Dumitrache, Tristan Miller, Jon Chamberlain, Barbara Plank, Edwin Simpson, and Massimo Poesio. 2021. Semeval-2021 task 12: Learning with disagreements. In *SemEval@ACL/IJCNLP*, pages 338–347. Association for Computational Linguistics.
- Matti Wiegmann, Magdalena Wolska, Christopher Schröder, Ole Borchardt, Benno Stein, and Martin Potthast. 2023. [Trigger Warning Assignment as a Multi-Label Document Classification Problem](#). In *Proceedings of the 61th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12113–12134, Toronto, Canada. Association for Computational Linguistics.
- Magdalena Wolska, Christopher Schröder, Ole Borchardt, Benno Stein, and Martin Potthast. 2022. [Trigger warnings: Bootstrapping a violence detector for fanfiction](#). *CoRR*, abs/2209.04409.
- 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.

A Figures and Tables

A.1 Dataset Evaluation

Racism	
Prompt	Provide a list of racist language that must be avoided in a story for people that are triggered by it. Return the list in json format.
Phrases in Set 1	nigger, chink, spic, gook, wetback, jap, darkie, kike, dago, ngger, chnk, spc, g**k, towelhead, curry muncher, redskin, half-breed, slant eyes
Phrases in Set 2	model minority, savage, primitive, barbaric, uncivilized, tribal, blackface, white supremacy, aryan race, master race, racial purity, colorblind, racism, articulate for, pretty for a, go back to your country, you're not like other, you're one of the good ones, I don't see color, all lives matter
Omitted	sand nigger, exotic, oriental, reverse racism, nip, speak english
Death	
Prompt	Provide a list of language relating to the death of others, character death, killing, corpses, or coping with loss or grief that must be avoided in a story for people that are triggered by it. Return the list in json format.
Phrases in Set 1	mourn, grieve, deal with loss, cope with grief, bury, inter, cremate, come to terms with, kill, murder, slay, weep, slaughter, sacrifice, take a life, end a life, cause death, corpse, dead body, lifeless form, cadaver, pass away, suffer loss, experience grief, navigate bereavement, struggle with mourning, manage sorrow, cope with death, die, cease, euthanize
Phrases in Set 2	choose to die, end suffering, peacefully pass away, undead, reanimate, come back to life, resurrect, revive, consume flesh, cannibalize, consume brains, devour flesh, tear into, consume blood, drink blood, suck blood, cannibalism, necromancy, flesh-eating, blood-sucking, risen from the dead
Omitted	cry, sob, lament, suffer, comfort, console, remember, honor, respect, accept, process, deal with, face, confront, put down, put to sleep, remains, lose, deal with deceased, struggle, fight, gasping, gasp, drown, sink, fade, weaken, wane, slip away, pass, depart, succumb, let go, release, end life, attack, bite, devour, consume, feast, eat, chew, gnaw, rip, tear, destroy, eliminate, terminate, annihilate, exterminate, rise again, prey upon, feast on, rip apart, consume human flesh, zombify, reanimate as a zombie, drain blood, vampirize, turn into a vampire, resurrection, zombification, vampirism, undeadness, death-like state, revival, post-mortem consumption, reanimation, expire

Table 3: **Warning Racism and Death.** Keyword selection for two warnings. We selected segments to be annotated based on topical keywords and phrases for each warning. The lists of **Phrases** were generated using gpt-3.5-turbo-0301 using the provided prompt. We manually removed keywords that better match a different warning (*chew, attack, . . .*), frequently lead to false positive matches (*let go, depart, lose, . . .*), and which had a large lexical overlap with non-omitted keywords (since our segment retrieval method applies stemming and partial matches). The remaining phrases were split into two sets for the out-of-distribution experiments. The keywords for all warnings can be found in the repository linked in the Abstract.

Warning	Majority Instances		Minority Instances		Majority Vote PR		Minority Vote PR	
	pos	neg	pos	neg	Set 1	Set 2	Set 1	Set 2
Violence	188	853	363 (35%)	678	107 (16 %)	81 (22 %)	205 (30 %)	158 (43 %)
Death	114	430	234 (43%)	310	79 (44 %)	35 (10 %)	122 (69 %)	112 (31 %)
War	102	725	282 (34%)	545	39 (19 %)	63 (10 %)	97 (48 %)	185 (30 %)
Abduction	132	379	273 (53%)	238	59 (23 %)	73 (28 %)	122 (48 %)	151 (57 %)
Racism	125	142	185 (69%)	82	66 (97 %)	59 (30 %)	67 (99 %)	118 (59 %)
Homophobia	138	175	216 (69%)	97	52 (40 %)	86 (47 %)	77 (60 %)	139 (76 %)
Misogyny	81	296	179 (47%)	198	57 (38 %)	24 (11 %)	90 (59 %)	89 (40 %)
Ableism	87	168	169 (66%)	86	43 (27 %)	44 (47 %)	92 (56 %)	77 (83 %)
Total	966	3,169	1,901 (46%)	2,234	501 (28 %)	465 (20 %)	872 (48 %)	1,029 (44 %)

Table 4: The number and ratio of positive and negative passages under majority and minority vote aggregation and the number and ratio (PR) of positively annotated examples across warnings. Shown are the PRs assuming Majority Vote (2 or more positives) and Minority Vote (1 or more positives), depending on if the keyword (used to retrieve the passage) was in set 1 or 2.

A.2 Classification Evaluation




Figure 6: Distribution and mean of model accuracy across all tested models. The coloring is for visual ease only.

B Model Training

All fine-tuned models are based on a ‘roberta-base’ checkpoint that was fine-tuned for fan fiction using masked language modeling (cf. B.1). We conducted a parameter sweep for each fine-tuned model on a 6th fold (cf. B.2). All fine-tuning models were trained on a single A100 40GB. All generative LLMs were prompted using the instructions shown in Figure 2. Section B.3 describes the prompt and our ablation study in detail. All used models are explained in detail in Section ??.

B.1 Language Modeling Fine-tuning of RoBERTa for Fan Fiction

We fine-tuned the ‘roberta-base’ checkpoint on fan fiction documents via masked language modeling using Huggingface’s ‘Trainer’ routine. As data, we extracted all English fan fiction documents from WTWC-22 that were marked as recommended for model training, i.e. where a trigger warning was assigned and where overly long, short, or unpopular works were removed (Wiegmann et al. (2023) list the precise parameters). All documents were split into training examples of ca. 450-500 words, respecting sentence and paragraph boundaries. We trained the checkpoint on the resulting 19 million examples for ca. 470,000 steps with a batch size of 32, the point of loss convergence on a 10,000 example hold-out validation set. As parameters, we used largely the standard parameters of the ‘Trainer’, with a random masking function, masking probability of 0.2, and an initial learning rate of $2e-5$.

B.2 Fine-tuning Parameter Sweep

We conducted a grid-based parameter sweep on the validation sample (cf. Section 5.1) for all fine-tuning strategy (binary, multi-label, multi-class, with and without extended training data) across three dimensions: in vs out-of-distribution, minority vs. majority vote aggregation, and learning rate within ($1e-5$, $2e-5$, $5e-5$). All models were trained for 20 epochs. The parameters did not vary between warnings. Table 6 shows the final parameter settings each model was trained on.

B.3 Prompt Ablation

Figure 2 shows the 5 parts of the instructions we used for both the annotators and as the prompt for the few-shot models: *Instruction*, *Passage*, *Persona*, *Definition*, and *Demonstrations* (one positive and one negative in that order). We counted model

Strategy	Dist.	Vote Agg.	LR	Acc.
binary	ood	majority	1e-5lr	0.62
binary	ood	minority	2e-5lr	0.64
binary	id	majority	1e-5lr	0.68
binary	id	minority	5e-5lr	0.74
binary extended	ood	minority	2e-5lr	0.62
binary extended	ood	majority	2e-5lr	0.62
binary extended	id	majority	2e-5lr	0.64
binary extended	id	minority	1e-5lr	0.71
multi-label	ood	majority	2e-5lr	0.58
multi-label	ood	minority	5e-5lr	0.57
multi-label	id	majority	2e-5lr	0.59
multi-label	id	minority	1e-5lr	0.57
multi-label extended	ood	majority	2e-5lr	0.62
multi-label extended	ood	minority	2e-5lr	0.58
multi-label extended	id	majority	1e-5lr	0.66
multi-label extended	id	minority	1e-5lr	0.57
multi-class	ood	majority	2e-5lr	0.59
multi-class	ood	minority	5e-5lr	0.64
multi-class	id	majority	2e-5lr	0.68
multi-class	id	minority	5e-5lr	0.84

Table 6: Overview of the different settings, the final parameter choices for each model, and the accuracy score (mean across all warnings) of that setting. The different tested learning rates varied between 0.01–0.09.

responses as positive when they contained ‘yes’ within the first 5 response tokens and negative if they contained ‘no’

For the prompt ablation, we queried all generative LLMs with open weights (excluding GPT) across all 4 test settings (aggregation and distribution), averaged the scores across all settings per model, and averaged the scores across all models to determine the best prompt template. We tested eight prompt template variations (cf. Table 7). Each variation started with *Instruction* and *Passage*, followed by all combinations of *Persona*, *Definition*, and *Demonstrations*, including neither. We manually selected the demonstrations that we rated as clear and representative for the respective trigger out of all unanimously annotated passages. We selected the prompt with all parts, which had with the highest average accuracy on the validation sample across all models and was also the one shown to annotators.

Prompt	Mistral 7B	Mixtral 8x7B	Llama 7B	Llama 13B	Mean Acc.
Instruction and Passage	0.72	0.74	0.52	0.50	0.62
+ Persona	0.71	0.74	0.57	0.57	0.65
+ Definition	0.72	0.75	0.54	0.58	0.65
+ Demonstrations	0.69	0.70	0.58	0.64	0.65
+ Persona + Definition	0.73	0.75	0.57	0.55	0.65
+ Persona + Demonstrations	0.68	0.74	0.64	0.64	0.68
+ Definition + Demonstrations	0.68	0.75	0.67	0.63	0.68
+ Persona + Definition + Demonstrations	0.68	0.75	0.68	0.59	0.68

Table 7: Overview of the accuracy on the validation sample of different candidate prompts across the tested models. All prompts ended with *Instruction* and *Passage* but varied in other parts.