

# WildFireCan-MMD: A Multimodal Dataset for Classification of User-Generated Content During Wildfires in Canada

Braeden Sherritt<sup>1</sup>, Isar Nejadgholi<sup>2</sup>, Efstratios Aivaliotis<sup>1</sup>,

Khaled Mslmani<sup>1</sup>, Marzieh Amini<sup>1</sup>

<sup>1</sup>Carleton University, Ottawa, Canada

<sup>2</sup>National Research Council Canada, Ottawa, Canada

BraedenSherritt@cmail.carleton.ca, isar.nejadgholi@nrc-cnrc.gc.ca, EfstratiosAivaliotis@cmail.carleton.ca,

KhaledMslmani@cmail.carleton.ca, MarziehAmini@cunet.carleton.ca

## Abstract

Rapid information access is vital during wildfires, yet traditional data sources are slow and costly. Social media offers real-time updates, but extracting relevant insights remains a challenge. In this work, we focus on multimodal wildfire social media data, which, although existing in current datasets, is currently under-represented in Canadian contexts. We present WildFireCan-MMD, a new multimodal dataset of posts on X from recent Canadian wildfires, annotated across twelve key themes. We evaluate zero-shot vision-language models on this dataset and compare their results with those of custom-trained and baseline classifiers. We show that while baseline methods and zero-shot prompting offer quick deployment, custom-trained models outperform them when labelled data is available. Our best-performing custom model reaches  $84.48 \pm 0.69\%$  f-score, outperforming VLMs and baseline classifiers. We also demonstrate how this model can be used to uncover trends during wildfires, through the collection and analysis of a large unlabeled dataset. Our dataset facilitates future research in wildfire response, and our findings highlight the importance of tailored datasets and task-specific training. Importantly, such datasets should be localized, as disaster response requirements vary across regions and contexts.

## 1 Introduction

Large-scale analysis of online interactions, using machine learning algorithms has enabled critical applications such as disaster response (Fauzi, 2023; Kumar et al., 2024b), mental health monitoring (Ghanadian et al., 2024; Dalal et al., 2024), detecting and combating hate speech in digital environments (Kiritchenko et al., 2021; Kumar et al., 2024a), pandemic discourse analysis (Imran et al., 2022) and sentiment and behavioral trend analysis (Ma et al., 2024; Curto et al., 2024). Among these applications, analysis of social media data

during natural disasters has gained increasing importance in recent years (Alam et al., 2018; Ofl et al., 2020; Alam et al., 2021; Dwarakanath et al., 2021; Algiriyyage et al., 2022; Hossain et al., 2022; Khattar and Quadri, 2022; Basit et al., 2023; Koshy and Elango, 2023; Rezk et al., 2023). Such data contains crucial details such as reports of casualties, infrastructure damage, and urgent needs of affected individuals, among other critical information. The effective use of the shared information can facilitate humanitarian aid efforts for more efficient management and faster recovery (Houston et al., 2015). This valuable data, however, is mixed in with noise, such as humorous content, advertisements, and other irrelevant information. Therefore, there is a critical need to develop an automated method for sifting through this information.

To foster research in this area, Alam et al. (2018, 2021) developed two multimodal datasets, Crisis-MMD and HumAID, both collected from Twitter, including several disasters such as earthquakes and hurricanes. However, the classes of these datasets are generalized across all types of disasters, meaning they are not tailored to specific disaster scenarios. As a result, the classification scheme is broad, and certain nuances or unique information that may be relevant to a particular type of disaster could be overlooked. Also, while the CrisisMMD dataset does contain wildfire data from the California wildfires of 2017, the proportion of this type of data compared to the rest is relatively small, only comprising 1589 samples, and is specific to the California wildfires of October 2017. The wildfire-related samples from Canada in the HumAID dataset (Alam et al., 2021) are limited to a single month in 2016, as well as another month of California samples from 2018. Therefore, wildfires in Canada remain underexplored in social media analysis research.

To address this limitation, we focus our work on gathering social media data specifically related to

Canadian wildfires, as they are becoming increasingly frequent and severe due to the adverse effects of climate change. Upon initial inspection of our Canadian wildfire-specific social media dataset, we observed that existing categorization schemes do not capture the specificity of this data. To address this, we devised a novel taxonomy of twelve categories reflecting the unique information needs and communication patterns of Canadian wildfire contexts. Using our developed taxonomy, we manually labeled 4,688 wildfire-related tweets from Canada and created **WildFireCan-MMD**, a new multimodal dataset of X posts from recent Canadian wildfires.

To benchmark the available classification models on WildFireCan-MMD, we evaluated the zero-shot classification capabilities of several Vision-Language Models (VLMs), custom deep-learning classifiers, as well as classical approaches, as a minimum achievable performance. Our best-performing classifier utilizes RoBERTa (Liu, 2019) and Vision Transformer (ViT) models to jointly process the text and image components of tweets. A custom transformer-based fusion module is designed to enable the modalities to interact deeply. While WildFireCan-MMD is sourced from data posted between 2022 and 2024, we collected an additional set of unlabeled tweets spanning seven years (2018-2024). Using our best-performing classifier, we analyze this unlabeled data and uncover trends and patterns of social media use during wildfires in Canada. The full set of labeled data (containing the tweet IDs and labels), as well as the experimental code, is available on Github (<https://github.com/Multimodal-Social-Media-Analysis/WildfireCanMMD>). The contributions of this study are summarized as:

- We collected 4,688 wildfire-related posts from X in Canada from 2022 to 2024 and devised a twelve-category taxonomy using unsupervised topic modeling and human analysis. We introduce **WildFireCan-MMD**, the first Canada-specific multimodal wildfire social-media dataset, annotated by three annotators.
- We propose a custom multimodal transformer model, trained on WildFireCan-MMD and compare it with several VLMs and classical methods. The fine-tuned custom classifier outperforms zero-shot VLMs significantly; our model achieves an f-Score of  $84.48 \pm 0.69\%$ .

- Additionally, we collected 46,279 posts from X from 2018 to 2024. We used our trained classifier to label them and reported the trends and patterns of social media use during wildfires in Canada. Results show that retrospective analysis using a classifier can uncover trends which relate to real situational updates.

## 2 Background and Literature Review

### 2.1 Social media analysis for disaster response

Significant research has been conducted on the role of social media during disasters. Fauzi (2023) conducted a bibliometric analysis of social media use in disaster response and offered a comprehensive overview of key studies and emerging research directions. Acikara et al. (2023) highlighted the potential of social media analytics in disaster response by leveraging users as “citizen sensors” to rapidly gather critical, time-sensitive information. Muniz-Rodriguez et al. (2020) focused on the public health aspects of social media use during emergencies, noting its effectiveness in spreading emergency warnings and assessing needs after a disaster. In other works, Dwarakanath et al. (2021), Bukar et al. (2020), and Mundottukandi et al. (2024) discussed the crucial role of social media in managing crises, pointing out how platforms like Facebook and Twitter support communication and coordination among affected individuals and communities. Furthermore, in a systematic review of the application of deep learning in disaster response, Al-giriyyage et al. (2022) discussed the successes and challenges faced in this area, proposed guidelines for future research, and highlighted the importance of using multimodal data in disaster scenarios. Our work is motivated by this literature as well as the literature shown in Appendix F which focuses this work on using machine learning models to triage Canada-specific multimodal social media data to inform disaster response and management.

### 2.2 Multimodal classification for disaster response

We review this literature based on the available datasets and previously explored algorithms.

**Datasets:** While many works utilize textual datasets sourced from social media to classify posts during disasters (Olteanu et al., 2014; Dong et al., 2021; Ochoa and Comes, 2021), it is now generally considered that multimodal approaches offer

stronger results (Ofli et al., 2020; Abavisani et al., 2020; Ramirez et al., 2025).

The study by Alam et al. (2018) released CrisisMMD, a collection of tweets (18k samples) with images from seven major natural disasters (including wildfires, earthquakes, floods, and hurricanes). CrisisMMD contains thousands of tweets labeled for categories like “infrastructure and utility damage” or “rescue volunteering or donation effort”.

Additionally, Alam et al. (2021) later released the HumAID dataset which consists of 77K human-labeled disaster tweets sampled from 24M tweets over 19 events (2016–2019). It is similarly annotated for emergency response categories (e.g. casualties, infrastructure damage).

There are several other multimodal datasets that exist, but most are variations or combinations of the CrisisLex (Olteanu et al., 2015) and CrisisMMD datasets (Wu et al., 2022). CrisisMMD has become a foundational resource for studying disaster-related communication patterns and the combined use of text and images. It remains the most prominent dataset in this domain, used in many recent studies (Ofli et al., 2020; Khattar and Quadri, 2022; Basit et al., 2023; Koshy and Elango, 2023; Rezk et al., 2023) and many others. Therefore, there is a notable gap in the availability of multimodal wildfire datasets.

In the context of Canada, where the most common and destructive natural disaster is wildfires, CrisisMMD is lacking in data. Although it contains 18k samples from other disasters such as hurricanes, earthquakes and floods, the amount of wildfire data is proportionally small (1589 samples). Additionally, this data is only from the California wildfires of October 2017. The wildfire-related samples from Canada in the HumAID dataset come from only a single month in 2016 (2,259 samples), as well as another month of California samples from 2018 (7,444 samples). Therefore, wildfires in Canada remain underexplored in social media analysis research. To address this limitation, this work focuses on gathering social media data related to Canadian wildfires, and proposes WildFireCan-MMD, a new Canadian wildfire social-media dataset with a wildfire-specific taxonomy 12 categories.

**Algorithms:** Previous work by Ofli et al. (2020) introduced a joint representation from two parallel deep learning architectures where one represents the text modality and the other represents the im-

age modality. Image and text modalities are vectorized through VGG16 and Word2Vec+CNN architectures respectively, followed by a dense layer for prediction. The work by Basit et al. (2023), expands on Ofli et al. (2020) by experimenting with multiple recent transformer structures, for both the text and image modalities, and compared them against the baseline reported in the previous work. They found that using a combination of ALBERT (Lan et al., 2020) as a text embedder, and DeiT (Touvron et al., 2021) as an image embedder, provided the best results for the eight-class classification task. Koshy and Elango (2023) implemented further experiments with text and image transformers and showed improvements by using bi-LSTMs instead of LSTMs. Furthermore, they have noted that multiplicative fusion strategies produce better results on a classification task when combined with the previously described methods. Specifically, they demonstrated that using the finetuned RoBERTa (Liu, 2019) model for text, Vision Transformer model for image, biLSTM+attention mechanism with multiplicative fusion, for prediction, achieves the best results for the eight-class classification task described by Alam et al. (2018).

Many architectures have been proposed in other works for classification on CrisisMMD, as seen in Appendix F. Most previous studies use a CNN for the image modality and then a BERT-based transformer for the text modality. Other works sometimes use ViT for the image modality. Concatenative fusion of the two modalities into a series of linear layers seems to be the most common approach, with other works using multiplicative or additive fusion. Cross-attention fusion of the two modalities seems to be another popular approach that seems to give promising results. The approach given by Pranesh (2022) proposes a transformer-like structure to fuse the two modalities. We build on the this previous work and explore the existing structures on the newly annotated data, WildFireCan-MMD.

In addition to fine-tuned structures, we explore using VLMs in a zero-shot setting (Zhang et al., 2024; Li et al., 2022). Specifically, we evaluate four VLMs, GPT-4o-mini (Achiam et al., 2023), LLaVA-v1.5-13B (Liu et al., 2023), Qwen2.5-VL7B (Bai et al., 2025), and SmolVLM-2B (Marafioti et al., 2025), to assess performance across a range of capabilities and computational requirements. GPT-4o-mini serves as a state-of-the-art, closed-source benchmark, while LLaVA and SmolVLM are open-source models that use visual instruction

tuning and are optimized for varying levels of hardware, reflecting ongoing efforts to build accessible, high-performing multimodal systems.

### 3 The WildFireCan-MMD Dataset

#### 3.1 Data Collection

To construct the dataset, X’s Pro-Tier paid API was utilized to collect a diverse set of posts, including images and text, surrounding the British Columbia and Alberta wildfires of 2022 and 2023, as well as the Jasper wildfires of 2024. The dataset comprises 4,688 unique image-text pairs gathered from the wildfire season (April-September) of each respective year. Data collection was conducted using a hashtag-based search strategy as seen in other works (Alam et al., 2018; Gupta and Hewett, 2020; Cruickshank and Carley, 2020). Table 1 presents the full queries used for data collection. To curate the hashtags, basic hashtags such as "#wildfire, #forestfire" were first used. Then, through an iterative process of adding new hashtags by collecting sample tweets, more relevant hashtags were discovered. This approach of using hashtags was necessary due to the limited availability of geotagged content, as most users on the platform do not enable precise geo-location. Instead, hashtags identified through iterative experimentation were found to be effective indicators of relevant content originating from Canada during wildfire events. The specificity of these hashtags ensured that the collected tweets were not only on-topic but also predominantly from users in the targeted geographic regions. The dataset includes the timestamp of when the tweet was posted, the user-provided location, and, in some cases, the specific region from which the tweet originated.

Year	Query
2022/2023	(#BCwildfire OR #BCfire OR #AB-Wildfire OR #albertawildfire OR #AB-Fire) -has:videos has:images lang:en -is:retweet -is:quote -is:reply
2024	(#BCwildfire OR #BCfire OR #ABWildfire OR #albertawildfire OR #ABFire OR #JasperStrong OR #JasperWildfire OR #JasperAB) -has:videos has:images lang:en -is:retweet -is:quote -is:reply

Table 1: The queries used for collecting the data for WildFireCan-MMD.

#### 3.2 Taxonomy Development

To examine the content of the tweets and discover the hidden latent semantics of the data, BERTopic

was employed to perform multimodal unsupervised topic modeling (Grootendorst, 2022). Text and image inputs were vectorized using the pre-trained “Clip-ViT-B-32”<sup>1</sup> model as suggested by BERTopic. The default parameter settings of BERTopic were used, which reduced the dimensionality to 5, while keeping the size of the local neighborhood at 15. Stop words were removed, and both unigrams and bigrams were considered as candidates of topic keywords. Using BERTopic on the dataset, the algorithm determined the number of topics to categorize the data into, resulting in 100 fine-grained topics. Hierarchical topic reduction was applied using Top2Vec, and the number of topics was gradually decreased. For each iteration, the representative text and visual samples of topics and keywords generated by BERTopic were qualitatively inspected. Once the number of topics was reduced to around 15, it became clear which topics were the strongest in the dataset, and the humanitarian utility of some of these topics began to appear. Some examples of useful topics included posts about infrastructure, weather reports, and reports on the actions of first responders. Less useful posts were also noticed, such as advertisements and political debates/complaints. Some examples of these topics, along with representative images and corresponding keywords generated by BERTopic, are presented in Appendix B.

The topics were inspected and 12 classes were decided upon. The full list of classes and their descriptions is given in Table 2. To ensure the reliability of the labeling process, the definitions and boundaries of each category were discussed among the research team and a shared understanding was reached. Through this discussion, it was determined that the classes are well-defined and grounded in observable cues (e.g., specific topics, hazards, or references in the text or image). This discussion, along with unsupervised topic modeling (BERTopic) were used to find strong topics which were then later assigned to a class. Also, the selection of classes was guided by the topics created by BERTopic, observations made through manual annotation, and the perceived relevance of classes to stakeholders in disaster scenarios. For example, ‘*Warnings & Status Updates*’, ‘*Reports of Actions of Responders*’, and ‘*Infrastructure*’ could be useful to first responders and emergency person-

<sup>1</sup><https://huggingface.co/sentence-transformers/clip-ViT-B-32>

Table 2: Categories of WildFireCan-MMD, their descriptions and total count of each class.

Category	Description	Count
Evacuees	Information relating to evacuees, their movements, needs, location.	252
<b>Smoke and Air Quality</b>	Tweets related to or showing signs of smoke or the current air quality.	1128
<b>General Information</b>	General info about the wildfire situation, such as total hectares burned.	170
<b>Preparedness</b>	Information for the general public to prepare themselves and their property.	264
<b>Weather Reports</b>	Information relating to the weather with a specific location mentioned.	296
<b>Warnings &amp; Status Updates</b>	Fire bans in certain areas, new information about a specific area, updates from officials.	669
<b>Reports of Actions of Responders</b>	Actions of responders within specific areas or times, including prescribed burns.	356
<b>Infrastructure</b>	Detours, road closures, damage to infrastructure (e.g., utility poles, highways), repairs by crews.	264
<b>Political</b>	Posts directed towards political figures or parties (excluding situation updates).	329
<b>Support</b>	Mental health and financial support, temporary housing for livestock.	178
<b>Insurance</b>	Information relating to insurance, employment insurance, and EI benefits.	158
<b>Advertisement</b>	Posts about food, restaurants, off-topic ads for services (e.g., apps, air purifiers, not insurance-related).	117
<b>Other</b>	No 'useful' information, focusing on images (e.g., scenery), general complaining, or irrelevant content.	507

nel. The ‘Evacuees,’ ‘General Information,’ and ‘Preparedness’ categories could be useful to the general public to be informed about the events happening around their locality. The ‘Political’ category could be useful to the media to gauge public opinion and policies on events surrounding the disaster. An ‘Other’ class is added to accommodate the samples that do not fit into any of the 12 classes.

Table 3: Annotation agreement statistics.

Metric	Agreement
Majority Agreement (2 same)	88.4%
Full Agreement (all same)	45.1%
Vote between all/annotator 1	78.7%
Vote between all/annotator 2	77.7%
Vote between all/annotator 3	65.8%
Cohen’s Kappa annotator 1-2	63.5%
Cohen’s Kappa annotator 1-3	50.6%
Cohen’s Kappa annotator 2-3	49.2%
Fleiss’ Kappa	54.3%

### 3.3 Data Labeling

Annotation was performed in Label Studio<sup>2</sup>, using the defined taxonomy presented in Table 2. Examples of labeled tweets are presented in Figure 1. Each image-text pair was given a single label. In cases where the label could differ depending on whether text or image was being considered, preference was given to the modality that provided greater value for that particular sample. For example, in Figure 1c the image might be labeled ‘Smoke & Air Quality’, but the text might be ‘Warnings & Status Updates’. In this case, ‘Warnings & Status Updates’ would be chosen as the label as the text contains more valuable information than the image alone. If neither contained detailed information about the time or location like the text does in this example, then the sample would be labeled as

‘Smoke & Air Quality’ as that is the only value that would otherwise be provided. The size of labeled classes is presented in Table 2.

Samples were annotated by three independent annotators, all authors of this paper. The annotators are male, 20-30 year old, fluent in English and reside in Canada. First, initial rounds of manual labeling on small subsets of the data gave a good understanding of what was being discussed and shown in the data, and confidence was gained in the decided classes. Over the course of labeling, samples were flagged that contained personal information. A vote between the labels from the three annotators decided the final labels. Annotation agreement scores can be seen in Table 3. A majority agreement was reached on 88.4% of the labels. Taking a vote between all annotators for the final label and comparing the result with each annotator, we can see which annotator was closest to the voted labels. From the results, annotators 1 & 2 had the best agreement scores. This is further demonstrated by Cohen’s Kappa metric between the different pairs of annotators. Fleiss’ Kappa metric however, indicates that a moderate level of agreement is achieved between the independent annotators. In cases of disagreement, the final label was chosen from the ‘expert’ annotator (Annotator 1), who had the most experience with the dataset collection and devising the taxonomy, and also had the highest agreement score with other annotators. Additional statistics can be found in Appendix A.

## 4 Classification Methodology

This section describes a range of classifiers evaluated for the 13-class (12 categories + ‘Other’) classification of WildFireCan-MMD. We investigate three categories of classification methods. First,

<sup>2</sup><https://labelstud.io/>



**Label:** Reports of Actions of Responders

(a) **Text:** PAFD is on the scene of a fire in Roger Creek Park. There is no threat to structures at this time. #portalberni #BCWildfire



**Label:** Infrastructure

(b) **Text:** The wildfires in #AB have burned down this schoolhouse. #ABWildfires



**Label:** Warnings & Status Updates

(c) **Text:** Wildfire updates, July 24: Fire reaches Jasper townsite | 'Significant loss has occurred': Parks Canada | Firefighters working to protect critical infrastructure #abfire

Figure 1: Examples of labeled posts from WildFireCan-MMD.

classical methods such as Decision Trees (DT), Gaussian Naive Bayes (GNB), K-Nearest Neighbours (KNN) and Support Vector Machines (SVM) are used to investigate a case where labeled data is available but computation is limited. Then, several VLMs such as LLaVA, Qwen, smolVLM, and GPT-4o-mini are employed as zero-shot classifiers. This is the right choice when no training data is available. Finally, we experimented with a fine-tuned transformer-based classifier to assess its performance in the presence of labeled data and moderate computational resources. Further information on classical methods and VLMs can be found in Appendices D and E.

**Architecture of custom classifier:** To determine which components to use in the selected architecture, we first experimented with various unimodal classifiers for each modality, as seen in Tables 4 & 5. These included CNNs and transformers for the image modality, as well as various BERT-based transformers for the text modality. After individually determining the best model for each modality (ViT for image with f1: 64.19%, and RoBERTA for text with f1: 77.98%), various fusion techniques found in the literature and included in Table F.1 were tested to improve the classification. The results of these multimodal experiments can be seen in Table 6. More information about these experiments, as well as an ablation study, can be found in Appendix G.

The best performing classifier consisted of two parallel encoders, an image and text encoder. The image encoder is ViT<sup>3</sup>, and the text encoder is

RoBERTa<sup>4</sup>. The output CLS token of the last hidden state from each are sent through separate projection layers to reduce the dimensionality from 768 to 512. The output is passed to the transformer fusion module consisting of two encoder layers, each with d\_model=512, nhead=8, dim\_feedforward=2048, dropout=0.2. A final classification head projects the dimension from 512 to 256, which is followed by ReLU activation and dropout=0.2, then finally from 256 to num\_classes. Figure H.1 shows the structure of this architecture.

**Training and Evaluation:** Three data splits with stratified sampling were used to train the model and report the results. The random seeds from splitting the data were kept for reproducibility. The hyperparameters used were a batch size of 8 and an image size of  $(384 \times 384 \times 3)$ . Text sequences are padded to a max length of 144. Adam optimizer is used with a learning rate of  $1e^{-5}$ , weight\_decay=0.01, and cross-entropy loss function. The model is fully trained after 10 epochs. The best model is saved based on f-weighted-average. The model was tested on the held-out test set (938 samples).

Table 4: Text-only classification results.

Text Model	Notes	F1-Weighted
'bert-base-cased'	mean-pooling	0.7779
'bert-base-cased'	CLS	0.7793
"bertweet-base"	mean-pooling	0.7900
"bertweet-base"	CLS	0.7927
"ModernBERT-base"	mean-pooling	0.7934
"ModernBERT-base"	CLS	0.7862
"roberta-base"	robertaforseq	0.7718
"roberta-base"	mean-pooling	0.7934
"roberta-base"	mean-pooling, half-frozen	0.7940
"roberta-base"	mean-pooling, custom vocab	0.7209
"roberta-base"	CLS	0.7798
"roberta-base"	CLS, half-frozen	<b>0.8008</b>

<sup>3</sup>google/vit-base-patch16-384'

<sup>4</sup>'roberta-base'

Table 5: Image-only classification results.

Image Model	Notes	F1-Weighted
DenseNet121	ImageNet weights	0.5733
InceptionV3	ImageNet weights	0.5804
ResNet50	ImageNet weights	0.5841
VGG16	ImageNet weights	0.5993
"facebook/deit-base-patch16-384"	CLS	0.6260
"google/vit-base-patch16-384"	half-frozen, CLS	0.6290
"google/vit-base-patch16-384"	mean-pooling	0.6404
"google/vit-base-patch16-384"	CLS	<b>0.6419</b>

## 5 Experimental Results

Tables 4 and 5 show the unimodal text/image models experimented with. All experiments were conducted using the same train/test split. Experiments were done on different final feature representations by using either mean pooling of the last hidden state or the CLS token. Freezing the first half of the encoder layers was also tried as this has been shown to increase accuracy (Ingle et al., 2022). All encoders had a linear layer appended for classification. For text, "roberta-base" performed best when using just the CLS token for classification and freezing the first half of the layers, as shown in Table 4. Later experimentation found that the unfrozen version performed better in the final classifier. For images, "vit-base-patch16-384" performed best when using just the CLS token for classification without any freezing. Many of the fusion experiments tested, as well as some architectures from other studies can be seen in Table 6. The same train/test split was used as before for unimodal classification. The best result, belonging to the selected classifier, is shown in bold. Table 7 shows the calculated f-scores for the VLMs, selected and baseline classifiers per class and as an average over all classes. Appendix I shows the confusion matrices of these classifiers. All metrics reported in these tables use 'weighted' F1-score to represent the class imbalance observed in the dataset.

The custom classifier achieved the highest weighted average F1-score (84.48%), outperforming both zero-shot vision-language models and traditional baseline classifiers across all classes. From the results in Table 7 and Appendix I, the '*Preparedness*' class had the lowest F1-score under the custom classifier ( $66.02 \pm 4.10\%$ ), likely due to high variability in both imagery and associated text. On the other hand, the '*Warnings & Status Updates*' class achieved the highest performance ( $93.19 \pm 2.02\%$ ), reflecting the strong visual and textual consistency across samples. Other high-performing classes include

'*Advertisement*' ( $89.37 \pm 1.71\%$ ) '*Weather Reports*' ( $89.01 \pm 4.51\%$ ), and '*Evacuees*' ( $88.55 \pm 3.20\%$ ), which may be more visually or textually distinct. In contrast, mid-range scores are seen in classes like '*Support*' ( $77.72 \pm 2.26\%$ ) and '*Political*' ( $83.73 \pm 0.29\%$ ), which often contain dense textual overlays that may not align well with the image modality.

## 6 Discussion

**Task-specific fine-tuning is critical:** Table 7 shows that classical models fine-tuned on labeled data generally outperform zero-shot VLMs, with the custom model achieving 23% higher f-score than VLMs. Despite a modest dataset, results affirm the value of targeted annotation and model tuning, echoing findings in literature presented in Appendix F.

**VLMs involve tradeoffs:** Zero-shot VLMs, while being a practical option when no labeled data is available, pose challenges in accuracy and cost. GPT-4o-mini offers strong results, but is costly and lacks transparency. Open-source models like LLaVA, Qwen, and smolVLM are more accessible but generally underperform, though Qwen is a notable exception, and trails GPT-4 by just 5%. Overall, VLMs may offer a solution when labeled data is unavailable, but their performance is comparable to traditional models such as SVM and KNN, which require training data but can be implemented with minimum computation.

**Future Work:** The dataset exhibits a significant class imbalance, and future work may involve steps to address this issue. Utilizing the unlabeled data paired with a labelled dataset in a semi-supervised setup is a promising avenue for future research (Sirbu et al., 2022; Sîrbu et al., 2025). Fine-tuning of VLMs is also a promising avenue which should be explored in future works (Duan et al., 2024; Zhai et al., 2024).

## 7 Analysis on Trends in Data

This section describes the continuation of the analysis from the work by Sherritt et al. (2025). Here, we present an analysis of an additional 46,279 posts from X, posted between 2018 and 2024, using the custom classifier described in Section 4. The collection methodology was the same as explained in section 3.1, where tweets were gathered using a curated set of wildfire-related hashtags specific to

Table 6: Multimodal classification results.

Text Model	Image Model	Fusion Type	Notes	F1
RoBERTa	ViT	Cross Attention	encoders pretrained individually, q=image k=text v=text	0.7855
RoBERTa	ViT	Cross Attention	encoders pretrained individually, q=text k=image v=image	0.6489
RoBERTa	Coca	Concatenative	COCCA/RoBERTa individually aligned w/ description of the label	0.1587
RoBERTa+BiLSTM	ViT	BiLSTM+BiLinear	Koshy's model (Koshy and Elango, 2023)	0.6149
CLIP	CLIP		"openai/clip-vit-base-patch32"	0.7072
VILT	VILT		"dandelin/vilt-b32-finetuned-nlvr2"	0.7638
RoBERTa	ViT	Concatenative	ViT/RoBERTa individually aligned w/ description of the label	0.6507
RoBERTa	ViT	Transformer	transformer fusion: $h\_dim$ 1024 and 0.1 dropout	0.8263
RoBERTa	ViT	Transformer	transformer fusion: $h\_dim$ 2048 and 0.2 dropout	<b>0.8364</b>
RoBERTa	ViT	Transformer	RoBERTa/ViT fully frozen	0.6164
RoBERTa	ViT	Concatenative	stopwords removed from text before RoBERTa	0.7870
RoBERTa	ViT	Concatenative	CLS into single classification layer	0.8249
RoBERTa x2	ViT	Concatenative	extract text from images with OCR	0.8102
ModernBERT	ViT	Transformer	transformer fusion: hidden dim 2048 and 0.2 dropout	0.8149
ModernBERT	ViT	Concatenative	CLS into single classification layer	0.7966
RoBERTa	ViT	Concatenative	mean-pooling of last hidden states into linear classification layer	0.7911
RoBERTa	ViT	Concatenative	roberta:mean, vit:cls	0.8028
RoBERTa	ViT	Transformer	transformer fusion: $h\_dim$ 2048 and 0.2 dropout RoBERTa:mean, ViT:cls	0.8223
RoBERTa	ViT	Cross Attention	same as proposed but with cross attention instead of transformer	0.8208

Table 7: F1-scores per-class for zero-shot VLMs, baseline and selected classifiers.

Class	Baseline Classifiers				Classifier	zero-shot VLMs			
	DT	GNB	KNN	SVM w/ PCA		smoVLM	LLaVA	Qwen	GPT-4o-mini
Evacuees	39.34±7.08	45.34±2.26	45.78±4.64	31.65±4.69	88.55±3.20	25.84	17.65	12.70	37.21
General Information	27.04±6.43	39.11±3.73	36.01±0.56	33.64±6.32	82.04±3.54	00.00	19.05	21.88	20.37
Preparedness	50.27±4.05	65.04±4.15	56.25±9.15	52.24±10.58	66.02±4.10	22.58	57.14	67.69	62.50
Weather Reports	51.48±6.89	61.22±3.86	63.70±3.44	65.37±2.89	89.01±4.51	00.00	64.52	50.00	62.39
Warnings & Status Updates	43.70±3.16	46.43±3.48	58.67±0.92	56.20±1.21	93.19±2.02	00.00	34.94	42.73	48.00
Reports of Actions of Responders	47.50±4.05	51.49±2.81	54.84±5.75	64.06±2.15	81.10±2.14	00.00	26.09	63.89	69.09
Infrastructure	55.13±8.21	67.51±6.28	69.16±6.58	66.01±8.15	87.88±1.86	00.00	36.36	50.00	37.50
Political	50.04±4.69	69.18±1.74	54.64±1.51	58.89±3.39	83.73±0.29	00.00	69.09	85.25	77.86
Insurance	86.66±1.52	91.25±1.31	90.57±3.62	92.09±1.12	87.01±1.22	93.75	88.52	90.91	86.15
Advertisement	48.80±6.41	49.56±7.07	65.92±6.43	66.11±3.23	89.37±1.71	00.00	08.33	45.16	73.17
Smoke and Air Quality	58.47±2.82	72.05±3.92	75.19±1.94	76.84±1.21	84.20±6.11	45.48	72.46	65.96	75.38
Support	47.99±10.26	60.76±4.66	47.53±6.10	43.86±6.57	77.72±2.26	05.86	27.27	43.84	56.52
Other	46.37±2.53	56.17±1.76	48.85±2.65	58.93±1.65	81.29±1.00	25.36	05.22	52.46	57.58
<b>F1 Weighted Average</b>	<b>50.88±1.11</b>	<b>60.77±1.45</b>	<b>61.22±1.23</b>	<b>62.21±0.44</b>	<b>84.48±0.69</b>	<b>19.86</b>	<b>44.69</b>	<b>55.94</b>	<b>61.17</b>

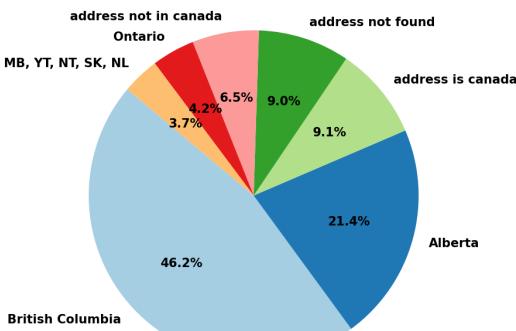


Figure 2: Distribution of collected tweets based on author location.

each year, selected based on known wildfire events and hotspot locations across affected regions. This was combined with a keyword-based search, and then the data was combined, and duplicates were removed. The full set of search queries can be seen in Appendix J.

The author locations were extracted from the

samples with the goal of obtaining the province-level region from which each sample was posted. These could be either automatically generated or user-provided. The queries were run through OSMNX<sup>5</sup>, which allowed for the automated search of cities or towns that users provided. If a user's location was not found, or not in Canada, it was labeled as such. The distribution of posts by province is shown in Figure 2, indicating that the majority of the data originates from British Columbia and Alberta. Additional information about this data, including a word cloud of the text, the frequency distribution of the posts by word count, can be found in Appendix K.

Using the custom classifier, predictions were obtained for the new data. This automated classification approach allowed for efficiency in processing the large volume of newly collected data without having to manually label each individual sample. This of course, will not uphold the same

<sup>5</sup><https://osmnx.readthedocs.io/en/stable/>

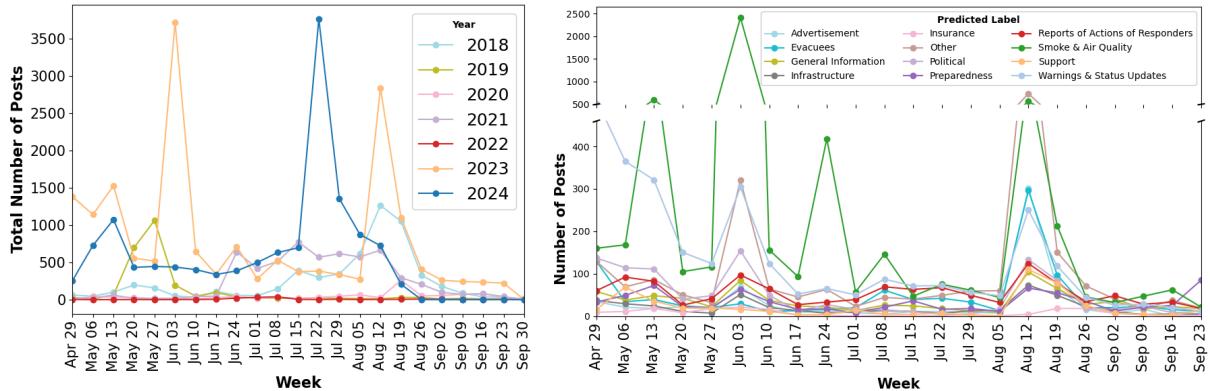


Figure 3: Weekly post distributions: (left) overall yearly trend, (right) class-wise distribution in 2023.

gold standard of labeling that a human annotator could achieve. However, while individual classifications may not always be accurate, general trends in the data can be analyzed and linked to real events.

Figure 3 (left) shows the total number of posts per week sorted by year. This demonstrates a significant event occurred the week of June 3, 2023, and July 22, 2024, due to several major fires in the country (Reuters Staff, 2023; CBC News, 2024). Figure 3 (right) illustrates weekly trends in predicted labels during 2023. Sharp spikes in *Smoke and Air Quality* and *Warnings & Status Updates* likely correspond to major wildfire events occurring in Alberta and BC. *Warnings & Status Updates* and *Evacuees* follow similar patterns around the week of Aug 12, which corresponds to the state of emergency declared as the city of West Kelowna was evacuated. Overall, at least 35,000 people were under evacuation order and another 30,000 under evacuation alert, as of August 19 (Williams, 2023). *Reports of Actions of Responders* and *Support* also rise during these peaks, reflecting coordinated emergency efforts.

## 8 Conclusion

Despite the existence of other natural disaster datasets from social media, these datasets often lack wildfire data from a significant period of time, typically spanning only a single month surrounding a particular event. Furthermore, even less of this data comes from Canada and is mainly restricted to the US context. To address this gap, we present WildFireCan-MMD, a wildfire-specific, multimodal dataset of social media posts curated from X, accompanied by a domain-focused tax-

onomy derived through multimodal topic modeling and manual annotation by multiple independent annotators. Unlike existing disaster datasets, WildFireCan-MMD is specifically designed for the Canadian wildfire context and provides a rich resource for training and evaluating specialized models. It also serves as a true test set for evaluating VLMs in a zero-shot setting. We provide the results of classifiers trained on this dataset from classical approaches to custom deep learning architectures and several VLMs, to showcase various scenarios in terms of data and computational resource availability. Our best results are obtained with an architecture that utilizes pre-trained encoders and transformer fusion for classification on this dataset, which highlights the importance of a labelled dataset. As climate-induced wildfires grow in frequency and complexity, this dataset and model support the development of robust tools for detection, coordination, and recovery.

## Acknowledgment

This work was funded by the National Program Office of the National Research Council Canada, under the New Beginning Initiative Program.

## Limitations

While this study introduces a novel wildfire-specific multimodal dataset and demonstrates the effectiveness of fine-tuned task-specific models, several limitations should be acknowledged.

The first limitation is the scope and size of the dataset. Although WildFireCan-MMD provides a targeted and realistic benchmark for wildfire-related content, its size remains modest compared to large-scale datasets used to train or evaluate deep

learning models. This limits the diversity of linguistic and visual expressions captured and may affect the generalizability of these results to other wildfire events or regions.

Platform and geographic bias might be present in our data. The dataset is sourced exclusively from the X platform and focuses on Canadian wildfires. This introduces both platform-specific and regional biases in language use, imagery, and public discourse. As a result, models trained on this dataset may not perform as well when applied to other social platforms or international wildfire contexts.

Moreover, zero-shot evaluation of VLMs introduces limitations. Prompt engineering for zero-shot setups is a heuristic process, and better prompts might yield different results.

Finally, there is a broader risk of misapplication. Deploying these models in real-world disaster contexts without appropriate oversight, interpretability mechanisms, or human-in-the-loop verification may lead to misclassification, misinformation amplification, or missed signals, potentially affecting emergency coordination and public trust.

## References

Mahdi Abavisani, Liwei Wu, Shengli Hu, Joel Tetreault, and Alejandro Jaimes. 2020. [Multimodal categorization of crisis events in social media](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14667–14677.

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Turgut Acikara, Bo Xia, Tan Yigitcanlar, and Carol Hon. 2023. Contribution of social media analytics to disaster response effectiveness: A systematic review of the literature. *Sustainability (Switzerland)*, 15(11):8860.

Mansi Agarwal, Maitree Leekha, Ramit Sawhney, and Rajiv Ratn Shah. 2020. [Crisis-dias: Towards multimodal damage analysis - deployment, challenges and assessment](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(01):346–353.

Firoj Alam, Ferda Ofli, and Muhammad Imran. 2018. [Crisismmd: Multimodal twitter datasets from natural disasters](#). In *12th International AAAI Conference on Web and Social Media, ICWSM 2018*.

Firoj Alam, Umair Qazi, Muhammad Imran, and Ferda Ofli. 2021. [Humaid: Human-annotated disaster incidents data from twitter with deep learning benchmarks](#). *Preprint, arXiv:2104.03090*.

Nilani Algiriyyage, Raj Prasanna, Kristin Stock, Emma E. H. Doyle, and David Johnston. 2022. [Multi-source multimodal data and deep learning for disaster response: A systematic review](#). *SN Computer Science*, 3.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025. [Qwen2.5-v1 technical report](#). *Preprint, arXiv:2502.13923*.

Mohammad Abdul Basit, Salman Ghufran Shaikh, Bashir Alam, and Zubaida Fatima. 2023. [Natural disaster tweets classification using multimodal data](#). In *EMNLP 2023 - 2023 Conference on Empirical Methods in Natural Language Processing, Proceedings*.

Umar Ali Bukar, Marzah A. Jabar, Fatimah Sidi, Rozi Nor Haizan Binti Nor, Salfarina Abdullah, and Mohamed Othman. 2020. [Crisis informatics in the context of social media crisis communication: Theoretical models, taxonomy, and open issues](#). *IEEE Access*, 8:185842–185869.

CBC News. 2024. [Wildfire near jasper, alta., now classified as 'being held'](#). *CBC News*. Accessed: 2025-05-15.

Iain Cruickshank and Kathleen Carley. 2020. [Characterizing communities of hashtag usage on twitter during the 2020 covid-19 pandemic by multi-view clustering](#). *Applied Network Science*, 5.

Georgina Curto, Svetlana Kiritchenko, Kathleen C Fraser, and Isar Nejadgholi. 2024. The crime of being poor: Associations between crime and poverty on social media in eight countries. In *Proceedings of the Sixth Workshop on Natural Language Processing and Computational Social Science (NLP+ CSS 2024)*, pages 32–45.

Sumit Dalal, Sarika Jain, and Mayank Dave. 2024. [Review of advancements in depression detection using social media data](#). *IEEE Transactions on Computational Social Systems*.

Shahid Shafi Dar, Mohammad Zia Ur Rehman, Karan Bais, Mohammed Abdul Haseeb, and Nagendra Kumar. 2025. [A social context-aware graph-based multimodal attentive learning framework for disaster content classification during emergencies](#). *Expert Systems with Applications*, 259:125337.

Adwaith Divakaran, Ashok K. Abishake, Siva V. Raghul, and Sivasankar Elango. 2022. [Enhancing multimodal disaster tweet classification using state-of-the-art deep learning networks](#). *Multimedia Tools and Applications*, 81(13):18483–18501. Copyright - © The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022; Last updated - 2024-11-13.

Zhijie Sasha Dong, Lingyu Meng, Lauren Christenson, and Lawrence Fulton. 2021. *Social media information sharing for natural disaster response*. *Natural Hazards*, 107.

Zhizhao Duan, Hao Cheng, Duo Xu, Xi Wu, Xiangxie Zhang, Xi Ye, and Zhen Xie. 2024. Cityllava: Efficient fine-tuning for vlms in city scenario. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7180–7189.

Lokabhiram Dwarakanath, Amirrudin Kamsin, Rasheed Abubakar Rasheed, Anitha Anandhan, and Liyana Shuib. 2021. Automated machine learning approaches for emergency response and coordination via social media in the aftermath of a disaster: A review. *IEEE Access*, 9:68917–68931.

Muhammad Ashraf Fauzi. 2023. *Social media in disaster management: review of the literature and future trends through bibliometric analysis*. *Natural Hazards*, 118.

Hamideh Ghanadian, Isar Nejadgholi, and Hussein Al Osman. 2024. Socially aware synthetic data generation for suicidal ideation detection using large language models. *IEEE Access*.

Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.

Shubham Gupta, Nandini Saini, Suman Kundu, and Debasis Das. 2024. Crisiskan: Knowledge-infused and explainable multimodal attention network for crisis event classification. *Preprint*, arXiv:2401.06194.

Vibhuti Gupta and Rattikorn Hewett. 2020. Real-time tweet analytics using hybrid hashtags on twitter big data streams. *Information*, 11(7).

Eftekhar Hossain, Mohammed Moshiul Hoque, Enamul Hoque, and Md. Saiful Islam. 2022. A deep attentive multimodal learning approach for disaster identification from social media posts. *IEEE Access*, 10:46538–46551.

J. Brian Houston, Joshua Hawthorne, Mildred F. Perreault, Eun Hae Park, Marlo Goldstein Hode, Michael R. Halliwell, Sarah E. Turner McGowen, Rachel Davis, Shivani Vaid, Jonathan A. Mcelderry, and Stanford A. Griffith. 2015. *Social media and disasters: A functional framework for social media use in disaster planning, response, and research*. *Disasters*, 39.

Muhammad Imran, Umair Qazi, and Ferda Ofli. 2022. Tbcov: Two billion multilingual covid-19 tweets with sentiment, entity, geo, and gender labels. *Data*, 7(1).

Digvijay Ingle, Rishabh Tripathi, Ayush Kumar, Kevin Patel, and Jithendra Vepa. 2022. Investigating the characteristics of a transformer in a few-shot setup: Does freezing layers in RoBERTa help? In *Proceedings of the Fifth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 238–248, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Anuradha Khattar and S. M. K. Quadri. 2022. Camm: Cross-attention multimodal classification of disaster-related tweets. *IEEE Access*, 10:92889–92902.

Svetlana Kiritchenko, Isar Nejadgholi, and Kathleen C Fraser. 2021. Confronting abusive language online: A survey from the ethical and human rights perspective. *Journal of Artificial Intelligence Research*, 71:431–478.

Rani Koshy and Sivasankar Elango. 2023. Multimodal tweet classification in disaster response systems using transformer-based bidirectional attention model. *Neural Computing and Applications*, 35.

S. Kotha, S. Haridasan, A. Rattani, A. Bowen, G. Rimington, and A. Dutta. 2022. Multimodal combination of text and image tweets for disaster response assessment. In *Data-driven Resilience Research 2022*, Leipzig, Germany.

Ankit Kumar, Jaspreet P. Singh, Yogesh K. Dwivedi, and Nripendra P. Rana. 2022. A deep multi-modal neural network for informative twitter content classification during emergencies. *Annals of Operations Research*, 319:791–822.

Gunjan Kumar, Jyoti Prakash Singh, and Amit Kumar Singh. 2024a. Detecting disseminators of hate speech by optimizing textual and user-based features. *IEEE Transactions on Computational Social Systems*.

Raghvendra Kumar, Ritika Sinha, Sriparna Saha, and Adam Jatowt. 2024b. Extracting the full story: A multimodal approach and dataset to crisis summarization in tweets. *IEEE Transactions on Computational Social Systems*.

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*.

Feng Li, Hao Zhang, Yi-Fan Zhang, Shilong Liu, Jian Guo, Lionel M Ni, PengChuan Zhang, and Lei Zhang. 2022. Vision-language intelligence: Tasks, representation learning, and large models. *arXiv preprint arXiv:2203.01922*.

Xukun Li and Doina Caragea. 2020. Improving disaster-related tweet classification with a multimodal approach. In *Proceedings of the 17th International Conference on Information Systems for Crisis Response and Management (ISCRAM)*, pages 893–902. Social Media for Disaster Response and Resilience.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. *Advances in Neural Information Processing Systems*, 36.

Yinhan Liu. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Zhiyuan Ma, Meiqi Pan, Yunfeng Hou, Ling Yang, and Wei Wang. 2024. Toward knowledge integration with large language model for end-to-end aspect-based sentiment analysis in social multimedia. *IEEE Transactions on Computational Social Systems*.

Sreenivasulu Madichetty, Sridevi M, and Sreekanth Madisetty. 2023. **A roberta based model for identifying the multi-modal informative tweets during disaster.** *Multimedia Tools and Applications*, 82(24):37615–37633. Copyright - © The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023. Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law; Last updated - 2024-12-14.

Sreenivasulu Madichetty, Sridevi Muthukumarasamy, and P. Jayadev. 2021. **Multi-modal classification of twitter data during disasters for humanitarian response.** *Journal of Ambient Intelligence and Humanized Computing*, 12:1–15.

Andrés Marafioti, Orr Zohar, Miquel Farré, Merve Noyan, Elie Bakouch, Pedro Cuenca, Cyril Zakka, Loubna Ben Allal, Anton Lozhkov, Nouamane Tazi, Vaibhav Srivastav, Joshua Lochner, Hugo Larcher, Mathieu Morlon, Lewis Tunstall, Leandro von Werra, and Thomas Wolf. 2025. Smolvlm: Redefining small and efficient multimodal models. *Preprint*, arXiv:2504.05299.

Mohammed S. Mundottukandi, Yusmadi Yah Jusoh, Noraini Che Pa, Rozi Nor Binti Haizan Nor, and Umar Ali Bukar. 2024. **Prioritizing factors in social media crisis communication for resilience enhancement using analytical hierarchy process.** *IEEE Access*, 12:54798–54815.

Kamalich Muniz-Rodriguez, Sylvia K Ofori, Lauren C Bayliss, Jessica S Schwind, Kadiatou Diallo, Manyun Liu, Jingjing Yin, Gerardo Chowell, and Isaac Chun-Hai Fung. 2020. Social media use in emergency response to natural disasters: a systematic review with a public health perspective. *Disaster medicine and public health preparedness*, 14(1):139–149.

Karla Saldana Ochoa and Tina Comes. 2021. A machine learning approach for rapid disaster response based on multi-modal data. the case of housing & shelter needs. *arXiv preprint arXiv:2108.00887*.

Ferda Ofli, Firoj Alam, and Muhammad Imran. 2020. Analysis of social media data using multimodal deep learning for disaster response. In *Proceedings of the International ISCRAM Conference*, volume 2020-May.

Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2014. Crisislex: A lexicon for collecting and filtering microblogged communications in crises. *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1):376–385.

Alexandra Olteanu, Sarah Vieweg, and Carlos Castillo. 2015. **What to expect when the unexpected happens: Social media communications across crises.** In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW ’15, page 994–1009, New York, NY, USA. Association for Computing Machinery.

Raj Pranesh. 2022. **Exploring multimodal features and fusion strategies for analyzing disaster tweets.** In *Proceedings of the Eighth Workshop on Noisy User-generated Text (W-NUT 2022)*, pages 62–68, Gyeongju, Republic of Korea. Association for Computational Linguistics.

Juan Ramirez, Chinasa Okolo, Mahima Hegde, Soumya Banerjee, Aayush Wadhwa, Rashmi Gangadharaiyah, and Niranjan Balasubramanian. 2025. **Using multimodal models for informative classification of disaster tweets.** In *Proceedings of the 3rd Workshop on Natural Language Processing for Digital Humanities (NLP4DH)*, pages 221–229, St. Julians, Malta. Association for Computational Linguistics.

Reuters Staff. 2023. **Fires rage in british columbia as more residents prepare for evacuations.** *Reuters*. Accessed: 2025-05-15.

Mariham Rezk, Noureldin Elmadany, Radwa K. Hamad, and Ehab F. Badran. 2023. **Categorizing crises from social media feeds via multimodal channel attention.** *IEEE Access*, 11:72037–72049.

Braeden Sherritt, Isar Nejadgholi, and Marzieh Amini. 2025. **Analysis of Canadian Wildfire Tweets Over Seven Years.** *Proceedings of the Canadian Conference on Artificial Intelligence*. [Https://caiac.pubpub.org/pub/nzt63ve8](https://caiac.pubpub.org/pub/nzt63ve8).

N. P. Shetty, Y. Bijalwan, P. Chaudhari, and 1 others. 2025. **Disaster assessment from social media using multimodal deep learning.** *Multimedia Tools and Applications*, 84:18829–18854.

Iustin Sirbu, Tiberiu Sosea, Cornelia Caragea, Doina Caragea, and Traian Rebedea. 2022. **Multimodal semi-supervised learning for disaster tweet classification.** In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2711–2723, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Iustin Sirbu, Robert-Adrian Popovici, Traian Rebedea, and Ștefan Trăusan-Matu. 2025. **Multihead average pseudo-margin learning for disaster tweet classification.** *Information*, 16(6).

Sumiko Teng and Emily Öhman. 2025. **Using multimodal models for informative classification of ambiguous tweets in crisis response.** In *Proceedings of the 5th International Conference on Natural Language Processing for Digital Humanities*, pages 265–271, Albuquerque, USA. Association for Computational Linguistics.

Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. 2021. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pages 10347–10357. PMLR.

Nia Williams. 2023. [Canada to deploy military in british columbia to tackle fast-spreading fires](#). *Reuters*. Accessed: 2025-05-14.

Xuehua Wu, Jin Mao, Hao Xie, and Gang Li. 2022. [Identifying humanitarian information for emergency response by modeling the correlation and independence between text and images](#). *Information Processing & Management*, 59(4):102977.

Simon Zhai, Hao Bai, Zipeng Lin, Jiayi Pan, Peter Tong, Yifei Zhou, Alane Suhr, Saining Xie, Yann LeCun, Yi Ma, and 1 others. 2024. Fine-tuning large vision-language models as decision-making agents via reinforcement learning. *Advances in neural information processing systems*, 37:110935–110971.

Duzhen Zhang, Yahan Yu, Jiahua Dong, Chenxing Li, Dan Su, Chenhui Chu, and Dong Yu. 2024. [MM-LLMs: Recent advances in MultiModal large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 12401–12430, Bangkok, Thailand. Association for Computational Linguistics.

## A Annotation

Figure A.1 shows the distribution of the labels over the classes by each independent annotator.

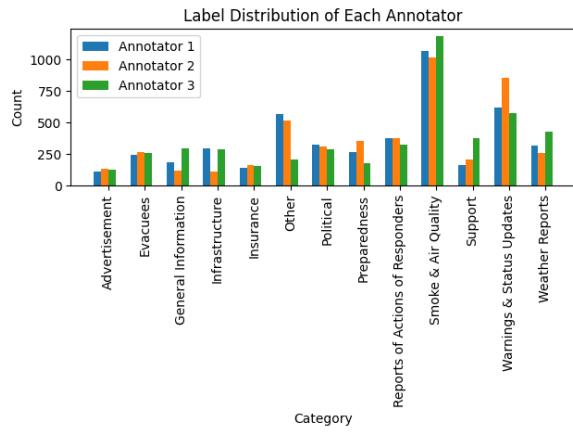


Figure A.1: Counts of labels by each annotator

## B Topic Modeling

Figure B.1 shows examples of topics extracted by the unsupervised BERTopic. We present representative images and corresponding keywords generated by BERTopic for each example. As seen in the captions of the image, some of the topics are not

useful and present non-coherent concepts, while others are useful and relevant to wildfires.

## C Prompts for Zero-Shot Classification

The full prompt used to evaluate the aforementioned VLMs is shown in Table C.1.

Table C.1: System and User prompt used for VLMs.

### System Prompt:

*"You are an assistant who is being given an image and text pair as a Twitter post, which was created during a natural disaster event. Your task is to use information from both the text and the image to decide which option the post should be labeled as. You must pay close attention to each option when deciding which label to use."*

### User Prompt:

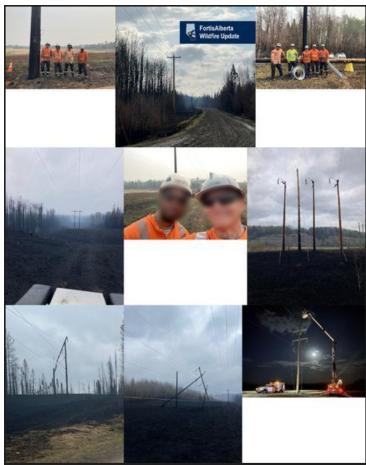
*"Which option should this post be labeled as?*

- A. Evacuees (information relating to evacuees, their movements, needs, location, etc)*
- B. General Information (GENERAL facts about the wildfire situation, hectares burned)*
- C. Preparedness (information for the general public to prepare themselves and property for wildfires)*
- D. Weather Reports (information relating to the weather, satellite imagery, radar imagery)*
- E. Warnings & Status Updates (warnings/updates to the public from authoritative bodies, fire bans, specific information relating to a certain time or area)*
- F. Reports of Actions of Responders (prescribed burns, responders responding to a specific location)*
- G. Infrastructure (road closures, damaged buildings or property, traffic)*
- H. Political (mentions of political or public figures or parties)*
- I. Insurance (mentions of insurance)*
- J. Advertisement (information about restaurants, food, apps or services)*
- K. Smoke & Air Quality (information about smoke or air quality, masks, breathing)*
- L. Support (information about financial, mental health, or other types of support for people)*
- M. Other (the post does not fit well in one of the previous categories)*

*You may only answer with the chosen option's letter."*

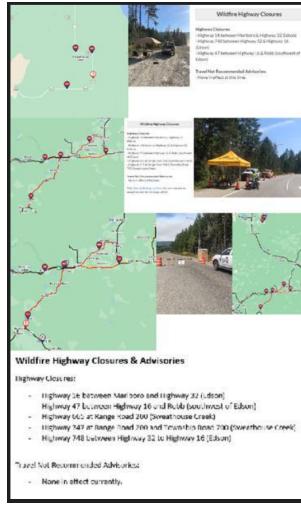
## D Baseline Classifiers

To establish a baseline, several classical methods were evaluated on the dataset, namely DT, GNB, KNN, SVM. For encoding data into numerical features for these methods, (pretrained, but not fine-tuned) RoBERTa and ViT were used to extract feature sets for each modality. The concatenated feature sets from the last hidden state were saved. Principle Component Analysis (PCA) was used on the feature sets to reduce the dimensionality from 768 to 250 for the SVM method as it improved the results. For DT, `class_weight='balanced'` was used. For KNN, `n_neighbors=1` was used. All other hy-



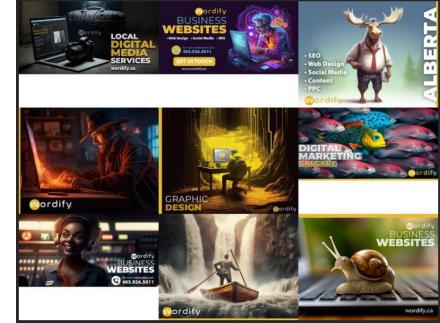
**Keywords:** [crews, safe, poles, damaged, continue, lines, working restore, power, conditions, restore]

(a) Useful



**Keywords:** [highway, closed, closures, abroad, detour, road, route, bchwy4, cameron, cameron lake]

(b) Useful



**Keywords:** ['yyknow', 'yycbusiness', 'yycbusiness yyknow', 'yyc yycbusiness', 'abfires abwildfire', 'calgary', 'startups', 'wordify', 'smallbusiness', 'smallbusiness startups']

(c) Non-useful

Figure B.1: Examples of useful/non-useful topics from topic modeling on our dataset.

perparameters remained default from Scikit-learn<sup>6</sup>. Three different random splits were used to train and test all methods to report a confidence interval.

## E VLMs as Zero-Shot Classifiers:

Four VLMs were considered for experimentation, GPT-4o-mini (Achiam et al., 2023), LLaVA (Liu et al., 2023), smolVLM (Marafioti et al., 2025), Qwen2.5-VL (Bai et al., 2025) and tested in a zero-shot setting. Classification prompts for all four VLMs were created. With each query to the model, a list of all categories and their descriptions are given as a prompt along with a data instance (text + image) to classify. The prompt specifies that the model’s task is to choose one of the provided categories and respond only with its corresponding letter. The system and user prompts can be seen in Appendix C and are the same for all VLMs. The same hyperparameters are used for all VLMs with temperature=0.1 num\_beams=1, max\_new\_tokens=1024. The same test set was used for all VLMs and consisted of a stratified split. This single split was discovered to be a good split from previous experimentation with baseline and custom classifiers, and is one of the same three splits used for these classifiers.

**SmolVLM-2B:** Being the smallest of the VLMs experimented with, with only 2.2B parameters and

<10GB of GPU ram, experiments were able to be performed locally without quantization.

**Qwen2.5-VL-7B:** Qwen is an intermediately-size VLM with around 7B parameters. Being too large for the hardware available, manual quantization using BitsandBytes was performed to shrink the precision down to 4-bit to be used locally.

**LLaVA-v1.5-13B:** Also considered an intermediately-sized option, the 13B 4-bit quantized version of LLaVA was used due to hardware constraints. Using a combination of a publicly available Colab notebook <sup>7</sup>, the author of which performed the quantization, and the original code from the open-source LLaVA repository <sup>8</sup>, the model was able to be used with Google Colab.

**GPT-4o-mini:** The largest of the VLMs tested, a paid API service had to be used for querying GPT-4o-mini <sup>9</sup>, and it is also the only closed-source model tested.

For each method, the responses are saved as a file where each line contains the model’s answer supplied as a letter from the options shown in Table C.1. This then allows for the calculation of metrics.

<sup>7</sup><https://github.com/camenduru/LLaVA-colab>

<sup>8</sup><https://github.com/haotian-liu/LLaVA>

<sup>9</sup>version:GPT-4o-mini-2024-07-18

<sup>6</sup><https://scikit-learn.org/stable/>

## F Related Architectures

Table F.1 briefly describes the architectures proposed by various authors for the CrisisMMD dataset. The purpose of this table is to demonstrate the literature review used to select the classifier proposed in section 4 and give an idea of the expected accuracy of such a classifier on this task of wildfire social-media data.

## G Ablation & Other Experiments

Table G.1 shows the ablation study performed on the selected model. (1) is the result of the selected model, (2) is the result of the model with a simpler classification head consisting of a single linear layer, (3) is the model with the proposed classification head but without the fusion module. Two other random splits were later used to support the result described in section 4.

## H Structure of Custom Classifier

Figure H.1 shows the architecture of the selected approach. This model is described in detail in Section 4.

## I Confusion Matrices

For deeper insights into the classifiers performances, the confusion matrices are shown in Figures I.1 and I.2.

## J Queries Used for Additional Data Collection

Shown here is the list of queries and keywords used to collect the additional unlabeled data discussed in Section 7.

### 2018

*Hashtag Query:* (#BCwildfire OR #britishcolumbiawildfire OR #BCfire OR #ABWildfire OR #albertawildfire OR #ABFire)

*Keyword Query:* (alberta OR british columbia OR Prince George OR Grande Praire OR Waterton OR Bulkley Nechako OR Nadina Lake OR Koote-  
nay OR Crowsnest Pass OR Medicine Lake OR Comstock Lake OR Tugwell Creek OR Sooke OR Nanaimo Lakes OR Tweedsmuir OR Johnny Creek OR Alkali Lake OR Lutz Creek OR Shovel Lake OR Nadina Lake OR Verdun Mountain OR Silver Lake OR Tommy Lakes OR Island Lake OR Chutanli Lake) (wildfire OR forest fire)

### 2019

*Hashtag Query:* (#ABWildfire OR #albertawildfire

OR #ABFire)

*Keyword Query:* (alberta OR calgary OR edson OR Fort McMurray OR Grande Prairie OR High Level OR Lac La Biche OR Whitecourt OR Steen River OR Chuckegg Creek OR Peace River OR Slave Lake OR Wood Buffalo National Park) (wildfire OR forest fire)

### 2020

*Hashtag Query:* (#BCwildfire OR #britishcolumbiawildfire OR #BCfire OR #ABwildfire OR #albertawildfire OR #ABFire OR #SKwildfire OR #saskatchewanwildfire OR #SKfire OR #YTwildfire OR #yukonwildfire OR #YTfire OR #NTwildfire OR #northwestterritorieswildfire OR #NTfire OR #NWTwildfire OR #NWTfire OR #MBwildfire OR #manitobawildfire OR #MBfire OR #ONwildfire OR #ontariowildfire OR #QCwildfire OR #quebecwildfire OR #QCfire)

*Keyword Query:* (british columbia OR alberta OR saskatchewan OR yukon OR northwest territories OR manitoba OR ontario OR quebec) (wildfire OR forest fire)

### 2021

*Hashtag Query:* (#MBwildfire OR #manitobawildfire OR #MBfire OR #ONwildfire OR #ontariowildfire OR #SKwildfire OR #saskatchewanwildfire OR #SKfire OR #pafire OR #ontariofire OR #manitobafire OR #saskatchewanfire)

*Keyword Query:* (manitoba OR ontario OR saskatchewan OR british columbia) (wildfire OR forest fire)

### 2022

*Hashtag Query:* (#YTwildfire OR #yukonwildfire OR #YTfire OR #yukonforestfire OR #NTwildfire OR #northwestterritorieswildfire OR #NTfire OR #NWTwildfire OR #NWTfire OR #nwtforestfire)

*Keyword Query:* (#yzf OR #nwt OR #Yellowknife OR Yukon OR Northwest Territories OR Whitehorse OR Yellowknife OR Dawson City OR Great Slave Lake OR Norman Wells OR Inuvik OR Watson Lake OR Hay River OR Fort Smith OR Tuktoyaktuk OR Behchoko) (wildfire OR forest fire) (-alaska -Eielson -CityofNorthPole)

### 2023

*Hashtag Query:* (#BCwildfire OR #britishcolumbiawildfire OR #BCfire OR #ABwildfire OR #albertawildfire OR #ABFire OR #SKwildfire OR #saskatchewanwildfire OR #SKfire OR #YTwildfire OR #yukonwildfire OR #YTfire OR #NTwildfire OR #northwestterritorieswildfire OR #NT-

Paper	Method	Task 1	Task 2	Task 3
Agar. 2020 (Agarwal et al., 2020)	InceptionV3+RCNN+AttentionFusion	99.00	99.00	97.00
Abav. 2020 (Abavisani et al., 2020)	DenseNet+BERT+CrossAttentionFusion	89.35	91.82	70.41
Ofli. 2020 (Ofli et al., 2020)	VGG16+CNN+Concat.Fusion	84.20	78.30	-
Xukun. 2020 (Li and Caragea, 2020)	AlexNet+GRU+Concat.Fusion	78.13	-	-
Madi. 2021 (Madichetty et al., 2021)	DenseNet+BERT+AverageFusion	74.85	-	-
Kumar. 2022 (Kumar et al., 2022)	VGG16+LSTM+Concat.Fusion	82.00	-	-
Pran. 2022 (Pranesh, 2022)	ResNet+RoBERTa+TransformerAttentionFusion	85.10	89.50	-
Sirbu. 2022 (Sirbu et al., 2022)	semi-supervised	82.00	87.20	88.70
Diva. 2022 (Divakaran et al., 2022)	DenseNet+RoBERTa+Concat.Fusion	89.26	85.87	-
Kotha. 2022 (Kotha et al., 2022)	RegNetY320+BERT+Concat.Fusion	89.63	86.59	-
Khat. 2022 (Khattar and Quadri, 2022)	VGG16+BiLSTM+CrossAttentionFusion	84.08	-	-
Madi. 2023 (Madichetty et al., 2023)	VGG16+RoBERTa+MultiplicativeFusion	90.21	-	-
Basit. 2023 (Basit et al., 2023)	DeiT+ALBERTa+Concat.Fusion	79.00	80.00	86.00
Koshy. 2023 (Koshy and Elango, 2023)	ViT+RoBERTa+BiLSTM+BilinearFusion	97.14	-	-
Gupta. 2024 (Gupta et al., 2024)	DenseNet+ELECTRA+CrossAttentionFusion	91.20	93.40	72.20
Shet. 2025 (Shetty et al., 2025)	ViT+RoBERTa+CrossAttentionFusion	91.25	89.63	68.72
Sirbu. 2025 (Sirbu et al., 2025)	semi-supervised	91.22	89.55	-
Teng. 2025 (Teng and Ohman, 2025)	VGG16+BERT+CrossAttentionFusion	81.00	-	-
Dar. 2025 (Dar et al., 2025)	(ResNet50+BERT)+(Graph)+IDEAFusion	98.23	90.13	-

Table F.1: Literature review of deep-learning approaches used by other works on the CrisisMMD dataset.

Ablation	Explanation	F1-Weighted
1	Proposed	83.64
2	Linear class. head	82.70
3	No fusion module	81.71

Table G.1: Ablation of model components.

## K Unlabeled Dataset

Figure 2 shows the distribution of post origins. A word cloud of the text data of the unlabeled dataset collected in Section 7, and also the frequency distribution of the posts by word count, are shown in Figure K.1.

fire OR #NWTwildfire OR #NWTfire OR #MB-wildfire OR #manitobawildfire OR #MBfire OR #ONwildfire OR #ontariowildfire OR #QCwildfire OR #quebecwildfire OR #QCfire OR #CanadaOn-Fire OR #CanadaWildfire OR #CanadaFires OR #CanadaIsOnFire)

*Keyword Query:* (ontario OR quebec OR saskatchewan OR british columbia OR manitoba OR northwest territories OR yukon OR alberta) (wildfire OR forest fire)

## 2024

*Hashtag Query:* (#JasperStrong OR #JasperWildfire OR #JasperAB OR #BCwildfire OR #british-columbiawildfire OR #BCfire OR #ABwildfire OR #albertawildfire OR #ABFire OR #CanadaOn-Fire OR #CanadaWildfire OR #CanadaFires OR #CanadaIsOnFire)

*Keyword Query:* (canada OR ontario OR quebec OR saskatchewan OR manitoba OR northwest territories OR yukon OR british columbia OR alberta OR jasper) (wildfire OR forest fire)

**Note:** Every query includes the suffix -has:videos has:images lang:en -is:retweet -is:quote -is:reply to filter for original English-language posts with images only.

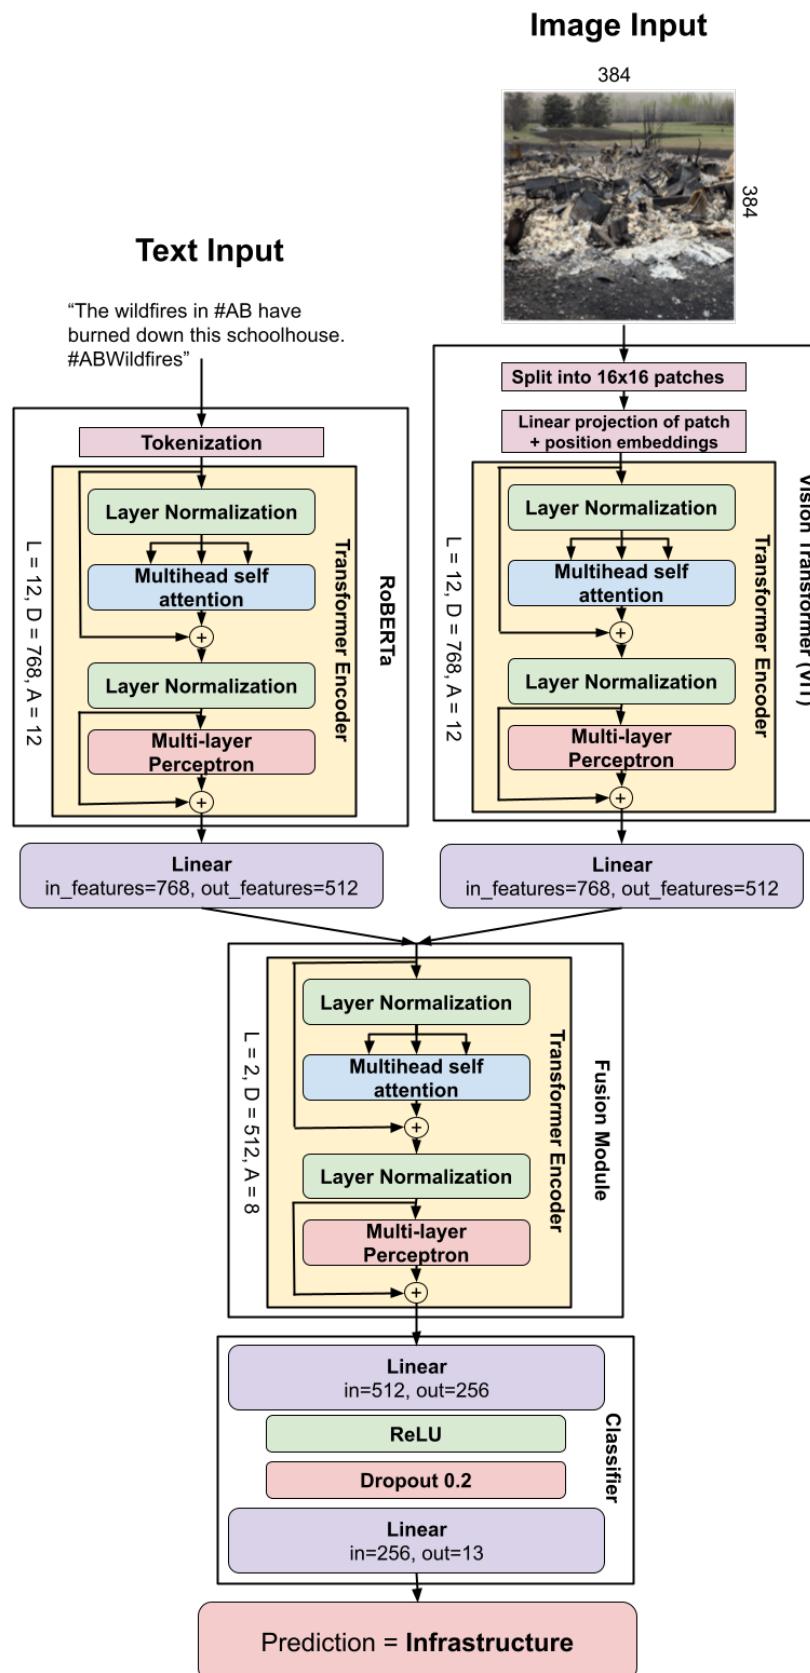


Figure H.1: Selected dual-encoder architecture utilizing transformer fusion module.

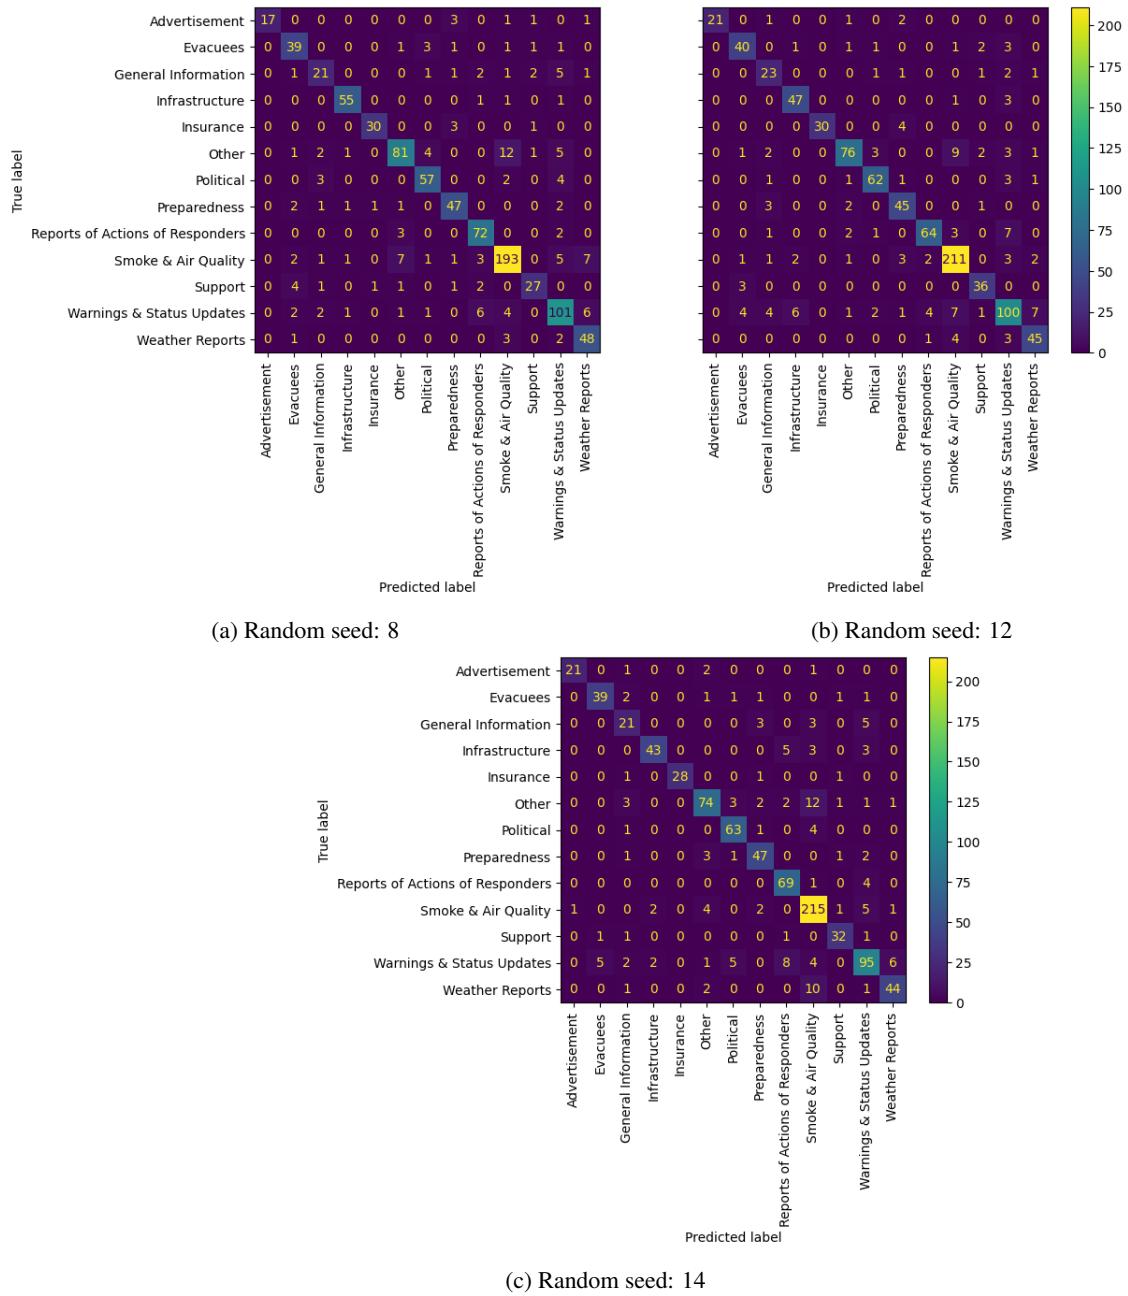


Figure I.1: Confusion matrices for selected classifier over the different train/test splits.

