

A Multi-Task Model for Sentiment Aided Stance Detection of Climate Change Tweets

Apoorva Upadhyaya, Marco Fisichella, Wolfgang Nejdl

L3S Research Center, Leibniz Universität Hannover, Hannover, Germany
{upadhyaya, mfishichella, nejdl}@l3s.de

Abstract

Climate change has become one of the biggest challenges of our time. Social media platforms such as Twitter play an important role in raising public awareness and spreading knowledge about the dangers of the current climate crisis. With the increasing number of campaigns and communication about climate change through social media, the information could create more awareness and reach the general public and policy makers. However, these Twitter communications lead to polarization of beliefs, opinion-dominated ideologies, and often a split into two communities of climate change deniers and believers. In this paper, we propose a framework that helps identify denier statements on Twitter and thus classifies the stance of the tweet into one of the two attitudes towards climate change (denier/believer). The sentimental aspects of Twitter data on climate change are deeply rooted in general public attitudes toward climate change. Therefore, our work focuses on learning two closely related tasks: Stance Detection and Sentiment Analysis of climate change tweets. We propose a multi-task framework that performs stance detection (primary task) and sentiment analysis (auxiliary task) simultaneously. The proposed model incorporates the feature-specific and shared-specific attention frameworks to fuse multiple features and learn the generalized features for both tasks. The experimental results show that the proposed framework increases the performance of the primary task, i.e., stance detection by benefiting from the auxiliary task, i.e., sentiment analysis compared to its uni-modal and single-task variants.

1 Introduction

Climate change is the burning crisis of our time, and it is happening even faster than we thought. A recent article on the BBC News website¹ states that many of the effects of global warming are now simply “irreversible” according to the latest assessment and that more than 40% of the world’s population is “at high risk” from climate. In fact, according to a report by the Intergovernmental Panel on Climate Change (IPCC), it is very likely that climate change is caused by man-made activities (Myhre et al. 2013). Despite the scientific consensus on the causes and main impacts of

climate change, it remains a controversial topic in public discourse. Therefore, understanding public perceptions plays a critical role in addressing climate change by increasing the public’s willingness to accept appropriate action on climate change (Shi, Visschers, and Siegrist 2015).

Recently, social media platforms such as Twitter have played an important role in raising public awareness of the current climate crisis and influencing public attitudes toward climate change (Lineman et al. 2015). Twitter conversations, however, are often influenced by the polarization of beliefs and, in the case of climate change, are divided into two competing groups, one that believes in climate change (Believers) and a second that is skeptical or denies that climate change is occurring (Disbelievers) (Jang and Hart 2015). The article published on the Euronews website² after the COP26 conference revealed that scientists have found that climate change deniers are not only skeptical about climate change, but have also led to the problem of delaying climate change by either shifting responsibility or eventually capitulating - the idea that it is not possible to prevent climate change, which often leads to the spread of misinformation (Zhou and Shen 2021). Therefore, it is important for government agencies, researchers, and technology companies to monitor such content on social media to identify and intervene in tweets from climate change deniers, which will help combat climate misinformation. This has motivated us to identify such content and understand public attitudes toward climate change by performing the important task of stance detection.

Stance detection is the task of automatically identifying the author’s point of view in relation to the target object (for/against/neutral). It has been used to identify social attitudes toward pressing issues (e.g., covid19 vaccination (Argyris et al. 2021), climate change, abortion, feminism (Mohammad et al. 2016)). In our study, we focus on climate change as a target and perform statement-level stance detection. The goal is to predict the attitude expressed in a single tweet, where the stance is for or against climate change.

Numerous works have classified tweets into favor or against the target using the SemEval 2016 benchmark dataset. The dataset contains 5 target topics, including climate change,

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://www.bbc.com/news/science-environment-60525591>

²<https://www.euronews.com/green/2021/11/18/climate-misinformation-is-getting-more-sophisticated-and-experts-say-cop26-progress-could->

with 29 denier and 335 believer tweets (Li and Caragea 2019; Wang et al. 2020). However, due to the small number of climate-specific tweets, these techniques do not focus on understanding the specific characteristics of climate change deniers and believers. This is because climate change deniers not only deny climate change, but also disagree with the solutions to combat climate change, which often leads to public negligence (Zhou and Shen 2021). Other works that deal with the classification of tweets on climate change lack a suitable architecture that efficiently performs the classification task (Kabaghe and Qin 2019). Therefore, in our work, we use a multi-tasking approach that uses different attention frameworks to classify the attitude of the climate change tweet into one of the two polarized classes (deniers/believers) to identify the denier statements on Twitter.

Sentiment analysis has helped various domains, like viral tweets about communal incidents (Upadhyaya and Chandra 2022) and many other tasks in a multi-task architecture. Studies on climate change have also justified the role of sentiment in conversations about climate change, either by assessing sentiment in tweets from climate change deniers or by examining the emotional impact of climate change data on Twitter (Dahal, Kumar, and Li 2019; El Barachi et al. 2021). Therefore, in our study, we leverage the sentiment analysis task to decipher the attitude of the tweets.

In addition, we use multiple inputs, i.e., combining tweet text and topic representations, to build a reliable classification model that helps identify the sentiment of the tweeter and determine the correct tweet attitude. Topic representations have helped in detecting fake news by providing more discriminatory power to the model used (Gautam, V, and Masud 2021). In our study, topic embedding provides a global context to a single tweet and thus more information and can circumvent the drawback of the short length of the tweet text to efficiently train the proposed system.

Our work focuses on learning two closely related tasks, stance detection and sentiment analysis of tweets on climate change. Stance detection is our main task, which is supported by sentiment analysis as an auxiliary task. We propose a multi-task framework that incorporates feature-specific and shared-specific attention frameworks to fuse multiple features and learn the features for both tasks. Our proposed approach is useful for government agencies and technology companies to detect the attitude of posts (deniers/believers) and curb the spread of such content that denies climate change and is false or misleading to combat climate misinformation.

We summarise the contributions of our work as follows: **(i)** We create a new dataset³ for the climate change domain consisting of tweets with the stance and sentiment labels which is beneficial for the research community. **(ii)** We illustrate the importance to consider the sentiment associated with the tweet while categorising the stance of tweet into favor(believers) or against(deniers) climate change. **(iii)** We propose a multi-task framework that jointly performs the stance detection (primary) and sentiment analysis (auxil-

³The dataset, code, and supplementary are available at <https://github.com/apoorva-upadhyaya/Climate-Change-Tweets>.

ary) tasks. We integrate feature-specific and shared-specific attention frameworks to integrate information across multiple features and shared tasks to learn features that optimise task performance. Experimental results indicate that the proposed framework increases performance of the primary task, i.e., stance detection by benefiting from the auxiliary task, i.e., sentiment analysis compared to its uni-modal and single-task variants.

2 Related Works

Climate change has become one of the greatest challenges of our time. Recently, (Upadhyaya et al. 2022) analyses the behavior of students on social media platforms related to climate change. Engaging the public is an important part of addressing climate change. Therefore, social media platforms like Twitter allow anyone to explore and report public viewpoints on the complex issue of climate change (Line-man et al. 2015; Dahal, Kumar, and Li 2019). However, debates and discussions on Twitter about climate change are widely associated with increasing polarization and are often divided into climate change believers and deniers (Jang and Hart 2015). In order to identify and understand public attitudes toward climate change, the task of identifying attitudes plays an important role.

Stance Detection Stance detection is about classifying the attitude that the author expresses towards a target object. The author may support the target object, reject it, or have a neutral stance. In our work, we focus on the climate change target, where opinions are either for climate change (believers) or against climate change (deniers). There are several climate change specific studies where the goal is to predict a user’s attitude (Chen, Zou, and Zhao 2019; Tyagi et al. 2020a). It has been found that multiple tweets from the same user can have different stance classes. In order not to miss the denier attitude of a single tweet that could interfere with the implementation of climate change policies, we focus on detecting statement-level stance detection, where the goal is to predict the stance described in a single tweet.

The stance detection of tweets has been studied in a variety of work on the popular SemEVAL 2016 dataset, which includes the 5 targets including climate change (with 364 climate change tweets) (Vychegzhanin and Kotelnikov 2021; Wang et al. 2020). However, in these previous studies, little attention was paid to understanding the characteristics of climate change denier and believer tweets in particular (only 29 climate change denier tweets in the SemEVAL 2016 dataset). One of the papers (Kabaghe and Qin 2019) classified tweets into three classes based on their attitude toward climate change: -1 (negative belief), 0 (neutral belief), and 1 (positive belief). However, the lack of an architectural framework led us to propose an efficient model that can efficiently classify a tweet’s attitude toward climate change into one of two categories (deniers/believers) in real-time, while taking advantage of the other task.

Sentiment Analysis Some of the works on stance recognition have emphasized the importance of sentiment (Wang et al. 2020), while some have cited the orthogonal relationship between stance and sentiment of the statement (Sen,

Sources	Denier Hashtags	Believer Hashtags
Tyagi et al. 2020a	ClimateHoax, YellowVests, Qanon	ClimateChange IsReal, ClimateAction Now, FactsMatter, ScienceMatters, ScienceIsReal
Most Used & Co-occur	GlobalWarming Hoax, ClimateChange Hoax, ClimateDenial, ClimateHoax	SaveClimate, ActOnClimate

Table 1: Denier & Believer Seed Hashtags

Flöck, and Wagner 2020). However, several works have focused on the sentimental aspects of climate change conversations and justified their role in climate change (Dahal, Kumar, and Li 2019). One recent work (El Barachi et al. 2021) proposes a real-time framework that uses sentiment and emotion analysis to provide meaningful insights into public opinion, and tested the model with tweets posted by Greta Thunberg and her followers on climate change. These studies motivated us to investigate the role of sentiment in classifying tweets on climate change.

Multi-Task Learning (MTL) is a learning paradigm that aims to learn multiple related tasks together in the hope of improving generalization performance for all tasks at hand. In our work, we focus on learning stance detection task with the help of sentiment analysis in a multi-tasking system. The study by (Li and Caragea 2019) uses sentiment to predict the stance through a multi-task learning model. Another work by (Chauhan, Kumar, and Ekbal 2019) uses sentiment as an auxiliary task to predict attitude. However, in our work, multiple features in the form of text and topic words are used to separate the task-dependent and independent feature spaces and perform both tasks simultaneously by using attention frameworks to focus on the most important feature representations and discarding the useless shared features that may affect the performance of both tasks.

3 Dataset

3.1 Data Collection Method

Previous works have used hashtags to identify the stances of different groups on social media, creating a method of tagging content by topic (Misra et al. 2016), we also use this explicit annotation quality of hashtags to collect a larger dataset. First, we select the hashtags for deniers and believers (Tyagi et al. 2020a) from the previous literature (see row 1 of Table 1). We then consider the two publicly available Twitter datasets on climate change, (i) tweet IDs collected from September 21, 2017 to May 17, 2019 using a set of climate change keywords available in the Harvard Dataverse (Littman and Wrubel 2019), and (ii) tweet IDs used in the work (Samantray and Pin 2019) collected from 2007 to 2019. We start by retrieving the tweet objects from

Algorithm 1: Label Propagation Algorithm

```

Input= Graph  $G$  with nodes  $n$  and edges  $e$  with  $e_{ij}$  as edge
weight between  $i \in n$  and  $j \in n$ 
initialize  $\gamma=50/100$  and  $i=0$ ;
for each  $n$  do
  define  $l = \text{integer}(i/\gamma)$ ;  $i+=1$ ;
  for each  $n$  do
    if  $n$  not labeled then
      compute  $t = \text{neighbors of } n$ ;
      compute  $t_l = \text{labeled neighbors of } n$ ;
      if  $t_l + l \geq t$  then
        initialize  $\text{score}, c$ ;
        for each  $t_i \in t$  do
           $\text{score}+ = \text{label } t_i * e_{nt_i}; c+=e_{nt_i}$ ;

```

these publicly available tweet IDs using the Tweepy API⁴. We analyzed the hashtags used in the collected tweets and found that the hashtags mentioned in row 2 of Table 1 are the most frequently used and co-occur with the denier and believer seed hashtags. We draw a random sample of 1000 tweets containing these most frequently used hashtags separately for both categories and identify 98% tweets as deniers and 99% tweets as believers. We conclude that the final set of seed hashtags (see Table 1) can be used to identify tweets from deniers and believers. Because the publicly available datasets contain tweets with a large time span that reflect Twitter trends and topics related to climate change and cover a wide range of audiences, they helped us identify additional relevant hashtags associated with climate change deniers and believers and enriched us with a final set of seed hashtags that can be used for further data collection related to climate change deniers and believers. We then select the unique tweets (excluding retweets) from the collected data that contain either the denier or believer hashtags. We obtain a total of 5,682 denier and 32,111 believer tweets after the filtering process based on seed hashtags. The collected dataset appears to be relatively small, with a smaller number of denier tweets, so we collect the real-time tweets from July 28 to December 26, 2021 using the live-streaming Tweepy API⁴ with the final set of seed hashtags. The tweets filtered out as 13,125 deniers and 60,430 believers according to the seed hashtags from various sources and are listed in Table 3 in Supplementary³.

3.2 Data Annotation

Stance Detection To validate the stances of the collected tweets provided by the hashtag self-annotation technique, we run a variant of label propagation algorithm (Tyagi et al. 2020a,b), which transfers the labels from the seed hashtags to other hashtags (refer Algorithm 1). The authors who provided the label propagation algorithm claim that their approach is similar to various other works (Weber, Garimella, and Batayneh 2013; Garimella et al. 2018). First, we weight the seed hashtags of believers by +1 and those of deniers by -1 (as in Table 1). We create a weighted hashtag*hashtag

⁴<https://docs.tweepy.org/en/stable/>

co-occurrence graph with all hashtags present in the data, where each node represents a hashtag and an edge is created between the hashtags that occur in the same tweet, where the weight of the edge is proportional to the frequency of their co-occurrence. The weights of the seed hashtags are then transferred to other hashtags as specified in the algorithm 1. The hashtag scores (believer hashtag = +1, denier hashtag = -1) are then arithmetically summed for all hashtags that occur in each tweet and then averaged. The final score is then used to classify tweets into deniers (score < 0) or believers (score > 0). In total, we found a set of 13, 125 denier and 60, 430 believer tweets. Three trained annotators drew a random sample of 1000 tweets from both categories and manually annotated them. To determine the consistency between the ratings of the annotators, we use the Fleiss-Kappa (Spitzer et al. 1967) measure and achieve an agreement score of 0.84, indicating that the annotations are of good quality. We consider the manually annotated tweets as the ground truth and compared them with the annotations found after the label propagation algorithm. We found that 98.40% of denier tweets belong to the denier category and 99.6% of believer tweets belong to the believer category. Therefore, to save time and cost, we consider the labels generated by the label propagation algorithm as the final labels for denier and believer tweets.

Sentiment Analysis We leverage weak supervision approach to annotate tweets for sentiment analysis. Similar to previous work (Singh et al. 2021), we use three sentiment classifiers to generate sentiment labels for each preprocessed text of the tweet, namely (i) VADER (Hutto and Gilbert 2014): a popular lexicon and rule-based sentiment analysis tool that relies on a dictionary to generate sentiment scores, (ii) TextBlob⁵: a Python-based library that provides an API for handling common NLP tasks such as sentiment analysis, POS tagging, etc., and (iii) NLTK⁶: a Python bundle that provides a collection of NLP algorithms such as sentiment analysis, NER, etc. This results in 3 labels (*positive, negative, neutral*) per tweet, from which a single label is finally selected as the sentiment expressed by the tweet based on the majority voting based ensemble method. The data statistics is mentioned in Table 2. Three trained annotators manually evaluated the labels for 1000 randomly selected tweets and obtained an inter-annotator agreement score of 0.81 using the Fleiss-Kappa measure. We consider the final annotations generated after inter-annotator agreement as the ground truth and compared them with the annotations generated using the weak supervision approach and found an accuracy of 97.6%. To save time and cost, we consider the annotations generated using the weak supervision approach for the sentiment analysis task.

3.3 Data Pre-processing

Data pre-processing is important because raw tweets without pre-processing are very unstructured and contain redundant and often problematic information that affects the performance of the model training and classification task.

⁵<https://textblob.readthedocs.io/en/dev/>

⁶<https://www.nltk.org/index.html>

Tweets	Negative	Positive	Neutral
Denier	60.2%	18.2%	21.6%
Believer	24.7%	46.2%	29.1%

Table 2: Data Statistics for Sentiment Task

Category	Tweet	Topic Words
Denier	CO2 is greening the planet and restoring the rainforest. Its almost like the planet is able to self regulate #ClimateHoax .	nonsense,hoax, myth,destroy, planet,hypocrisy
Believer	Great format and read #climate #ClimateActionNow	warm,hot, emergency, possible,crisis, urgent,warming

Table 3: Examples with stance, tweet text and topic words

Text We remove mentions, URLs, punctuation, spaces, and unwanted characters such as RT (retweet), CC (carbon copy), and stopwords from the tweet text. We use ekphrasis (Baziotis, Pelekis, and Doukeridis 2017) to extract hashtags by segmenting long strings into their individual words. For further text processing, we use the Python toolkit NLTK⁶. The NLTK-based tokenizer is used to tokenize tweets. All words are converted to lowercase letters. Then, we reduce the inflected words by applying the NLTK Wordnet lemmatizer, and then apply PorterStemmer for stemming.

Topic We first remove the seed hashtags used for data collection, otherwise topics created can be biased towards the hashtags. The tweet text is pre-processed using the procedure described above. In this study, we use BERTopic modeling, which uses transformer-based embeddings to create easily interpretable topics and their distributions (Grootendorst 2022). This modeling technique has recently gained popularity and provided promising results in previous studies (Anwar et al. 2021), therefore we focus on using BERTopic that detect semantic similarity and integrate topics with pre-trained contextual representations. The tweet text is then fed into the BERTopic library with the *calculate_probabilities=True*, which creates topics from the data and assigns a probability score to each created topic for each tweet sample in the data. We select the m-most similar topics for each tweet sample, where each topic is represented by the top ‘p’ topic words.

Role of Topic Words as Feature In Table 3, we present two samples from the dataset that illustrate the importance of considering topic words along with the tweet text for the analysis tasks. It can be observed that the tweet text alone is not helpful in identifying the stance of the tweet in the given samples. However, the addition of topic words gives the tweet more context and information that helps in correctly predicting the denier or believer stance of the tweet. These examples show that the presence of complementary information in the form of topic words aids the process of stance detection.

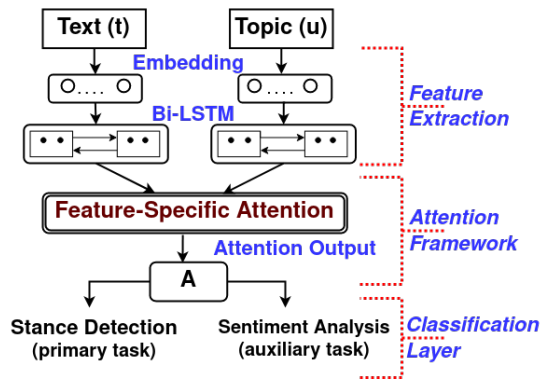


Figure 1: Architectural diagram of the Shared-Only Multi-Task (SO-MT) Framework.

4 Methodology

In this section, we outline the working of our proposed multi-task model for the stance detection (primary task) and the sentiment analysis (auxiliary task). The proposed model consists of the following components : *Feature Extraction*, *Attention Framework*, and *Classification Layer*. The key factor of multi-task learning is the sharing of features across the tasks. Therefore, we first describe the two variants of the model depending on the framework of multi-task learning and then explain the model components in detail.

4.1 Variants of the Proposed Model

We depict the following 2 variants of the multi-task model depending on the sharing of features across the tasks: (i)*Shared-Only Multi-Task Model (SO-MT)* : In the SO-MT model, we use single shared layers for the feature extractor and the attention framework to extract features for all tasks, as shown in Fig. 1. The single shared output of the attention framework is then used as an input to the classification layer, which produces separate outputs for the stance and sentiment tasks. This model focuses on the task-invariant features and ignores the fact that some features are task-dependent.

(ii)*Shared-Private Multi-Task Model (SP-AMT)*: In the model SP-AMT, we have two separate feature spaces for the two tasks that capture task-specific features, and a single shared feature space that captures the task-invariant features (refer Fig. 2). The final features are the concatenation of the features from the private space and the shared space, which are then fed into the classification layer to generate the output for both tasks.

The input and output of each model component for both variants (Figures 1 and 2) are mentioned in the following subsection 4.2. We now describe each of the model components in detail.

4.2 Components of the Model

Feature Extraction

Text: Each tweet text ‘T’ contains n_t number of words, where the embedding of each word w_1, \dots, w_{n_t} is acquired from BERT (Devlin et al. 2019) with dimension(d) = 768. We obtain final embedding for each tweet text as $T \in$

$R^{n_t \times d}$. Since Bi-LSTM has shown excellent performance in text classification due to its ability to learn long-term dependencies and incorporate past and future context information without retaining duplicate information (Zhou et al. 2016), we use Bi-LSTM to sequentially encode the embedded input text representations. The embedded text is fed to Bi-LSTM with dimension d_l , which learns the long-term context dependent semantic features into hidden states. The final hidden matrix of text is $H_t \in R^{n_t \times 2d_l}$.

Topic: Each tweet is tagged with the number of m-most similar topics based on the probability score assigned to each of the topics, where each topic is represented by the top ‘p’ topic words created using BERTopic (described in Section 3.3). Here we represent the topic feature as ‘U’ containing a set of n_u words ($n_u \leq m \times p$), where the representation of each word w_1, \dots, w_{n_u} are obtained from BERT with $d = 768$. We obtain final embedding for each tweet topic as $U \in R^{n_u \times d}$. This representation of topic is then fed to the Bi-LSTM layer with d_l that sequentially encodes these representation, and gives the final hidden matrix of topic $H_u \in R^{n_u \times 2d_l}$.

Attention Framework Attention mechanism has been used as an important component across a wide range of NLP models (Bahdanau, Cho, and Bengio 2014). Typically, the attention layer concentrates on the relevant part of the input and extracts the most important information from the input. We apply the attention framework similar to (Vaswani et al. 2017), in which the authors consider an attention function as a mapping to a set of queries, keys, and values. To obtain queries, keys, and values for the final feature representations, we pass the hidden matrix output from the Bi-LSTM layer of text (H_t) and topic (H_u), respectively, through three fully connected layers of dimension d_a . There are two triplets of query, key, and value for text (Q_t, K_t, V_t) and topic (Q_u, K_u, V_u) in the model SO-MT, while we have a total of four triplets for the model SP-AMT, forming two pairs of two triplets each for text and topic, which are used for stance detection ($(Q_{td}, K_{td}, V_{td}), (Q_{ud}, K_{ud}, V_{ud})$) and sentiment ($(Q_{ts}, K_{ts}, V_{ts}), (Q_{us}, K_{us}, V_{us})$) tasks respectively. Fig. 3 visually shows the two attention frameworks used in our model: *feature-specific attention* and *shared-specific attention*. The lower part of the figure shows how different queries, keys, and values are encoded to obtain self-attention and inter-attention, which form the two sub-modules of feature-specific attention. The upper part of the figure shows the connections between queries, keys, and values to achieve shared-specific attention. In the following, we describe in detail the attention mechanisms used in our study.

Feature-Specific Attention We apply two types of attention to the features to capture the most informative parts of them. Fig. 3 (bottom) shows the visual representation and connections of feature-specific attention. Feature-specific attention is further divided into Self Attention (SA) and Inter Attention (IA).

Self Attention (SA) We use Self Attention (SA) to relate different positions of a single sequence of say tweet text or topic to quantify the most important part of that sequence

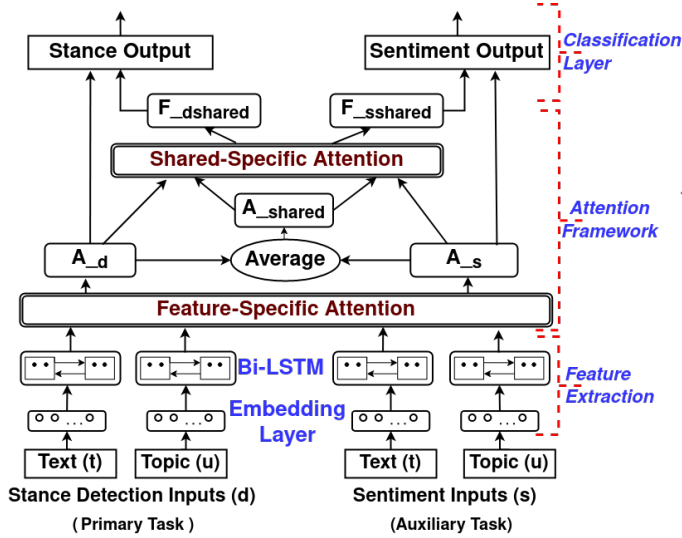


Figure 2: Architecture of the Shared-Private Multi-Task variant with Attention Framework (SP-AMT)

(Vaswani et al. 2017).

$$SA_j = \text{softmax}(Q_j K_j^T) V_j \quad (1)$$

SA scores are calculated using the equation 1, where $SA_t \in R^{n_t \times d_a}$, and $SA_u \in R^{n_u \times d_a}$. Here, two SA scores are computed for SO-MT, while four such SA scores are required for SP-AMT model variant ($SA_{td}, SA_{ud}, SA_{ts}, SA_{us}$) (as shown with black dotted arrow connections in Fig. 3 (bottom)).

Inter Attention (IA) We find out the Inter Attention (IA) scores to learn the interdependence between different features. IA scores are determined using below equations where query of one feature is intervened with key and value of the other. IA scores help to reveal the significant contributions amongst different inputs to learn optimal features for both tasks. The equations that represent the IA scores for text and topic (IA_{tu}, IA_{ut}) are:

$$IA_{tu} = \text{softmax}(Q_t K_u^T) V_u, \quad (2)$$

$$IA_{ut} = \text{softmax}(Q_u K_t^T) V_t, \quad (3)$$

where $IA_{tu} \in R^{n_t \times d_a}$, and $IA_{ut} \in R^{n_u \times d_a}$. IA equations are represented graphically with orange and royal red dotted arrows in Fig. 3 (bottom) part. The SA and IA scores are then concatenated finally, where A is directly used for shared-only (SO) variant (refer Fig. 1) while average of attention vector (A_{shared}) specific to stance and sentiment tasks (A_d, A_s) is used for Shared-private (SP) variant of model (mentioned in Fig. 3 (bottom)).

$$A = \text{concat}(SA_t, SA_u, IA_{tu}, IA_{ut}), \quad (4)$$

$$A_d = \text{concat}(SA_{td}, SA_{ud}, IA_{tud}, IA_{utd}), \quad (5)$$

$$A_s = \text{concat}(SA_{ts}, SA_{us}, IA_{tus}, IA_{uts}), \quad (6)$$

$$A_{shared} = \text{Average}(A_d, A_s) \quad (7)$$

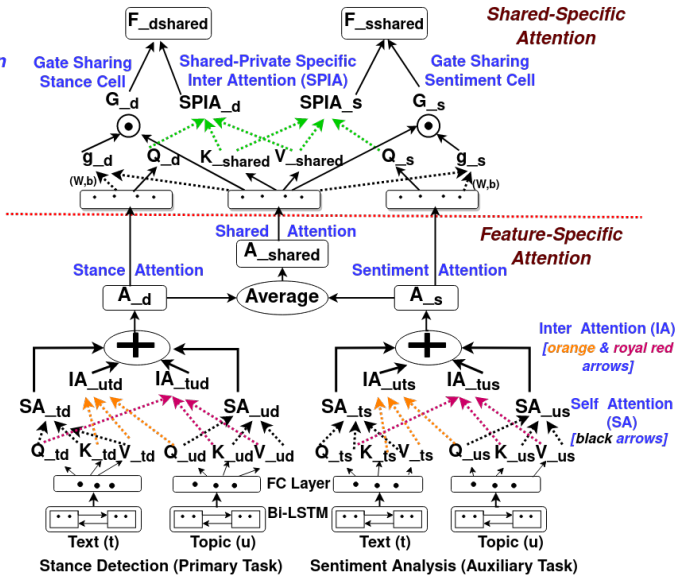


Figure 3: Attention Framework : (Top) Shared-Specific Attention; (Bottom) Feature-Specific Attention

Shared-Specific Attention Some of the works mentioned in Section 2 focused on the orthogonal relationship between sentiment and stance detection. Although climate change deniers and proponents dominate in sentiment, there are a few examples where sentiment does not match attitude/stance (see Table 2). (Wu et al. 2019) also mentioned the disadvantage of the shared-private model of multi-task learning, explaining that the shared space usually mixes some task-relevant features, which makes learning different tasks difficult. To solve the above problems, we use the (Wu et al. 2019) inspired shared-specific attention, which filters out the useless features that interfere with the model prediction and only pay attention to the selected features from the shared layer that lead to the correct predictions of the SP-AMT model (see Fig. 2). Next, we describe the sub-modules to achieve the desired result (as shown in Fig. 3 (top))

Gate Sharing Cell We use similar approach used by authors (Wu et al. 2019) where a single gate mechanism removes the useless shared features from shared layer. We first express the cell with reference to stance detection task. The stance specific, and shared attention scores (A_d, A_{shared}) from equations (5,7) are passed through dense layers with d_s units. The weights and biases are captured when passing A_d through dense layer which are used for A_{shared} , and is expressed as gated sharing cell.

$$g_d = \sigma(W_d \cdot A_{shared} + b_d) \quad (8)$$

where, $W_d \in R^{n_t(d_s) \times n_t(d_s)}$ and $b_d \in R^{1 \times n_t(d_s)}$. Similar equations are followed for sentiment task also, hence, we do not iterate here. The final output of the shared features for both the task after filtering will be represented as :

$$G_d = g_d \odot A_{shared}, \quad (9)$$

$$G_s = g_s \odot A_{shared} \quad (10)$$

where \odot denotes element-wise multiplication. Fig. 3 (top) shows the connections of gate sharing cell for stance detection and sentiment tasks.

Shared-Private Specific Inter Attention (SPIA) We use similar concept of the inter attention of feature-specific attention (equations 2, 3). We capture the important shared features relevant to the specific task, by using query matrix of the particular task (stance/sentiment) and keys and values of the shared task. The attention vectors (A_d, A_s) are passed through fully connected layers with d_s units to create Q_d, Q_s for stance and sentiment tasks, while A_{shared} is passed through dense layer to generate K_{shared} , and V_{shared} . The equations are represented visually with green dotted arrows in Fig. 3 (top) part.

$$SPIA_d = softmax(Q_d K_{shared}^T) V_{shared}, \quad (11)$$

$$SPIA_s = softmax(Q_s K_{shared}^T) V_{shared} \quad (12)$$

Fusion The final output of the shared layer is the fusion of the output of the gated cell and shared-private specific inter attention. Recently, fusion technique with absolute difference and element-wise product is found to be effective in (Mou et al. 2015).

$$C_1 = [G_d; SPIA_d; G_d - SPIA_d; G_d \odot SSIA_d]; \quad (13)$$

$$C_2 = [G_s; SPIA_s; G_s - SPIA_s; G_s \odot SSIA_s]; \quad (14)$$

$$F_{dshared} = \tanh(W_{fd} \cdot C_1 + b_{fd}), \quad (15)$$

$$F_{sshared} = \tanh(W_{fs} \cdot C_2 + b_{fs}) \quad (16)$$

Classification Layer The final representation of the tweet obtained (A_d, A_s) is passed through separate outputs for stance and sentiment tasks (for SO-MT model (Fig. 1)), however individual task specific tweet representations along with the shared layer representations are passed through two output channels, subjected to polarisation ($A_d, F_{dshared}$) and sentiment ($A_s, F_{sshared}$) tasks for SP-AMT Model (Fig. 2). The task specific and shared loss are used as

$$L_{total} = L_{task} + \lambda L_{shared} \quad (17)$$

where $\lambda = 0.5$ is a hyper-parameter (Liu, Qiu, and Huang 2017).

5 Experiment

5.1 Datasets

We evaluate our model performance and compare with the other baselines on the two datasets:

Climate Change Data The details of the data collection and statistics are covered in Section 3.

SemEval is provided in the SemEval-2016 shared task 6.A on tweet stance detection (Mohammad et al. 2016). Each tweet is in favor, against or neutral corresponding to one of the five targets: *Atheism*, *Climate Change is a Real Concern*, *Feminist Movement*, *Hillary Clinton*, and *Legalization of Abortion*(*Abortion*). There has been several works that use this benchmark dataset for stance classification.

5.2 Set-Up

We use the python-based library Keras⁷ at various stages of our implementations. For the experiments, we perform stratified k-fold cross-validation on our dataset, oversample the minority class (deniers) in the k-1 training data using the sklearn resampling technique, and report the averaged scores and standard deviation (over 5 folds) for the accuracy and F1 scores. We select $m = 5$ and $p = 10$, which fits our dataset well, where m denotes the number of most similar topics and each topic contains p number of words for the topic feature of each tweet. In the feature extraction sub-module, Bi-LSTM (d_l) with 100 memory cells is used. The dimensions d_a and d_s of the fully connected layers used in the attention framework to extract queries, keys and values for feature-specific and shared-specific attention (refer Section 4.2) are used with 100 units each. The stance and sentiment output channels contain 2 and 3 output neurons, respectively. The loss functions binary cross-entropy and categorical cross-entropy are used for the stance and sentiment output channels, respectively. The experiments are run on an NVIDIA GeForce GTX 1080Ti GPU and the models are optimized using Adam optimizer with a learning rate of 0.0001. All these values are selected using TPE in the hyperopt python library (Bergstra et al. 2013) and after a thorough sensitivity analysis of the parameters that minimise the loss functions.

5.3 Baseline Techniques

We compare our proposed approach with the following baselines on our climate change dataset:

Logistic Regression (Argyris et al. 2021): uses logistic regression with Count Vectorizer feature extraction method to classify vaccine-related tweets into provaccine, antivaccine, and neutral stances.

ESD (Vychegzhnin and Kotelnikov 2021): The authors form a relevant feature set using an ensemble of feature selection methods and propose the model ESD by selecting an optimal ensemble of classifiers. They evaluate the performance of the model using the UKP Sentential Argument Mining Corpus and the SemEval-2016 dataset.

HAN (Wang et al. 2020): is a hierarchical attention neural model, focusing on different features such as document, sentiment, dependency, and argument representations. The model is evaluated on SemEval-2016 and H&N14 dataset.

AT-JSS-LEX (Li and Caragea 2019): is a multi-task framework for stance detection with sentiment analysis as auxiliary task. The attention mechanism of the model is guided by target-specific attention along with sentiment and stance lexicons.

MNB (Kabaghe and Qin 2019): Multinomial naive bayes performed better with respect to other models proposed in the study, to classify tweets into positive, negative or neutral beliefs towards climate change.

DNN (Chen, Zou, and Zhao 2019): Deep Neural Network (DNN) is used as a classifier to identify users who either believe or deny climate change based on the content of tweets. The model’s performance is assessed on the real-time collection of climate change twitter data.

⁷<https://keras.io/>

Model	Single-Task Stance Detection				Single-Task Sentiment			
	Text		Text+Topic		Text		Text+Topic	
	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1
Shared-only (SO)	76.73/2.48	70.75/2.11	79.56/1.10	75.72/1.08	71.61/0.45	70.11/0.79	75.82/2.03	74.63/2.41
SO + Self Attn. (SA)	80.29/0.55	74.24/0.31	82.88/0.78	77.56/0.87	73.49/0.92	72.57/1.27	77.68/1.27	76.59/1.66
SO + Inter Attn. (IA)	81.33/1.04	73.35/1.41	83.13/1.42	76.04/1.11	73.68/3.20	72.23/3.14	77.93/2.01	74.11/1.98
SO + SA + IA (Feature-Specific Attn.)	82.15/0.05	76.78/0.81	85.01/1.05	80.03/2.42	76.21/0.84	74.95/1.02	80.81/1.29	79.28/1.44

Table 4: Results (Avg./Std.dev.) of the Single-Task Stance Detection and Sentiment tasks for various combinations

Model	Multi-Task Stance Detection + Sentiment							
	Stance Detection				Sentiment			
	Text		Text+Topic		Text		Text+Topic	
	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1
Shared-only (SO)	81.16/1.89	77.01/2.61	84.86/1.13	81.66/3.28	75.44/1.23	74.40/1.75	80.49/2.05	79.19/2.8
Shared-Private (SP)	84.64/1.95	81.53/2.55	87.47/2.01	84.24/2.04	77.93/2.09	74.11/2.19	85.28/0.91	84.58/1.02
Shared-Private(SP)+ Feature-Sp. Attn.	86.49/2.29	81.67/3.21	91.31/1.06	85.93/1.22	81.46/0.71	80.71/0.39	85.46/0.9	84.73/1.10
Shared-Private(SP)+ Shared-Sp. Attn.	88.10/1.04	84.09/1.39	92.29/1.32	86.12/1.4	81.67/0.61	80.89/0.32	87.59/0.45	87.07/0.39
Shared-Private + Attention Framework (Feature-Sp.Attn.+ Shared-Sp.Attn.) [SP-AMT]	89.99/2.67	86/2.02	93.95/1.27	90.24/1.16	84.60/0.44	83.98/0.81	89.08/1.01	88.48/1.60

Table 5: Results (Avg./Std.dev.) of the Multi-Task Stance Detection and Sentiment tasks for various combinations

SVM-ngram (Sobhani, Mohammad, and Kiritchenko 2016): is trained on word and character n-grams features for stance detection task on SemEval 2016 dataset. The model surpassed the best model in SemEval-2016 competition.

We evaluate our model performance on *SemEval 2016* dataset and contrast with the the **ESD**, **HAN**, **AT-JSS-LEX**, and **SVM-ngram** methods as described above.

6 Result and Analysis

In this section, we investigate the performance of the proposed approach. We first compare different single-task and multi-task variants and then compare them with the state-of-the-art methods mentioned in Section 5.3. We also analyze the importance of each feature and the different variants of the attention framework. We report all the results of the five-fold cross-validation (mean and standard deviation of accuracy and F1 score) for the different combinations of the proposed system. The Tables 4 and 5 illustrate the results of the single-task and the various combinations of the proposed multi-task models for both the stance detection and sentiment tasks. It is evident that the addition of topic words consistently improves the performance of the models. This improvement means that the proposed architecture makes very effective use of the interaction between input features. This shows the importance of incorporating multiple features for various analysis tasks.

6.1 Comparison amongst Single-Task and Multi-Task Framework

From Tables 4 and 5, the multi-task variants perform better than the single-task variants by achieving an average macro

F1 score of 90.24 and 88.48 for the stance detection (primary) and sentiment analysis (auxiliary) tasks respectively. The results show that the sentiment and stance tasks improve each other’s performance when learned together. The single stance detection task is able to correctly label some tweets from deniers and believer that contain predominantly negative and positive sentiments, respectively (examples 2 and 6 from Table 6). However, examples 1, 4, and 5 from Table 6 clearly show that the stance task, together with the sentiment analysis task, is able to unambiguously identify denier and believers tweets with the corresponding less dominant positive, neutral, and negative sentiment polarities. As stated earlier, we consider sentiment analysis as an auxiliary task that supports the main task, i.e., stance detection. However, we report the performance of the sentiment task for the proposed model for both single-task and multi-task frameworks in the tables 4 and 5 to illustrate the impact of the main task on the auxiliary task and to show that the multiple features in the form of tweet text and topic words, as well as the attention framework, also benefit the sentiment classification task. However, we do not make explicit efforts to improve the model performance on the auxiliary task.

6.2 Comparison amongst Different Multi-Task Frameworks

Table 5 shows the improvement of the multi-task framework from the shared-only variant (Fig. 1) to the shared-private multi-task model (SP-AMT) (Fig. 2). The inclusion of feature-specific and shared-specific attention frameworks helped the multi-task models focus on the important parts of the features and effectively discard the useless shared features, resulting in a 7.40% increase in accuracy and a

No.	Tweet	Sentiment	True	Predicted Stance	Predicted Stance + Sentiment
1.	My family support Oil and Gas! #ClimateHoax	positive	denier	believer	denier
2.	Once again brainwashing #kids to push the green tax agenda, under the #globalwarminghoax umbrella! Stop	negative	denier	denier	denier
3.	His Green BS policies will send us back to the Dark Ages. #ClimateChangeHoax	negative	denier	believer	denier
4.	ClimateHoax The climate has fluctuated since the time of creation, and nothing those people will do can change that one way or the other	neutral	denier	believer	denier
5.	And yet, there are those who deny climate change?? Ice-shelves breaking off, heat waves, etc.#Sad #ScienceIsReal	negative	believer	denier	believer
6.	Have you seen this? Its very moving. We definitely need more #ClimateAction	positive	believer	believer	believer
7.	For those adamant that global warming is real, THIS is Today in Alaska. Four inches of snow overnight, and still coming down! ClimateChange	neutral	denier	believer	believer
8.	I am glad you went by plane. Way better for our climate instead of zoommeetings...	positive	believer	denier	denier
9.	Over 60° today. Over 6" of snow tomorrow. But yeah, #climatechange is total bullshit right?	negative	believer	denier	denier

Table 6: Tweets with true and predicted labels for single and multi-task models [bold indicates incorrect predictions]

Model	Training Time (secs)	Mean Accuracy
Single Task Best (SO + SA + IA) [Table 4]	870	85.01
Multi Task Variants [Refer Table 5]		
Shared-only (SO)	918	84.86
Shared-Private (SP)	1218	87.47
Shared-Private (SP) + Feature-Specific Attn.	1419	91.31
Shared-Private (SP) + Shared-Specific Attn.	1506	92.29
Shared-Private + Attention Framework (Feature-Sp. Attn. + Shared-Sp. Attn.) [SP-AMT]	1791	93.95

Table 7: Training time of Different Text + Topic Models

12.46% increase in F1 score. Furthermore, in Table 7 we give the training times of the best performing single-task model and different variants of the multi-task model for 20 epochs to analyze the additional time required by the best performing multi-task model with text and topic as input features (SP-AMT) compared to other variants. As can be seen from Table 7, SP-AMT requires about 15 more minutes (approximately twice the time) to achieve a 10.51% improvement in accuracy compared to the best performing single-task model, while SP-AMT requires 9.5 more minutes to achieve a 7.4% improvement in performance compared to the shared-private (SP) multi-task variant without any attention framework.

All results reported here are statistically significant as we performed a t-test at the 5% significance level (Welch 1947) against the null hypothesis, which states that the mean accuracy/F1 score of all the multi-task variants is more when compared to the the best performing proposed model SP-AMT (Shared-Private + Feature-Specific Attention +Shared-Specific Attention) [refer Table 5]. If the p-

value is significant ($p < 0.05$), we reject the null hypothesis. Our best performing proposed model outperforms all the other multi-task variants while meeting statistical significance under t-tests ($p < 0.05$). For the confidence analysis, we also report the p-values and t-test statistics of all the multi-task variant models compared to the best performing model in Tables 1 and 2 in Supplementary³.

6.3 Comparison with the Baseline Methods

In Table 8, we report the results for the baseline methods by re-implementing them on the Climate Change dataset (Section 3). It is observed that our proposed multi-task approach SP-AMT outperforms the SOTA approaches in terms of accuracy and F1 score. Our best performing model achieved better results compared to ESD (Vychegzhanin and Kotelnikov 2021) and HAN (Wang et al. 2020). This highlights that the shared-private multi-tasking approach takes advantage of task-specific and invariant features to improve classification task performance. Although the AT-JSS-LEX (Li and Caragea 2019) model was implemented with a multi-tasking approach, our model performs better because it keeps the task-dependent and task-independent feature spaces separate and removes the useless shared features that hinder task performance of the stance detection, demonstrating the importance of the shared-specific attention framework. It is also observed that the methods that use sentiment features (ESD, HAN, and AT-JSS-LEX) perform better than the other baselines. This proves the importance of the proposed sentiment analysis approach for climate change. Our best performing single-task polarization framework (Table 4) also outperforms MNB, DNN, and SVM-ngram approaches. This justifies the benefits of using topic words in addition to tweet text and feature-specific attention framework to improve the performance of the model. Consequently, we also performed a comparative analysis of the proposed multi-tasking approach SP-AMT with the state-of-the-art models (SOTA) on the SemEval 2016

Model	Accuracy	F1-score
Prop. SP-AMT(Fig. 2)	93.95	90.24
LR	81.48	81.00
ESD	89.65	85.11
HAN	89.47	86.00
AT-JSS-LEX	88.02	84.01
MNB	85.44	78.08
DNN	84.61	76.23
SVM-ngram	85.55	66.33

Table 8: Results of Proposed Framework SP-AMT with baselines on our Climate Change Dataset

dataset. The model is trained with three polarized classes (*Favour*, *Against*, *None*) and the metrics (F_{avg} , $MacF_{avg}$) are evaluated according to the procedure defined in (Li and Caragea 2019). Table 9 shows that our approach outperforms other methods with an overall $MacF_{avg}$ value of 66.84. Our proposed framework performs better in the climate (C), feminism (F), and abortion (AB) domains, while the F_{avg} values are comparable in the atheism (AT) and Hillary (H) domains, showing that our framework generalises well in different domains.

6.4 Error Analysis

We perform an in-depth error analysis to understand where the proposed model has faltered. These are the following scenarios: (i) The climate change dataset is an imbalanced dataset with a high proportion of believer tweets, resulting in low F1 values compared to accuracy. Although we applied oversampling to partially counter this problem, even finer categories of believers can be identified and labeled, which can be beneficial for the model to learn different classes with a clear separation of distribution in tweets, such as “tweet conveys causes of climate change” or “tweet believes in human-caused climate change”. (ii) We determine the frequency of unigrams and bigrams extracted using TF-IDF and find that some of the denier’s tweets containing either rarely used keywords or keywords frequently used in believers’ tweets were misclassified. For example, the denier’s tweet in example 7 (Table 6) contains words such as *real*, *snow overnight*, which are most commonly found in believers’ tweets and confuse the model and lead to an incorrect prediction. (iii) We investigated that of the total misclassified denier tweets, 35.7% of the tweets contained sarcasm to express their denial. Of the sarcastic denier tweets, 50.16% of the tweets have positive sentiment, 31.70% have neutral sentiment, and the rest have negative sentiment, while 25.78% of the misclassified believer tweets have sarcastic labels (examples 8 and 9 of Table 6). The labeling of sarcasm is based on the majority vote of three trained annotators with an inter-rater agreement of 0.78, calculated with the Fleiss-Kappa measure. This motivated us to investigate the presence of sarcasm in climate change tweets to further improve the performance of the model as a part of our future work.

7 Discussion and Implications

In this section, we discuss the implications of our work. We observed that in our novel curated dataset, a high degree of

Model	AT F_{avg}	C F_{avg}	F F_{avg}	H F_{avg}	AB F_{avg}	Mac F_{avg}
SP-AMT (Fig. 2)	69.5	63.5	63.2	67.5	70.5	66.84
ESD	66.64	43.82	62.85	67.79	64.94	61.20
HAN	70.53	49.56	57.50	61.23	66.16	61.00
AT-JS S-LEX	69.22	59.18	61.49	68.33	68.41	65.33
SVM-ngram	65.19	42.35	57.46	58.63	66.42	58.01

Table 9: Results of Proposed Framework SP-AMT with baselines on SemEval 2016 Dataset

negativity is prevalent in the tweets of climate deniers. Theoretically, our study describes a new dimension by incorporating sentiment information for stance detection of climate change tweets. It introduces a framework that uses tweet text and topic words to extract useful features. We empirically tested and validated the role of sentiment in detecting the attitude of a tweet. These results contribute to theoretical knowledge in the domain.

Climate change deniers are not only skeptical about climate change but also emphasize the disadvantages of all measures proposed to combat climate change and abandon the idea that it is not possible to prevent climate change. This often leads to the spread of misinformation, resulting in a delay in the implementation of effective climate change mitigation measures (Zhou and Shen 2021). Since our work is dedicated to classifying Twitter content into climate change deniers or believers, the proposed approach can be useful for government agencies, researchers, and tech companies that monitor such content on social media to identify and intervene tweets from climate change deniers. The input feature of the proposed model, such as the tweet text, is available as soon as the user posts something. However, the topic feature can be extracted by performing topic modeling for a collection of tweets after a fixed interval, e.g., after every 5, 10, 15, or t minutes of duration. Therefore, our proposed framework can be used in a real-time environment by interested agencies and authorities to classify social media content into one of the two polarized classes. Moreover, the performance of the model on the SemEval dataset shows the applicability of our approach in different domains. These findings contribute to the practical implications of our work and justify the usefulness of our approach.

8 Conclusion and Future Work

In this work, we investigated the role of sentiment in classifying stance of the tweets related to climate change. We curate a novel dataset that includes annotations for both stance detection and sentiment analysis tasks, which will be useful to the research community in exploring other needed classification tasks. We propose a shared-private multi-task framework for the optimization of stance detection task benefiting from the sentiment analysis (auxiliary task). The proposed module uses feature-specific and shared-specific attention to fuse multiple features and learn useful and relevant private and shared features for both tasks. The results

show that multi-tasking increased the performance of the stance detection task compared to its uni-modal and single-task variants. Although we examined the performance of the proposed approach in detecting attitudes in the domain of climate change, the performance of the model on the SemEval dataset shows that it is much more broadly applicable beyond the domain of climate change, suggesting that our framework can generalize well in different domains. Future work will attempt to analyse what other aspects of natural language, such as sarcasm, aspect-based sentiment, and emotion recognition, might help to more accurately classify attitudes toward climate change. The inclusion of other modality encodings such as images, emoji, and advanced architectures will also be the subject of our future work.

Ethical Statement

The data in this paper comes from publicly available user-generated online content. Although we focus on identifying the attitude of a tweet based on the tweet text rather than individual user characteristics, such data poses risks for targeting the authors of the tweets and problems with their privacy. To mitigate these issues and comply with the terms of use, we are committed to protecting individual privacy and therefore avoid sharing personally identifiable content. The dataset that is made publicly available consists only of tweet IDs and annotations.

Acknowledgments

This work was partly funded by the SoMeCliCS project under the Volkswagen Stiftung and Niedersächsisches Ministerium für Wissenschaft und Kultur, by the European Commission for the research project eXplainable Artificial Intelligence in Healthcare Management (xAIM), agreement no. INEA/CEF/ICT/A2020/2276680, and by the European Commission for the research project “SoBigData++” under the Horizon 2020 program with grant agreement number 871042.

References

Anwar, A.; Ilyas, H.; Yaqub, U.; and Zaman, S. 2021. Analyzing QAnon on Twitter in Context of US Elections 2020: Analysis of User Messages and Profiles Using VADER and BERT Topic modeling. In *DG. O2021: The 22nd Annual International Conference on Digital Government Research*, 82–88.

Argyris, Y. A.; Monu, K.; Tan, P.-N.; Aarts, C.; Jiang, F.; and Wiseley, K. A. 2021. Using Machine Learning to Compare Provacine and Antivaccine Discourse Among the Public on Social Media: Algorithm Development Study. *JMIR public health and surveillance*, 7(6): e23105.

Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Baziotis, C.; Pelekis, N.; and Doulkeridis, C. 2017. Datasets at semeval-2017 task 4: Deep lstm with attention for message-level and topic-based sentiment analysis. In *Proceedings of the 11th international workshop on semantic evaluation (SemEval-2017)*, 747–754.

Bergstra, J.; Yamins, D.; Cox, D. D.; et al. 2013. Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms. In *Proceedings of the 12th Python in science conference*, volume 13, 20. Citeseer.

Chauhan, D. S.; Kumar, R.; and Ekbal, A. 2019. Attention based shared representation for multi-task stance detection and sentiment analysis. In *International Conference on Neural Information Processing*, 661–669. Springer.

Chen, X.; Zou, L.; and Zhao, B. 2019. Detecting climate change deniers on twitter using a deep neural network. In *Proceedings of the 2019 11th International Conference on Machine Learning and Computing*, 204–210.

Dahal, B.; Kumar, S. A.; and Li, Z. 2019. Topic modeling and sentiment analysis of global climate change tweets. *Social Network Analysis and Mining*, 9(1): 1–20.

Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, 4171–4186.

El Barachi, M.; AlKhatib, M.; Mathew, S.; and Oroumchian, F. 2021. A novel sentiment analysis framework for monitoring the evolving public opinion in real-time: Case study on climate change. *Journal of Cleaner Production*, 127820.

Garimella, K.; Morales, G. D. F.; Gionis, A.; and Mathioudakis, M. 2018. Quantifying controversy on social media. *ACM Transactions on Social Computing*, 1(1): 1–27.

Gautam, A.; V, V.; and Masud, S. 2021. Fake News Detection System Using XLNet Model with Topic Distributions: CONSTRAINT@AAAI2021 Shared Task. In Chakraborty, T.; Shu, K.; Bernard, H. R.; Liu, H.; and Akhtar, M. S., eds., *Combating Online Hostile Posts in Regional Languages during Emergency Situation - First International Workshop, CONSTRAINT 2021, Collocated with AAI 2021*, volume 1402 of *Communications in Computer and Information Science*, 189–200. Springer.

Grootendorst, M. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794*.

Hutto, C.; and Gilbert, E. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the International AAI Conference on Web and Social Media*, volume 8.

Jang, S. M.; and Hart, P. S. 2015. Polarized frames on “climate change” and “global warming” across countries and states: Evidence from Twitter big data. *Global Environmental Change*, 32: 11–17.

Kabaghe, C.; and Qin, J. 2019. Classifying Tweets Based on Climate Change Stance. *Training*, 66(60.9): 61–6.

Li, Y.; and Caragea, C. 2019. Multi-task stance detection with sentiment and stance lexicons. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 6299–6305.

- Lineman, M.; Do, Y.; Kim, J. Y.; and Joo, G.-J. 2015. Talking about climate change and global warming. *PLoS one*, 10(9): e0138996.
- Littman, J.; and Wrubel, L. 2019. Climate Change Tweets Ids. *Harvard Dataverse*, <https://doi.org/10.7910/DVN/5QCCUU>.
- Liu, P.; Qiu, X.; and Huang, X. 2017. Adversarial Multi-task Learning for Text Classification. In Barzilay, R.; and Kan, M., eds., *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, 1–10.
- Misra, A.; Ecker, B.; Handleman, T.; Hahn, N.; and Walker, M. 2016. NLDS-UCSC at SemEval-2016 Task 6: A Semi-Supervised Approach to Detecting Stance in Tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. Association for Computational Linguistics.
- Mohammad, S.; Kiritchenko, S.; Sobhani, P.; Zhu, X.; and Cherry, C. 2016. Semeval-2016 task 6: Detecting stance in tweets. In *Proceedings of the 10th international workshop on semantic evaluation (SemEval-2016)*, 31–41.
- Mou, L.; Men, R.; Li, G.; Xu, Y.; Zhang, L.; Yan, R.; and Jin, Z. 2015. Natural language inference by tree-based convolution and heuristic matching. *arXiv preprint arXiv:1512.08422*.
- Myhre, G.; Shindell, D.; Bréon, F.-M.; Collins, W.; Fuglestedt, J.; Huang, J.; Koch, D.; Lamarque, J.-F.; Lee, D.; Mendoza, B.; et al. 2013. Climate change 2013: the physical science basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. K., Tignor, M., Allen, SK, Boschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, PM, Cambridge University Press Cambridge, United Kingdom and New York, NY, USA.
- Samantray, A.; and Pin, P. 2019. Credibility of climate change denial in social media. *Palgrave Communications*, 5(1): 1–8.
- Sen, I.; Flöck, F.; and Wagner, C. 2020. On the reliability and validity of detecting approval of political actors in tweets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1413–1426.
- Shi, J.; Visschers, V. H.; and Siegrist, M. 2015. Public perception of climate change: The importance of knowledge and cultural worldviews. *Risk Analysis*, 35(12): 2183–2201.
- Singh, A.; Saha, S.; Hasanuzzaman, M.; and Dey, K. 2021. Multitask learning for complaint identification and sentiment analysis. *Cognitive Computation*, 1–16.
- Sobhani, P.; Mohammad, S.; and Kiritchenko, S. 2016. Detecting stance in tweets and analyzing its interaction with sentiment. In *Proceedings of the fifth joint conference on lexical and computational semantics*, 159–169.
- Spitzer, R. L.; Cohen, J.; Fleiss, J. L.; and Endicott, J. 1967. Quantification of agreement in psychiatric diagnosis: A new approach. *Archives of General Psychiatry*, 17(1): 83–87.
- Tyagi, A.; Babcock, M.; Carley, K. M.; and Sicker, D. C. 2020a. Polarizing tweets on climate change. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, 107–117. Springer.
- Tyagi, A.; Field, A.; Lathwal, P.; Tsvetkov, Y.; and Carley, K. M. 2020b. A computational analysis of polarization on Indian and Pakistani social media. In *International Conference on Social Informatics*, 364–379. Springer.
- Upadhyaya, A.; and Chandra, J. 2022. Spotting flares: The vital signs of the viral spread of tweets made during communal incidents. *ACM Transactions on the Web*, 16(4): 1–28.
- Upadhyaya, A.; Pfeiffer, C.; Astappiev, O.; Marenzi, I.; Lenzer, S.; Nehring, A.; and Fisichella, M. 2022. How learnweb can support science education research on climate change in social media. In *12th International Conference on Methodologies and Intelligent Systems for Technology Enhanced Learning*. Springer.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Vychezhnanin, S.; and Kotelnikov, E. 2021. A New Method for Stance Detection Based on Feature Selection Techniques and Ensembles of Classifiers. *IEEE Access*, 9: 134899–134915.
- Wang, Z.; Sun, Q.; Li, S.; Zhu, Q.; and Zhou, G. 2020. Neural Stance Detection With Hierarchical Linguistic Representations. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28: 635–645.
- Weber, I.; Garimella, V. R. K.; and Batayneh, A. 2013. Secular vs. islamist polarization in egypt on twitter. In *Proceedings of the 2013 IEEE/ACM international conference on advances in social networks analysis and mining*, 290–297.
- Welch, B. L. 1947. The generalization of ‘STUDENT’S’ problem when several different population variances are involved. *Biometrika*, 34(1-2): 28–35.
- Wu, L.; Rao, Y.; Jin, H.; Nazir, A.; and Sun, L. 2019. Different absorption from the same sharing: Sifted multi-task learning for fake news detection. In *Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4636–4645.
- Zhou, P.; Qi, Z.; Zheng, S.; Xu, J.; Bao, H.; and Xu, B. 2016. Text Classification Improved by Integrating Bidirectional LSTM with Two-dimensional Max Pooling. In Calzolari, N.; Matsumoto, Y.; and Prasad, R., eds., *COLING 2016, 26th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, December 11-16, 2016, Osaka, Japan*, 3485–3495. ACL.
- Zhou, Y.; and Shen, L. 2021. Confirmation Bias and the Persistence of Misinformation on Climate Change. *Communication Research*, 00936502211028049.