

M2SA: Multimodal and Multilingual Model for Sentiment Analysis of Tweets

Gaurish Thakkar¹, Sherzod Hakimov², Marko Tadić¹

¹Faculty of Humanities and Social Sciences, University of Zagreb

²Computational Linguistics, University of Potsdam

{gthakkar, marko.tadic}@ffzg.hr, first.last@uni-potsdam.de

Abstract

In recent years, multimodal natural language processing, aimed at learning from diverse data types, has garnered significant attention. However, there needs to be more clarity when it comes to analysing multimodal tasks in multi-lingual contexts. While prior studies on sentiment analysis of tweets have predominantly focused on the English language, this paper addresses this gap by transforming an existing textual Twitter sentiment dataset into a multimodal format through a straightforward curation process. Our work opens up new avenues for sentiment-related research within the research community. Additionally, we conduct baseline experiments utilising this augmented dataset and report the findings. Notably, our evaluations reveal that when comparing unimodal and multimodal configurations, using a sentiment-tuned large language model as a text encoder performs exceptionally well.

Keywords: sentiment analysis, multilingual, multimodal

1. Introduction

Social media platforms serve as conduits for the dissemination of information. Tweets have emerged as a trendy medium through which individuals communicate and express their ideas and opinions. Twitter (aka X) is widely used by researchers as a prominent social media platform for engaging in micro-blogging and fostering interactions. Sentiment analysis (Pang and Lee, 2005) is a well-studied topic in natural language processing. The topic has received consideration in both unimodal and multimodal contexts. The proliferation of social media platforms, including Twitter and YouTube, has led to a common practice of assessing content using several modalities (You et al., 2016; Yu et al., 2020). This approach offers additional context through spoken, non-verbal, and auditory aspects. The primary focus in many domains of natural language processing (NLP) often revolves around higher-resourced languages. However, the challenge of processing lower-resourced languages remains unresolved.

The process of annotating supervised datasets for natural language processing (NLP) tasks is a labour-intensive endeavour requiring significant investment of time, financial resources, and effort. Recently, several shared tasks, including SemEval (Nakov et al., 2013a; Ghosh et al., 2015), have introduced tasks aimed at identifying the polarity of tweets, categorising them into predetermined classes. All of the datasets for the shared tasks are accompanied by labels considered the gold standard. Another point to take into account here is that previous approaches (Raffel et al., 2019; Xie et al., 2020; Cliche, 2017) focused on text-only,

while posts shared on social media sometimes include images, videos, etc. Approaches incorporating multimodal information (Poria et al., 2016b; Cheema et al., 2021; Poria et al., 2016a) for the classification of sentiment are predominantly focused on the English language.

This paper presents a straightforward approach for enhancing pre-existing publicly accessible datasets to conduct multimodal (image & text) sentiment analysis on Twitter called M2SA (Multimodal Multilingual Sentiment Analysis). We have collected existing datasets in 21 languages where each annotated post includes both text and image with the annotated labels being either *positive*, *negative*, or *neutral*. We then trained a multimodal model that combines image and text embedding features to classify the target labels.

Our contributions are as follows:

- We engage in curating, enriching, and analysing pre-existing Twitter sentiment datasets in 21 different languages.
- The pre-trained model architectures use a fusion of textual information and visual features, utilising large language models for text encoding and image encoding.
- The study examines the effects of utilising machine translation instances in the context of lower-resourced languages.

All resources (pre-trained models, datasets) and the source code are shared publicly¹. The subse-

¹<https://github.com/cleopatra-itn/M2SA-multimodal-multilingual-sentiment-analysis>

quent sections of the paper are structured in the following manner: Section 2 provides an overview of the existing literature and research in the field. Section 3 provides a comprehensive overview of the processes involved in data collection, enrichment, and the statistical characteristics of the dataset. The methodology for classification is outlined in Section 4. The experimental setup and results are outlined in Sections 5 and 6. The paper concludes in Section 7.

2. Related Work

Poria et al. (2016b) presented a framework that uses CNNs to extract features from multimodal data's visual and textual modalities. The visual features are extracted using a CNN model that has been previously trained, such as VGG16 or ResNet-50. A CNN model trained on a massive corpus of text data is used to extract the textual features. The combined extracted features from the visual and textual modalities are then fed into an MKL classifier. The MKL classifier discovers the optimal combination of kernels for distinguishing between distinct emotions or sentiments. Poria et al. (2016a) used both feature and decision-level fusion methods to merge affective information extracted from multiple modalities. Cheema et al. (2021) evaluated various embedding features from both text and visual content. Huang et al. (2023) proposed a new framework for multimodal sentiment analysis in realistic environments, with two main components: a module for multimodal word refinement and a module for cross-modal hierarchical fusion. Baecchi et al. (2016) employed a strategy that uses a skip-gram neural network to extract features from the text mode. Image-specific features are extracted using a denoising autoencoder (Vincent et al., 2010) neural network. The denoising autoencoder network is taught to reconstruct an image from its corrupted version. The extracted features from the text and image modalities are then concatenated and fed to an SVM classifier. In addition to considering the modelling strategies for the sentiment analysis task, it is essential to identify the available benchmarking datasets. English contains a substantial quantity of multimodal datasets on sentiment and emotion analysis (Go et al., 2009; Mohammad et al., 2018). While the TweetEval (Barbieri et al., 2020) examines the application of large language models to seven tasks in Twitter, including emotion, emoji, irony sentiment, and others, the test set is monolingual. The paper authored by Garg et al. (2022) provides a comprehensive exposition on diverse multimodal datasets, encompassing the domain of multimodal sentiment analysis.

3. Multimodal Multilingual Sentiment Analysis (M2SA)

According to our investigation, numerous datasets are available for conducting both unimodal and multimodal sentiment analysis. The conversion of unimodal datasets, particularly those derived from Twitter, to a multimodal format has been limited. The fundamental hypothesis underlying our utilisation of the unimodal dataset posits that, given its gold annotation, the Twitter dataset can be linked to an image that has not been previously examined or employed in the context of multimodal sentiment classification. Thus, we present our contribution in this field called *M2SA* (Multimodal Multilingual Sentiment Analysis).

3.1. Data Collection

An initial step in initiating the enrichment process involves conducting a manual search of pre-existing Twitter sentiment datasets. We do not target any other social media datasets to process them uniformly and keep them from a single source. This search is conducted through the utilisation of search engines and data repositories such as HuggingFace Datasets², European Language Grid³, and GitHub⁴. To retrieve the dataset, a search is conducted using specific keywords such as **twitter sentiment analysis dataset**, **social media sentiment analysis dataset**, and **twitter sentiment shared tasks**. Next, the compiled list of datasets undergoes the process of querying tweet information using the Twitter API. The JSON format is used to store the text and images associated with each individual tweet in a dataset. The initially collected datasets are then subjected to manual checking to exclude tasks unrelated to sentiment analysis. Lastly, label transformations were applied on class labels to convert them from a five-class to a three-class format in cases where they did not initially possess three distinct sentiment categories: positive, negative, and neutral. The preliminary investigation yielded approximately 100 datasets in multiple languages. However, the final version of our dataset consisted of only 56 distinct datasets, encompassing 21 different languages. This reduction in the number of datasets was primarily due to the absence of tweet IDs linked to the corresponding text in most datasets. Table 1 presents a comprehensive overview of various languages and their respective datasets that were collected.

²<https://huggingface.co/datasets>

³<https://live.european-language-grid.eu>

⁴<https://github.com/>

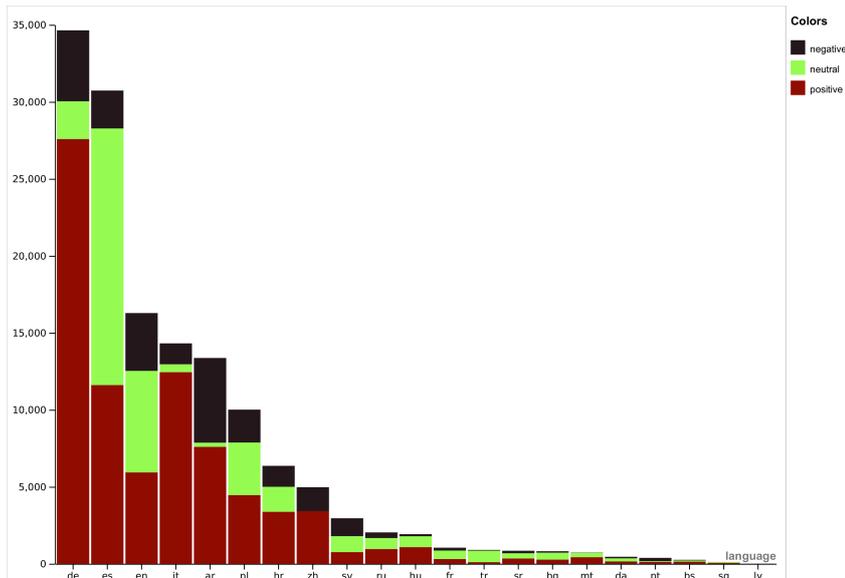


Figure 1: The dataset’s distribution across different languages.

3.2. Preprocessing

Preprocessing social media texts is imperative due to their inherent informality and noise. The preprocessing steps are delineated as follows:

- Removal of all black and white images.
- Tweet normalisation for USERS, URLs and HASHTAGS i.e., replace @ElonMusk → <user_1>..., URLs → <URL_1>, #tweet → <hashtag>tweet<hashtag/>
- Filtering of tweets with text content less than five characters, not accounting for USER and URL tags.
- Deduplication is performed using tweet IDs.
- Checking if the same tweet ID has more than one label assigned and employing a majority vote when needed.
- Filtering of tweets with corrupted or no images or with images of less than 200 × 200 pixels size.
- Checking the language tag in the tweet JSON and see if it matches the target language.
- Translation of English tweets for lower-resourced languages using the NLLB⁵ machine translation (MT) model.

The complete preprocessed dataset is structured according to a schema that can be described as follows:

⁵<https://huggingface.co/facebook/nllb-200-3.3B>

- tweetid: unique identifier for the tweet.
- normalised-text: text obtained after applying preprocessing steps.
- language: the language of the text.
- translated-text: text in the target language obtained using the NLLB model.
- image-paths: list of images associated with the tweet.
- label: POSITIVE|NEGATIVE|NEUTRAL

3.3. Dataset

Figure 1 illustrates the comprehensive distribution of datasets across different classes, encompassing 21 languages. The final dataset consists of 143K data points.

The dataset contains the following languages: Arabic-ar (Nakov et al., 2013b; Yin et al., 2021), Bulgarian-bg (Mozetič et al., 2016), Bosnian-bs (Mozetič et al., 2016), Danish-da (Pauli et al., 2021), German-de (Rei et al., 2016), English-en (Nakov et al., 2013c; Ghosh et al., 2015; Nakov et al., 2013d,a; Vanzo et al., 2014; Castellucci et al., 2015; Tayebi Arasteh et al., 2021), Spanish-es (Adrián, 2016-05-25; Santamaría et al., 2022; Agüero-Torales et al., 2021; Villena-Román and Garcia-Morera, 2013; Román et al., 2015; Vilares et al., 2015; Montejo-Ráez and Díaz-Galiano, 2016; Cámara et al., 2018; Díaz-Galiano et al., 2019; García-Vegaa et al., 2020), French-fr (Vukotić et al., 2015), Croatian-hr (Babić et al., 2021), Hungarian-hu (Mozetič et al., 2016), Italian-it (Moctezuma et al., 2016), Maltese-mt (Cortis and Davis, 2021), Polish-pl (Mozetič et al., 2016), Portuguese-pt (Patrick

Lang	Dataset name
ar	SemEval-2017
ar	TM-Senti@ar
bg	Twitter-15@Bulgarian
bs	Twitter-15@Bosnian
da	AngryTweets
de	xLiMe@German, Twitter-15@German, TM-Senti@de
en	SemEval-2013-task2, SemEval-2015, SemEval-2016
en	CB COLING2014 vanzo
en	CB IJCOL2015 ENG castellucci
en	RETWEET
es	xLiMe@spanish
es	Copres14
es	mavis@tweets
es	Twitter-15@Spanish
es	JOSA corpus
es	TASS 2018, 2019, 2020
es	TASS 2012, 2013, 2014, 2015
fr	DEFT 2015
hr	InfoCoV-Senti-Cro-CoV-Twitter
hr	Twitter-15@Croatian
hu	Twitter-15@Hungarian
it	CB IJCOL2015 ITA castellucci
it	xLiMe@Italian
it	sentipolc16
it	TM-Senti@it
lv	Latvian tweet corpus
mt	Malta-Budget-2018, 2019, 2020
pl	Twitter-15@Polish
pt	Twitter-15@Portuguese
pt	Brazilian tweet@tweets
ru	Twitter-15@Russian
sq	Twitter-15@Albanian
sr	doiserbian@tweet
sr	Twitter-15@Serbian
sv	Twitter-15@Swedish
tr	BounTi Turkish
zh	TM-Senti@zh-ids

Table 1: Languages and their corresponding dataset names

et al., 2022), Russian-ru (Mozetič et al., 2016), Serbian-sr (Ljajić and Marovac, 2018), Swedish-sv (Mozetič et al., 2016), Turkish-tr (Mutlu and Özgür, 2022), Chinese-zh (Yin et al., 2021), Latvian-lv (Muischnek and Müürisep, 2018) and Albanian-sq (Mozetič et al., 2016). Languages with more data points, such as German, Spanish, English, Italian, Arabic, and Polish, possess dataset instances exceeding 10,000, whereas other languages exhibit 5,000 or fewer instances of text and images. The dataset exhibits an average token count ranging from 4.25 to 5.94 words, separating each token

by a space. One observable pattern is that tweets classified as positive tend to be more likely to be accompanied by images than other categories. The diagram also indicates an imbalance in the datasets across the languages.

4. Methodology

4.1. Problem Definition

In the task of unimodal sentiment analysis, the model receives a sequence X_m as input, where m represents the length of the sequence. The model then produces a single class as output, which belongs to a closed set consisting of positive, negative, or neutral sentiments. In the context of multimodal sentiment analysis, the model receives input from multiple modalities denoted as $X_m^1 \dots X_n^2$, and the output is equivalent to that of unimodal sentiment analysis. The objective of the models is to extract features from the input vectors and acquire the ability to classify sentiment accurately.

The model architecture of the overall sentiment classification system is depicted in Figure 2. We examine distinct computational scenarios, encompassing the analysis of textual data alone and the integration of both textual and visual information, to classify the sentiment expressed in tweets. In the context of unimodal textual experiments, the models employed include Multilingual-BERT (Devlin et al., 2019a), XLM-RoBERTa (Conneau et al., 2020a), and XLMR-SM, a fine-tuned model specifically designed for sentiment analysis. In the context of multimodal systems, pre-trained vision models (CLIP and DINOv2) are employed as feature extractors. The combined textual and visual features are modelled using a concatenation operator.

In the context of language processing, datasets pertaining to a specific language are regarded as a cohesive entity. The dataset containing train, validation, and test sets is utilised directly within their respective sets. In this scenario, if there is no distinct set available, we will partition it into train (85%), test (10%), and validation (5%) sets manually.

Given that the text has already undergone processing, no additional processing has been applied to it. The input text undergoes tokenisation, during which it is padded and truncated according to the maximum length supported by the language models. The 'input id' and 'attention mask' are passed into the Language Model (LLM) to extract textual features for each instance in the dataset. Regarding image modality, the images undergo preprocessing through an image preprocessor linked to the corresponding vision models. The image preprocessor's output is subsequently inputted into the vision encoder. The concatenation of the output from the text and vision encoder is subsequently

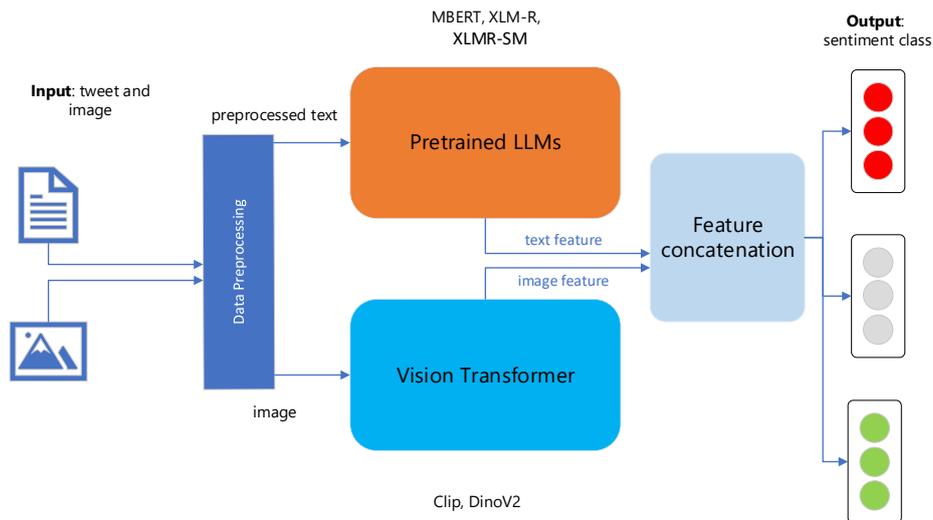


Figure 2: Model architecture of the M2SA.

projected onto a linear layer, which is then followed by a softmax layer for the purpose of classification. The models undergo separate fine-tuning processes using a combined dataset comprising samples from multiple languages. Furthermore, given the limited amount of data available for languages such as Latvian and Albanian, we opt to employ translation techniques to convert the existing text from another dataset into these languages with fewer than 10,000 tweets. Prior to the translation process, a language detection procedure using an existing model⁶ is executed. This process aims to accurately identify the source text, as the dataset contains text from various languages. Machine translation models rely on providing source and target language codes to perform translations effectively. In the context of the dataset, a language detection process is conducted to classify the instances into their respective source languages. Subsequently, the machine translation pipeline receives each grouped set along with the corresponding source and target language codes. The subsequent subsections discuss the model architectures and pertinent details associated with the models.

4.2. Text Encoders

To evaluate the model's efficacy in the absence of visual features, we conducted a standard fine-tuning process utilising the Transformer model's output. Specifically, we employed contextualised sentence embedding, which consists of a 768-element vector. The vector is subsequently fed into a fully connected (FC) layer consisting of three neurons, which is accompanied by a softmax layer for clas-

sification. The subsequent text models were employed as encoders for textual data.

4.2.1. Multilingual-BERT (M-BERT)

BERT (Devlin et al., 2019b) is a bidirectional transformer pre-trained with the masked language modelling (MLM) and next sentence prediction (NSP) objectives on the top 104 languages with the largest Wikipedia. This model is chosen due to its multilingual nature.

4.2.2. XLM-RoBERTa (XLM-R)

The XLM-R model (Conneau et al., 2020b) is a large multilingual language model trained on 2.5 TB of filtered Common Crawl data containing 100 languages. The model was trained with the Masked Language Modelling (MLM) objective, with 15% of the input words masked. The model has been shown to perform really well on downstream tasks when fine-tuned for supervised tasks. XLM-R can understand the input's language solely based on the input IDs without having to use language tensors. This model has been proven to improve the M-BERT scores in various tasks.

4.2.3. XLM-RoBERTa-Sentiment-Multilingual (XLMR-SM)

The XLMR-SM model (Antypas et al., 2022) is a fine-tuned version of XLM-T (Barbieri et al., 2022) on the tweet sentiment multilingual dataset (all), which consists of text from the following languages: Arabic, English, French, German, Hindi, Italian, Portuguese, and Spanish. The XLM-T model has been pre-trained on approximately 198 million multilingual tweets. We introduce this model to study the ef-

⁶papluca/xlm-roberta-base-language-detection

fect of the presence of sentiment in the pre-trained encoder. Since this model is based on XLM-R, which is trained on tweets and fine-tuned on sentiment datasets, it should perform better at the classification task.

4.3. Vision Encoders

The vision encoders divide an image into fixed-size segments and turn them into a sequence that the model can interpret. The encoder analyses the links between these image patches to capture the image's overall meaning, much like transformers do with text. The visual features obtained from the vision models are combined with those obtained from the Transformers text models. The combined output from the encoders is projected into the latest shared space and fine-tuned on a supervised dataset. We employed the following vision encoder models:

4.3.1. CLIP

The CLIP model, as described in the paper by Radford et al. (2021), is a multimodal framework that combines visual and linguistic information. The CLIP model utilises a transformer architecture, specifically the Vision Transformer (ViT), to extract visual features. Additionally, it employs a causal language model to acquire text features. Consequently, the textual and visual attributes are subsequently mapped onto a latent space with equivalent dimensionality. Computing the similarity score involves calculating the dot product between the projected image and text features.

4.3.2. DINOv2

The DINOv2 (Oquab et al., 2023) model is a self-supervised learning approach that builds upon the DINO framework proposed by Caron et al. (2021). The dataset utilised for pre-training purposes is meticulously curated to encompass a diverse range of images sourced from various domains and platforms, including but not limited to natural images, social media images, and product images. This ensures that the acquired features can be applied to diverse practical scenarios.

5. Experimental Setup

In this section, we provide details about the implementation and configurations that we used to train the model architecture.

5.1. Implementation

The neural network's implementation is founded on the PyTorch library. The pre-trained models in

the HuggingFace model hub are utilised through direct API calls. All monolingual models employed a batch size of 8 and a learning rate of $3e^{-5}$. All experiments were conducted using an NVIDIA V100 GPU with a memory capacity of 16GB. The translation module employed the NLLB-200-3.3B model (Costa-jussà et al., 2022), which encompasses all the languages in the dataset that are considered lower-resourced. All multilingual models employed a learning rate of $5e^{-5}$.

5.2. Model Configurations

We used the following configurations to train the model architecture and evaluate the results:

- **Unimodal vs. Multimodal:** First, we experiment with training the unimodal model by using only the text. In another configuration, we train the model using tweets' image and text content. Such a model considers both modalities and predicts the sentiment label jointly.
- **Original data vs. Inclusion of translations:** In one configuration, we used only the extracted tweets as input for the text encoder to train the model. As shown in Figure 1, not all languages within the curated dataset possess many instances that can be utilised for training purposes. Therefore, the original tweets are machine-translated from English into the target language, and we combine the original text with the translations to train the models for lower-resourced languages.
- **Monolingual vs. Multilingual:** In the monolingual setting, we train separate models for each language using data only from the respective language (either the original data or the addition of translations). In the multilingual setting, the data for all languages are merged, and we train a single model for all languages.

6. Evaluation

In this section, we analyse the outcomes produced by the aforementioned configurations. Additionally, we proceed to analyse the obtained results. The training and early-stopping of the train set are conducted based on the loss observed on the validation set. Final scoring is performed using the test set. The experiments were conducted with five different random seeds (42, 123, 777, 2020, 31337), and the resulting macro F1 scores were reported.

6.1. Results

The results (F1-score) for the model configurations are given in Table 2.

Lang	M	X	M+C	X-SM	M+D	X-SM+C	X-SM	M+C	X-SM+C
	monolingual					multilingual			
ar	57.3	64.6	53.6	66.5	25.1	69.1	41.0	61.3	72.7
bg	51.9	38.0	53.7	63.1	11.1	60.5	53.5	57.8	60.8
bs	62.4	57.0	60.5	64.4	35.4	66.5	40.3	63.1	67.9
da	48.8	34.5	46.9	66.9	21.9	59.1	55.1	57.8	75.2
de	68.7	89.1	69.4	90.1	10.7	89.6	56.3	75.3	92.9
en	34.1	18.8	30.4	36.2	6.6	33.0	64.2	52.2	53.7
es	46.5	22.6	36.9	51.6	8.0	46.4	61.4	49.4	59.6
fr	51.1	40.2	50.9	64.5	18.5	64.9	65.8	41.0	51.5
hr	58.5	28.7	56.4	64.6	25.7	55.9	40.5	57.7	63.4
hu	50.9	43.1	50.5	62.5	17.8	66.3	47.3	56.1	63.7
it	40.3	29.8	24.0	55.8	4.4	60.2	54.4	56.6	63.1
mt	60.3	60.3	60.0	68.3	11.9	62.0	35.9	44.0	56.8
pl	67.8	45.3	46.2	68.7	12.7	69.5	51.2	63.8	72.3
pt	67.2	48.1	51.8	64.3	29.5	74.6	48.3	52.8	61.8
ru	65.5	43.9	70.6	73.1	27.1	75.3	64.9	65.7	82.3
sr	42.6	23.4	38.1	49.7	21.6	43.8	48.7	49.9	65.3
sv	68.2	43.0	59.2	73.1	28.7	73.3	54.5	66.0	80.2
tr	45.9	32.1	44.4	49.6	11.6	49.4	47.9	41.3	47.8
zh	57.6	98.9	64.9	99.0	26.3	98.4	43.9	68.7	98.4
lv	22.6	19.0	24.8	22.0	21.5	18.1	76.8	52.4	61.6
sq	20.7	20.7	20.5	20.5	7.8	20.5	33.7	43.5	45.4
bg_mt	26.1	23.5	25.8	23.5	9.1	29.4			
bs_mt	17.3	19.0	15.6	18.5	9.1	20.6			
da_mt	20.7	20.7	20.7	24.5	15.0	24.7			
fr_mt	23.1	23.1	23.1	25.8	13.6	23.4			
hr_mt	34.0	25.4	28.9	34.9	16.5	46.9			
hu_mt	28.7	21.2	22.8	28.6	10.3	28			
mt_mt	30.1	18.6	20.9	43.8	12.0	26.3			
pt_mt	16.4	8.7	10.5	22.9	23.4	21.9			
ru_mt	41.3	17.8	28.8	46.9	23.7	45.6			
sr_mt	18.8	18.8	18.6	25.5	17.8	23.0			
sv_mt	31.7	17.3	24.3	54.6	19.8	34.7			
tr_mt	33.7	30.8	32.5	31.5	13.8	30.8			
zh_mt	38.2	66.7	38.0	78.1	25.3	85.4			

Table 2: F1 comparison of models using visual and textual features. M: M-BERT, C: CLIP, X: XLM-Roberta, X-SM: XLM-RoBERTa-Sentiment-Multilingual, D: DINOv2. {lang}_mt: it refers to the model that uses data from original tweets and their translations for that specific lower-resourced language. The value included within a cell containing model headers signifies the model’s performance on the test set for the specific language indicated by the lang column. Monolingual training involves the use of data from a single language, whereas multilingual training involves the incorporation of training data from multiple languages. The best result for each language is highlighted in bold.

Unimodal vs. Multimodal: In terms of using textual features to train unimodal models, we can observe that, on average, textual features from XLM-RoBERTa-Sentiment-Multilingual yielded higher F1-scores than Multilingual-BERT or XLM-RoBERTa-base. When we combine both modalities to train multimodal models, we can observe that the combination of XLM-RoBERTa-Sentiment-Multilingual with CLIP (X-SM+C) demonstrated superior performance compared to other multimodal models. The unimodal models of the Bulgarian, Danish, German, Croatian, Maltese, and Chinese languages exhibit superior performance

compared to their multimodal counterparts. In contrast, the multimodal model demonstrated superior performance for the remaining higher-resourced languages.

Original data vs Inclusion of translated text: In the context of lower-resourced languages, the utilisation of machine-translated instances sourced from higher-resourced languages, such as English, did not yield significant performance improvements. In the context of the Chinese language, including translated instances resulted in a decline in the overall performance. We hypothesise that, in contrast to product and movie reviews, which encom-

pass comprehensive contextual information as a cohesive entity, one single tweet lacks the wider contextual frame. Consequently, the translation of the original language is of lower quality and results in a modification of the overall meaning.

Monolingual vs Multilingual: When compared with monolingual models' results, on average, training a single model for all languages yielded the best performance for 17 languages, where results for Croatian, Hungarian, Maltese, Portuguese are higher with monolingual models. It suggests that providing a single model instead of 21 language-specific models is adequate for many languages of interest in this paper. Regarding modality for multilingual configuration, the combination of XLM-RoBERTa-Sentiment-Multilingual with CLIP (X-SM+C) yielded the best performance across many languages. Thus, we can confirm that the model trained with the configuration of multimodal and multilingual achieved the best score for the sentiment analysis of tweets that include both text and image content.

Figure 3 displays (on right) the average F1-scores for each language and for each combination of pre-trained models (on left). In the first subplot, it is evident that X-SM+C exhibits superior performance across all languages, with XLM-RoBERTa-Sentiment-Multilingual (X-SM) following closely behind. These findings also suggest the significance of pre-trained models, particularly those that are highly specialised or domain-specific in the context of sentiment tasks. In the second subplot, we observe that languages such as Chinese, Russian, Swedish, and German have overall better scores on all the trained models.

6.2. Error Analysis

A manual inspection was conducted on the predictions generated by the best performing unimodal and multimodal model. The errors observed in the model can be classified into the following categories:

Missing Context: The tweets exhibited a level of ambiguity that required the application of external knowledge about the world in order to determine the polarity of the messages. Given that tweets often only capture a fragment of a larger conversation and lack the necessary background context, it can be argued that these tweets require additional information beyond the presented text in order to accurately classify their content. The majority of incorrect predictions for unimodal and multimodal can be placed in this category.

Disputable: It is important to note that not all labels present in numerous datasets can be regarded as definitive ground truth, particularly in the case of (Mozetič et al., 2016), which has been previously identified as having noise and exhibiting

low inter-rater agreement (Rasooli et al., 2018). It is our contention that the identification of these instances with noisy labels should be accomplished through the utilisation of established frameworks such as Northcutt et al. (2021). This observation suggests that there remains significant potential for enhancement and validates the efficacy of the collaborative assessment of multimodal data.

Figurative Language: Although the multimodal features help in majority of the case, the models cannot comprehend cases such as sarcasm. In this case, the textual model predicts a neutral class while the multimodal model predicts a positive class, despite the fact that the original class from the dataset is negative.

In Figure 4, we show a few examples from X-SM+C where the multilingual model predicts the correct label and the unimodal makes an incorrect classification. In the example (c), the tweet contains the text "Wishing Prince George a very Happy Birthday! Mum & Dad may not be looking forward to the terrible <number>'s, but we are!" is classified as negative by the text model, but the multimodal multimodal multilingual model correctly predicts it as positive.

7. Conclusion

This paper presents the model architecture trained on the dataset extracted from various sources for multimodal sentiment classification in a multilingual context. To achieve the objective, this study employed a straightforward methodology to enhance an existing unimodal dataset from Twitter, transforming it into a multimodal one. Numerous models have been trained utilising textual data and a combination of textual and visual modalities. The primary conclusion drawn from this study is that incorporating sentiment knowledge into transformers-based models enhances the accuracy of tweet sentiment classification. The efficacy of the same model settings varies across different languages. Training a single model for all languages multilingual and multimodal data yielded the best performance across many languages. In our prospective endeavours, we intend to utilise tweets devoid of images that underwent filtration during the preprocessing phase. We aim to augment the existing dataset by incorporating additional languages. One potential avenue for advancing research is using translated datasets derived from languages other than the target language.

Limitations

The performance of pre-trained models in highly specialised or domain-specific tasks may be limited due to the broad coverage of topics in their training

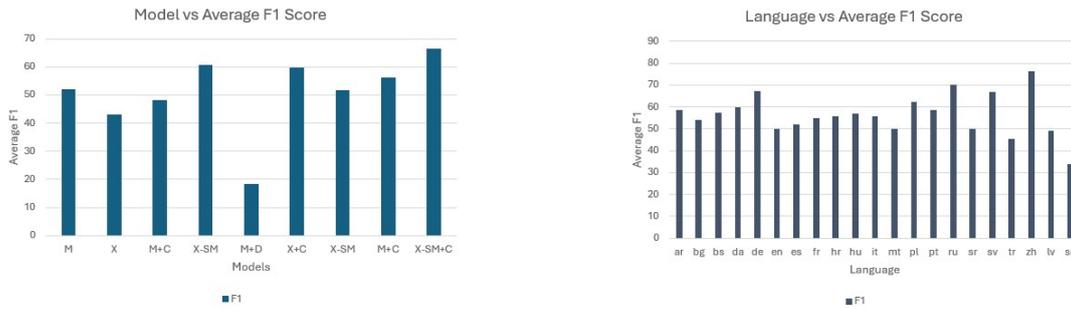


Figure 3: The (left) plot illustrates the averaged F1-score across various models. The (right) plot illustrates the averaged F1-score across languages.

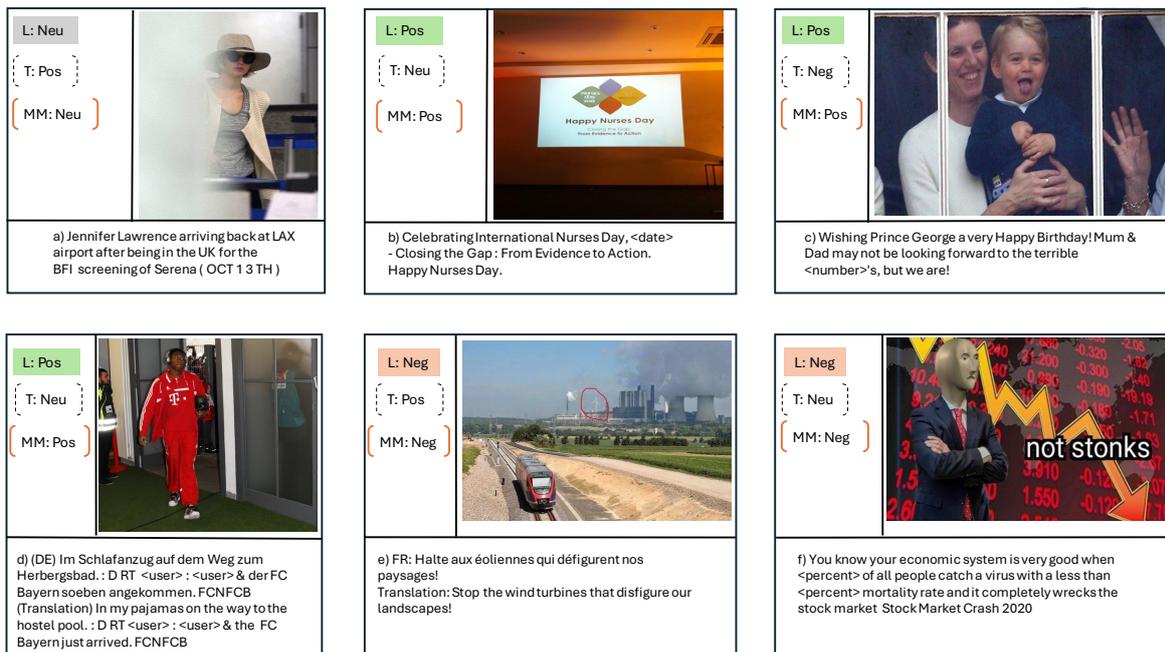


Figure 4: Examples from multilingual multimodal model predicts the correct label and text-only model fail

data. The pre-trained models learn from the data they are trained on, which can result in the introduction of any inherent biases in the training data. This bias can affect model outputs, particularly when the data do not represent all demographic, cultural, or social groups. The sentiment datasets used contain a bias towards a particular topic, which was incorporated by the annotators when the datasets were labelled.

Acknowledgements

This work was partially funded by the EU Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie grant agreement no. 812997 (CLEOPATRA ITN). This work was partially funded from the European Union's Horizon Europe Research and Innovation Programme un-

der Grant Agreement No 101070631 and from the UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee (Grant No 10039436).

8. Bibliographical References

Cerón-Guzmán Jhon Adrián. 2016-05-25. [A sentiment analysis model of spanish tweets. case study: Colombia 2014 presidential election.](#)

Marvin Agüero-Torales, David Vilares, and Antonio López-Herrera. 2021. [On the logistical difficulties and findings of jopara sentiment analysis.](#) In *Proceedings of the Fifth Workshop on Computational Approaches to Linguistic Code-Switching*, pages

- 95–102, Online. Association for Computational Linguistics.
- Dimosthenis Antypas, Asahi Ushio, Jose Camacho-Collados, Leonardo Neves, Vitor Silva, and Francesco Barbieri. 2022. Twitter Topic Classification. In *Proceedings of the 29th International Conference on Computational Linguistics*, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Karlo Babić, Milan Petrović, Slobodan Beliga, Sanda Martinčić-Ipšić, Mihaela Matešić, and Ana Meštrović. 2021. Characterisation of covid-19-related tweets in the croatian language: Framework based on the cro-cov-csebert model. *Applied Sciences*, 11(21).
- Claudio Baccchi, Tiberio Uricchio, Marco Bertini, and Alberto Del Bimbo. 2016. A multimodal feature learning approach for sentiment analysis of social network multimedia. *Multimedia Tools and Applications*, 75:2507–2525.
- Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. TweetEval: Unified benchmark and comparative evaluation for tweet classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, Online. Association for Computational Linguistics.
- Francesco Barbieri, Luis Espinosa Anke, and Jose Camacho-Collados. 2022. XLM-T: Multilingual language models in Twitter for sentiment analysis and beyond. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 258–266, Marseille, France. European Language Resources Association.
- Eugenio Martínez Cámara, Yudián Almeida-Cruz, Manuel Carlos Díaz-Galiano, Sulian Estévez-Velarde, Miguel Ángel García Cumberas, Manuel García Vega, Yoan Gutiérrez, Arturo Montejo-Ráez, Andrés Montoyo, Rafael Muñoz, Alejandro Piad-Morffis, and Julio Villena-Román. 2018. Overview of TASS 2018: Opinions, health and emotions. In *Proceedings of TASS 2018: Workshop on Semantic Analysis at SEPLN, TASS@SEPLN 2018, co-located with 34th SEPLN Conference (SEPLN 2018), Sevilla, Spain, September 18th, 2018*, volume 2172 of *CEUR Workshop Proceedings*, pages 13–27. CEUR-WS.org.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. 2021. Emerging properties in self-supervised vision transformers. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 9630–9640. IEEE.
- Giuseppe Castellucci, Andrea Vanzo, Danilo Croce, and Roberto Basili. 2015. Context-aware models for twitter sentiment analysis. *IJCoL vol. 1, n. 1 december 2015: Emerging Topics at the First Italian Conference on Computational Linguistics*, page 69.
- Gullal S. Cheema, Sherzod Hakimov, Eric Müller-Budack, and Ralph Ewerth. 2021. A fair and comprehensive comparison of multimodal tweet sentiment analysis methods. In *MMPT@ICMR2021: Proceedings of the 2021 Workshop on Multimodal Pre-Training for Multimedia Understanding, Taipei, Taiwan, August 21, 2021*, pages 37–45. ACM.
- Mathieu Cliche. 2017. Bb_twtr at semeval-2017 task 4: Twitter sentiment analysis with cnns and lstms. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 573–580.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020a. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020b. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Keith Cortis and Brian Davis. 2021. A Dataset of Multidimensional and Multilingual Social Opinions for Malta’s Annual Government Budget.
- Marta R Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. 2022. No language left behind: Scaling human-centered machine translation. *ArXiv preprint*, abs/2207.04672.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019a. 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis,

- Minnesota. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019b. [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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Manuel Carlos Díaz-Galiano, Manuel García Vega, Edgar Casasola, Luis Chiruzzo, Miguel Ángel García Cumbreñas, Eugenio Martínez Cámara, Daniela Moctezuma, Arturo Montejo-Ráez, Marco Antonio Sobrevilla Cabezudo, Eric Sadit Tellez, et al. 2019. Overview of tass 2019: One more further for the global spanish sentiment analysis corpus. In *IberLEF@ SEPLN*, pages 550–560.
- Manuel García-Vegaa, Manuel Carlos Díaz-Galiano, Miguel Á García-Cumbreñas, Flor Miriam Plaza del Arcoa, Arturo Montejo-Ráeza, Salud María Jiménez-Zafraa, Eugenio Martínez Cámarab, César Antonio Aguilarc, Marco Antonio, Sobrevilla Cabezudod, et al. 2020. Overview of tass 2020: Introducing emotion detection. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2020) co-located with 36th Conference of the Spanish Society for Natural Language Processing (SEPLN 2020), Málaga, Spain*, pages 163–170.
- Muskan Garg, Seema Wazarkar, Muskaan Singh, and Ondřej Bojar. 2022. [Multimodality for NLP-centered applications: Resources, advances and frontiers](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6837–6847, Marseille, France. European Language Resources Association.
- Aniruddha Ghosh, Guofu Li, Tony Veale, Paolo Rosso, Ekaterina Shutova, John Barnden, and Antonio Reyes. 2015. [SemEval-2015 task 11: Sentiment analysis of figurative language in Twitter](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 470–478, Denver, Colorado. Association for Computational Linguistics.
- Alec Go, Richa Bhayani, Lei Huang, and Riloff Ellen. 2009. Twitter sentiment analysis. *arXiv preprint arXiv:0903.4294*.
- Ju Huang, Pengtao Lu, Shuifa Sun, and Fangyi Wang. 2023. Multimodal sentiment analysis in realistic environments based on cross-modal hierarchical fusion network. *Electronics*, 12(16):3504.
- Adela Ljajić and Ulfeta Marovac. 2018. [Improving sentiment analysis for twitter data by handling negation rules in the serbian language](#). *Computer Science and Information Systems*, 16:13–13.
- Daniela Moctezuma, Circuito Tecnopolo Norte No, Eric S Tellez, Mario Graff, Sabino Miranda-Jiménez, and Circuito Tecnopolo Sur. 2016. On the performance of b4msa on sentipolc'16. In *of the Final Workshop 7 December 2016, Naples*, page 200.
- Saif Mohammad, Svetlana Kiritchenko, Mohammad Salameh, and John Paul Dredze. 2018. Multimodal twitter sentiment analysis. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3377–3389.
- Arturo Montejo-Ráez and Manuel Carlos Díaz-Galiano. 2016. Participación de sinai en tass 2016. In *TASS@ SEPLN*, pages 41–45.
- Igor Mozetič, Miha Grčar, and Jasmina Smailović. 2016. [Twitter sentiment for 15 european languages](#). Slovenian language resource repository CLARIN.SI.
- K Muischnek and K Müürisepp. 2018. Latvian tweet corpus and investigation of sentiment analysis for latvian. In *Human Language Technologies—The Baltic Perspective: Proceedings of the Eighth International Conference Baltic HLT 2018*, volume 307, page 112. IOS Press.
- Mustafa Melih Mutlu and Arzucan Özgür. 2022. [A dataset and BERT-based models for targeted sentiment analysis on Turkish texts](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 467–472, Dublin, Ireland. Association for Computational Linguistics.
- Preslav Nakov, Sara Rosenthal, Zornitsa Kozareva, Veselin Stoyanov, Alan Ritter, and Theresa Wilson. 2013a. [SemEval-2013 task 2: Sentiment analysis in Twitter](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 312–320, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Preslav Nakov, Sara Rosenthal, Zornitsa Kozareva, Veselin Stoyanov, Alan Ritter, and Theresa Wilson. 2013b. [SemEval-2013 task 2: Sentiment analysis in Twitter](#). In *Second Joint Conference*

- on *Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 312–320, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Preslav Nakov, Sara Rosenthal, Zornitsa Kozareva, Veselin Stoyanov, Alan Ritter, and Theresa Wilson. 2013c. [SemEval-2013 task 2: Sentiment analysis in Twitter](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 312–320, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Preslav Nakov, Sara Rosenthal, Zornitsa Kozareva, Veselin Stoyanov, Alan Ritter, and Theresa Wilson. 2013d. [SemEval-2013 task 2: Sentiment analysis in Twitter](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 312–320, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Curtis Northcutt, Anish Athalye, and Jonas Mueller. 2021. [Pervasive label errors in test sets destabilize machine learning benchmarks](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1. Curran.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. 2023. [Dinov2: Learning robust visual features without supervision](#). *ArXiv preprint*, abs/2304.07193.
- Bo Pang and Lillian Lee. 2005. [Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 115–124, Ann Arbor, Michigan. Association for Computational Linguistics.
- Hallan Patrick, Fernando Paulo Belfo, and António Trigo. 2022. Brazilian tweets classified for sentiment analysis.
- Amalie Brogaard Pauli, Maria Barrett, Ophélie Lacroix, and Rasmus Hvingelby. 2021. [DaNLP: An open-source toolkit for Danish natural language processing](#). In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 460–466, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.
- Soujanya Poria, Erik Cambria, Newton Howard, Guang-Bin Huang, and Amir Hussain. 2016a. [Fusing audio, visual and textual clues for sentiment analysis from multimodal content](#). *Neurocomputing*, 174:50–59.
- Soujanya Poria, Iti Chaturvedi, Erik Cambria, and Amir Hussain. 2016b. [Convolutional mkl based multimodal emotion recognition and sentiment analysis](#). In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pages 439–448.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Colin Raffel, Noam M. Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Mohammad Sadegh Rasooli, Noura Farra, Axinia Radeva, Tao Yu, and Kathleen McKeown. 2018. Cross-lingual sentiment transfer with limited resources. *Machine Translation*, 32:143–165.
- Luis Rei, Simon Krek, and Dunja Mladenić. 2016. [xLiMe twitter corpus XTC 1.0.1](#). Slovenian language resource repository CLARIN.SI.
- Julio Villena Román, Eugenio Martínez Cámara, Janine García Morera, and Salud M Jiménez Zafra. 2015. Tass 2014-the challenge of aspect-based sentiment analysis. *Procesamiento del Lenguaje Natural*, 54:61–68.
- Lucia Prieto Santamaría, Juan Manuel Tuñas, Diego Fernández Peces-Barba, Almudena Jaramillo, Manuel Cotarelo, Ernestina Menasalvas, Antonio Conejo Fernández, Amalia Arce, Angel Gil de Miguel, and Alejandro Rodríguez González. 2022. [Influenza and measles-mmr: two case study of the trend and impact of vaccine-related twitter posts in spanish during 2015-2018](#). *Human Vaccines & Immunotherapeutics*, 18(1):1–16. PMID: 33662222.
- Soroosh Tayebi Arasteh, Mehrpad Monajem, Vincent Christlein, Philipp Heinrich, Anguelos Nicolaou, Hamidreza Naderi Boldaji, Mahshad Lotfinia, and Stefan Evert. 2021. [How will your tweet be received? predicting the sentiment polarity of tweet replies](#). In *Proceedings of the 2021*

IEEE 15th International Conference on Semantic Computing (ICSC), pages 370–373, Laguna Hills, CA, USA.

Andrea Vanzo, Danilo Croce, and Roberto Basili. 2014. [A context-based model for sentiment analysis in Twitter](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2345–2354, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.

David Vilares, Yerai Doval, Miguel A Alonso, and Carlos Gómez-Rodríguez. 2015. Lys at tass 2015: Deep learning experiments for sentiment analysis on spanish tweets. In *TASS@ SEPLN*, pages 47–52.

Julio Villena-Román and Janine Garcia-Morera. 2013. Tass 2013—workshop on sentiment analysis at sepln 2013: An overview. In *Proceedings of the TASS workshop at SEPLN*, pages 112–125.

Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol. 2010. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *J. Mach. Learn. Res.*, 11:3371–3408.

Vedran Vukotić, Vincent Claveau, and Christian Raymond. 2015. Irisa at deft 2015: supervised and unsupervised methods in sentiment analysis. In *DeFT, Défi Fouille de Texte, joint à la conférence TALN 2015*.

Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. 2020. Unsupervised data augmentation for consistency training. *Advances in neural information processing systems*, 33:6256–6268.

Wenjie Yin, Rabab Alkhalifa, and Arkaitz Zubiaga. 2021. [TM-Senti](#).

Quanzeng You, Liangliang Cao, Hailin Jin, and Jiebo Luo. 2016. [Robust visual-textual sentiment analysis: When attention meets tree-structured recursive neural networks](#). In *Proceedings of the 2016 ACM Conference on Multimedia Conference, MM 2016, Amsterdam, The Netherlands, October 15-19, 2016*, pages 1008–1017.

Wenmeng Yu, Hua Xu, Fanyang Meng, Yilin Zhu, Yixiao Ma, Jiele Wu, Jiyun Zou, and Kaicheng Yang. 2020. [CH-SIMS: A Chinese multimodal sentiment analysis dataset with fine-grained annotation of modality](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3718–3727, Online. Association for Computational Linguistics.