



HAL
open science

Optimal Strategies to Perform Multilingual Analysis of Social Content for a Novel Dataset in the Tourism Domain

Maxime Masson, Rodrigo Agerri, Christian Sallaberry, Marie-Noelle Bessagnet, Annig Le Parc Lacayrelle, Philippe Roose

► To cite this version:

Maxime Masson, Rodrigo Agerri, Christian Sallaberry, Marie-Noelle Bessagnet, Annig Le Parc Lacayrelle, et al.. Optimal Strategies to Perform Multilingual Analysis of Social Content for a Novel Dataset in the Tourism Domain. 2023. hal-04326748

HAL Id: hal-04326748

<https://hal.science/hal-04326748>

Preprint submitted on 6 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimal Strategies to Perform Multilingual Analysis of Social Content for a Novel Dataset in the Tourism Domain

Maxime Masson^{a,b,*}, Rodrigo Agerri^b, Christian Sallaberry^a,
Marie-Noelle Bessagnet^a, Annig Le Parc Lacayrelle^a, Philippe Roose^a

^a*LIUPPA, E2S, University of Pau and Pays Adour (UPPA), Av. de
l'Université, Pau, 64000, France*

^b*HiTZ Center - Ixa, University of the Basque Country UPV/EHU, M. Lardizabal
1, Donostia-San Sebastian, 20080, Basque Country, Spain*

Abstract

The rising influence of social media platforms in various domains, including tourism, has highlighted the growing need for efficient and automated natural language processing (NLP) approaches to take advantage of this valuable resource. However, the transformation of multilingual, unstructured, and informal texts into structured knowledge often poses significant challenges.

In this work, we evaluate and compare few-shot, pattern-exploiting and fine-tuning machine learning techniques on large multilingual language models (LLMs) to establish the best strategy to address the lack of annotated data for 3 common NLP tasks in the tourism domain: (1) Sentiment Analysis, (2) Named Entity Recognition, and (3) Fine-grained Thematic Concept Extraction (linked to a semantic resource). Furthermore, we aim to ascertain the quantity of annotated examples required to achieve good performance in those 3 tasks, addressing a common challenge encountered by NLP researchers in the construction of domain-specific datasets.

Extensive experimentation on a newly collected and annotated multilin-

*Corresponding author

Email addresses: maxime.masson@univ-pau.fr (Maxime Masson),
rodrigo.agerri@ehu.eus (Rodrigo Agerri), christian.sallaberry@univ-pau.fr
(Christian Sallaberry), marie-noelle.bessagnet@univ-pau.fr
(Marie-Noelle Bessagnet), annig.lacayrelle@univ-pau.fr (Annig Le Parc Lacayrelle),
philippe.roose@univ-pau.fr (Philippe Roose)

gual (French, English, and Spanish) dataset composed of tourism-related tweets shows that current few-shot learning techniques allow us to obtain competitive results for all three tasks with very little annotation data: 5 tweets per label (15 in total) for Sentiment Analysis, 10% of the tweets for *location detection* (around 160) and 13% (200 approx.) of the tweets annotated with *thematic concepts*, a highly fine-grained sequence labeling task based on an inventory of 315 classes.

This comparative analysis, grounded in a novel dataset, paves the way for applying NLP to new domain-specific applications, reducing the need for manual annotations and circumventing the complexities of rule-based, ad hoc solutions.

Keywords: Tourism, Few-Shot Learning, Large Language Models, Multilinguality, Computational Social Science, Natural Language Processing

1. Introduction

Social media platforms have become essential channels for sharing opinions and experiences about tourist practices and itineraries. Twitter, in particular, has become a popular medium for people to spontaneously share their thoughts and recommendations, making it a valuable source of information and user-generated content (UGC) for the tourism industry [1] (such as tourism offices, destination marketing organizations - DMOs, etc.).

However, analyzing vast volumes of social media data can be challenging [2], especially when it comes to extracting structured knowledge from unstructured text. Consequently, tourism stakeholders often delegate the task of knowledge extraction to researchers, who rely on Natural Language Processing techniques. NLP offers a powerful set of techniques for processing and analyzing text data and is often used to address three common knowledge extraction tasks: Sentiment Analysis, Named Entity Recognition (NER) for locations, and Fine-Grained Thematic Concept Extraction [3, 4, 5].

Recently, NLP techniques based on deep learning and large language models (LLMs) have emerged. These techniques offer several advantages over traditional rule-based approaches. Deep learning-based approaches can adapt to changing language patterns and structures [6], ensuring a more dynamic and up-to-date analysis. In addition, they can handle vast amounts of data efficiently. However, to achieve optimal results in domain-specific

applications, language models must be fine-tuned. Consequently, researchers often face two recurrent challenges: (1) determining which NLP strategy is most suitable for a given domain (e.g., which technique, which LLM, etc.), and (2) discerning how many domain-specific examples are truly necessary to achieve competitive NLP results. As annotating datasets is both costly and time-consuming, researchers strive to keep the annotation work to a minimum while maintaining high-quality results.

In this paper, we present a comparative study on the data requirements for achieving competitive performance in NLP tasks within the **tourism domain**. We hypothesize that among the existing language models and training techniques, some will be better suited for this specific domain, especially given the unique context of social media characterized by short informal texts, frequent grammatical errors, and the presence of emojis. We focus on the three common knowledge extraction tasks mentioned above: Sentiment Analysis (text classification), NER for Locations (sequence labeling), and Fine-Grained Thematic Concept Extraction (sequence labeling). Specifically, we investigate which NLP strategies are best to keep manual data annotation to a minimum and to avoid cumbersome and costly rule-based ad hoc approaches. To enable this study, we contribute a new manually annotated and multilingual (French, English, Spanish) dataset extracted from social media which includes annotations regarding sentiment, locations and touristic thematic concepts.

The main contributions are the following: (i) we present a **novel dataset** of tourism tweets. This dataset is multilingual (French, English, and Spanish) and has been manually annotated at the text level with *sentiment* labels and at the token level with *locations* and *thematic concepts* linked to a fine-grained tourism thesaurus (WTO Thesaurus of Tourism and Leisure Activities [7]) which, to the best of our knowledge, is the first dataset of its kind; (ii) we perform a **comparative study** on rule-based, fine-tuning and few-shot learning techniques with the aim of establishing which of the methods is the most efficient for each NLP task (Sentiment Analysis, NER for locations and Fine-Grained Thematic Concept Extraction) on social media data for the tourism domain; (iii) Additionally, we experiment with various numbers of examples and dataset sampling methods to determine **how many annotated examples are really needed** to achieve competitive results for each tasks, using different training techniques and large language models (LLMs).

Experiments using the new dataset show that current few-shot learning techniques allow to obtain competitive results for all three tasks with very

little annotation data: 5 tweets per label (15 in total) for Sentiment Analysis, 10% of the tweets for *location detection* (around 160 tweets) and 13% of the tweets annotated with *thematic concepts* (around 200 tweets). Interesting results also point out that few shot prompting for sequence labeling [8] seems to be particularly effective for highly fine-grained tasks (more than 300 classes) even with a limited number of examples.

These results provide a promising solution to apply NLP to new domain-specific applications keeping the manual annotation requirements low while avoiding complex rule based ad hoc solutions. The code and annotated data will be made publicly available in the future to facilitate the reproducibility of the results and encourage research on this particular topic.

In summary, we believe that the findings of our paper may be useful not only for researchers interested in the tourism domain, but for any application that requires NLP analysis in a domain-specific scenario when no annotated data is available or when complex and costly ad hoc rule-based approaches are to be avoided.

The rest of the paper is structured as follows. In Section 2, we provide an overview of the NLP techniques based on deep learning used in tourism. Section 3 covers the construction and annotation of the dataset. Section 4 describes the experimental setup. In Section 5, we present a comparative analysis of fine-tuning and few-shot approaches for three common NLP tasks. Furthermore, the results and limitations of the experiment are discussed in Section 6. Finally, Section 7 provides some concluding remarks and offers some insight into the future application of the work presented here.

2. Related Work

In the rapidly evolving field of Natural Language Processing (NLP), one of the most significant advancements has been the advent of pre-trained large language models (LLMs). These models serve as foundational architectures trained on vast corpora, capturing a broad spectrum of linguistic structures, nuances, and knowledge [6]. As a result, they offer a significant boost in performance and generalization across a myriad of NLP tasks, from text classification to question-answering systems.

Pre-trained language models like BERT (Bidirectional Encoder Representations from Transformers) [9] or XLM-RoBERTa [10] are trained on large amounts of textual data, sometimes encompassing terabytes of information

from diverse sources such as books, articles, and websites. Through this extensive training process, these models are equipped to understand intricate relationships between words, phrases, and larger text constructs. They capture semantic, syntactic and contextual signals, which provide them with a powerful ability to generate and understand language [11, 12].

2.1. Previous work on social media for the tourism domain

A popular approach in NLP is to **fine-tune** pre-trained language models for domain-specific downstream tasks (see Table 1). This results in altering the model weights to adapt it to the new task. For instance, pre-trained language models have been fine-tuned for text classification tasks, including spam detection in hotel reviews [13] and sentiment analysis in touristic reviews [14, 15] or reviews about sustainable transport [16], leading to improved accuracy. In NER, fine-tuning of pre-trained language models has been employed to extract location information from a tourism corpus [17, 18]. Furthermore, pre-trained language models have demonstrated promising results in thematic concept extraction, such as identifying travel-related themes and topics from tourism texts [19].

Table 1: Application of Fine-Tuning in the Tourism Domain

Ref.	Year	Objective	Label
[13]	2021	Spam Detection (Hotel Reviews)	Text
[14]	2022	Sentiment Analysis (Tourism Reviews)	Text
[15]	2021	Sentiment Analysis (Tourism Reviews)	Text
[16]	2021	Sentiment Analysis (Transport)	Text
[20]	2018	Classification of Basque Users	Text
[17]	2022	Location Extraction (Touristic Corpus)	Token
[18]	2020	Location Extraction (Touristic Corpus)	Token
[19]	2021	Travel Themes Identification	Token

2.2. Addressing the lack of domain-specific annotated data

Apart from fine-tuning techniques, **few-shot** learning techniques have emerged as a novel approach for NLP tasks. This approach enables a language model to learn from a few labeled examples, which is especially useful when a considerable amount of labeled data is unavailable [21]. Few-shot methods often use prompting to guide pre-trained language models towards specific tasks, rather than updating the model weights. Pattern-Exploiting

Training (PET) is a semi-supervised few-shot training approach for natural language processing tasks related to text classification. It combines the idea of providing a pre-trained language model with task descriptions in natural language with a cloze-style phrase generation approach to help the model understand the task [22]. For example, to classify movie reviews based on the predominant sentiment they express, the model would be prompted with the query: “*The movie was ?*” The model would then try to predict the blank, choosing from options such as outstanding (positive) or terrible (negative).

Several recent studies have also explored methods (e.g., EntLM [8]) to replace complex templates used in few-shot prompting for sequence labeling tasks such as Named Entity Recognition [23, 8]. These approaches aim to simplify the process of generating task-specific queries and reduce the reliance on manual template construction. In both cases, few-shot approaches have shown their superiority over fine-tuning methods to obtain competitive results with very little annotation data.

The effectiveness of pre-trained language models highly depends on the diversity and quantity of the training data and the specificity of the task.

Therefore, it is crucial to determine how much training data must be annotated to get competitive results. An alternative to few-shot learning to address the lack of annotated data for NLP tasks on social media is the use of multilingual LLMs to perform data augmentation via machine translation. The idea is to translate into multiple languages the training data from a source language for which there is training data available, and then use the translated versions to perform data augmentation during training. This approach has been tested for many sequence labeling [24, 25] and classification tasks [26, 27] in many genres of text, including sentiment analysis in social media [28]. Although interesting, cross-lingual transfer approaches are based on transferring *existing* annotations from (at least) a source to a given target language(s). However, our starting point is the absence of *any annotated data in any language* for tourist locations, *fine-grained tourist concept detection* and sentiment analysis for the tourist domain.

2.3. Existing Annotated Resources

While annotated data for the tourism domain is non-existent, there are several existing annotated corpora, from other domains, which could be used for experimentation, as highlighted in Table 2. Here, we compare a selection of existing annotated datasets along different criteria: (1) the source of the data collection, (2) the types of annotations available, (3) the languages

covered by the dataset, and (4) the method used for generating annotations (manual by humans, semi-automatic or automatic).

For example, the **ESTER** corpus [29] is a comprehensive collection of French radio transcripts and **AnCora** [30], is a multilevel annotated corpus (*mostly from newspaper*) for Catalan and Spanish. Both of these resources are annotated for Named Entity Recognition (*such as persons, locations, organizations*). In terms of social media-specific resources, the **Broad Twitter Corpus** (BTC) [31] includes coarse-grained NER annotations, while **Sentiment140** [32], **STS-Gold** [33], and many other datasets developed as part of shared evaluation tasks at **SemEval** [34, 35, 3, 36, 37], are used for sentiment analysis. Other corpora include the **MultiWOZ** dialogue dataset [38], the **Stanford NLI** dataset [39] for text inference, and the **Heldugazte** corpus [20], which assists in categorizing tweets as formal or informal.

These datasets are extensive but broad and often focus on English only, therefore lacking the necessary contextual information relevant to the tourism domain. Most importantly, we could not find any public dataset annotated for *fine-grained thematic concept extraction* in the tourism domain. Taking this into account, we decided to build our own custom annotated dataset.

Table 2: Comparison of Existing Annotated Dataset

Dataset	Source	Annotations	Lang.	Type
ESTER [29]	Radio	Named Entities	FR	Manual
AnCora [30]	Newspapers	Named Entities	ES, CA	Semi
BTC [31]	Twitter	Named Entities	EN	Manual
Sent140 [32]	Twitter	Sentiment	EN	Auto
STS-Gold [33]	Twitter	Sentiment	EN	Manual
GSC [16]	Reviews	Sentiment	EN	Auto
SemEval [34]	Twitter	Sentiment	EN	Manual
MultiWOZ [38]	Humans	Dialogue States	EN	Manual
SNLI [39]	Humans	Inference Pairs	EN	Manual
Heldugazte [20]	Twitter	Formal/Informal	EU	Auto

3. Dataset Building and Annotation

In this section, we describe the creation of a novel multilingual dataset consisting of tourism-related tweets annotated for three important NLP tasks for tourist applications: (1) **Sentiment Analysis**, (2) **Named Entity**

Recognition for Locations, and (3) **Fine-grained Thematic Concept Extraction** (based on the WTO Thesaurus on Tourism and Leisure Activities [7]).

3.1. Collection

The dataset was collected from Twitter using the Academic API¹ and the collection process was carried out using a novel methodology that we have designed for dataset building. This methodology is both generic, iterative and incremental, [40]. Several iterations were carried out, each of them with successive filtering corresponding to the dataset target *dimensions*:

- **Spatial:** French Basque Coast area (*defined by spatial coordinates for geotagged tweets or a list of toponyms*).
- **Temporal:** Summer of 2019 (*21 Jun – 21 Sept*). Based on tweets' timestamps.
- **Thematic:** Tourism domain (*defined by the World Tourism Organization Thesaurus on Tourism and Leisure*).

To prevent excessive noise, each iteration was followed by human feedback to adjust and balance the filters. For example, a large number of tweets related to the G7 summit, which took place in the region that year, were collected because they contained tourism-related keywords; these were subsequently blacklisted using hashtag and keyword exclusion rules. Similarly, we excluded professional and institutional users because we are primarily interested in understanding the behaviors and feelings of individual tourists, not analyzing promotional or institutional content.

The final dataset is made up of 27,379 tweets, of which 2,961 tweets from 624 users were selected to be annotated and used for experimentation. The tweets in this subset (2,961) were manually checked to confirm that they were *about tourism* and *from tourist visitors* (unlike the tweets from tourism professionals, or from news outlets speaking about tourism, etc.). The dataset is multilingual and includes an unbalanced variety of tweets in English, French, and Spanish, which reflects the reality of the use of social

¹<https://developer.twitter.com/en/use-cases/do-research/academic-research> (discontinued in April 2023)

media on the French Basque Coast. Table 3 provides the language distribution of the dataset and the splits created for experimentation (60% train, 20% dev, 20% test). Splits were generated maintaining a balance between the number of users and languages in each set.

Table 3: Breakdown of the Collected Dataset by Language – Tweets (Users)

	All	French	English	Spanish
Train	1,662 (503)	1,297 (391)	283 (129)	82 (32)
Dev	619 (300)	450 (213)	99 (66)	70 (31)
Test	680 (431)	401 (273)	102 (100)	177 (93)

3.2. Sentiment Annotation

The annotation process of those 2,961 tweets was carried out in a semi-automatic manner, following the procedure depicted in Figure 1.

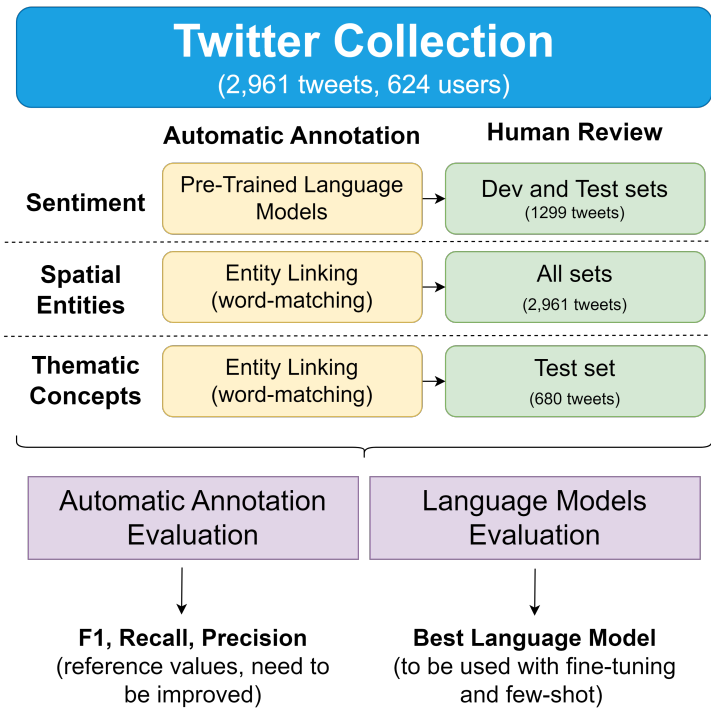


Figure 1: Dataset Building and Annotation Process

Firstly, to assist human annotators, the 1,299 tweets in the development and test splits underwent a process of **automatic annotation** using the

5 language models listed in Table 4. Subsequently, they were manually reviewed. Each tweet was assigned to two annotators to evaluate agreement (*Cohen’s kappa coefficient*) and ensure the quality of the annotations. We achieved $\kappa = 0.79$ for French tweets, $\kappa = 0.75$ for Spanish tweets, and $\kappa = 0.67$ for English tweets, which corresponds to a strong agreement. Any disagreements were resolved through collaborative discussion.

Table 4: Accuracy of Available Sentiment Language Models on manually annotated test data.

Sentiment Models:	[41]	[42]	[43]	[44]	[45]
	French				
Global	0.56	0.45	0.43	0.47	0.82
Positive	0.34	0.14	0.11	0.95	0.82
Negative	0.06	0.11	0.00	0.28	1.00
Neutral	0.97	0.97	1.00	0.00	0.74
	Spanish				
Global	0.71	0.64	0.61	0.34	0.83
Positive	0.31	0.09	0.03	1.00	0.84
Negative	0.00	0.00	0.00	0.29	0.43
Neutral	1.00	1.00	0.98	0.00	0.87
	English				
Global	0.81	0.81	0.71	0.66	0.80
Positive	0.75	0.75	0.59	1.00	0.72
Negative	0.75	0.50	0.75	0.50	0.75
Neutral	0.94	0.97	0.94	0.00	0.97

The next step was to evaluate the performance of the 5 language models used to automatically label the tweets with respect to the human annotations. Table 4 shows that XLM-T Sentiment, fine-tuned with multilingual sentiment analysis data from various domains different to tourism [45], outperformed, on average for the three languages, any other approach. Following this, we automatically annotated using XLM-T Sentiment the training split which was then manually revised for its use in the experimentation.

3.3. Locations and Thematic Concepts

Although sentiment analysis is a text classification task in which each tweet is assigned a polarity label, we are also interested in identifying lo-

cations and thematic concepts relevant to the tourism domain. These two tasks are addressed as sequence labeling problems.

Before experimenting with supervised approaches, we implemented a basic rule-based *word-matching* approach as a baseline for *locations* and *thematic concepts*. Locations were matched using 625 local toponyms extracted from Open Street Map (*cities, POIs, landmarks, etc.*) while thematic concepts were matched using their label and synonyms in the WTO thesaurus of tourism (which contains 1,494 touristic concepts).

Tweet preprocessing (lowercase, removal of URLs, hashtag splitting, decomposing hashtags to find concepts or toponyms in them) was performed to facilitate word-matching. We applied this algorithm to annotate the train, dev and test splits. Automatic annotations were then manually corrected by human annotators: for *locations* all train, dev, and test sets. For *thematic concepts* the *word-matching* algorithm detected 315 concept classes for the full dataset (out of the 1,494 concepts included in the WTO ontology), making it a highly fine-grained sequence labeling task. Due to this fact, we just revised the *thematic concepts* for the test set, as annotating 315 concept classes is a complex task requiring a large human effort. Finally, inter-annotator agreement was calculated on a subset of 100 random tweets. For location entities: $\kappa = 0.91$ for exact matches, when all tokens forming an entity are precisely the same (e.g., New (B-LOC), - (I-LOC), York (I-LOC)). On the other hand, $\kappa = 0.93$ for partial matches, when an entity is mostly recognized but has missing or extra tokens (e.g., New (B-LOC), - (O), York (O)). Both values indicate a near-perfect consensus.

The F1 results of evaluating *word-matching* baseline on the test set are reported in Table 5. These results will serve as a baseline to compare with the different supervised approaches in the experimental section.

Table 5: Performance of the Word-Matching Algorithm

NER for Locations			
	Recall	Precision	F1-score
Location Exact Match	0,692	0,722	0,707
Location Partial Match	0,780	0,814	0,796

Fine-grained Thematic Concept Extraction

	Recall	Precision	F1-score
Concept Exact Match	0,746	0,952	0,836
Concept Partial Match	0,747	0,953	0,837

From the results in Table 5, it can be observed that, for locations, results are not that good, particularly in terms of recall. However, the *word-matching* algorithm performs remarkably well on fine-grained thematic concept extraction, especially in terms of precision. Although this constitutes a rather strong baseline, the recall remains comparatively low, which means that many thematic concepts remain undetected by the system. Furthermore, developing such as *word-matching* algorithm is a rather complex and time-consuming exercise which we would like to ideally avoid for any domain-specific new application.

Thus, our main objective is now to establish whether deep-learning supervised techniques based on multilingual LLMs can match or improve over the *word-matching* algorithm while keeping the amount of manual annotation to a minimum, especially for *fine-grained thematic concept extraction*. Thus, in addition to standard fine-tuning approaches, it is of particular interest investigating techniques based on few-shot learning, where the aim is to generate competitive taggers using only a very small amount of labeled data.

Table 6: Dataset Annotations (following human review)

Set	Tweets	Locations	Concepts	Sentiments
Train	1,662	4,030	3,841	787 (+) 191 (-) 684 (=)
Dev	619	1,419	1,337	271 (+) 82 (-) 266 (=)
Test	680	1,679	1,844	299 (+) 93 (-) 288 (=)
Global	2,961	7,128	7,022	1,357 (+) 366 (-) 1,238 (=)

The data used for experimentation is the one generated by manually revising the annotations for the three tasks, as described above. Table 6

provides a detailed description of the dataset including the total number of tweets (2,961) and the number of annotations per task. Note that for the sentiment column in Table 6, the colors and symbols used represent each type of polarity (- for negative annotations, + for positive and = for neutral).

In summary, the goal is to establish how much labeled data do we actually need to obtain competitive performance by comparing the *fine-tuning* and *few-shot* approaches with respect to the *word-matching* algorithm, and which of these approaches is the most efficient in terms of performance and human effort.

4. Experimental Setup

The experimentation is focused on the three tasks presented above: *Sentiment Analysis* (Section 4.1), *NER for Locations* and *Fine-grained Thematic Concept Extraction* (Section 4.2).

These experiments leverage the tourist dataset described in the previous section (*2,961 multilingual tweets including 1,662 for training*) and summarized in Table 6. To study few-shot learning approaches, the dataset is sampled using two different methods:

- **k-shot sampling:** In this method, we selected a specific number of examples for each tweet or token label from the training set. We performed few-shot using the following k-values: 5, 10, 20, 30, 40, 50, and 100 examples per label. For sentiment analysis, we used the PET sampling method [22] while for *locations* and *fine-grained thematic concept extraction*, we apply the EntLM method with default parameters [8].
- **Percentage sampling:** For Sequence labeling (*locations* and *fine-grained thematic concepts*) we also experiment with sampling on percentages of tweets rather than labels, as k-shot does. We successively used 5%, 10%, 20%, 30%, 40%, 50%, 60%, 70%, 90%, and 100% of the training set, while trying to maintain the original label classes distribution, including the O labels (the O label is assigned to tokens which are neither location nor thematic entities), which the k-shot sampling method does not contemplate.




Figure 2: Experimental Setup of the Comparative Study

Figure 2 provides an overview of our experimental setup with the models, sampling methods and learning techniques used for each task. We will go into the details in the next two subsections.

4.1. Text Classification – Sentiment Analysis

Based on the results reported by Table 4, we chose XLM-T for experiments on Sentiment Analysis. XLM-T is based on XLM-RoBERTa [10], further pre-trained on a corpus of 198 million tweets for 15 languages [45]. This version of the model is specifically designed to handle the unique characteristics of tweets and social media posts, such as their limited length, informal language, and the presence of emojis and hashtags. More specifically, we use two variants of the XLM-T model:

- The **base version** [45] (XLM-T), which is a masked language model (MLM).
- XLM-T previously fine-tuned specifically for *sentiment analysis* [45] (XLM-T SENTIMENT). This sentiment variant has been fine-tuned using 24,264 tweets in eight different languages, including French, English, and Spanish. However, it is important to note that these tweets cover a wide range of topics which do not include tourism.

These two models are then fine-tuned or use in a few-shot prompting approach with various dataset sizes using the two sampling methods described above.

- **Fine-Tuning:** optimal hyperparameters were found through grid search. For XLM-T 8 batch size, 2e-5 learning rate, weight decay 0.01; for XLM-T SENTIMENT 32 batch, 1e-5 learning rate and decay: 0.1.
- **Pattern-Exploiting Training (PET) [22]:** an approach to perform few-shot learning via prompting for text classification tasks. It turns sentiment analysis into a cloze task, where the model fills in the blanks in text, enabling the learning of task-specific knowledge from minimal data. For instance, in our specific use case, the cloze statement is formulated as: “*The dominant sentiment expressed in the following text: [Tweet] is ?*”. The language model will then attempt to fill in the “?” with the appropriate sentiment from a list of possible labels: *positive, negative, or neutral*.

By comparing these two deep learning methods and assessing their effectiveness with varying amounts of annotated data, we aim to learn any insights about the minimum data requirements for achieving reliable sentiment analysis results in the tourism domain. This would allow us to establish which method requires less annotated data to obtain competitive performance.

4.2. Sequence labeling – Locations and Fine-Grained Thematic Concept Extraction

The setup for sequence labeling tasks like named entity recognition for *locations* and *fine-grained thematic concept extraction* also aims to compare *fine-tuning* with *few-shot* learning approaches. However in this case: (i) we use *EntLM* [8] for few-shot in sequence labeling tasks; (ii) in addition to XLM-T, we include two more models, XLM-RoBERTa and multilingual BERT [9], the latter being the model used by *EntLM* in its default setting, and (iii) we also compare with the baseline established by the *word-matching* algorithm, as presented in Table 5. Grid search for hyperparameter fine-tuning was 8 for batch-size, 5-e5 learning rate and 0.1 weight decay.

The primary objective of this comparison is to ascertain the minimum amount of annotated data required to justify transitioning from costly and manual word-matching approaches to more advanced deep learning methods for each task, respectively. More specifically, we seek to determine the tipping point at which the benefits of employing deep learning techniques outweigh their data requirements. This is particularly true for a highly complex task such as *fine-grained thematic concept extraction* which involves labeling 315

different concept classes, and for which developing *word-matching* algorithms or manually annotating data are highly inefficient and expensive methods.

Experiments were conducted on servers equipped with Tesla V100 (32GB VRAM) GPUs, Intel Xeon Gold 6226R CPUs (2.90 GHz) and 256 GB of RAM. The language models were used through the HuggingFace Transformers API [46].

As it is customary, for sentiment analysis, we report accuracy results, while for sequence labeling we use the usual F1 micrometric calculated at the span level as defined in the CoNLL 2002 shared task [47]. All reported results are the average of three randomly initialized runs.

5. Results

5.1. Sentiment Analysis

The results of the sentiment analysis are reported in Figure 3, consisting of annotating each tweet as *positive*, *negative* or *neutral*. The x-axis represents the number of training examples used (*in this case, of tweets*), while the y-axis indicates the accuracy scores achieved using two different techniques: fine-tuning (F-T) and Pattern-Exploiting Training (PET, prompt-based) using the two language models, namely, XLM-T and XLM-T SENTIMENT.




Figure 3: Sentiment Analysis with k-shot sampling

The most noteworthy aspect from the results is that fine-tuning XLM-T SENTIMENT clearly outperforms any other approach in a few-shot setting. In contrast to previous work [22], this suggests that fine-tuning on a large multilingual dataset for sentiment analysis, even with texts from different domains, dramatically helps improve the results in domain-specific tourist data, substantially improving the results in few-shot settings over prompting methods such as PET. In fact, the fine-tuned XLM-T SENTIMENT model reaches optimal results with just 10 examples and similar performance to using the full training data with just a 5-shot training.

Among the approaches that use only our domain-specific training data, PET consistently outperforms fine-tuning XLM-T until we reach 100 examples. This highlights the effectiveness of PET, which is able to achieve high accuracy even with a limited number of examples. However, PET still falls short of the performance achieved by fine-tuning XLM-T SENTIMENT, which benefits from being pre-trained on a large out-of-domain corpus of thousands of tweets.

Summarizing, the observed patterns provide two significant takeaways: (1) when using an already fine-tuned sentiment model such as XLM-T SENTIMENT, **a dataset containing as few as 10 examples is sufficient to achieve reliable sentiment analysis performance in the tourism domain**, adding more examples does not seem to improve accuracy significantly and (2) when employing a base masked language model (MLM) such as XLM-T, **PET appears to be a preferable choice for few-shot scenarios**, given that close to optimal performance can be reached with 50 examples.

5.2. Named Entity Recognition (NER) for Locations

Figures 4 and 5 show the performance of Named Entity Recognition (NER) for Locations with the two sampling methods (*k-shot* and *percentage*). The fine-tuned versions of three models: XLM-T, XLM-RoBERTa (XLM-R), and mBERT were compared to EntLM, a few-shot method for sequence labeling tasks using multilingual BERT as backbone.




Figure 4: NER for Locations – k-shot Sampling




Figure 5: NER for Locations – Percentage Sampling

If we compare the results using the two different samplings (Figures 4 and 5), it can be observed that when using the full training data, all four methods obtain comparable results. However, in a few-shot setting, the EntLM method requires fewer labeled data. In other words, the fine-tuned multilingual BERT does not start to outperform EntLM until 30% of the training data is used. Overall, fine-tuning is only competitive when using percentage sampling, which we hypothesized might be due to the low amount of O tokens generated by the k-shot sampling. In contrast, EntLM behaves rather robustly when using a smaller number of examples with both sampling methods.

In our case, an interesting finding of our work is that both fine-tuned multilingual BERT and EntLM are clearly superior to the *word-matching* algorithm when using only approximately 13% of the tweets for training (around 200 tweets). Intriguingly, the models not adapted to Twitter (multilingual BERT and XLM-RoBERTa) demonstrate better performance than those specifically trained on Twitter data.

5.3. Fine-Grained Thematic Concept Extraction

Perhaps it is in the evaluation of *fine-grained thematic concept extraction*, shown in Figures 6 and 7, where the few-shot learning approach for sequence labeling clearly makes its mark. For a task which involves detecting and classifying sequences into a predetermined inventory of 315 classes, EntLM performs very competitively. Thus, with just five examples per class (5-shot setting), it obtains a 0.760 F1 score, equalling the *word-matching* algorithm’s results with just a 50-shot training. These scores indicate a strong ability to accurately identify touristic concepts, as reflected by the high precision values spanning from 0.80 to 0.913.




Figure 6: Fine-grained Thematic Concept Extraction – k-shot Sampling

Although overall results with *word matching* were similar, EntLM was slightly superior in terms of recall while being slightly worse in precision. Still, EntLM’s performance shows great promise to avoid costly manual annotation effort or complex development of rule-based algorithms for domain-specific fine-grained sequence labeling tasks. EntLM’s results are perhaps

magnified by the very poor results obtained by the fine-tuned approaches in both data sampling scenarios, which indicates the difficulty of learning good sequence taggers for fine-grained tasks.




Figure 7: Fine-grained Thematic Concept Extraction – Percentage Sampling

Having conducted these experiments, we will now delve into the insights gained and discuss the key findings that emerged. Additionally, we will address the potential limitations and biases that might have affected the results.

6. Discussion and limitations

In our comparative analysis of few-shot learning and fine-tuning for three NLP tasks in the tourism domain, namely sentiment analysis, named entity recognition (NER) for locations, and fine-grained thematic concept extraction, we have gained valuable insights into the data requirements and performance of these techniques.




Figure 8: Number of Label Words for the Most Frequent Thematic Concepts Found in the Tweets

For **Sentiment Analysis**, our findings suggest that a model previously fine-tuned on a large out-of-domain dataset for the same downstream text classification task can outperform prompt-based approaches even in few-shot settings. If such extra data is not available for our target task, then addressing the task by means of few-shot learning techniques (such as PET) has demonstrated to be the best approach to avoid costly manual annotation work.

Regarding **Named Entity Recognition (NER) for Locations**, it is a sequence labeling task with a single class but many label words associated with it (995 different location names). This results in low representativeness of label words which does not benefit EntLM’s approach. Still, for both sampling methods EntLM was competitive, improving over the results of *word matching* using from 10% of the data upwards (around 450 annotated tweets). In fact, overall results demonstrate the superiority of percentage sampling for few-shot learning in sequence labeling tasks.

We also explored the idea of improving our dataset by combining it with other existing corpora from other domains, such as AnCora [30] (*Spanish*) and ESTER [29] (*French*), both already annotated with location entities. However, this experiment did not lead to any significant improvements in the F1 score. Upon merging the three datasets, the F1-score saw only a minor increase, rising from 0.808 to 0.835 for XLM-T and from 0.821 to 0.830 for XLM-R. The limited improvement could be attributed to the fact that these corpora are not specifically designed for social media. ESTER consists of radio broadcast transcripts, while AnCora is comprised of newspaper texts. Consequently, they lack the contextual information pertinent to the tourism

domain.

In the case of **Fine-Grained Thematic Concept Extraction**, it is a different and more complex sequence labeling task which involves a large inventory of classes (315 concepts instantiated out of the 1,494 from the WTO tourism thesaurus), each having very few instance label words that are highly representative of the classes to which they refer. Figure 8 shows the low count of unique label words for the thematic concepts most often found in the tweets. As we expected, fine-tuning did not yield any satisfactory results regardless of the amount of data used or the sampling method employed. We believe this is due to the low number of examples per class. However, a major finding of this paper is that the EntLM few-shot learning technique demonstrated solid performance for this task, even when trained with a limited number of examples on a very small percentage of the available data. This robustness could be attributed to the model’s ability to generalize effectively from a smaller set of label words which are associated to each of the classes, namely, the *thematic concepts*. This result is particularly useful for practical applications where manually annotated data is usually very scarce.

Our experiments provide valuable insights but have some limitations. Firstly, we intentionally focused on three common NLP tasks within the tourism domain, ensuring in-depth understanding in this area. Although our findings are specific to these tasks, additional studies can broaden the applicability to other domains. Second, we worked with a curated dataset of 2,961 tweets. This limitation in dataset size was purposefully chosen to maintain focus, but larger and more diverse datasets might offer further perspectives. Lastly, while our dataset has a diverse language representation, it predominantly features French tweets, mirroring the tourist demographics of the French Basque Coast region.

7. Concluding Remarks and Future Work

In this paper we present a comparative study on few-shot learning for sentiment analysis, location detection and fine-grained thematic concept extraction for social media data in the tourism domain.

The main objective of this work is to establish the best strategy to obtain competitive results while keeping to a minimum any costly manual annotation and/or the development of complex and cumbersome rule-based approaches, especially for complex tasks such as *fine-grained thematic concept*

extraction.

In order to be able to do the experimentation, we provide a novel tourist-specific multilingual (French, English, and Spanish) dataset annotated for the tasks which are the subject of our study. This dataset, one of its kind, will be made publicly available in the future to facilitate further research on this particular topic and the reproducibility of the results.

Results show that current few-shot learning techniques allows to obtain competitive results for all three tasks with a very little amount of annotation: 5 tweets per label (15 in total) for Sentiment Analysis, 10% of the tweets for *location detection* (around 160 tweets) and 13% of the tweets annotated with *thematic concepts* (around 200 tweets). We believe that these findings are helpful not only for the development of our own application, but also for other domain-specific applications, which may require NLP analysis as a model enrichment. However, further research on analogous NLP tasks in different domains would be needed to validate the generalizability of our results in various domains, other than tourism. We believe these findings offer value not only for NLP researchers focusing on tourism but also for applications needing domain-specific NLP analysis, especially when there's a lack of annotated data or when one wishes to exclude ad hoc rule-based methods.

Our project's next step involves presenting the NLP results to stakeholders in the tourism industry. To achieve this, we have initiated the development of multidimensional, dynamic dashboards based on the output generated by the NLP processing modules discussed in this paper. These dashboards are designed to highlight the frequency, relationships, and trajectories of extracted entities, categorized spatially (cities, POIs), temporally (date, time periods), and thematically (touristic concepts), within a specified social media corpus, all correlated with contextual data such as sentiment and engagement levels.

References

- [1] B. Zeng, R. Gerritsen, What do we know about social media in tourism? a review, *Tourism management perspectives* 10 (2014) 27–36.
- [2] D. Maynard, K. Bontcheva, D. Rout, Challenges in developing opinion mining tools for social media, in: *Workshop Programme*, p. 15.

- [3] S. Rosenthal, P. Nakov, S. Kiritchenko, S. Mohammad, A. Ritter, V. Stoyanov, SemEval-2015 task 10: Sentiment analysis in Twitter, in: Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015), Association for Computational Linguistics, Denver, Colorado, 2015, pp. 451–463. doi:10.18653/v1/S15-2078. URL <https://aclanthology.org/S15-2078>
- [4] X. Liu, H. Chen, W. Xia, Overview of named entity recognition, Journal of Contemporary Educational Research 6 (5) (2022) 65–68.
- [5] S. Fu, D. Chen, H. He, S. Liu, S. Moon, K. J. Peterson, F. Shen, L. Wang, Y. Wang, A. Wen, et al., Clinical concept extraction: a methodology review, Journal of biomedical informatics 109 (2020) 103526.
- [6] B. Min, H. Ross, E. Sulem, A. P. B. Veyseh, T. H. Nguyen, O. Sainz, E. Agirre, I. Heinz, D. Roth, Recent advances in natural language processing via large pre-trained language models: A survey, arXiv preprint arXiv:2111.01243 (2021).
- [7] World Tourism Organization, Thesaurus on tourism and leisure activities 2001, World Tourism Organization, Madrid, Spain, 2002.
- [8] R. Ma, X. Zhou, T. Gui, Y. Tan, L. Li, Q. Zhang, X. Huang, Template-free prompt tuning for few-shot NER, in: Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, 2022, pp. 5721–5732. doi:10.18653/v1/2022.naacl-main.420.
- [9] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional transformers for language understanding, CoRR abs/1810.04805 (2018). arXiv:1810.04805. URL <http://arxiv.org/abs/1810.04805>
- [10] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, V. Stoyanov, Unsupervised cross-lingual representation learning at scale, CoRR abs/1911.02116 (2019). arXiv:1911.02116. URL <http://arxiv.org/abs/1911.02116>

- [11] C. D. Manning, K. Clark, J. Hewitt, U. Khandelwal, O. Levy, Emergent linguistic structure in artificial neural networks trained by self-supervision, *Proceedings of the National Academy of Sciences* 117 (2020) 30046–30054. doi:<https://doi.org/10.1073/pnas.1907367117>.
- [12] O. Toporkov, R. Agerri, On the role of morphological information for contextual lemmatization, *arXiv preprint 2302.00407* (2023).
- [13] M. Crawford, T. M. Khoshgoftaar, Using inductive transfer learning to improve hotel review spam detection, in: *2021 IEEE 22nd International Conference on Information Reuse and Integration for Data Science (IRI)*, IEEE, 2021, pp. 248–254.
- [14] M. P. Enríquez, J. A. Mencía, I. Segura-Bedmar, Transformers approach for sentiment analysis: Classification of mexican tourists reviews from tripadvisor (2022).
- [15] J. Vásquez, H. Gómez-Adorno, G. Bel-Enguix, Bert-based approach for sentiment analysis of spanish reviews from tripadvisor., in: *IberLEF@SEPLN*, 2021, pp. 165–170.
- [16] A. Serna, A. Soroa, R. Agerri, Applying deep learning techniques for sentiment analysis to assess sustainable transport, *Sustainability* 13 (4) (2021). doi:[10.3390/su13042397](https://doi.org/10.3390/su13042397). URL <https://www.mdpi.com/2071-1050/13/4/2397>
- [17] I. Bouabdallaoui, F. Guerouate, S. Bouhaddour, C. Saadi, M. Sbihi, Named entity recognition applied on moroccan tourism corpus, *Procedia Computer Science* 198 (2022) 373–378.
- [18] X. Cheng, W. Wang, F. Bao, G. Gao, Mtner: A corpus for mongolian tourism named entity recognition, in: *Machine Translation: 16th China Conference, CCMT 2020, Hohhot, China, October 10-12, 2020, Revised Selected Papers 16*, Springer, 2020, pp. 11–23.
- [19] C. Chantrapornchai, A. Tunsakul, Information extraction on tourism domain using spacy and bert, *ECTI Transactions on Computer and Information Technology* 15 (1) (2021) 108–122.

- [20] J. F. de Landa, R. Agerri, Social analysis of young basque-speaking communities in twitter, *Journal of Multilingual and Multicultural Development* 0 (2021) 1–15.
- [21] S. Kadam, V. Vaidya, Review and analysis of zero, one and few shot learning approaches, in: *Intelligent Systems Design and Applications: 18th International Conference on Intelligent Systems Design and Applications (ISDA 2018) held in Vellore, India, December 6-8, 2018, Volume 1*, Springer, 2020, pp. 100–112.
- [22] T. Schick, H. Schütze, Exploiting cloze questions for few shot text classification and natural language inference, *arXiv preprint arXiv:2001.07676* (2020).
- [23] L. Wang, R. Li, Y. Yan, Y. Yan, S. Wang, W. Wu, W. Xu, Instructionner: A multi-task instruction-based generative framework for few-shot ner, *arXiv preprint 2203.03903* (2022).
- [24] I. García-Ferrero, R. Agerri, G. Rigau, Model and data transfer for cross-lingual sequence labelling in zero-resource settings, in: *Findings of the Association for Computational Linguistics: EMNLP 2022*, Association for Computational Linguistics, 2022, pp. 6403–6416.
- [25] I. Garcia-Ferrero, R. Agerri, G. Rigau, T-projection: High quality annotation projection for sequence labeling tasks, *arXiv preprint 2212.10548* (2022).
- [26] M. Artetxe, G. Labaka, E. Agirre, Translation artifacts in cross-lingual transfer learning, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020, pp. 7674–7684.
- [27] M. Artetxe, V. Goswami, S. Bhosale, A. Fan, L. Zettlemoyer, Revisiting machine translation for cross-lingual classification, *arXiv preprint 2305.14240* (2023).
- [28] V. Barriere, A. Balahur, Improving sentiment analysis over non-English tweets using multilingual transformers and automatic translation for

- data-augmentation, in: Proceedings of the 28th International Conference on Computational Linguistics, International Committee on Computational Linguistics, 2020, pp. 266–271. doi:10.18653/v1/2020.coling-main.23.
- [29] S. Galliano, E. Geoffrois, G. Gravier, J.-F. Bonastre, D. Mostefa, K. Choukri, Corpus description of the ester evaluation campaign for the rich transcription of french broadcast news., in: LREC, Citeseer, 2006, pp. 139–142.
- [30] M. Taulé, M. A. Martí, M. Recasens, Ancora: Multilevel annotated corpora for catalan and spanish., in: Lrec, 2008.
- [31] L. Derczynski, K. Bontcheva, I. Roberts, Broad Twitter corpus: A diverse named entity recognition resource, in: Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers, Osaka, Japan, 2016, pp. 1169–1179.
- [32] A. Go, R. Bhayani, L. Huang, Twitter sentiment classification using distant supervision, CS224N project report, Stanford 1 (12) (2009) 2009.
- [33] H. Saif, M. Fernandez, Y. He, H. Alani, Evaluation datasets for twitter sentiment analysis: a survey and a new dataset, the sts-gold (2013).
- [34] P. Nakov, S. Rosenthal, Z. Kozareva, V. Stoyanov, A. Ritter, T. Wilson, 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), Association for Computational Linguistics, Atlanta, Georgia, USA, 2013, pp. 312–320.
URL <https://aclanthology.org/S13-2052>
- [35] S. Rosenthal, A. Ritter, P. Nakov, V. Stoyanov, SemEval-2014 task 9: Sentiment analysis in Twitter, in: Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014), Association for Computational Linguistics, Dublin, Ireland, 2014, pp. 73–80. doi:10.3115/v1/S14-2009.
URL <https://aclanthology.org/S14-2009>
- [36] P. Nakov, A. Ritter, S. Rosenthal, F. Sebastiani, V. Stoyanov, SemEval-2016 task 4: Sentiment analysis in Twitter, in: Proceedings of the 10th

- International Workshop on Semantic Evaluation (SemEval-2016), Association for Computational Linguistics, San Diego, California, 2016, pp. 1–18. doi:10.18653/v1/S16-1001.
URL <https://aclanthology.org/S16-1001>
- [37] S. Rosenthal, N. Farra, P. Nakov, SemEval-2017 task 4: Sentiment analysis in Twitter, in: Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017), Association for Computational Linguistics, Vancouver, Canada, 2017, pp. 502–518. doi:10.18653/v1/S17-2088.
URL <https://aclanthology.org/S17-2088>
- [38] P. Budzianowski, T.-H. Wen, B.-H. Tseng, I. Casanueva, S. Ultes, O. Ramadan, M. Gašić, MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling, in: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 5016–5026. doi:10.18653/v1/D18-1547.
URL <https://aclanthology.org/D18-1547>
- [39] S. R. Bowman, G. Angeli, C. Potts, C. D. Manning, A large annotated corpus for learning natural language inference, arXiv preprint arXiv:1508.05326 (2015).
- [40] M. Masson, C. Sallaberry, R. Agerri, M.-N. Bessagnet, P. Roose, A. Le Parc Lacayrelle, A domain-independent method for thematic dataset building from social media: The case of tourism on twitter, in: Web Information Systems Engineering–WISE 2022: 23rd International Conference, Biarritz, France, November 1–3, 2022, Proceedings, Springer, 2022, pp. 11–20.
- [41] F. Barbieri, J. Camacho-Collados, L. Neves, L. Espinosa-Anke, Tweet-eval: Unified benchmark and comparative evaluation for tweet classification, arXiv preprint arXiv:2010.12421 (2020).
- [42] J. M. Pérez, J. C. Giudici, F. Luque, pysentimiento: A python toolkit for sentiment analysis and socialnlp tasks (2021). arXiv:2106.09462.
- [43] Seethal, Sentiment analysis generic dataset, https://huggingface.co/Seethal/sentiment_analysis_generic_dataset, accessed on 23th March 2023.

- [44] J. Hartmann, M. Heitmann, C. Siebert, C. Schamp, More than a feeling: Accuracy and application of sentiment analysis, *International Journal of Research in Marketing* 40 (1) (2023) 75–87. doi:<https://doi.org/10.1016/j.ijresmar.2022.05.005>.
- [45] F. Barbieri, L. Espinosa Anke, J. Camacho-Collados, XLM-T: Multilingual language models in Twitter for sentiment analysis and beyond, in: *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, European Language Resources Association, Marseille, France, 2022, pp. 258–266.
URL <https://aclanthology.org/2022.lrec-1.27>
- [46] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, A. Rush, Transformers: State-of-the-art natural language processing, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, Association for Computational Linguistics, 2020, pp. 38–45. doi:10.18653/v1/2020.emnlp-demos.6.
- [47] E. F. Tjong Kim Sang, Introduction to the CoNLL-2002 shared task: Language-independent named entity recognition, in: *COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002)*, 2002.