

UNIVERSIDADE FEDERAL FLUMINENSE

SÉRGIO DE MELO BARRETO JUNIOR

SENTIMENT ANALYSIS IN TWEETS:
AN ASSESSMENT STUDY FROM
CLASSICAL TO MODERN TEXT
REPRESENTATION MODELS

NITERÓI

2021

UNIVERSIDADE FEDERAL FLUMINENSE

SÉRGIO DE MELO BARRETO JUNIOR

**SENTIMENT ANALYSIS IN TWEETS:
AN ASSESSMENT STUDY FROM
CLASSICAL TO MODERN TEXT
REPRESENTATION MODELS**

Dissertação de Mestrado apresentada ao Programa de Pós-Graduação em Computação da Universidade Federal Fluminense como requisito parcial para a obtenção do Grau de Mestre em Computação. Área de concentração: Ciência da Computação

Orientador:

ALEXANDRE PLASTINO DE CARVALHO

Coorientadora:

ALINE MARINS PAES CARVALHO

NITERÓI

2021

Ficha catalográfica automática - SDC/BEE
Gerada com informações fornecidas pelo autor

B273s Barreto junior, Sérgio de Melo
Sentiment analysis in tweets: An assessment study from classical to modern text representation models / Sérgio de Melo Barreto junior ; ALEXANDRE PLASTINO DE CARVALHO, orientador ; ALINE MARINS PAES CARVALHO, coorientadora. Niterói, 2021.
133 f.

Dissertação (mestrado)-Universidade Federal Fluminense, Niterói, 2021.

DOI: <http://dx.doi.org/10.22409/PGC.2021.m.13122652722>

1. Sentiment analysis. 2. Text representations. 3. Language models. 4. Natural language processing. 5. Produção intelectual. I. CARVALHO, ALEXANDRE PLASTINO DE, orientador. II. CARVALHO, ALINE MARINS PAES, coorientadora. III. Universidade Federal Fluminense. Instituto de Computação. IV. Título.

CDD -

SÉRGIO DE MELO BARRETO JUNIOR

SENTIMENT ANALYSIS IN TWEETS:
AN ASSESSMENT STUDY FROM CLASSICAL TO MODERN TEXT
REPRESENTATION MODELS

Dissertação de Mestrado apresentada ao Programa de Pós-Graduação em Computação da Universidade Federal Fluminense como requisito parcial para a obtenção do Grau de Mestre em Computação. Área de concentração: Ciência da Computação

Aprovada em Julho de 2021.

BANCA EXAMINADORA



Prof. D.Sc. ALEXANDRE PLASTINO DE CARVALHO -
Orientador, PGC/UFF



Profª. D.Sc. ALINE MARINS PAES CARVALHO -
Coorientadora, PGC/UFF



Profª. D.Sc. VIVIANE PEREIRA MOREIRA, UFRGS



Prof. D.Sc. JOSÉ VITERBO FILHO, PGC/UFF

Niterói

2021

Acknowledgments

The author would like to thank the Brazilian research agencies FAPERJ and CNPq for their financial support. The author would also like to thank his family and his fiancée for all the support that made the whole journey possible.

Resumo

Com o crescimento das mídias sociais, como o Twitter, muitos dados gerados pelos usuários surgem diariamente. Os textos curtos publicados no Twitter – os tweets – têm ganhado atenção significativa como uma fonte rica de informações para auxiliar na tomada de decisão. No entanto, as características inerentes dos tweets, como o tamanho reduzido e o estilo linguístico informal e ruidoso, continuam sendo um desafio para muitas tarefas de processamento de linguagem natural (PLN), incluindo a análise de sentimentos. A classificação de sentimentos em textos é abordada principalmente por classificadores baseados em aprendizado de máquina. Para transformar tweets em entradas baseadas em vetores para alimentar classificadores de sentimento, a literatura tem adotado representações de palavras de naturezas distintas. A geração das representações vetoriais são obtidas desde por métodos simples baseados em contagem, como bag-of-words, até por métodos mais sofisticados que geram embeddings, como o BERTweet, construído sobre a notável arquitetura BERT. No entanto, a maioria dos estudos concentra-se principalmente na avaliação desses modelos de representação usando apenas um número reduzido de bases de dados. Apesar do progresso feito nos últimos anos na construção de modelos de linguagem, ainda há uma lacuna em relação a uma avaliação robusta de embeddings aplicados à análise de sentimento em tweets. Além disso, enquanto o ajuste do modelo visando a tarefa final tem se tornado popular, pouca atenção tem sido dada aos ajustes com base no estilo linguístico específico dos textos. Neste contexto, este estudo realiza uma avaliação dos modelos de linguagem existentes para a análise de sentimento expresso em tweets, usando uma rica coleção de 22 bases de dados de domínios distintos e cinco algoritmos de classificação. A avaliação inclui representações vetoriais estáticas e contextualizadas, totalizando 280 diferentes combinações de estratégias para extrair sentimentos de textos. Os contextos são identificados a partir de modelos de autoencoder baseados em Transformer, que também são ajustados com base na tarefa intermediária de modelo de linguagem mascarada, usando uma vasta quantidade de estratégias.

Palavras-chave: análise de sentimento, representação de texto, modelos de linguagem, processamento de linguagem natural, Twitter

Abstract

With the growth of social medias, such as Twitter, plenty of user-generated data emerge daily. The short texts published on Twitter – the tweets – have earned significant attention as a rich source of information to guide many decision-making processes. However, their inherent characteristics, such as the informal and noisy linguistic style, remain challenging to many natural language processing (NLP) tasks, including sentiment analysis. Sentiment classification is tackled mainly by machine learning-based classifiers. The literature has adopted word representations from distinct natures to transform tweets to vector-based inputs to feed sentiment classifiers. The representations come from simple count-based methods, such as bag-of-words, to more sophisticated ones, such as BERTweet, built upon the trendy BERT architecture. Nevertheless, most studies mainly focus on evaluating those models using only a small number of datasets. Despite the progress made in recent years in language modelling, there is still a gap regarding a robust evaluation of induced embeddings applied to sentiment analysis on tweets. Furthermore, while fine-tuning the model from downstream tasks is prominent nowadays, less attention has been given to adjustments based on the specific linguistic style of the data. In this context, this study fulfils an assessment of existing language models in distinguishing the sentiment expressed in tweets by using a rich collection of 22 datasets from distinct domains and five classification algorithms. The evaluation includes static and contextualized representations, assessing in total 280 different combinations of strategies for extracting sentiments from text. Contexts are assembled from Transformer-based autoencoder models that are also fine-tuned based on the masked language model task, using a plethora of strategies.

Keywords: sentiment analysis, text representations, language models, natural language processing, Twitter

List of Figures

2.1	BERT model architecture [14]	11
-----	------------------------------	----

List of Tables

3.1	Characteristics of the Twitter sentiment datasets ordered by size (Total column)	13
3.2	Summary of the advantages and disadvantages of the classification algorithms adopted in this study	14
3.3	Characteristics of the static pretrained embeddings ordered by the number of dimensions	16
4.1	Accuracies (%) achieved using the SVM classifier	18
4.2	F_1 -macro scores (%) achieved by evaluating static representation using the SVM classifier	19
4.3	Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by evaluating each static representation on the 22 datasets, for each classification algorithm, in terms of accuracy	20
4.4	Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by evaluating each static representation on the 22 datasets, for each classification algorithm, in terms of F_1 -macro	20
4.5	Top 10 average rank results achieved for each combination Model-Classifier	23
4.6	Tail 10 average rank results achieved for each combination Model-Classifier	23
4.7	Average Rank results achieved for each Embedding, evaluating static representations	24
4.8	Average Rank results achieved for each Classifier, evaluating static representations	24
4.9	Best results achieved for each dataset	25
5.1	Accuracies and F_1 -macro scores (%) achieved by evaluating Transformer-Autoencoder language models using the SVM classifier	27

5.2	Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by evaluating each Transformer-Autoencoder model on the 22 datasets, for each classification algorithm	28
5.3	Average rank position results achieved for each combination Model-Classifier	29
5.4	Average rank position results achieved for each Embedding	29
5.5	Average rank position results achieved for each Classifier	29
5.6	Best results achieved for each dataset by evaluating the combination of language model and classifier	30
5.7	Percentage of vocabulary's tokens of language model in the line that is also in the vocabulary's tokens of language model in the column.	31
5.8	Top 10 average rank position results achieved for each combination Model-Classifier by evaluating Transformer-Autoencoder model and static embeddings	33
5.9	Average rank position results achieved for each Embedding evaluating Transformer-Autoencoder model and static embedding	33
5.10	Best results achieved for each dataset by evaluating Transformer-Autoencoder and static models	34
6.1	Average classification accuracies (%) achieved by fine-tuning BERT with different samples of generic unlabeled tweets, using the SVM classifier . . .	37
6.2	Average classification accuracies (%) achieved by fine-tuning RoBERTa with different samples of generic unlabeled tweets, using the SVM classifier	38
6.3	Average classification accuracies (%) achieved by fine-tuning BERTweet with different samples of generic unlabeled tweets, using the SVM classifier	39
6.4	Average F_1 -macro scores (%) achieved by fine-tuning BERT with different samples of unlabeled tweets, using the SVM classifier	40
6.5	Average F_1 -macro scores (%) achieved by fine-tuning RoBERTa with different samples of unlabeled tweets, using the SVM classifier	41
6.6	Average F_1 -macro scores (%) achieved by fine-tuning BERTweet with different samples of unlabeled tweets, using the SVM classifier	42

6.7	Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by each classifier when fine-tuning the Transformer-Autoencoder models with different samples of unlabeled tweets in terms of accuracy	43
6.8	Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by each classifier when fine-tuning the Transformer-Autoencoder models with different samples of unlabeled tweets in terms of F_1 -macro	44
6.9	Best results achieved for each dataset by fine-tuning the Transformer-based models with different samples of generic tweets	45
6.10	Top 10 average rank position results achieved for each combination Model-Classifer by evaluating Transformer-Autoencoder models	46
6.11	Top 10 average rank position results achieved for each Embedding evaluating Transformer-Autoencoder model	47
6.12	Average rank position results achieved for each Classifier evaluating Transformer-Autoencoder model	47
6.13	Average rank position results achieved for BERT model and its tuned models	49
6.14	Average rank position results achieved for RoBERTa model and its tuned models	49
6.15	Average rank position results achieved for BERTweet model and its tuned models	50
7.1	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the SVM classifier and BERT . . .	53
7.2	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the SVM classifier and RoBERTa	54
7.3	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the SVM classifier and BERTweet	55
7.4	Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by each classifier when fine-tuning the Transformer-Autoencoder models using strategies InData, LOO, and AllData in terms of accuracy and F_1 -macro	56

7.5	Best results achieved for each dataset by fine-tuning the Transformer-Autoencoder models using strategies InData, LOO, and AllData	57
7.6	Top 10 average rank position results achieved for each combination Model-Classifier by evaluating Transformer-Autoencoder model	58
7.7	Average position results achieved for each Embedding	59
7.8	Average rank position results achieved for each Classifier	59
7.9	Average rank position results achieved for BERT models	60
7.10	Average rank position results achieved for RoBERTa models	60
7.11	Average rank position results achieved for BERTweet models	61
7.12	Top 10 average rank position results achieved for each combination Model-Classifier by evaluating all assessed model in this study	62
7.13	Tail 10 average rank position results achieved for each combination Model-Classifier by evaluating all assessed model in this study	63
7.14	Top 10 average rank position results achieved comparing all assessed models in this study	63
7.15	Tail 10 average rank position results achieved comparing all assessed models in this study	64
A.1	Accuracies (%) achieved using the RF classifier	73
A.2	F_1 -macro scores (%) achieved by evaluating static representation using the RF classifier	74
A.3	Accuracies (%) achieved using the MLP classifier	74
A.4	F_1 -macro scores (%) achieved by evaluating static representation using the MLP classifier	75
A.5	Accuracies (%) achieved using the LR classifier	75
A.6	F_1 -macro scores (%) achieved by evaluating static representation using the LR classifier	76
A.7	Accuracies (%) achieved using the XGB classifier	76
A.8	F_1 -macro scores (%) achieved by evaluating static representation using the XGB classifier	77

A.9	Accuracies and F_1 -macro scores (%) achieved by evaluating Transformer-Autoencoder language models using the RF classifier	78
A.10	Accuracies and F_1 -macro scores (%) achieved by evaluating Transformer-Autoencoder language models using the MLP classifier	79
A.11	Accuracies and F_1 -macro scores (%) achieved by evaluating Transformer-Autoencoder language models using the LR classifier	80
A.12	Accuracies and F_1 -macro scores (%) achieved by evaluating Transformer-Autoencoder language models using the XGB classifier	81
A.13	Average classification accuracies (%) achieved by fine-tuning BERT with different samples of generic unlabeled tweets, using the RF classifier	82
A.14	Average F_1 -macro scores (%) achieved by fine-tuning BERT with different samples of unlabeled tweets, using the RF classifier	83
A.15	Average classification accuracies (%) achieved by fine-tuning BERT with different samples of generic unlabeled tweets, using the LR classifier	84
A.16	Average F_1 -macro scores (%) achieved by fine-tuning BERT with different samples of unlabeled tweets, using the LR classifier	85
A.17	Average classification accuracies (%) achieved by fine-tuning BERT with different samples of generic unlabeled tweets, using the MLP classifier . . .	86
A.18	Average F_1 -macro scores (%) achieved by fine-tuning BERT with different samples of unlabeled tweets, using the MLP classifier	87
A.19	Average classification accuracies (%) achieved by fine-tuning BERT with different samples of generic unlabeled tweets, using the XGB classifier . . .	88
A.20	Average F_1 -macro scores (%) achieved by fine-tuning BERT with different samples of unlabeled tweets, using the XGB classifier	89
A.21	Average classification accuracies (%) achieved by fine-tuning RoBERTa with different samples of generic unlabeled tweets, using the RF classifier .	90
A.22	Average F_1 -macro scores (%) achieved by fine-tuning RoBERTa with different samples of unlabeled tweets, using the RF classifier	91
A.23	Average classification accuracies (%) achieved by fine-tuning RoBERTa with different samples of generic unlabeled tweets, using the LR classifier .	92

A.24 Average F_1 -macro scores (%) achieved by fine-tuning RoBERTa with different samples of unlabeled tweets, using the LR classifier	93
A.25 Average classification accuracies (%) achieved by fine-tuning RoBERTa with different samples of generic unlabeled tweets, using the MLP classifier	94
A.26 Average F_1 -macro scores (%) achieved by fine-tuning RoBERTa with different samples of unlabeled tweets, using the MLP classifier	95
A.27 Average classification accuracies (%) achieved by fine-tuning RoBERTa with different samples of generic unlabeled tweets, using the XGB classifier	96
A.28 Average F_1 -macro scores (%) achieved by fine-tuning RoBERTa with different samples of unlabeled tweets, using the XGB classifier	97
A.29 Average classification accuracies (%) achieved by fine-tuning BERTweet with different samples of generic unlabeled tweets, using the RF classifier	98
A.30 Average F_1 -macro scores (%) achieved by fine-tuning BERTweet with different samples of unlabeled tweets, using the RF classifier	99
A.31 Average classification accuracies (%) achieved by fine-tuning BERTweet with different samples of generic unlabeled tweets, using the LR classifier	100
A.32 Average F_1 -macro scores (%) achieved by fine-tuning BERTweet with different samples of unlabeled tweets, using the LR classifier	101
A.33 Average classification accuracies (%) achieved by fine-tuning BERTweet with different samples of generic unlabeled tweets, using the MLP classifier	102
A.34 Average F_1 -macro scores (%) achieved by fine-tuning BERTweet with different samples of unlabeled tweets, using the MLP classifier	103
A.35 Average classification accuracies (%) achieved by fine-tuning BERTweet with different samples of generic unlabeled tweets, using the XGB classifier	104
A.36 Average F_1 -macro scores (%) achieved by fine-tuning BERTweet with different samples of unlabeled tweets, using the XGB classifier	105
A.37 Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the LR classifier and BERT	106

A.38	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the RF classifier and BERT	107
A.39	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the MLP classifier and BERT	108
A.40	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the XGB classifier and BERT	109
A.41	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the LR classifier and RoBERTa	110
A.42	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the RF classifier and RoBERTa	111
A.43	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the MLP classifier and RoBERTa	112
A.44	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the XGB classifier and RoBERTa	113
A.45	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the LR classifier and BERTweet	114
A.46	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the RF classifier and BERTweet	115
A.47	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the MLP classifier and BERTweet	116
A.48	Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the XGB classifier and BERTweet	117

List of Abbreviations and Acronyms

NLP	: Natural Language Processing;
BERT	: Bidirectional Encoder Representations from Transformers;
RoBERTa	: Robustly Optimized BERT;
BERTweet	: Bidirectional Encoder Representations from Transformers trained with Tweets;
VSM	: Vector space model;
BoW	: Bag of Word;
TF-IDF	: Term Frequency-Inverse Document Frequency;
CBOW	: Continuous Bag-Of-Word;
OOV	: Out-of-vocabulary;
ULMFit	: Universal Language Model Fine-tuning;
STLR	: Slanted triangular learning rates;
RNN	: Recurrent Neural Network;
GPT	: Generative Pre-Training Transformer;
LSTM	: Long-Short-Term-Memory Recurrent Neural Network;
BiLSTM	: Bi-directional Long-Short-Term-Memory Recurrent Neural Network;
ELMO	: Embeddings from Language Models;
SVM	: Support Vector Machine;
XGB	: XGBoost;
MLP	: Multi-layer Perceptron;
RF	: Random Forest ;
LR	: Logistic Regression;

Contents

1	Introduction	1
2	Literature Review	6
3	Experimental Methodology	12
3.1	Datasets	12
3.2	Experimental Protocol	13
3.3	Computational Experiments Details	15
4	Evaluation of Static Text Representations	17
5	Evaluation of the Transformer-based Text Representations	26
6	Fine-tuning Transformer-based Models Using a Large Collection of English Tweets	35
7	Fine-tuning Transformer-based Models Using Sentiment Datasets	51
8	Conclusions and Future Works	65
	References	68
	Appendix A – Extra Experimental Results	73
A.1	Static Text Representation	73
A.2	Transformer-based Text Representation	78
A.3	Transformer-based Text Representation tuned With a Large Collection of English Tweets	82

A.4 Transformer-based Text Representation Tuned With Sentiment Datasets	106
---	-----

Chapter 1

Introduction

In recent years, the use of social media networks, such as Twitter¹, has been growing exponentially. It is estimated that about 500 million tweets – the short informal messages sent by Twitter users – are published daily². Unlike other text styles, tweets have an informal linguistic style, misspelled words, careless use of grammar, URL links, user mentions, hashtags, and more. Due to these inherent characteristics, discovering patterns from tweets represents a challenge and opportunities for machine learning and natural language processing (NLP) tasks, such as sentiment analysis.

Sentiment analysis is the field of study that analyzes people’s opinions, sentiments, appraisals, attitudes, and emotions toward entities and their attributes expressed in written text [30]. Usually, one reduces the sentiment analysis task to find out the polarity classification, i.e., whether they carry a positive or negative connotation. One of the biggest challenges concerning the sentiment classification of tweets is that people often express their sentiments and opinions using a casual linguistic style, resulting in the presence of misspelled words and the careless use of grammar. Consequently, the automated analysis of tweets’ content requires machines to build a deep understanding of natural text to deal effectively with its informal structure [3]. However, before discovering patterns from text, it is essential to define a more fundamental step: how automatic methods can numerically represent textual content.

Vector space models (VSMs) [21] are one of the earliest and most common strategies adopted in text classification literature to allow machines to deal with texts and their structures. The VSM represents each document in a corpus as a point in a vector space. Points that are close together in this space are semantically similar, and points that are far

¹<http://www.twitter.com>

²<https://www.dsayce.com/social-media/tweets-day/>

apart are semantically distant [55]. The firsts VSM approaches are count-based methods, such as Bag-of-Words (BoW) [33] and BoW with TF-IDF [33]. Although VSMs have been extensively used in the literature, they cannot deal with the curse of dimensionality. More clearly, considering the inherent characteristics of tweets, a corpus of tweets may contain different spellings for each unique word leading to an extensive vocabulary, making the vector representation of those tweets very large and often sparse.

To tackle the curse of dimensionality inherent from BOW-based approaches, in the last years it has become a standard practice to learn dense vectors to represent words and texts, the *embeddings*. Methods such as Word2Vec [53], Fast-Text [54], and others [2, 20, 51, 58] have been used with relative success to address a plethora of NLP tasks. Nevertheless, in general, the performance of such techniques is still unsatisfactory to solve sentiment analysis from tweets, taking into account the dynamic vocabulary used by Twitter users to express themselves. Specifically, in tweets, the ironic and sarcastic content expressed in a limited space, regularly out of context and informal, makes even more challenging to retrieve meaning from the words. Such attributes may degrade the performance of traditional word embeddings methods if not handled properly. In this context, contextualized word representations have recently emerged in the literature, aiming at allowing the vector representation of words to adapt to the context they appear. Contextual embedding techniques, including ELMo [42] and Transformer-based autoencoder methods, such as BERT [14], RoBERTa [31], and BERTweet [39], capture not only complex characteristics of word usage, such as syntax and semantics, but also how the word usage vary across linguistic contexts. Those methods have achieved state-of-the-art results on various NLP tasks, including sentiment analysis [4, 12, 22, 1].

Much effort in recent language modeling research is focused on scalability issues of existing word embedding methods. On this basis, inductive transfer learning strategies and pre-trained embedding models have gained important application in the literature, especially when the amount of labeled data to train a model is relatively small. With that, models obtained from the aforementioned contextual embeddings methods are rarely trained from scratch but are instead fine-tuned from models pre-trained on datasets with a huge amount of texts [27, 42, 45]. Pre-trained models reduce the use of computational resources and tend to increase the classification performance of several NLP tasks, sentiment analysis included.

Despite the successful achievements in developing efficient word representation methods in NLP literature, there is still a gap regarding a robust evaluation of existing lan-

guage models applied to the sentiment analysis task on tweets. Most studies are mainly focused on evaluating those models for different NLP tasks using only a small number of datasets [43, 28, 31, 26, 58]. Our main goal is to identify appropriate embeddings-based text representations for the sentiment analysis of English tweets in this study. For this purpose, we evaluate embeddings of different natures, including: i) static embeddings learned from generic texts [2, 34, 35, 41]; ii) static embeddings learned from datasets of Twitter sentiment analysis [5, 8, 20, 41, 51, 58]; iii) contextualized embeddings learned from transformer-based autoencoders with generic texts with no adjustments [14, 31]; iv) contextualized embeddings learned from Transformer-based autoencoders with a dataset of tweets with no adjustments [39]; v) contextualized embeddings fine-tuned to the tweets language; and vi) contextualized embeddings fine-tuned to the tweets *sentiment* language. In all assessments, we use a representative set of twenty-two sentiment datasets [11] as input to five classifiers to evaluate the predictive performance of the embeddings. To the best of our knowledge, there is no previous study that has conducted such a robust evaluation regarding language models of several flavors and a large number of datasets, assessing 280 different strategies of extracting sentiment from short texts, such as tweets. In order to identify the most appropriate text embeddings, we conduct this study to answer the following four research questions.

RQ1. Which static embeddings are the most effective in the sentiment classification of tweets? Our motivation to evaluate those models is that many state-of-the-art deep learning models can require a lot of computational power, such as memory usage and storage. Thus, running those models locally on some devices may be difficult for mass-market applications that depend on low-cost hardware. To overcome this limitation, embeddings generated by language models can be gathered by simply looking up at the embedding table to achieve a static representation of textual content. We intend to assess how these static representations work and which are the most appropriate in this context. We answer this research question by evaluating a rich set of text representations from the literature [2, 5, 8, 14, 20, 34, 35, 39, 41, 51, 58, 59]. To achieve a good overview of the static representations, we conduct an experimental evaluation in the sentiment analysis task with five different classifiers and 22 datasets.

RQ2. Considering state-of-the-art Transformer-based autoencoder models, which are the most effective in the sentiment classification of tweets? Regarding recent advances in language modeling, Transformer-based architectures have achieved state-of-the-art performances in many NLP tasks. Specifically, BERT [14] is the first method that successfully uses the encoders components of the Transformer architecture [57] to learn contextualized

embeddings from texts. Shortly after that, RoBERTa [31] was introduced by Facebook as an extension of BERT that uses an optimized training methodology. Next, BERTweet [39] was proposed as an alternative to RoBERTa for NLP tasks focusing on tweets. While RoBERTa was trained on traditional English texts, such as Wikipedia, BERTweet was trained from scratch using a massive corpus of 850M English tweets. In this context, to answer this research question, we conduct an experimental evaluation of BERT, RoBERTa, and BERTweet models in the sentiment analysis task with five different classifiers and 22 datasets to obtain a comprehensive analysis of their predictive performances. By evaluating these models we may obtain a robust overview of the Transformer-based autoencoder representations that better fit tweet’s style.

RQ3. Can Transformer-based autoencoder models benefit from a fine-tuning procedure with a large set of English tweets? One of the benefits of pre-trained language models, such as the Transformer-based models exploited in this study, is the possibility to adjust the language model to a specific domain by applying a fine-tuning procedure. We aim at assessing whether the sentiment analysis of tweets can benefit from fine-tuning BERT, RoBERTa and BERTweet language models with a vast, generic, and unlabeled set of around 6.7M English tweets from distinct domains. To that, we fine-tuned the pre-trained language model using the intermediate masked-language model task. Besides, considering that the fine-tuning procedure can be a very data-intensive task that may demand a lot of computational power, in addition to the large corpus of 6.7M tweets, we use in the fine-tuning process nine other samples with different sizes, varying from 500 to 1.5M tweets. We conduct an experimental evaluation with all models in the sentiment analysis task with five different classifiers and 22 datasets as in the previous questions.

RQ4. Can Transformer-based autoencoder models benefit from a fine-tuning procedure with tweets from sentiment analysis datasets? Although using unlabeled generic tweets to adjust a language model seems to be promising regarding the availability of data, we believe that the fine-tuning procedure may benefit from the sentiment information that tweets from labeled datasets contain. In this context, we aim at identifying whether fine-tuning models with positive and negative tweets can boost the sentiment classification of tweets. We perform this evaluation by assessing three distinct strategies in order to simulate three real-world situations, as follows. In the first strategy, we use a specific sentiment dataset itself as the target domain dataset to fine-tune a language model. The second strategy simulates the case where a collection of general sentiment dataset is available to fine-tune a language model. In the third and last strategy, we combine the two previous situations. In short, we put together tweets from a target dataset and from

a collection of sentiment datasets in the fine-tuning procedure. Finally, we present a comparison between the predictive performances achieved by these three evaluations and the fine-tuned models evaluated in RQ3. As in the previous questions, we conduct the experiments with five different classifiers and 22 datasets.

In summation, given the large number of language models exploited in this study, our main contributions are: (i) a comparative study of a rich collection of publicly available static representations generated from distinct deep learning methods, and with different dimensions, vocabulary size, and from various kinds of corpora; (ii) an assessment of state-of-the-art contextualized language models from the literature, that is, Transformer-based autoencoder models, including BERT, RoBERTa, and BERTweet; (iii) an evaluation of distinct strategies for fine-tuning Transformer-based autoencoder language models; and (iv) a general comparison over static, Transformer-based autoencoder, and fine-tuned language models, aiming at determining the most suitable ones for detecting the sentiment expressed in tweets³.

In order to present our contributions, we organized this work as follows. Chapter 2 presents a literature review related to the language models examined in this study. In Chapter 3, we describe the experimental methodology we followed in the computational experiments, which are reported in Chapters 4, 5, 6, and 7, responding the four research question, respectively. Finally, in Chapter 8, we present the conclusions and directions for future research.

³The code and detailed results from our investigation are publicly available at https://github.com/MeLL-UFF/tuning_sentiment

Chapter 2

Literature Review

Sentiment analysis is an automated process used to predict people’s opinions, sentiments, evaluations, appraisals, attitudes, and emotions towards entities such as products, services, organizations, individuals, issues, events, topics, and their attributes [30]. Recently, sentiment analysis has been recognized as a suitcase research problem [9], which involves solving different NLP classification sub-tasks, including sarcasm, subjectivity, and polarity detection, which is the focus of this study.

Pioneer works in the sentiment classification of tweets mainly focused on the polarity detection task, which aims at categorizing a piece of text as carrying a positive or negative connotation. For example, Go et al. [24] define sentiment as a personal positive or negative feeling. There, they used unigrams as features to train different machine learning classifiers, using tweets with emoticons as training data. The unigram model, or bag-of-words (BoW), is the most basic representation in text classification problems.

Over the years, different techniques have been developed in NLP literature in an effort to make natural language easily triable by computers. Vector Space Models (VSMs) [21] are one of the earliest strategies used to represent the knowledge extracted from a given corpus. Earlier approaches to build VSMs are grounded on count-based methods, such as BoW [56] with TF-IDF (Term Frequency-Inverse Document Frequency) [49] representation, which measures how important a word is to a document, relying on its frequency of occurrence in a corpus.

The BoW model, which assumes word order is not important, is based on the hypothesis that the frequencies of words in a document tend to indicate the relevance of the document to a query [21]. This hypothesis expresses the belief that a column vector in a term-document matrix captures an aspect of the meaning of the corresponding document

or phrase. Precisely, Let X be a term-document matrix. Suppose the document collection contains n documents and m unique terms. The matrix X will then have m rows (one row for each unique term in the vocabulary) and n columns (one column for each document). Let w_i be the i -th term in the vocabulary and let d_j be the j -th document in the collection. The i -th row in X is the row vector $x_{i:}$ and the j -th column in X is the column vector $x_{:j}$. The row vector $x_{i:}$ contains n elements, one element for each document, and the column vector $x_{:j}$ contains m elements, one element for each term. Suppose X is a simple matrix of frequencies, then the element x_{ij} in X is the frequency of the i -th term w_i in the j -th document d_j [56].

Such a simple way of creating numeric representations from texts have motivated early studies in detecting the sentiment expressed in tweets [6, 24, 40]. However, though widely adopted, this kind of feature representation leads to the curse of dimensionality due to the large number of uncommon words tweets contain [47].

Thus, with the revival and success of neural-based learning techniques, several methods that learn dense real-valued low dimensional vectors to represent words have been proposed, such as Word2Vec [53], FastText [54], and GloVe [41]. Word2Vec [53] is one of the pioneer models to become popular taking advantage from the development of neural networks over the years. Word2Vec is actually a software package composed of two distinct implementations of language-models, both based on a feed-forward neural architecture, namely Continuous Bag-Of-Words (CBOW) and Skip-gram. The CBOW model aims at predicting a word given its surrounding context words. Conversely, the Skip-gram model predicts the words in the surrounding context given a target word. Both architectures consist of input, a hidden layer and an output layer. The input layer has the size of the vocabulary and encodes the context by combining the one-hot vector representations of surrounding words of a given target word. The output layer has the same size as the input layer and contains a one-hot vector of the target word obtained during the training. However, one of the main disadvantages of those models is that they usually struggle to deal with out-of-vocabulary (OOV) words, i.e., words that have not been seen in the training data before. To address this weakness, more complex approaches have been proposed, such as FastText [54].

FastText [54] is based on the Skip-gram model [53], still it considers each word as a bag of character n -grams, which are contiguous sequences of n characters from a word, including the word itself. A dense vector is learned to each character n -gram and the dense vector associated to a word is taken from the sum of those representations. Thus,

FastText can deal with different morphological structure of words that covers the words not seen in the training phase, i.e., OOV words. For that reason, FastText is also able to deal with tweets, considering the huge number of uncommon and unique words in this kind of text.

Going to another direction, the GloVe model [41] attempts at making efficient use of statistics of word occurrences in a corpus to learn better word representations. In [41], Pennington et al. present a model that rely on the insight that ratios of co-occurrences, rather than raw counts, encode semantic information about pair of words. This relationship is used to derive a suitable loss function for a log-linear model, which is then trained to maximize the similarity of every word pair, as measured by the ratios of co-occurrences. Given a probe word, the ratio can be small, large or equal to one depending on their correlations. This ratio gives hints on the relations between three different words. For example, given a probe word and two others w_i and w_j , if the ratio is large, the probe word is related to w_i but not w_j .

In general, methods for learning word embeddings deal well with the syntactic context of words but ignore the potential sentiment they carry. In the context of sentiment analysis, words with similar syntactic structure but opposite sentiment polarity, such as *good* and *bad*, are usually mismapped to neighbouring word vectors. To address this issue, Tang et al. [51] proposed the Sentiment-Specific Word Embedding model (SSWE), which encodes the sentiment information in the embeddings. Specifically, they developed neural networks that incorporate the supervision from sentiment polarity of texts in their loss function. To that, they slide the window of n-gram across a sentence, and then predict the sentiment polarity based on each n-gram with a shared neural network. In addition to SSWE, other methods have been proposed in order to improve the quality of word representations in sentiment analysis, by leveraging the sentiment information in the training phase, such as DeepMoji [20], Emo2Vec [58], and EWE [2].

The aforementioned word embedding models have been used as standard components in most sentiment analysis methods. However, they pre-compute the representation for each word independently from the context they are going to appear. This static nature of these models results in two problems: (i) they ignore the diversity of meaning each word may have, and (ii) they suffer from learning long-term dependencies of meaning. Different from those static word embedding techniques, *contextualized embeddings* are not fixed, adapting the word representation to the context it appears. Precisely, at training time, for each word in a given input text, the learning model analyzes the context, usually

using sequence-based models, such as recurrent neural networks (RNNs), and adjusts the representation of the target word by looking at the context. These context-awareness embeddings are actually the internal states of a deep neural network trained in an self-supervised setting. Thus, the training phase is carried out independently from the primary task on an extensive unlabeled data. Depending on the sequence-based model adopted, these contextualized models can be divided into two main groups, namely RNN-based [43] and Transformers-based [29, 31, 39, 28, 57].

Transfer learning strategies have also been emerging to improve the quality of word representation, such as ULMFit (Universal Language Model Fine-tuning) [27]. ULMFit is an effective transfer learning method that can be applied to any NLP task, and introduces key techniques for fine-tuning a language model, consisting of three stages, described as follows. First, the language model is trained on a general-domain corpus to capture generic features of the language in different layers. Next, the full language model is fine-tuned on the target task data using discriminative fine-tuning and slanted triangular learning rates (STLR) to learn task-specific features. Lastly, the model is fine-tuned on the target task using gradual unfreezing and STLR to preserve low-level representations and to adapt high-level ones.

Fine-tuning techniques enable the development and availability of pre-trained contextualized language models using massive amounts of data. For example, Peters et al. [42] introduced ELMo (Embeddings from Language Models), a deep contextualized model for word representation. ELMo comprises a Bi-directional Long-Short-Term-Memory Recurrent Neural Network (BiLSTM) to combine a forward model, looking at the sequence in the traditional order, and a backward model, looking at the sequence in the reverse order. ELMo is composed of two layers of BiLSTM sequence encoder responsible for capturing the semantics of the context. Besides, some weights are shared between the two directions of the language modeling unit and there is also a residual connection between the LSTM layers to accommodate the deep connections without the gradient vanishing issue. ELMo also makes use of the character-based technique for computing embeddings. Therefore, it benefits from the characteristics of character-based representations to avoid OOV words.

Although ELMo is more effective as compared to static pre-trained models, its performance may be degraded when dealing with long texts, exposing a trade-off between efficient learning by gradient descent and latching on information for long periods [7].

Transformer-based language models, on the other hand, have been proposed to solve the gradient propagation problems described in [7]. Compared to RNNs, which process

the input sequentially, Transformers work in parallel, which brings benefits when dealing with large corpora. Moreover, while RNNs by default process the input in one direction, Transformers-based models can attend to the context of a word from distant parts of a sentence and pay attention to the part of the text that really matters, using self-attention [57].

The OpenAI Generative Pre-Training Transformer model (GPT) [45] is one of the first attempts to learn representations using Transformers. It encompasses only the decoder component of the Transformer architecture with some adjustments, discarding the encoder part. Therefore, instead of having a source and a target sentence for the sequence transduction model, a single sentence is given to the decoder. GPT's objective function targets at predicting the next word given a sequence of words, as a standard language modeling goal. To comply with the standard language model task, while reading a token, GPT can only attend to previously seen tokens in the self-attention layers. This setting can be limiting for encoding sentences, since understanding a word might require processing the ones coming after it in the sentence.

Devlin et al. [14] addressed the unidirectional nature of GPTs by presenting an strategy called BERT (Bidirectional Encoder Representations from Transformers) that, as the name says, encodes sentences by looking them at both directions. BERT is also based on the Transformer architecture but, contrary to the GPT, it is based on the encoder component of that architecture. The essential improvement over GPT is that BERT provides a solution for making Transformers bidirectional by applying masked language models, which randomly masks some percentage of the input tokens, and the objective is to predict those masked tokens based on their context. Also, in [14], they use a next sentence prediction task for predicting whether two text segments follow each other. Figure 2.1 presents a representation of the BERT model architecture. The language model receives as input a pair of sentences separated by the special tokens [CLS] and [SEP], as shown in the figure, and corrupted by the masks. The model seeks to reconstruct the sentences and learn the relationship between the sentences. All those improvements have made BERT to achieve state-of-the-art results in various NLP tasks when it was published.

Later, Liu et al. [31] proposed RoBERTa (Robustly optimized BERT approach), achieving even better results than BERT. RoBERTa is an extension of BERT with some modifications, such as: (i) training the model for a longer period of time, with bigger batches, over more data, (ii) removing the next sentence prediction objective, (iii) training on longer sequences, and (iv) dynamically changing the masking pattern applied to

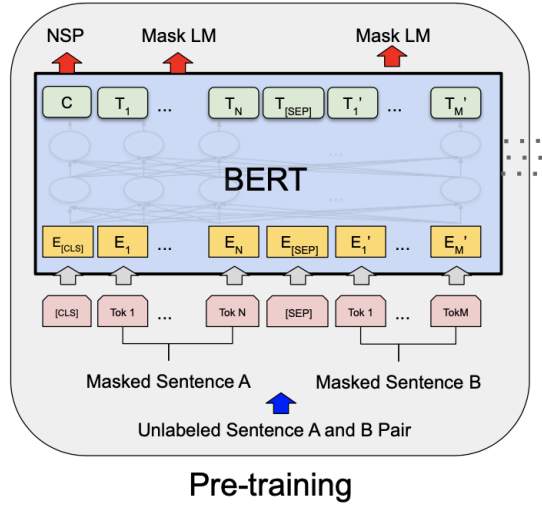


Figure 2.1: BERT model architecture [14]

the training data.

Recently, Nguyen et al. [39] introduced BERTweet, an extension of RoBERTa trained from scratch with tweets. BERTweet has also the same architecture as BERT, but it is trained using the same RoBERTa pre-training procedure instead. BERTweet consumes a corpus of 850M English tweets, which is a concatenation of two corpora. The first corpus contains 845M English tweets from the Twitter Stream dataset and the second one contains 5M English tweets related to the COVID-19 pandemic. In [39], the proposed BERTweet model outperformed RoBERTa baselines in some tasks on tweets, including sentiment analysis.

As far as we know, most studies in language modeling focus on designing new effective models in order to improve the predictive performance of distinct NLP tasks. For example, Devlin et al. [14] and Liu et al. [31] have respectively introduced BERT and RoBERTa, which achieved state-of-the-art results in many NLP tasks. Nevertheless, they did not evaluate the performance of such methods on the sentiment classification of tweets. Nguyen et al. [39], on the other hand, used only a unique generic collection of tweets when evaluating their BERTweet strategy. In this context, we fulfill a robust evaluation of existing language models from distinct natures, including static representations, Transformer-based autoencoder models, and fine-tuned models, by using a significant set of 22 datasets of tweets from different domains and sizes. In the following chapters, we present the assessment of such models.

Chapter 3

Experimental Methodology

This section presents the experimental methodology we followed in this work. We begin by describing, in Section 3.1, the twenty-two benchmark datasets used to evaluate the different language models we investigate in this study. In Section 3.2, we present the experimental protocol we followed. Then, in Section 3.3, we describe the computational experiments reported in Chapters 4, 5, 6, and 7.

3.1 Datasets

We used a large set of twenty-two English datasets [10] to assess the effectiveness of the distinct language models described in Section 2¹. Table 3.1 summarizes the main characteristics of these datasets, namely the abbreviation we use when reporting the experimental results to save space (*Abbrev.* column), the domain they belong (*Domain* column), number of positive tweets (*#pos.* column), proportion of positive tweets (*%pos.* column), number of negative tweets (*#neg.* column), proportion of negative tweets (*%neg.* column), and the total number of tweets (*Total* column).

Those datasets have been extensively used in the literature of Twitter sentiment analysis and we believe they provide a diverse scenario in evaluating embeddings of tweets in the sentiment classification task, regarding a variety of domains, sizes, and class balance. For example, while datasets SemEval13, SemEval16, SemEval17, and SemEval18 contain generic tweets, other datasets, such as *iphone6*, *movie*, and *archeage*, contain tweets of a particular domain. Also, the datasets vary a lot in size, with some of them containing only dozens of tweets, such as *irony* and *sarcasm*. We believe that this diverse and

¹The datasets are publicly available at <https://github.com/joncarv/air-datasets>

Table 3.1: Characteristics of the Twitter sentiment datasets ordered by size (Total column)

Dataset	Abbrev.	Domain	#pos.	%pos.	#neg.	%neg.	Total
irony [25]	iro	Irony	22	34%	43	66%	65
sarcasm [25]	sar	Sarcasm	33	46%	38	54%	71
aisopos ²	ntu	Generic	159	57%	119	43%	278
SemEval-Fig ³	S15	Irony/Metaphors	47	15%	274	85%	321
sentiment140 [24]	stm	Generic	182	51%	177	49%	359
person [13]	per	Towards a Person	312	71%	127	29%	439
hobbit [32]	hob	Movies	354	68%	168	32%	522
iphone6 [32]	iph	Products	371	70%	161	30%	532
movie [13]	mov	Movies	460	82%	101	18%	561
sanders ⁴	san	Business	570	47%	654	53%	1,224
Narr [38]	Nar	Generic	1,739	60%	488	40%	1,227
archeage [32]	arc	Games	724	42%	994	58%	1,718
SemEval18 [36]	S18	Equity Evaluation Corpus	865	47%	994	53%	1,859
OMD [15]	OMD	Presidential Debate	710	37%	1,196	63%	1,906
HCR [50]	HCR	Health Care Reform	539	28%	1,369	72%	1,908
STS-gold [48]	STS	Generic	632	31%	1,402	69%	2,034
SentiStrength [52]	SSt	Generic	1,340	59%	949	41%	2,289
Target-dependent [18]	Tar	Celebrities	1,734	50%	1,733	50%	3,467
Vader [23]	vad	Generic	2,897	69%	1,299	31%	4,196
SemEval13 ⁵	S13	Generic	3,183	73%	1,195	23%	4,378
SemEval17 [46]	S17	Generic	2,375	37%	3,972	63%	6,347
SemEval16 [37]	S16	Generic	8,893	73%	3,323	27%	12,216

large collection of datasets may help drawing more concise and robust conclusions on the effectiveness of distinct language models in the sentiment analysis task.

3.2 Experimental Protocol

To assess the effect of different kinds of language models in the polarity classification task, we follow the protocol of first extracting the features from the several vector-based language representation mechanisms (BOW, static embeddings, contextualized embeddings, fine-tuned embeddings). Next, those features compose the input attribute space for five distinct classifiers, namely Support Vector Machine (SVM), Logistic Regression (LR), Random Forest (RF), XGBoost (XBG), and Multi-layer Perceptron (MLP). We adopted scikit-learn’s⁶ implementations of those machine learning algorithms. Although we have used the default parameters in most of the cases, it is important to mention that we set the class balance parameter for SVM, LR, and RF (*class_weight = balanced*). Also, for LR, we set the maximum number of iterations to 500 (*max_iter = 500*) and the solver parameter to *liblinear*. Moreover, for MLP, we set the number of hidden layers to 100. Table 3.2 shows a summary of the classification algorithms used in this study, remarking their characteristics. We aim at determining which language models are the

⁶<https://scikit-learn.org>

Table 3.2: Summary of the advantages and disadvantages of the classification algorithms adopted in this study

Classifier	Advantages	Disadvantages
SVM	<ul style="list-style-type: none"> – Works well when there is a clear margin of separation between classes – More effective in high dimensional spaces – Memory efficient 	<ul style="list-style-type: none"> – Underperforms when the number of features for each data point exceeds the number of training data samples
LR	<ul style="list-style-type: none"> – Easy to interpretate – Provides association directions 	<ul style="list-style-type: none"> – Non-linear problems cannot be solved – Needs that independent variables are i.i.d – Linearly related to the log odds
RF	<ul style="list-style-type: none"> – No feature scaling required – Handles non-linear parameters efficiently 	<ul style="list-style-type: none"> – Long training period – Difficult to interpret
XGB	<ul style="list-style-type: none"> – Highly scalable/parallelizable – Quick to execute 	<ul style="list-style-type: none"> – More likely to overfit – Many parameters
MLP	<ul style="list-style-type: none"> – Learns non-linear and complex relationships 	<ul style="list-style-type: none"> – Difficult to interpret parameters – Convergence of the weights can be very slow

most effective ones in Twitter sentiment analysis by leveraging classifiers from distinct natures, thus examining how they deal with the peculiarities of each evaluated model. Furthermore, it is important to note that we do not aim at establishing the best classifier for the sentiment analysis task, which may require a specific study and additional computational experiments.

Preprocessing is the first step in many text classification problems and the use of appropriate techniques can reduce noise hence improving classification effectiveness [19]. As this manuscript’s main goal is to evaluate the performance of different models of tweet representation, the preprocessing step is simple so that the focus is on the language models and classifiers. Thus, for each tweet in a given dataset, we only replace URLs by the token *someurl*, user mentions by the token *someuser*, and all tokens were lowercased.

In the experimental evaluation, the predictive performance of the sentiment classification is measured in terms of accuracy and F_1 -macro. Precisely, for each evaluated dataset, the accuracy of the classification was computed as the ratio between the number of correctly classified tweets and the total number of tweets, following a stratified ten-fold cross-validation. F_1 -macro was computed as the unweighted average of the F_1 -score for the positive and negative classes.

Moreover, all experiments were performed by using Tesla P100-SXM2 GPU within Ubuntu operational system, running in a machine with Intel(R) Xeon(R) CPU E5-2698

v4 processor.

3.3 Computational Experiments Details

In the next sections, we evaluate a significant collection of vector-based textual representations attempting to answer the research questions introduced in Chapter 1. Specifically, we conduct a comparative study of vector-based language representation models from distinct natures, including Bag of Words, as a classic baseline, static representations and representations induced from Transformer-based autoencoder models, by fine-tuning or not the intermediate masked language task, in order to acknowledge their effectiveness in the polarity classification of English tweets. These language representation models are incrementally evaluated throughout Chapters 4, 5, 6, and 7.

In Chapter 4, we begin by analyzing the predictive performance of the static representations, which include 13 pretrained embeddings from the literature, as shown in Table 3.3, as well as the classical BOW with TF-IDF representation schema. Regarding the static embeddings described in Table 3.3, we have selected representations trained on distinct kinds of texts (Corpus column) and built from different architectures (Architecture column), from feedforward neural networks to Transformer-based ones. The $|D|$ and $|V|$ columns refer to the dimension and vocabulary size of each pretrained embedding, respectively. Although the most usual way of employing embeddings trained from Transformer-based architectures is running the text through the model to obtain contextualized representations, here we first investigate how these models behave when the experimental protocol is the same as earlier embeddings models: pretrained embeddings are collected from the embeddings layer and are the input of the classifiers.

Next, in Chapter 5, we present an evaluation of state-of-the-art Transformer-based autoencoder models, including BERT [14], RoBERTa [31], and BERTweet [39]. In this evaluation, for each assessed dataset, we represent their tweets as the average of the concatenation of the last four layers for each token representation of the models. For the sake of simplicity, the Transformer-based autoencoder models assessed in this study are referred to hereafter as Transformer-based models.

Lastly, in Chapters 6 and 7, we evaluate the effectiveness of fine-tuning the aforementioned Transformer-based models regarding the intermediate masked-language task in two different ways: (i) by using a huge collection of unlabeled, or non-sentiment, tweets, and (ii) by using tweets from sentiment datasets.

Table 3.3: Characteristics of the static pretrained embeddings ordered by the number of dimensions

Embedding	$ D $	$ V $	Architecture	Corpus
SSWE [51]	50	137K	Feed-forward network	Twitter (10M tweets)
Emo2Vec [58]	100	1.2M	Convolutional network	Twitter (1.9M tweets)
GloVe-TWT [41]	200	1.2M	log-bilinear model	Twitter (27B tokens)
DeepMoji [20]	256	50K	Recurrent network	Twitter (1B tweets)
EWE [2]	300	183K	Recurrent network	Amazon reviews (200K reviews)
GloVe-WP [41]	300	400K	log-bilinear model	Wikipedia/Gigaword (6B tokens)
fastText [34]	300	1M	Feed-forward network	Wikipedia/web pages/news (16B tokens)
w2v-GN [35]	300	3M	Feed-forward network	Google news (100B tokens)
w2v-Edin [8]	400	259K	Feed-forward network	Twitter (10M tweets)
w2v-Araque [5]	500	57K	Feed-forward network	Twitter (1.28M tweets)
BERT [14]	768	30K	Transformers	BooksCorpus [59]/Eng. Wiki (3.3B words)
RoBERTa (RoB) [59]	768	50K	Transformers	5 datasets [31] (161GB)
BERTweet (BTWT) [39]	768	64K	Transformers	Twitter (850M tweets)

In Chapter 6, regarding the non-sentiment fine-tuning approach, we adopted the general purpose collection of unlabeled tweets from the Edinburgh corpus [44], which contains 97M tweets in multiple languages. Tweets written in languages other than English were discarded, resulting in a final corpus of 6.7M English tweets, which was then used to fine-tune BERT, RoBERTa, and BERTweet. In addition to the entire corpus of 6.7M tweets, we used nine other samples with different sizes, varying from 500 to 1.5M tweets. Specifically, we generated samples containing 500 (0.5K), 1K, 5K, 10K, 25K, 50K, 250K, 500K, and 1.5M non-sentiment tweets.

Conversely, in Chapter 7, we evaluated the sentiment fine-tuning procedure using positive and negative tweets from the twenty-two benchmark datasets described in Table 3.1. For this purpose, we used each dataset once as the target dataset, while the others were used as the source datasets. More clearly, for each assessed dataset, referred to as the target dataset, we explored three distinct strategies to fine-tune the masked-language model: (i) by using only the tweets from the target sentiment dataset itself, (ii) by using the tweets from the remaining 21 datasets, and (iii) by using the entire collection of tweets from the 22 datasets, including the tweets from the target dataset.

Chapter 4

Evaluation of Static Text Representations

The computational experiments conducted in this section aim at answering the research question RQ1, as follows:

RQ1. Which static embeddings are the most effective in the sentiment classification of tweets?

We answer this question by assessing the predictive power of the 13 pretrained embeddings described in Table 3.3. These embeddings were generated from distinct neural networks architectures, with different dimensions and vocabulary size, and trained on various kinds of corpora. Recall that by static embeddings we mean that the features are gathered from the embeddings layer working as a look-up table of tokens. In addition to the pretrained embeddings, we evaluate the BoW model with the TF-IDF representation, which is the most basic text representation used in Twitter sentiment analysis and text classification tasks in general. For all tweet representation, we take the average of all tokens representation of the tweet.

We begin by evaluating the predictive performance of the static representations for each classification algorithm presented in Table 3.2. We report the computational results in detail for SVM as an example of this evaluation (refer to Appendix A for the detailed evaluation for each classifier). Tables 4.1 and 4.2 show the results achieved by using each static representation to train an SVM classifier, in terms of classification accuracy and unweighted F_1 -macro, respectively. The boldfaced values indicate the best results, and the last three lines show the total number of wins for each static representation (#wins row), as well as a ranking of the results (rank sums and position rows). Precisely, for each dataset, we assign scores, from 1.0 to 14.0, to each assessed representation (each

Table 4.1: Accuracies (%) achieved using the SVM classifier

Dataset	w2v GN	GloVe WP	Fast Text	EWE	GloVe TWT	BERT static	TF- IDF	w2v Araque	w2v Edin	SSWE	Emo2 Vec	Deep Moji	RoB static	BTWT static
iro	70.95	69.52	67.86	52.14	51.67	60.24	64.52	63.10	66.43	70.48	78.33	61.90	70.48	39.29
sar	68.75	68.75	61.61	68.75	67.32	70.54	54.64	71.79	68.93	87.32	73.04	67.50	56.07	57.68
ntu	79.15	88.47	73.04	70.54	73.70	83.47	74.40	77.71	93.15	93.53	81.28	91.73	89.91	81.98
S15	87.88	83.20	81.95	88.80	83.50	89.72	87.23	87.54	89.41	77.56	82.88	88.16	90.34	87.52
stm	85.22	83.02	83.56	83.01	83.56	81.06	81.59	82.71	87.74	80.75	85.24	81.34	84.94	71.59
per	83.37	76.75	76.77	74.69	73.56	73.32	80.18	74.93	79.03	72.44	76.09	76.98	80.40	69.23
hob	89.66	85.24	83.51	90.22	79.88	91.75	93.29	91.37	90.04	76.99	87.18	91.57	92.90	90.03
iph	78.00	75.57	76.13	72.17	73.88	79.33	80.07	80.46	78.01	76.49	81.57	78.56	80.08	76.51
mov	82.90	76.66	76.65	75.95	71.67	83.96	84.67	83.78	84.14	78.09	83.07	81.29	85.38	79.35
san	82.59	78.50	80.55	80.54	80.38	81.37	83.33	81.46	83.25	80.14	81.12	80.80	83.57	78.84
Nar	84.51	83.94	84.03	83.78	85.33	83.46	80.36	84.18	88.67	88.84	88.34	88.83	87.21	80.61
arc	85.45	83.59	85.16	84.05	85.27	86.61	87.54	85.16	86.67	79.63	83.35	86.09	87.43	84.87
S18	80.37	78.48	80.69	80.26	80.20	84.56	82.09	75.79	82.36	81.28	80.96	78.97	86.50	79.51
OMD	83.79	82.21	81.42	77.55	77.38	84.15	79.22	77.70	82.84	76.50	75.76	77.02	85.10	82.37
HCR	74.78	72.58	72.64	71.59	73.21	76.36	80.24	73.95	75.94	67.29	70.54	73.21	78.30	73.27
STS	85.94	83.19	85.74	84.71	84.90	86.97	83.97	86.92	87.02	88.99	86.19	88.69	87.86	83.09
SSt	77.98	75.10	76.98	77.76	77.24	78.51	73.48	76.28	79.56	79.77	84.93	79.64	80.21	73.00
Tar	83.67	81.92	83.82	83.18	82.81	83.44	82.35	81.40	83.39	79.18	82.98	82.90	84.42	80.36
Vad	88.25	84.51	87.01	86.44	85.96	87.51	83.27	85.30	87.73	85.70	85.94	86.77	89.32	81.82
S13	81.52	78.53	79.49	78.39	78.69	81.59	80.52	79.19	80.31	81.04	87.80	81.89	83.60	77.64
S17	88.61	86.34	88.29	87.22	87.71	88.47	87.95	84.53	87.96	81.63	85.60	85.47	89.03	86.01
S16	84.50	82.51	84.50	83.15	83.87	85.27	85.85	81.37	84.39	78.51	82.29	82.79	86.00	81.69
#wins	1	0	0	0	0	0	3	0	1	4	4	0	9	0
rank sums	116.5	221.0	183.5	212.0	216.0	117.0	155.0	185.5	95.0	200.5	154.0	152.5	56.5	245.0
position	3.0	13.0	8.0	11.0	12.0	4.0	7.0	9.0	2.0	10.0	6.0	5.0	1.0	14.0

column), in ascending order of accuracy (F_1 -macro), where the score 1.0 is assigned to the representation with the highest accuracy (F_1 -macro). Thus, low score values indicate better results. When two assessed representation has the same performance, we take an average of their scores. If two assessed representations achieve the best performance, they will receive a score of 1.5 $((1+2)/2)$. Finally, we sum up the assigned scores obtained in each dataset for each assessed representation to calculate rank sums. With the rank sum of each assessed representation, we rank the rank-sum result from the best (1) to the worst (14), calculating the rank position.

As we can see in Tables 4.1 and 4.2, RoBERTa (RoBstatic column) achieved the best performance in nine out of the 22 datasets in terms of accuracy, in 11 out of the 22 datasets in terms of F_1 -macro, and was ranked first in the overall evaluation (position row). Regarding the number of wins (#wins row), we can note that Emo2Vec and SSWE achieved the second best results, reaching the best performance in four out of the 22 datasets for both accuracy and F_1 -macro. However, regarding the overall evaluation (position row), w2v-Edin and w2v-GN were ranked among the top three best static representations along with RoBERTa, in terms of accuracy. Regarding F_1 -macro, the top three best static representations were RoBERTa, w2v-Edin and BERT (BERT-static column).

Tables 4.3 and 4.4 show a summary of the results by evaluating each static representation on the 22 datasets, for each classification algorithm. Each cell indicates the number

Table 4.2: F_1 -macro scores (%) achieved by evaluating static representation using the SVM classifier

Dataset	w2v GN	GloVe WP	Fast Text	EWE	GloVe TWT	BERT static	TF- IDF	w2v Araque	w2v Edin	SSWE	Emo2 Vec	Deep Moji	RoB static	BTWT static
iro	60.03	64.46	60.78	47.51	49.40	54.84	39.11	48.35	58.24	67.04	72.17	53.46	61.16	35.48
sar	66.00	67.75	60.22	67.55	63.32	69.29	51.68	69.97	66.70	86.93	70.72	64.72	50.28	52.69
ntu	78.79	87.81	72.46	70.19	73.02	83.15	71.55	76.72	92.99	93.29	80.99	91.53	89.42	81.22
S15	69.91	65.74	68.13	75.02	70.93	77.04	58.51	66.33	77.62	67.16	70.24	73.65	78.48	75.38
stm	85.18	82.94	83.46	82.96	83.51	80.96	81.46	82.60	87.67	80.71	85.18	81.28	84.87	71.50
per	80.58	74.31	74.86	73.03	71.80	70.85	72.27	71.59	76.97	70.46	74.16	74.46	77.30	67.53
hob	88.56	83.89	82.25	89.26	78.71	90.92	91.73	90.52	88.95	75.49	86.03	90.73	92.08	89.05
iph	75.96	74.31	74.96	71.00	72.63	77.34	71.55	77.92	76.18	74.93	79.97	76.62	78.00	74.70
mov	73.42	68.28	68.39	67.23	64.17	73.72	58.15	73.55	76.15	71.60	76.06	72.97	73.86	68.37
san	82.38	78.24	80.28	80.32	80.03	81.20	82.99	81.33	83.10	79.99	80.98	80.60	83.41	78.43
Nar	84.08	83.50	83.76	83.42	85.02	82.93	79.32	83.72	88.38	88.48	88.01	88.45	86.82	80.06
arc	85.14	83.27	84.84	83.80	84.87	86.37	87.11	84.90	86.34	78.88	82.84	85.61	87.19	84.56
S18	80.21	78.31	80.51	80.03	79.91	84.44	81.55	75.44	82.15	81.13	80.86	78.74	86.40	79.24
OMD	82.50	80.91	80.03	76.00	75.81	82.70	76.57	76.01	81.36	74.86	74.28	75.71	83.85	80.63
HCR	70.97	68.96	69.78	68.26	70.44	72.94	72.16	69.17	72.38	62.99	66.96	68.79	74.55	69.74
STS	84.12	81.29	83.99	82.93	83.26	84.99	79.56	84.95	85.20	87.62	84.49	87.08	85.95	80.67
SSt	77.69	74.90	76.69	77.59	77.07	78.20	72.49	75.89	79.27	79.49	84.68	79.36	79.83	72.74
Tar	83.67	81.91	83.81	83.18	82.80	83.43	82.33	81.39	83.38	79.13	82.96	82.87	84.42	80.32
Vad	86.68	82.62	85.33	84.82	84.23	85.76	78.48	83.38	86.14	84.08	84.28	85.14	87.80	79.66
S13	78.43	75.83	76.82	75.52	76.10	78.67	72.24	75.95	77.64	78.40	85.59	79.16	80.40	74.63
S17	88.02	85.67	87.69	86.61	87.05	87.75	86.9	83.52	87.30	80.51	84.89	84.69	88.37	85.13
S16	81.58	79.52	81.62	80.15	81.10	82.48	81.68	78.17	81.63	75.82	79.54	80.00	83.18	78.62
#wins	1	0	0	0	0	0	0	0	2	4	4	0	11	0
rank sums	123.5	215.0	168.0	206.0	202.0	111.0	212.0	198.0	89.0	193.0	143.5	153.0	56.0	240.0
position	4.0	13.0	7.0	11.0	10.0	3.0	12.0	9.0	2.0	8.0	5.0	6.0	1.0	14.0

of wins, the rank sums, and the rank position achieved by the related static representation (each line) used to train the corresponding classifier (each column). The *Total* column indicates the total number of wins, the total rank sums, and the total rank position, i.e., the sum of the rank positions presented in each cell for each assessed model. Moreover, in the total column, we underline the top three best overall results in terms of total rank position.

Regarding the overall evaluation (Total column), from Tables 4.3 and 4.4, we can see that although Emo2Vec achieved the highest total number of wins (i.e., 27 wins in terms of accuracy, and 29 wins in terms of F_1 -macro), w2v-Edin was ranked as the best overall model, achieving the lowest total rank position for both accuracy (22.0) and F_1 -macro (21.0). Nevertheless, considering each classifier (each column), we can note that RoBERTa achieved the best performance when used to train LR, SVM, and MLP, for both accuracy and F_1 -macro. Conversely, Emo2Vec achieved the best overall results when used to train RF and XGB classifiers. Analyzing the overall results in terms of the total rank position (Total column), we observe that Emo2Vec and w2v-GN, along with w2v-Edin, are ranked as the top three best static representations. These results suggest that w2v-Edin, Emo2Vec, and w2v-GN are well-suited static representations for Twitter sentiment analysis.

In the previous evaluations, we analyzed the predictive performance achieved by each

Table 4.3: Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by evaluating each static representation on the 22 datasets, for each classification algorithm, in terms of accuracy

Representation	LR	SVM	MLP	RF	XGB	Total
w2v-GN	3/116.5/4.0	1/116.5/3.0	1/151.5/6.0	1/159.0/6.0	3/125.5/4.0	9/669.0/ <u>23.0</u>
GloveWP	0/149.0/7.0	0/221.0/13.0	0/212.5/11.0	0/207.0/11.0	0/195.5/10.0	0/985.0/52.0
FastText	0/192.5/10.0	0/183.5/8.0	2/148.5/5.0	0/190.0/9.0	0/151.5/6.0	2/866.0/38.0
EWE	1/101.5/3.0	0/212.0/11.0	0/142.0/4.0	0/170.5/7.0	0/153.5/8.0	1/779.5/33.0
GloveTW	1/97.0/2.0	0/216.0/12.0	1/152.0/7.0	1/105.0/3.0	2/123.0/3.0	5/693.0/27.0
SSWE	4/152.0/8.0	4/200.5/10.0	1/234.0/14.0	6/79.5/2.0	4/153.0/7.0	19/819.0/41.0
TF-IDF	5 /146.5/6.0	3/155.0/7.0	1/225.0/13.0	3/116.0/5.0	2/215.5/13.0	14/858.0/44.0
DeepMoji	0/174.0/9.0	0/152.5/5.0	1/167.5/8.0	0/171.0/8.0	1/144.0/5.0	2/809.0/35.0
w2v-Araque	0/204.0/11.0	0/185.5/9.0	0/221.0/12.0	0/202.5/10.0	0/214.5/12.0	0/1027.5/54.0
w2v-Edin	0/220.5/12.0	1/95.0/2.0	2/100.5/2.0	1/105.5/4.0	2/106.5/2.0	6/628.0/ 22.0
Emo2Vec	4/120.0/5.0	4/154.0/6.0	2/192.0/10.0	10 / 46.0 / 1.0	7 / 98.5 / 1.0	27 / 610.5 / <u>23.0</u>
BERT-static	0/264.0/13.0	0/117.0/4.0	3/117.5/3.0	0/235.5/12.0	0/207.5/11.0	3/941.5/43.0
RoBERTa-static	5 / 82.5 / 1.0	9 / 56.5 / 1.0	8 / 74.0 / 1.0	0/244.0/13.0	1/167.5/9.0	23/624.5/25.0
BERTweet-static	0/290.0/14.0	0/245.0/14.0	0/172.0/9.0	0/278.5/14.0	0/254.0/14.0	0/1239.5/65.0

Table 4.4: Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by evaluating each static representation on the 22 datasets, for each classification algorithm, in terms of F_1 -macro

Representation	LR	SVM	MLP	RF	XGB	Total
w2v-GN	3/113.0/4.0	1/123.5/4.0	1/146.5/5.0	1/161.5/6.0	3/129.0/4.0	9/673.5/ <u>23.0</u>
GloveWP	0/152.0/6.5	0/215.0/13.0	1/213.0/11.0	0/208.0/11.0	0/195.5/10.0	1/983.5/51.5
FastText	0/189.0/10.0	0/168.0/7.0	1/149.5/6.0	0/200.0/9.0	0/153.0/7.0	1/859.5/39.0
EWE	1/102.5/3.0	0/206.0/11.0	0/144.0/4.0	0/165.5/7.0	0/160.5/8.0	1/778.5/33.0
GloveTW	1/96.0/2.0	0/202.0/10.0	1/152.0/7.0	0/109.5/5.0	1/123.0/3.0	3/682.5/27.0
SSWE	5 /152.0/6.5	4/193.0/8.0	1/228.0/13.0	6/70.0/2.0	4/135.0/5.0	20/778.0/34.5
TF-IDF	3/166.0/8.0	0/212.0/12.0	1/243.0/14.0	4/106.5/4.0	3/211.5/11.0	11/939.0/49.0
DeepMoji	0/169.0/9.0	0/153.0/6.0	1/164.0/8.0	0/173.5/8.0	1/143.0/6.0	2/802.5/37.0
w2v-Araque	0/200.0/11.0	0/198.0/9.0	0/222.0/12.0	0/204.5/10.0	0/216.5/12.0	0/1041.0/54.0
w2v-Edin	0/221.0/12.0	2/89.0/2.0	2/95.0/2.0	0/105.0/3.0	2/110.0/2.0	6/620.0/ 21.0
Emo2Vec	4/117.5/5.0	4/143.5/5.0	2/191.0/10.0	11 / 42.0 / 1.0	8 / 81.5 / 1.0	29 / 575.5 / <u>22.0</u>
BERT-static	0/261.0/13.0	0/111.0/3.0	3/112.5/3.0	0/234.0/12.0	0/217.5/13.0	3/936.0/44.0
RoBERTa-static	5 / 82.0 / 1.0	11 / 56.0 / 1.0	8 / 75.5 / 1.0	0/245.5/13.0	0/177.0/9.0	24/636.0/25.0
BERTweet-static	0/289.0/14.0	0/240.0/14.0	0/174.0/9.0	0/284.5/14.0	0/257.0/14.0	0/1244.5/65.0

representation for one classification algorithm at a time, focusing on the individual contribution of the text representations in the performance on the final task. Next, we investigate the classification performance of the final sentiment analysis process, that is, the combination of text representation and classifier. Considering that the final classification is a combination of both representation and classifier, an appropriate choice of the classification algorithm may affect the performance of a text representation. For this purpose, we present an overall evaluation of all possible combinations of text representations and classification algorithms, examining them as pairs $\{text\ representation, classifier\}$. More clearly, we evaluate the classification effectiveness of 70 possible combinations of text representations and classifiers (14×5) on the 22 datasets of tweets. Table 4.5 presents the top ten results in terms of the average rank position and 4.6 presents the ten worst average rank position. Specifically, for each dataset, we calculate a rank of the 70 combinations and then average the rank position of each combination over the 22 datasets. From Table 4.5, we can note that the best overall results were achieved by using RoBERTa to train an SVM classifier for both accuracy and F_1 -macro. Also, w2v-Edin + SVM and RoBERTa + MLP appear in the top three results along with RoBERTa + SVM. By Table 4.6, we can notice the high-frequency of RF in the pair Model-Classifier.

Tables 4.7 and 4.8 show a summary of the results for each text representation and classifier, respectively, from best to worst, in terms of the average rank position. As we can observe, Emo2Vec, RoBERTa, and w2v-Edin appear in the top three, being the representations that achieved the best overall performances. Among the classifiers, we can note that SVM and MLP seem to be good choices in Twitter sentiment analysis regarding the usage of static text representations. Conversely, RF achieved the worst overall performance across all evaluations.

In addition to the individual assessment of text representations and classifiers presented in Tables 4.7 and 4.8, Table 4.9 shows the best results achieved for each dataset. We can see that RoBERTa achieved the highest accuracies in seven out of the 22 datasets, and highest F_1 -macro scores in nine out of the 22 datasets. Furthermore, as highlighted in Table 4.5, RoBERTa + SVM achieved the best performances in six out of the 22 datasets in terms of accuracy, and in eight out of the 22 datasets in terms of F_1 -macro.

The top three static representations identified in the previous experiments, i.e., RoBERTa, w2v-Edin, and Emo2Vec, are very different from each other. While w2v-Edin and Emo2Vec were trained from scratch on tweets, RoBERTa was trained on traditional English texts.

The better performance of Emo2Vec and w2v-Edin can be caused by the inclusion of

the sentiment analysis task in its training process. We also have other models built in this same strategy and trained from scratch with tweets, such as Deepmoji and SSWE, respectively, the seventh and eighth position by Table 4.7. The Emo2Vec better performance may be a result of its multi-task learning approach. Considering another model with the same architecture of w2v-Edin and also trained from scratch with tweets, the differential performance between w2v-Edin and w2v-Araque (the fourteenth position by Table 4.7) may lie in its volume of training data (w2v-Araque: 1.28M, and W2V-Edin: 10M) and the vocabulary size (w2v-Araque: 57K and w2v-Edin: 259K).

However, among these, RoBERTa is the only Transformer-based model, which holds state-of-the-art performance in capturing context and semantics of terms from texts. Furthermore, regarding w2v-Edin, although it was trained with a more straightforward architecture (feedforward neural network) as compared to others, its training parameters were optimized for the emotion detection task on tweets [8], which may have helped determining the sentiment expressed in tweets.

Surprisingly, as shown in Table 4.7, BERTweet achieved the worst overall performance among all assessed text representations, despite having been trained using the same state-of-the-art Transformer-based architecture as RoBERTa *while yet using tweets*. One possible explanation for this behavior is that BERTweet training procedure limits the representation of its training tweets to 60 tokens only, while RoBERTa uses a limit of 512 tokens. For that reason, we believe that RoBERTa model is able to capture more semantic information to the tokens from its training vocabulary as compared to BERTweet when one collects the token representation from the embeddings layer.

Finally, regarding research question RQ1, we can highlight and suggest that: (i) disregarding the classification algorithms, Emo2Vec, w2v-Edin, and RoBERTa seem to be well-suited representations for determining the sentiment expressed in tweets, and (ii) considering the combination of text representations and classifiers, RoBERTa + SVM achieved the best overall performance, which may represent a good choice for Twitter sentiment analysis in hardware-restricted environments, since the cost here is most due to the classifier induction.

Table 4.5: Top 10 average rank results achieved for each combination Model-Classifier

Representation	Classifier	Accuracy avg. rank pos.	Representation	Classifier	F_1 -macro avg. rank pos.
RoBERTa-static	SVM	9.32	RoBERTa-static	SVM	8.59
RoBERTa-static	MLP	11.57	W2V-Edin	SVM	9.39
W2V-Edin	SVM	14.50	RoBERTa-static	MLP	12.52
W2V-Edin	MLP	15.36	BERT-static	SVM	14.70
BERT-static	MLP	16.68	W2V-GN	SVM	14.95
W2V-GN	SVM	17.80	W2V-Edin	MLP	15.55
BERT-static	SVM	19.02	RoBERTa-static	LR	16.02
Emo2Vec	Xgb	21.23	BERT-static	MLP	17.48
W2V-GN	MLP	22.50	GloVe-TWT	LR	17.91
fastText	MLP	23.20	Emo2Vec	SVM	18.05

Table 4.6: Tail 10 average rank results achieved for each combination Model-Classifier

Representation	Classifier	Accuracy avg. rank pos.	Representation	Classifier	F_1 -macro avg. rank pos.
DeepMojj	RF	51.41	EWE	RF	56.86
BERTweet-static	Xgb	52.14	DeepMojj	RF	57.43
fastText	RF	52.75	BERTweet-static	LR	57.45
GloVe-WP	RF	54.11	W2V-GN	RF	58.02
W2V-Araque	RF	54.68	fastText	RF	60.36
BERT-static	RF	57.18	W2V-Araque	RF	61.00
RoBERTa-static	RF	57.95	GloVe-WP	RF	61.16
BERT-static	LR	60.86	BERT-static	RF	63.91
BERTweet-static	RF	61.02	RoBERTa-static	RF	64.11
BERTweet-static	LR	66.09	BERTweet-static	RF	67.32

Table 4.7: Average Rank results achieved for each Embedding, evaluating static representations

Representation	Accuracy avg. rank pos.	Representation	F_1 -macro avg. rank pos.
Emo2Vec	26.35	Emo2Vec	25.34
RoBERTa-static	27.93	RoBERTa-static	29.06
W2V-Edin	29.55	W2V-Edin	30.05
W2V-GN	30.05	W2V-GN	31.04
GloVe-TWT	32.00	GloVe-TWT	31.55
EWE	34.50	SSWE	33.40
SSWE	34.78	EWE	34.15
DeepMoji	35.93	DeepMoji	34.98
TF-IDF	36.21	fastText	36.34
fastText	37.00	BERT-static	39.38
BERT-static	39.43	GloVe-WP	39.43
GloVe-WP	40.25	TF-IDF	40.25
W2V-Araque	43.15	W2V-Araque	42.85
BERTweet-static	49.88	BERTweet-static	49.18

Table 4.8: Average Rank results achieved for each Classifier, evaluating static representations

Classifier	Accuracy avg. rank pos.	Classifier	F_1 -macro avg. rank pos.
MLP	26.28	SVM	23.07
SVM	28.30	MLP	26.69
XGB	35.83	LR	31.27
LR	39.17	XGB	41.33
RF	47.92	RF	55.15

Table 4.9: Best results achieved for each dataset

Dataset	Accuracy	Classifier	Representation	F_1 -macro	Classifier	Representation
iro	78.81	LR	Emo2Vec	75.87	LR	Emo2Vec
sar	87.50	LR	SSWE	87.19	LR	SSWE
ntu	95.30	MLP	w2v-Edin	95.19	MLP	w2v-Edin
S15	90.35	LR	TF-IDF	78.48	SVM	RoBERTa-static
stm	87.74	SVM	w2v-Edin	87.67	SVM	w2v-Edin
per	83.83	MLP	w2v-GN	80.58	SVM	w2v-GN
hob	94.82	MLP	BERT-static	94.05	MLP	BERT-static
iph	84.39	MLP	GloVe-TWT	81.15	MLP	GloVe-TWT
mov	88.78	XGB	Emo2Vec	77.86	MLP	fastText
san	84.71	MLP	TF-IDF	84.56	MLP	TF-IDF
Nar	89.00	LR	SSWE	88.58	LR	SSWE
arc	87.60	MLP	RoBERTa-static	87.29	MLP	RoBERTa-static
S18	86.50	SVM	RoBERTa-static	86.40	SVM	RoBERTa-static
OMD	85.10	SVM	RoBERTa-static	83.85	SVM	RoBERTa-static
HCR	80.24	SVM	TF-IDF	74.55	SVM	RoBERTa-static
STS	89.08	LR	SSWE	87.70	LR	SSWE
SST	85.06	LR	Emo2Vec	84.77	LR	Emo2Vec
Tar	84.42	SVM	RoBERTa-static	84.42	SVM	RoBERTa-static
Vad	89.32	SVM	RoBERTa-static	87.80	SVM	RoBERTa-static
S13	88.24	XGB	Emo2Vec	85.59	LR	Emo2Vec
S17	89.03	SVM	RoBERTa-static	88.37	SVM	RoBERTa-static
S16	86.00	SVM	RoBERTa-static	83.18	SVM	RoBERTa-static

Chapter 5

Evaluation of the Transformer-based Text Representations

In this chapter, we address the research question RQ2, as follows:

RQ2. Considering state-of-the-art Transformer-based autoencoder models, which are the most effective in the sentiment classification of tweets?

To answer that question, we conduct a thorough evaluation of the widely used BERT and RoBERTa models and the BERT-based transformer trained from scratch with tweets, namely, BERTweet. These models represent a set of the most recent Transformer-based autoencoder language modeling techniques that have achieved state-of-the-art performance in many NLP tasks. While BERT is the first Transformer-based autoencoder model to appear in the literature, RoBERTa is an evolution of BERT with improved training methodology, due to the elimination of the Next Sentence Prediction task, which may fit NLP tasks on tweets considering they are limited in size and self-contained in context. Moreover, by evaluating BERTweet we analyze the performance of a Transformer-based model trained from scratch on tweets.

In this set of experiments, we give an example tweet as input to the transformer model and concatenate its last four layers to be the token representation and the tweet representation is the average of the tokens representation. Next, those representations collected from the whole dataset are given as input to the learning classifier method together with the labels of the tweets. Finally, the learned classifier is employed to perform the evaluation. In this way, we once again follow the feature extraction plus classification strategy but now using the contextualized embedding from each tweet.

Table 5.1 presents the classification results when using the SVM classifier in terms

Table 5.1: Accuracies and F_1 -macro scores (%) achieved by evaluating Transformer-Autoencoder language models using the SVM classifier

Dataset	Accuracy			F_1 -macro		
	RoBERTa	BERT	BERTweet	RoBERTa	BERT	BERTweet
iro	46.67	71.43	69.52	31.0	66.51	61.22
sar	57.86	76.07	61.96	46.23	75.63	59.24
ntu	78.45	87.02	91.03	78.18	86.73	90.86
S15	86.31	90.03	91.59	73.33	77.58	82.06
stm	84.12	89.14	90.25	84.04	89.11	90.23
per	72.66	83.13	83.14	71.18	80.73	81.34
hob	69.15	83.52	83.13	68.62	81.76	81.53
iph	75.96	81.58	83.65	74.65	79.63	81.97
mov	74.35	84.14	86.47	68.61	77.58	80.55
san	83.66	85.54	89.87	83.43	85.47	89.81
Nar	89.73	91.6	95.35	89.48	91.34	95.22
arc	88.18	87.25	90.16	88.0	87.02	89.99
S18	86.28	87.25	88.97	86.07	87.16	88.87
OMD	82.16	85.62	87.36	81.2	84.71	86.4
HCR	76.67	78.61	79.82	72.89	74.75	76.22
STS	89.48	90.46	93.56	88.11	89.16	92.65
SSt	84.01	85.19	86.76	83.83	84.87	86.53
Tar	84.83	85.64	86.93	84.81	85.62	86.92
Vad	87.73	89.63	90.56	86.24	88.18	89.28
S13	84.26	86.62	88.15	82.0	84.21	86.13
S17	90.61	91.54	92.56	90.08	91.03	92.08
S16	87.62	88.77	90.72	85.5	86.59	88.86
#wins	0	3	19	0	3	19
rank sums	65.0	42.0	25.0	65.0	42.0	25.0
position	3.0	2.0	1.0	3.0	2.0	1.0

of accuracy and F_1 -macro, and Table 5.2 shows a summary of the complete evaluation regarding all classifiers. As in previous chapter, to limit the number of tables in the manuscript, we only report the computational results in detail for the LR classifier as an example of this evaluation (refer to Appendix A for the detailed evaluation). From Table 5.1, we can note that BERTweet achieved the best results in 18 out of the 22 datasets for both accuracy and F_1 -macro. Similarly, regarding all classifiers, Table 5.2 shows that BERTweet outperformed BERT and RoBERTa by a significant difference in terms of the total number of wins for both accuracy and F_1 -macro.

Next, we present an overall analysis of using BERT, RoBERTa, and BERTweet models to train each one of the five classification algorithms, examining them as pairs $\{language\ model, classifier\}$. Table 5.3 presents the average rank position across all 15 possible combinations (3 language models \times 5 classification algorithms), from best to worst, as

Table 5.2: Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by evaluating each Transformer-Autoencoder model on the 22 datasets, for each classification algorithm

Embedding	LR	SVM	MLP	RF	XGB	Total
ACCURACY						
BERT	2/45.0/2.0	0/65.0/3.0	3/47.0/2.0	4/46.5/2.0	3/43.5/2.0	12/247.0/11.0
RoBERTa	2/60.0/3.0	3/42.0/2.0	2/57.0/3.0	5/52.5/3.0	0/63.0/3.0	12/274.5/14.0
BERTweet	18/27.0/1.0	19/25.0/1.0	17/28.0/1.0	15/33.0/1.0	20/25.5/1.0	89/138.5/5.0
F_1-MACRO						
BERT	2/45.0/2.0	0/65.0/3.0	3/47.0/2.0	3/48.0/2.0	3/44.0/2.0	11/249.0/11.0
RoBERTa	2/60.0/3.0	3/42.0/2.0	2/57.0/3.0	5/52.5/3.0	0/61.0/3.0	12/272.5/14.0
BERTweet	18/27.0/1.0	19/25.0/1.0	17/28.0/1.0	15/31.5/1.0	19/27.0/1.0	88/138.5/5.0

explained in Chapter 4. We can observe that BERTweet combined with LR, MLP, and SVM classifiers achieved the best overall performances for both accuracy and F_1 -macro. Conversely, using RF to train the Transformer-based embeddings seems to harm the performance of the models.

Tables 5.4 and 5.5 show a summary of the results for each model and classifier, respectively, from best to worst, in terms of the average rank position. From Table 5.4, we can see that BERTweet achieved the best overall classification effectiveness and was ranked first. Also, RoBERTa and BERT achieved comparable overall performances for both accuracy and F_1 -macro. Regarding the classifiers, as shown in Table 5.5, MLP and LR achieved rather comparable performances and were ranked as the top two best classifiers regarding the Transformer-based models, followed by SVM, XGB, and RF.

Regarding the results achieved for each dataset, Table 5.6 presents the best results in terms of accuracy and F_1 -macro. As we can notice, BERTweet outperformed BERT and RoBERTa in 17 out of the 22 datasets in terms of accuracy and in 18 out of the 22 datasets in terms of F_1 -macro. These results may confirm that Twitter sentiment classification benefits most from contextualized language models trained from scratch on Twitter data. Unlike BERT and RoBERTa, which were trained on traditional English texts, BERTweet was trained on a huge amount of 850M tweets. This fact may have helped BERTweet on learning the specificities of tweets, such as their morphological and semantic characteristics.

For a better understanding of the results, we present an analysis of the difference between the vocabulary embedded in the assessed models. For this purpose, Table 5.7 highlights the number of tokens shared between BERT, RoBERTa, and BERTweet. In

Table 5.3: Average rank position results achieved for each combination Model-Classifier

Model	Classifier	Accuracy avg. rank pos.	Model	Classifier	F_1 -macro avg. rank pos.
BERTweet	LR	2.55	BERTweet	LR	2.32
BERTweet	MLP	2.68	BERTweet	MLP	3.00
BERTweet	SVM	4.16	BERTweet	SVM	3.55
RoBERTa	MLP	5.05	RoBERTa	LR	4.98
RoBERTa	LR	5.68	RoBERTa	MLP	5.39
BERT	MLP	6.23	BERT	SVM	5.75
BERT	SVM	6.86	BERT	MLP	6.61
BERTweet	Xgb	7.34	BERT	LR	7.30
BERT	LR	7.73	BERTweet	Xgb	8.36
RoBERTa	Xgb	9.57	RoBERTa	Xgb	10.18
BERT	Xgb	11.61	RoBERTa	SVM	10.80
BERTweet	RF	11.95	BERT	Xgb	11.77
RoBERTa	SVM	12.05	BERTweet	RF	12.48
RoBERTa	RF	12.98	RoBERTa	RF	13.55
BERT	RF	13.57	BERT	RF	13.98

Table 5.4: Average rank position results achieved for each Embedding

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
BERTweet	5.74	BERTweet	5.94
RoBERTa	9.06	RoBERTa	8.98
BERT	9.20	BERT	9.08

Table 5.5: Average rank position results achieved for each Classifier

Classifier	Accuracy avg. rank pos.	Classifier	F_1 -macro avg. rank pos.
MLP	4.65	LR	4.86
LR	5.32	MLP	5.00
SVM	7.69	SVM	6.70
XGB	9.51	XGB	10.11
RF	12.83	RF	13.33

Table 5.6: Best results achieved for each dataset by evaluating the combination of language model and classifier

Dataset	Accuracy	Classifier	Model	F_1 -macro	Classifier	Model
iro	80.48	LR	BERT	73.08	LR	BERT
sar	76.07	SVM	BERT	75.63	SVM	BERT
ntu	91.03	LR	BERTweet	90.86	SVM	BERTweet
S15	92.23	LR	BERTweet	83.33	LR	BERTweet
stm	90.79	LR	RoBERTa	90.75	LR	RoBERTa
per	87.69	LR	BERTweet	85.29	LR	BERTweet
hob	87.93	LR	RoBERTa	86.29	LR	RoBERTa
iph	87.59	MLP	BERTweet	84.72	MLP	BERTweet
mov	89.47	MLP	RoBERTa	82.12	LR	BERTweet
sand	91.17	MLP	BERTweet	91.11	MLP	BERTweet
Nar	95.60	MLP	BERTweet	95.43	MLP	BERTweet
arc	90.74	MLP	BERTweet	90.51	MLP	BERTweet
S18	88.97	SVM	BERTweet	88.87	SVM	BERTweet
OMD	87.36	SVM	BERTweet	86.40	SVM	BERTweet
HCR	81.55	XGB	BERTweet	76.77	LR	BERTweet
STS	93.90	LR	BERTweet	92.98	LR	BERTweet
SSt	86.76	SVM	BERTweet	86.53	SVM	BERTweet
Tar	86.93	SVM	BERTweet	86.92	SVM	BERTweet
Vad	90.80	LR	BERTweet	89.38	LR	BERTweet
S13	89.61	LR	BERTweet	87.37	LR	BERTweet
S17	92.56	SVM	BERTweet	92.08	SVM	BERTweet
S16	91.03	LR	BERTweet	89.05	LR	BERTweet

Table 5.7: Percentage of vocabulary’s tokens of language model in the line that is also in the vocabulary’s tokens of language model in the column.

	BERT $V = 30K$	RoBERTa $V = 50K$	BERTweet $V = 64K$
BERT	–	61	62
RoBERTa	41	–	71
BERTweet	32	55	–

other words, we show the number of tokens (in %) embedded in the models presented in each row that are also included in the models presented in each column, i.e., the intersection between their vocabularies. For example, regarding BERT (first row), we can see that 61% of its tokens can be found on RoBERTa (second column). The information below each model name in the columns refers to their vocabulary size (number of embedded tokens). It is possible to note that only 32% of the 64K tokens from BERTweet vocabulary (i.e., about 20K tokens) can be found in BERT. It means that, when compared to BERT, BERTweet contains about 44K ($64 - 20$) specific tokens extracted from tweets. Similarly, 55% of the tokens embedded in BERTweet (i.e., about 35K tokens) can be found in RoBERTa, meaning that BERTweet holds about 29K ($64 - 35$) specific tokens from tweets that are not included in RoBERTa. As a matter of fact, analyzing the tokens embedded in BERTweet, we find some specific tokens, such as “Awww”, “hahaha”, “broo”, and other internet expressions and slang that social media users often use to express themselves. While creating representations for these tokens is straightforward in BERTweet, BERT and RoBERTa need to do some extras steps. Specifically, when BERT and RoBERTa do not find a token in their vocabularies, they split the token into subtokens until all of them are found. For example, the token “hahaha” would be split into “ha”, “ha”, and “ha” to represent the original token. This analysis points out that this particular vocabulary, combined with a language model that was trained focused on learning the intrinsic structure of tweets, is the responsible for the BERTweet language model’s best performance on tweet sentiment classification.

In this context, regarding RQ2, we believe BERTweet is an effective language modeling technique in distinguishing the sentiment expressed in tweets. Also, regarding the classifiers, in general, MLP and LR seem to be good choices when using Transformer-based models.

Different from static representation, when we used only the embedding layer of the 13

language models, in this chapter, we use the whole language model: the tweet goes from the embedding layer up to the last layer to be transformed in a vector representation. Attempting to understand the benefits from using the whole language model (embedding layer and language model), we compare the predictive performance of Transformer-based models evaluated in this chapter against all the static representations assessed in Chapter 4. Table 5.8 presents the top ten results across all 85 possible combinations of models and classifiers (17 models \times 5 classification algorithms), and Table 5.9 shows an overall evaluation of the models, from best to worst, in terms of the average rank position. In addition, Table 5.10 shows the best results achieved for each dataset.

From Tables 5.8 and 5.9 we can notice that the Transformer-based BERTweet model outperformed all other models and was ranked first in both evaluations. Also, Table 5.9 shows that the Transformed-based models achieved the best overall results against all static models and were ranked as the top three best representations. Furthermore, from Table 5.10, the Transformer-based BERTweet model achieved the best overall classification effectiveness in 16 out of the 22 datasets in terms of accuracy and in 17 out of the 22 datasets in terms of F_1 -macro.

These results point out that learning language model parameters is essential in distinguishing the sentiment expressed in tweets. Static representations may lose a lot of relevant information considering they ignore the diversity of meaning that words may have depending on the context they appear. In contrast, Transformer-based models benefit from learning how to encode the context information of a token in an embedding.

Table 5.8: Top 10 average rank position results achieved for each combination Model-Classifer by evaluating Transformer-Autoencoder model and static embeddings

Model	Classifier	Accuracy avg. rank pos.	Model	Classifier	F_1 -macro avg. rank pos.
BERTweet	LR	6.20	BERTweet	LR	5.91
BERTweet	MLP	6.32	BERTweet	MLP	6.55
RoBERTa	MLP	10.27	BERTweet	SVM	9.14
BERT	MLP	10.43	BERT	SVM	10.00
BERTweet	SVM	10.45	RoBERTa	MLP	10.89
RoBERTa	LR	12.23	RoBERTa	LR	10.91
BERT	LR	12.91	BERT	MLP	10.91
BERT	SVM	13.11	BERT	LR	12.07
BERTweet	XGB	17.02	RoBERTa-static	SVM	17.32
RoBERTa-static	SVM	18.66	W2V-Edin	SVM	18.75

Table 5.9: Average rank position results achieved for each Embedding evaluating Transformer-Autoencoder model and static embedding

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
BERTweet	14.84	BERTweet	17.38
BERT	20.86	BERT	23.42
RoBERTa	23.00	RoBERTa	24.94
Emo2Vec	37.49	Emo2Vec	36.09
RoBERTa-static	39.65	RoBERTa-static	40.37
W2V-Edin	41.50	W2V-Edin	41.54
W2V-GN	42.37	2V-GN	43.02
GloVe-TWT	44.64	GloVe-TWT	43.55
SSWE	46.75	SSWE	44.78
EWE	47.21	EWE	46.17
DeepMoji	48.35	DeepMoji	46.79
TF-IDF	49.09	fastText	48.95
fastText	50.18	GloVe-WP	51.55
BERT-static	52.30	BERT-static	51.67
GloVe-WP	53.01	TF-IDF	53.00
W2V-Araque	56.25	W2V-Araque	55.45
BERTweet-static	63.52	BERTweet-static	62.34

Table 5.10: Best results achieved for each dataset by evaluating Transformer-Autoencoder and static models

Dataset	Accuracy	Classifier	Model	F_1 -macro	Classifier	Model
iro	80.48	LR	BERT	75.87	LR	Emo2Vec
sar	87.50	LR	SSWE	87.19	LR	SSWE
ntu	95.30	MLP	w2v-Edin	95.19	MLP	w2v-Edin
S15	92.23	LR	BERTweet	83.33	LR	BERTweet
stm	90.79	LR	RoBERTa	90.75	LR	RoBERTa
per	87.69	LR	BERTweet	85.29	LR	BERTweet
hob	94.82	MLP	BERT-static	94.05	MLP	BERT-static
iph	87.59	MLP	BERTweet	84.72	MLP	BERTweet
mov	89.47	MLP	RoBERTa	82.12	LR	BERTweet
san	91.17	MLP	BERTweet	91.11	MLP	BERTweet
Nar	95.60	MLP	BERTweet	95.43	MLP	BERTweet
arc	90.74	MLP	BERTweet	90.51	MLP	BERTweet
S18	88.97	SVM	BERTweet	88.87	SVM	BERTweet
OMD	87.36	SVM	BERTweet	86.40	SVM	BERTweet
HCR	81.55	XGB	BERTweet	76.77	LR	BERTweet
STS	93.90	LR	BERTweet	92.98	LR	BERTweet
SST	86.76	SVM	BERTweet	86.53	SVM	BERTweet
Tar	86.93	SVM	BERTweet	86.92	SVM	BERTweet
Vad	90.80	LR	BERTweet	89.38	LR	BERTweet
S13	89.61	LR	BERTweet	87.37	LR	BERTweet
S17	92.56	SVM	BERTweet	92.08	SVM	BERTweet
S16	91.03	LR	BERTweet	89.05	LR	BERTweet

Chapter 6

Fine-tuning Transformer-based Models Using a Large Collection of English Tweets

In this chapter, we aim at performing computational experiments in order to answer the research question RQ3, stated as follows:

RQ3. Can the fine-tuning of Transformer-based autoencoder models using a large set of English tweets improve the sentiment classification performance?

To answer this research question, we evaluate the classification effectiveness of BERT, RoBERTa, and BERTweet language models fine-tuned with tweets from a corpus of 6.7M unlabeled, or generic unlabeled, tweets, as described in Chapter 3.3. Precisely, we use this set of tweets to fine tuning the model weights using the intermediate masked language model task as the training objective with the probability of 15% to (randomly) mask tokens in the input. We also compare the fine-tuning results of such models against those achieved by using the original weights of the Transformer-based models, as presented in Chapter 5, in order to analyze whether the adjustment of the models via fine-tuning improves the predictive performance of the sentiment classification.

In general, the performance of the fine-tuned models is very sensitive to different random seeds [17]. For that reason, all the results presented in this chapter are the average of three executions using different seeds (12,34,56) to account for the sensitivity of the fine-tuning process regarding different seeds [16].

The first part of the experiments reported in this chapter consists in determining whether the predictive performance of the Transformer-based models are affected by the fine-tuning procedure using tweets from corpora of different sizes. For this purpose, in

addition to the entire Edinburgh corpus of 6,657,700 tweets (around 6.7M tweets), we used nine other smaller samples of tweets with different sizes, varying from 500 to 1.5M tweets. Specifically, we generated samples containing 0.5K, 1K, 5K, 10K, 25K, 50K, 250K, 500K, and 1.5M generic unlabeled tweets. In the fine-tuning processes, we performed three training epochs, except for the tuned models with 6.7M tweets, when we used one epoch, as there was a degradation of some models, such as BERTweet. In all fine-tuning process, all layers are unfrozen. Regarding the batch size, we use the available hardware capacity of eight instances per device. We used a learning rate of $5e-5$ with a linear scheduler and Adam optimizer with beta1 equal to 0.9, beta2 , 0.999, and epsilon , $1e-8$. We also use a max gradient of 1 and with no weight decay.

Tables 6.1, 6.2, 6.3, 6.4, 6.5, and 6.6 present the average classification accuracies and F_1 -macro scores, respectively, when fine-tuning the Transformer-based models with the different samples of tweets generated from the Edinburgh corpus. These results were achieved by using the SVM classifier (refer to Appendix A for the detailed evaluation of each classifier). Regarding the variance in performance across the different seeds, the mean and maximum standard deviations are 0.05% and 0.5% for both accuracy and F_1 -macro, respectively.

Note that BERT was most benefited when fine-tuned with samples of 250K tweets (position row), for both accuracy and F_1 -macro. RoBERTa achieved the best overall results when fine-tuned with samples of 1.5M and 250K tweets, in terms of accuracy and F_1 -macro, respectively. On the other hand, BERTweet benefited from smaller samples, achieving higher overall predictive performances when fine-tuned with samples of 25K and 5K tweets in terms of accuracy and F_1 -macro, respectively. This is an expected result as BERTweet is already trained from scratch from tweets. As we are fine-tuning the language model task, BERT and RoBERTa seems to require more samples to accommodate the Twitter-based vocabulary into the weights' model.

Next, we analyze the overall performance of the fine-tuned Transformer-based models for each classification algorithm. Tables 6.7 and 6.8 summarize the results. Regarding the variance across the different seeds, the mean and maximum standard deviations are 0.2% and 0.7% in terms of accuracy, and 0.26% and 0.98% in terms of F_1 -macro.

Interestingly, from Tables 6.7 and 6.8, we can note that when fine-tuning a language model to fit a specific type of text, such as tweets, applying large corpora does not guarantee better predictive performances. Specifically, the best overall results (Total column) were achieved when fine-tuning BERT, RoBERTa, and BERTweet models with

Table 6.1: Average classification accuracies (%) achieved by fine-tuning BERT with different samples of generic unlabeled tweets, using the SVM classifier

Dataset	BERT									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	73.81	74.05	72.38	66.19	76.43	73.57	76.67	61.19	62.07	56.74
sar	68.93	70.18	74.46	67.32	73.04	77.32	74.46	70.36	71.31	63.69
ntu	88.84	86.65	84.18	85.97	83.08	85.62	84.89	92.09	93.88	94.38
S15	90.35	89.4	90.02	90.64	88.78	90.03	87.22	87.23	85.96	87.13
stm	89.14	89.14	88.29	89.41	89.13	88.86	89.13	88.3	88.57	87.01
per	84.72	82.91	85.42	83.14	82.91	82.01	81.1	81.32	79.5	76.16
hob	83.9	85.24	86.39	86.78	86.18	84.85	84.29	83.72	84.15	80.89
iph	81.02	81.39	81.58	82.15	81.02	81.2	82.15	81.21	80.34	78.96
mov	82.72	83.43	84.49	85.2	84.5	85.03	85.91	83.26	82.41	82.71
san	86.19	86.19	85.87	87.01	86.27	87.42	87.66	87.0	86.98	87.6
Nar	91.93	91.44	92.42	91.52	91.6	91.28	92.5	92.83	93.67	93.5
arc	86.85	88.07	88.94	88.19	87.49	88.01	87.95	89.58	88.13	88.5
S18	87.25	86.82	86.45	86.45	86.5	86.61	86.5	86.82	87.09	86.0
OMD	85.68	85.99	85.26	85.15	84.73	84.31	85.79	84.84	85.43	84.58
HCR	78.56	78.82	78.4	78.35	78.35	79.5	78.82	77.98	77.32	76.56
STS	90.36	90.51	90.41	90.81	90.22	91.05	91.79	92.23	91.89	91.29
SSst	84.58	85.06	85.37	84.58	84.01	83.79	84.97	84.8	84.84	84.43
Tar	85.67	85.67	85.93	86.01	85.9	85.84	85.9	85.69	85.69	85.4
Vad	89.75	89.66	89.63	89.44	89.56	90.3	90.28	89.51	90.5	90.0
S13	86.52	86.66	87.64	87.05	87.39	86.96	87.23	86.98	87.07	86.32
S17	91.08	91.11	91.05	91.15	90.61	90.92	90.63	90.18	89.85	89.56
S16	88.45	88.34	89.24	88.83	88.83	88.82	88.9	88.29	88.01	87.93
#wins	1	1	4	5	0	2	3	2	2	1
rank sums	125.5	112.0	101.0	101.5	132.5	117.0	88.0	132.0	129.5	171.0
position	6.0	4.0	2.0	3.0	9.0	5.0	1.0	8.0	7.0	10.0

Table 6.2: Average classification accuracies (%) achieved by fine-tuning RoBERTa with different samples of generic unlabeled tweets, using the SVM classifier

RoBERTa										
Dataset	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	46.67	46.67	46.67	48.1	46.67	46.67	46.67	46.67	45.24	46.74
sar	60.71	59.29	65.0	57.86	62.14	60.71	65.0	62.14	63.57	64.7
ntu	81.69	81.31	79.15	81.34	82.06	85.63	87.08	90.67	91.75	90.41
S15	83.18	83.5	84.75	86.93	84.75	85.38	86.93	88.79	87.24	87.03
stm	85.23	86.35	89.7	85.24	87.75	88.04	88.03	87.48	90.81	86.35
per	69.92	69.68	70.38	68.33	71.74	71.28	72.66	74.26	76.32	77.0
hob	73.76	71.44	73.56	74.32	72.41	77.4	77.38	77.2	77.97	77.82
iph	78.41	78.41	77.46	78.96	78.78	79.71	79.89	80.09	79.15	79.02
mov	71.14	76.84	81.64	80.76	78.44	78.09	79.52	80.4	81.12	81.05
san	84.39	84.23	85.21	85.78	85.13	86.27	85.86	87.17	88.23	86.16
Nar	91.04	92.18	92.58	92.25	91.2	92.58	92.74	92.83	93.07	93.88
arc	88.88	88.53	89.0	87.72	88.48	89.23	88.59	89.64	89.47	88.57
S18	87.9	87.15	87.15	87.09	86.61	87.52	87.14	87.47	86.61	86.62
OMD	82.79	83.58	83.63	84.0	83.53	83.47	83.79	82.27	82.74	82.55
HCR	76.52	77.25	77.2	76.73	76.1	77.78	76.41	77.77	75.94	77.74
STS	89.62	90.71	91.39	90.95	91.84	92.08	92.08	92.52	92.92	92.0
SSt	84.67	86.28	85.93	85.98	85.67	86.06	85.8	86.02	86.02	85.22
Tar	85.43	85.67	86.36	85.78	85.95	85.41	85.52	85.23	85.98	84.33
Vad	88.25	89.56	89.82	89.73	89.3	89.63	89.85	90.4	91.11	89.87
S13	85.59	85.54	86.18	86.57	85.68	85.93	86.23	85.54	87.03	85.89
S17	90.77	91.0	91.37	91.29	91.07	91.15	91.02	90.86	90.63	89.68
S16	88.37	88.4	88.85	89.01	88.96	88.26	88.74	88.96	89.01	87.95
#wins	1	1	3	2	0	1	0	3	7	2
rank sums	174.0	161.0	111.0	124.0	148.0	105.5	101.5	92.0	78.5	114.5
position	10.0	9.0	5.0	7.0	8.0	4.0	3.0	2.0	1.0	6.0

Table 6.3: Average classification accuracies (%) achieved by fine-tuning BERTweet with different samples of generic unlabeled tweets, using the SVM classifier

Dataset	BERTweet									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	68.81	68.1	71.19	73.81	74.76	79.76	71.67	70.48	67.78	66.43
sar	64.82	69.11	67.68	73.21	73.21	68.93	70.36	70.36	72.26	71.79
ntu	89.95	91.4	91.03	92.82	93.17	92.09	93.16	93.16	94.38	92.45
S15	90.34	92.22	90.04	92.53	92.23	91.6	90.03	90.04	89.4	89.09
stm	90.52	91.92	93.04	91.36	91.07	91.9	92.19	93.02	91.45	90.52
per	84.28	84.04	85.64	86.77	86.79	85.41	86.1	85.87	83.14	81.32
hob	85.06	85.25	86.78	86.78	86.78	87.92	86.6	86.77	87.03	86.21
iph	84.04	83.84	83.47	85.35	84.78	84.22	83.47	85.55	82.22	82.34
mov	85.93	87.71	88.6	88.95	88.42	89.49	88.78	88.42	87.3	87.89
san	90.03	90.69	91.42	90.93	91.18	91.67	91.01	91.01	89.46	90.03
Nar	95.93	96.33	96.58	96.41	96.58	96.41	96.58	95.92	95.71	94.87
arc	90.05	90.39	90.63	90.34	91.09	90.45	91.21	90.98	91.12	90.74
S18	89.99	89.78	90.26	90.48	89.89	89.83	89.94	89.03	88.26	87.9
OMD	88.09	88.93	88.99	87.67	88.25	88.51	88.88	88.09	87.84	86.88
HCR	80.24	80.5	81.23	81.18	81.24	80.6	80.61	78.98	78.25	78.09
STS	94.35	95.23	95.08	94.99	94.84	95.13	94.49	94.15	93.97	94.3
SSt	87.64	88.16	88.73	89.6	88.99	88.03	87.51	87.59	86.98	87.68
Tar	87.02	87.83	87.68	87.54	87.68	87.6	87.28	86.79	86.38	85.84
Vad	90.8	92.42	92.56	92.66	92.23	92.59	92.71	92.18	92.11	92.09
S13	89.4	90.0	90.25	89.84	90.0	88.88	89.38	88.85	88.76	88.44
S17	92.75	93.37	93.35	93.49	93.4	92.97	92.75	92.01	91.31	91.16
S16	90.7	91.38	91.98	91.35	91.54	91.36	91.18	90.6	90.38	89.55
#wins	0	2	4	4	2	4	2	1	1	0
rank sums	165.0	116.5	84.5	82.0	70.0	94.5	103.5	135.0	169.0	190.0
position	8.0	6.0	3.0	2.0	1.0	4.0	5.0	7.0	9.0	10.0

Table 6.4: Average F_1 -macro scores (%) achieved by fine-tuning BERT with different samples of unlabeled tweets, using the SVM classifier

Dataset	BERT									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	71.8	59.35	57.57	58.71	55.4	58.55	62.92	49.72	55.0	49.95
sar	56.06	49.58	59.97	58.0	64.49	68.15	57.05	56.54	66.52	57.29
ntu	82.27	83.73	80.3	80.78	80.5	80.67	83.24	90.02	86.99	89.62
S15	65.69	62.12	60.62	65.57	63.82	64.1	64.6	61.86	62.76	66.05
stm	84.92	83.48	84.66	85.18	81.3	83.72	84.61	83.17	82.8	86.29
per	71.38	77.86	76.33	75.95	76.83	75.75	76.14	72.98	71.0	71.03
hob	79.14	80.33	80.63	78.85	83.11	80.32	81.6	80.79	82.09	81.83
iph	78.71	80.53	79.09	78.76	81.91	79.35	76.93	79.86	77.63	74.78
mov	60.89	60.98	64.29	60.29	57.75	58.49	63.56	66.66	60.49	67.71
san	85.75	85.4	84.66	85.03	83.38	85.66	86.45	85.5	85.44	85.46
Nar	88.81	88.71	88.52	89.31	88.58	88.79	89.79	90.29	90.81	91.36
arc	88.63	87.78	87.81	87.74	88.51	87.85	88.48	89.52	88.76	88.76
S18	83.83	84.66	83.8	83.52	83.42	83.31	83.04	83.57	83.31	84.03
OMD	82.77	82.03	80.68	79.83	81.66	79.66	82.59	81.9	81.81	80.56
HCR	72.62	72.99	71.85	72.68	72.31	71.46	72.48	71.03	69.89	68.45
STS	84.29	85.24	84.42	85.58	84.63	85.35	87.2	87.12	86.67	87.81
SSt	82.87	82.27	82.04	81.74	80.77	82.1	82.94	81.86	82.62	81.13
Tar	83.18	83.96	84.22	84.19	83.0	83.9	83.64	83.32	83.59	83.71
Vad	85.11	84.66	85.11	85.59	84.7	85.31	85.3	85.47	85.83	85.85
S13	82.45	82.15	82.67	80.76	81.98	82.26	82.84	82.43	82.35	81.75
S17	89.12	88.76	88.98	88.47	88.66	88.69	88.71	87.77	87.73	87.68
S16	84.65	84.86	84.59	85.0	84.64	84.63	84.64	84.08	83.75	83.64
#wins	3	3	1	1	2	1	3	2	0	6
rank sums	105.5	112.0	129.5	126.0	146.5	131.5	95.5	120.0	128.0	115.5
position	2.0	3.0	8.0	6.0	10.0	9.0	1.0	5.0	7.0	4.0

Table 6.5: Average F_1 -macro scores (%) achieved by fine-tuning RoBERTa with different samples of unlabeled tweets, using the SVM classifier

Dataset	RoBERTa									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	54.77	61.69	47.99	58.14	57.59	67.22	42.02	54.39	53.42	54.93
sar	53.42	56.62	59.94	73.89	64.55	54.6	68.13	49.36	50.89	67.67
ntu	83.63	83.39	75.86	81.64	79.9	84.13	88.92	88.63	91.87	87.65
S15	67.46	63.75	74.11	66.24	62.03	73.59	67.08	69.26	61.61	66.92
stm	86.32	86.0	85.21	85.77	87.43	86.61	86.58	84.32	83.77	83.63
per	72.78	73.45	75.6	75.1	76.31	74.48	77.13	74.92	74.8	76.17
hob	76.79	79.35	77.42	76.73	80.1	79.58	80.82	80.22	78.26	82.06
iph	77.2	79.94	80.96	82.27	81.08	79.67	75.66	79.93	78.21	78.86
mov	64.06	70.45	73.38	66.84	68.1	66.22	69.41	72.07	68.51	68.91
san	86.56	87.22	86.8	88.14	86.99	87.8	86.24	87.12	87.73	86.33
Nar	89.57	91.0	92.06	90.61	90.4	91.92	91.4	91.76	91.21	91.42
arc	89.53	88.65	89.29	89.06	89.84	89.56	90.22	89.88	88.94	88.84
S18	85.46	87.06	85.68	86.27	86.86	86.33	86.7	85.68	84.43	84.61
OMD	83.23	83.07	84.4	83.69	83.75	82.83	83.73	81.68	80.76	81.73
HCR	71.23	74.87	72.89	72.13	72.09	73.59	72.17	73.06	71.71	70.58
STS	87.5	88.8	88.7	88.46	88.84	89.61	89.71	89.14	88.95	88.6
SSt	83.38	83.23	84.71	83.78	84.34	84.21	84.43	84.5	84.8	82.36
Tar	83.84	84.13	85.14	84.47	85.0	84.51	83.78	83.49	84.79	83.27
Vad	85.85	85.72	86.89	86.75	86.44	87.19	88.04	88.13	87.98	86.75
S13	82.03	82.75	83.57	83.82	83.63	83.73	84.38	83.54	83.73	82.31
S17	89.72	89.98	90.07	90.17	89.69	90.2	90.21	90.01	89.31	87.6
S16	85.29	85.99	86.44	86.27	86.59	86.3	86.31	86.06	86.4	84.46
#wins	0	2	5	3	2	1	5	1	2	1
rank sums	171.0	132.0	100.5	118.5	105.0	92.5	87.0	110.5	136.5	156.5
position	10.0	7.0	3.0	6.0	4.0	2.0	1.0	5.0	8.0	9.0

Table 6.6: Average F_1 -macro scores (%) achieved by fine-tuning BERTweet with different samples of unlabeled tweets, using the SVM classifier

Dataset	BERTweet									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	54.36	61.9	53.1	49.09	51.4	53.28	62.24	66.75	53.0	42.85
sar	49.87	53.96	59.51	53.78	62.33	62.79	54.39	55.69	56.94	55.57
ntu	83.16	89.22	85.29	87.38	85.11	87.31	85.77	88.22	87.64	86.64
S15	71.56	77.95	73.8	69.48	73.47	70.08	63.77	63.21	61.07	58.37
stm	85.2	85.99	86.81	85.72	85.73	88.27	85.71	87.71	87.42	86.28
per	82.43	81.37	77.92	81.86	78.86	77.46	77.25	78.55	77.76	74.77
hob	77.56	84.41	81.24	81.87	80.0	81.5	79.35	80.79	80.85	78.1
iph	81.95	82.07	84.28	80.85	80.71	80.71	81.08	82.04	78.76	81.09
mov	71.91	73.18	76.03	75.34	73.91	78.82	70.84	72.74	73.3	70.2
san	88.62	89.79	90.55	90.87	88.88	88.46	88.7	88.29	88.19	86.71
Nar	91.39	94.35	93.48	93.12	92.72	92.78	92.88	92.83	92.26	91.6
arc	89.52	90.66	91.59	90.76	90.81	90.04	90.35	90.4	89.33	90.62
S18	87.25	88.09	88.26	87.55	87.99	87.03	87.35	85.32	85.07	84.41
OMD	83.01	83.39	85.56	84.26	83.75	83.14	85.0	83.56	82.72	82.69
HCR	73.08	72.69	74.28	74.31	73.16	73.97	74.0	73.57	70.16	69.09
STS	89.3	90.18	90.85	90.69	90.47	89.82	89.35	89.98	88.91	88.98
SSt	84.74	85.88	86.41	86.7	84.54	84.48	83.79	85.15	83.59	84.18
Tar	85.13	85.54	85.8	85.31	85.34	84.3	84.21	84.64	83.67	82.8
Vad	85.68	87.98	88.26	87.9	87.32	87.14	87.89	87.45	87.55	86.75
S13	83.74	84.73	85.43	84.33	83.94	83.48	83.49	83.42	82.51	81.36
S17	90.76	91.27	91.79	91.77	91.71	90.88	90.74	90.09	88.76	88.9
S16	87.16	88.07	88.43	88.06	87.94	87.26	87.4	86.3	86.11	85.17
#wins	1	4	10	3	0	3	0	1	0	0
rank sums	153.0	74.0	53.0	83.0	108.5	122.5	136.0	121.0	169.0	190.0
position	8.0	2.0	1.0	3.0	4.0	6.0	7.0	5.0	9.0	10.0

Table 6.7: Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by each classifier when fine-tuning the Transformer-Autoencoder models with different samples of unlabeled tweets in terms of accuracy

ACCURACY						
Sample	LR	SVM	MLP	RF	XGB	Total
BERT						
0.5k	3/127.5/5.5	1/125.5/6.0	1/129.0/7.0	1/153.0/10.0	5 /108.0/2.0	11/643.0/30.5
1k	0/134.0/7.0	1/112.0/4.0	1/139.5/8.0	3/114.0/3.5	2/113.5/3.0	7/613.0/25.5
5k	3/115.0/4.0	4/101.0/2.0	3/114.5/5.0	1/120.5/6.0	1/128.0/8.0	12/579.0/25.0
10k	0/143.5/9.0	5 /101.5/3.0	1/142.0/10.0	4 /119.0/5.0	2/125.5/7.0	12/631.5/34.0
25k	0/136.0/8.0	0/132.5/9.0	0/141.0/9.0	1/134.0/8.0	2/146.0/10.0	3/689.5/44.0
50k	0/146.0/10.0	2/117.0/5.0	2/121.5/6.0	3/113.5/2.0	0/131.5/9.0	7/629.5/32.0
250k	0/127.5/5.5	3/88.0/1.0	3/101.5/1.0	2/73.5/1.0	2/97.5/1.0	10/ 488.0/9.5
500k	1/110.5/3.0	2/132.0/8.0	2/108.0/4.0	1/114.0/3.5	2/119.5/5.0	8/584.0/23.5
1.5M	4/96.0/2.0	2/129.5/7.0	3/107.0/3.0	1/131.5/7.0	1/122.0/6.0	11/586.0/25.0
6.7M	10/74.0/1.0	1/171.0/10.0	6 /106.0/2.0	4 /137.0/9.0	5 /118.5/4.0	26 /606.5/26.0
RoBERTa						
0.5k	1/140.0/9.0	1/174.0/10.0	0/165.5/9.0	0/171.5/9.0	0/173.0/10.0	2/824.0/47.0
1k	2/137.0/8.0	1/161.0/9.0	2/143.0/8.0	0/165.0/8.0	2/130.5/7.0	7/736.5/40.0
5k	3/92.0/1.0	3/111.0/5.0	0/99.5/3.5	4/104.0/5.0	4/100.0/ 4.0	14/506.5/18.5
10k	1/125.0/7.0	2/124.0/7.0	4 /111.5/6.0	1/120.0/7.0	3/118.5/6.0	11/599.0/33.0
25k	4 /103.0/3.0	0/148.0/8.0	2/107.5/5.0	1/104.5/6.0	2/98.0/3.0	9/561.0/25.0
50k	4 /100.5/2.0	1/105.5/4.0	3/85.5/1.0	4/85.0/2.0	1/97.0/2.0	13/ 473.5/11.0
250k	0/124.0/6.0	0/101.5/3.0	3/131.5/7.0	8/77.5/1.0	4/88.5/1.0	15/523.0/18.0
500k	3/113.5/5.0	3/92.0/2.0	2/99.5/3.5	3/100.5/4.0	2/109.5/5.0	13/515.0/19.5
1.5M	3/109.0/4.0	7/78.5/1.0	3/91.5/2.0	1/98.5/3.0	2/136.0/8.0	16 /513.5/18.0
6.7M	0/166.0/10.0	2/114.5/6.0	2/175.0/10.0	0/183.5/10.0	1/159.0/9.0	5/798.0/45.0
BERTweet						
0.5k	1/143.0/7.0	0/165.0/8.0	0/167.0/9.0	0/174.5/8.0	0/152.5/8.0	1/802.0/40.0
1k	5 /78.5/2.0	2/116.5/6.0	3/95.0/4.0	0/99.0/4.0	4/76.5/2.0	14/465.5/18.0
5k	5/69.5/1.0	4 /84.5/3.0	3/75.5/1.0	10/42.5/1.0	10/56.0/1.0	32/328.0/7.0
10k	2/92.0/3.0	4 /82.0/2.0	4/95.5/5.0	2/72.5/3.0	4/81.0/3.0	16/423.0/16.0
25k	1/95.0/4.0	2/70.0/1.0	3/78.5/2.0	5/71.0/2.0	0/112.0/4.0	11/426.5/13.0
50k	2/110.0/5.0	4/94.5/4.0	6 /91.0/3.0	2/117.0/5.5	3/119.5/6.0	17/532.0/23.5
250k	2/114.5/6.0	2/103.5/5.0	0/128.0/6.0	0/117.0/5.5	0/138.0/7.0	4/601.0/29.5
500k	0/162.0/8.0	1/135.0/7.0	0/150.0/7.0	2/138.0/7.0	1/117.0/5.0	4/702.0/34.0
1.5M	0/172.5/9.0	1/169.0/9.0	0/174.0/10.0	0/176.0/9.0	0/171.0/9.0	1/862.5/46.0
6.7M	1/173.0/10.0	0/190.0/10.0	3/155.5/8.0	0/202.5/10.0	0/186.5/10.0	4/907.5/48.0

samples of 250K, 50K, and 5K tweets, respectively, for both accuracy and F_1 -macro.

Regarding the results achieved for each dataset, Table 6.9 shows the best predictive performances in terms of accuracy and F_1 -macro. We can see that BERTweet achieved the best results for most datasets when fine-tuned with fewer number of tweets. More specifically, BERTweet outperformed the other models when fine-tuned with samples varying from 1K to 25K tweets in 14 out of the 22 datasets for both accuracy and F_1 -macro.

As in previous chapters, we also present an overall evaluation of combining all fine-tuned models and classifiers across the 22 datasets, in terms of the average rank position. Table 6.10 shows the top ten results among all 150 possible combinations ($3 \text{ models} \times 10$

Table 6.8: Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by each classifier when fine-tuning the Transformer-Autoencoder models with different samples of unlabeled tweets in terms of F_1 -macro

F_1 -MACRO						
Sample	LR	SVM	MLP	RF	XGB	Total
BERT						
0.5k	3/128.0/6.0	1/127.0/6.0	1/132.0/7.0	0/155.0/10.0	3/105.5/2.0	8/647.5/31.0
1k	1/140.5/8.0	1/113.0/4.0	1/141.0/9.0	3/113.5/3.0	3/112.0/3.0	9/620.0/27.0
5k	2/120.0/4.0	4/99.5/2.0	3/115.5/5.0	1/118.5/5.0	1/129.5/8.0	11/583.0/24.0
10k	0/144.5/9.0	5 /104.5/3.0	1/146.5/10.0	1/124.5/6.0	1/126.0/6.0	8/646.0/34.0
25k	0/140.0/7.0	0/131.0/8.0	0/139.0/8.0	1/136.5/8.0	2/146.5/10.0	3/693.0/41.0
50k	0/148.5/10.0	2/115.0/5.0	2/119.0/6.0	4/106.0/2.0	1/131.5/9.0	9/620.0/32.0
250k	0/122.0/5.0	4/ 92.0/1.0	3/ 96.5/1.0	6/68.0/1.0	3/ 95.5/1.0	16/ 474.0/9.0
500k	1/108.0/3.0	2/132.0/9.0	2/113.5/4.0	1/116.0/4.0	2/120.0/5.0	8/589.5/25.0
1.5M	4/87.5/2.0	2/129.5/7.0	3/104.0/3.0	1/135.0/7.0	0/128.0/7.0	10/584.0/26.0
6.7M	11/71.0/1.0	1/166.5/10.0	6 /103.0/2.0	3/137.0/9.0	6 /115.5/4.0	27 /593.0/26.0
RoBERTa						
0.5k	1/140.0/9.0	1/175.0/10.0	0/163.0/9.0	0/170.5/9.0	0/171.0/10.0	2/819.5/47.0
1k	3/132.5/8.0	1/160.0/9.0	2/142.0/8.0	0/165.0/8.0	2/132.0/7.0	8/731.5/40.0
5k	3/ 90.0/1.0	3/114.0/5.0	0/102.5/3.0	4/105.5/5.0	5 /100.5/3.0	15/512.5/17.0
10k	1/125.5/6.0	1/129.0/7.0	3/112.0/6.0	1/118.5/7.0	3/118.5/6.0	9/603.5/32.0
25k	3/103.5/3.0	0/151.0/8.0	2/109.0/5.0	1/108.5/6.0	2/105.0/4.0	8/577.0/26.0
50k	4 /99.0/2.0	1/106.5/4.0	4/85.5/1.0	4/83.0/2.0	1/92.5/2.0	14/ 466.5/11.0
250k	0/128.0/7.0	0/98.5/3.0	3/126.0/7.0	8/71.0/1.0	5/87.0/1.0	16/510.5/19.0
500k	2/115.0/5.0	3/88.0/2.0	1/104.0/4.0	4/102.0/4.0	1/110.5/5.0	11/519.5/20.0
1.5M	4/108.5/4.0	8/72.0/1.0	4/87.0/2.0	0/101.0/3.0	2/136.5/8.0	18 /505.0/18.0
6.7M	0/168.0/10.0	4/116.0/6.0	2/179.0/10.0	0/185.0/10.0	1/156.5/9.0	7/804.5/45.0
BERTweet						
0.5k	1/142.0/7.0	0/166.5/8.0	0/169.0/9.0	0/174.0/8.0	1/153.0/8.0	2/804.5/40.0
1k	7 /79.5/2.0	2/112.0/6.0	3/89.0/3.0	0/99.5/4.0	4/74.0/2.0	16/454.0/17.0
5k	4/ 71.0/1.0	5 /80.0/3.0	3/ 74.0/1.0	12/39.5/1.0	10/53.0/1.0	34/317.5/7.0
10k	3/89.5/3.0	5 /79.0/2.0	4/95.0/5.0	1/76.5/3.0	3/83.0/3.0	16/423.0/16.0
25k	1/94.0/4.0	1/ 73.5/1.0	3/77.0/2.0	5/71.0/2.0	0/108.5/4.0	10/424.0/13.0
50k	2/110.5/5.0	4/96.5/4.0	6/94.0/4.0	2/117.0/ 6.0	3/122.5/6.0	17/540.5/25.0
250k	2/116.5/6.0	1/106.0/5.0	0/126.5/6.0	0/116.0/5.0	0/136.0/7.0	3/601.0/29.0
500k	0/163.0/8.0	0/135.5/7.0	1/151.5/7.0	2/135.5/7.0	1/121.0/5.0	4/706.5/34.0
1.5M	0/173.0/10.0	1/170.0/9.0	0/175.0/10.0	0/178.0/9.0	0/169.0/9.0	1/865.0/47.0
6.7M	2/171.0/9.0	0/191.0/10.0	2/159.0/8.0	0/203.0/10.0	0/190.0/10.0	4/914.0/47.0

Table 6.9: Best results achieved for each dataset by fine-tuning the Transformer-based models with different samples of generic tweets

Dataset	Accuracy	Classifier	Model	F_1 -macro	Classifier	Model
iro	82.30	MLP	BERTweet-50K	75.87	LR	BERT-500K
sar	77.32	SVM	BERT-50K	75.85	SVM	BERT-50K
ntu	94.38	SVM	BERTweet-1.5M	94.26	SVM	BERTweet-1.5M
S15	94.18	MLP	BERTweet-1K	86.22	MLP	BERTweet-1K
stm	93.04	SVM	BERTweet-5K	93.02	SVM	BERTweet-5K
per	89.51	LR	BERTweet-10K	87.53	LR	BERTweet-10K
hob	89.83	MLP	BERTweet-50K	88.30	MLP	BERTweet-50K
iph	88.16	MLP	RoBERTa-25K	85.85	MLP	RoBERTa-25K
mov	93.29	MLP	BERTweet-50K	88.27	LR	BERTweet-50K
san	91.83	LR	BERTweet-10K	91.77	LR	BERTweet-10K
Nar	97.04	MLP	BERTweet-1K	96.91	MLP	BERTweet-1K
arc	92.08	LR	BERTweet-25K	91.92	LR	BERTweet-25K
S18	90.48	SVM	BERTweet-10K	90.40	SVM	BERTweet-10K
OMD	88.99	SVM	BERTweet-5K	88.17	SVM	BERTweet-5K
HCR	82.18	XGB	RoBERTa-1K	78.18	LR	BERTweet-250K
STS	95.38	MLP	BERTweet-50K	94.59	MLP	BERTweet-50K
SSt	89.60	SVM	BERTweet-10K	89.36	SVM	BERTweet-10K
Tar	87.83	SVM	BERTweet-1K	87.82	SVM	BERTweet-1K
Vad	92.80	LR	BERTweet-1K	91.64	LR	BERTweet-1K
S13	90.70	LR	BERTweet-5K	88.59	LR	BERTweet-5K
S17	93.49	SVM	BERTweet-10K	93.07	SVM	BERTweet-10K
S16	91.98	SVM	BERTweet-5K	90.30	SVM	BERTweet-5K

Table 6.10: Top 10 average rank position results achieved for each combination Model-Classifer by evaluating Transformer-Autoencoder models

Model	Classifier	Accuracy avg. rank pos.	Model	Classifier	F_1 -macro avg. rank pos.
BERTweet-5K	LR	11.95	BERTweet-5K	LR	11.18
BERTweet-5K	MLP	14.05	BERTweet-25K	SVM	13.43
BERTweet-25K	MLP	14.64	BERTweet-10K	SVM	13.95
BERTweet-25K	LR	16.14	BERTweet-10K	LR	14.20
BERTweet-50K	MLP	16.43	BERTweet-25K	LR	14.32
BERTweet-1K	MLP	16.77	BERTweet-1K	LR	15.11
BERTweet-10K	LR	16.82	BERTweet-5K	MLP	15.95
BERTweet-25K	SVM	17.02	BERTweet-25K	MLP	16.11
BERTweet-1K	LR	17.68	BERTweet-50K	LR	16.80
BERTweet-10K	SVM	17.93	BERTweet-50K	SVM	17.43

samples of tweets \times 5 classification algorithms). As we can see in Table 6.10, fine-tuned BERTweet embeddings achieved the best overall performances when used to train LR, MLP, and SVM, mastering the top ten results. Also, note that by using LR, MLP, and SVM, BERTweet outperformed all other models when fine-tuned with samples containing 50K tweets or less.

Tables 6.11 and 6.12 show the top ten results among all fine-tuned models and a summary of the results for each classifier, from best to worst, respectively, in terms of the average rank position. From Table 6.11, we can notice that all BERTweet fine-tuned models (0.5K, 1K, 5K, 10K, 25K, 50K, 250K, 500K, 1.5M, and 6.7M) were ranked in the top ten results. Furthermore, neither BERT nor RoBERTa appear in the top results, even when they are fine-tuned with the entire corpus of 6.7M tweets. RoBERTa appears only in the top 24 accuracy score with an average rank of 37.02 tuned with 50K tweets and combined MLP classifier and in top 28 F_1 -macro score with an average rank of 37.27 tuned with 50K tweets and combined LR classifier. BERT appears only in the top 56 accuracy score with an average rank of 66.05 tuned with 1.5M tweets and combined MLP classifier and in top 51 F_1 -macro score with an average rank of 60.77 tuned with 6.7M tweets and combined LR classifier. Among the classifiers, as we can see in Table 6.12, MLP and LR achieved the best predictive performances and were ranked as the top two best classifiers. Conversely, RF was ranked as the worst classifier.

From all previous evaluations, we can note that as the size of the samples increases, the fine-tuning procedure seems to be less effective. It may be due to the adjustment of

Table 6.11: Top 10 average rank position results achieved for each Embedding evaluating Transformer-Autoencoder model

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
BERTweet-5K	37.76	BERTweet-5K	40.13
BERTweet-10K	40.81	BERTweet-10K	43.27
BERTweet-25K	41.24	BERTweet-25K	43.80
BERTweet-1K	43.09	BERTweet-1K	44.78
BERTweet-50K	45.22	BERTweet-50K	47.42
BERTweet-250K	48.65	BERTweet-250K	50.19
BERTweet-500K	53.96	BERTweet-500K	55.59
BERTweet-0.5K	58.00	BERTweet-0.5K	59.59
BERTweet-1.5M	66.63	BERTweet-1.5M	66.48
BERTweet-6.7M	71.09	BERTweet-6.7M	70.46

Table 6.12: Average rank position results achieved for each Classifier evaluating Transformer-Autoencoder model

Classifier	Accuracy avg. rank pos.	Classifier	F_1 -macro avg. rank pos.
MLP	48.52	LR	48.71
LR	53.17	MLP	49.92
SVM	70.74	SVM	60.83
XGB	84.90	XGB	93.16
RF	120.18	RF	124.87

the weights of the models' layers during the back-propagation process. Considering that the fine-tuning procedure consists in unfreezing the entire model obtained previously and adjusting their weights with the new data, the original model and the semantic and syntactic knowledge learned in its layers are changed. In that case, we believe that after some training iterations, the adjustment of the weights starts to damage the original knowledge embedded in the models' layers. The aforementioned conclusion may further explain why BERTweet achieved improved classification performance by using smaller samples of tweets as compared to BERT and RoBERTa. Our hypothesis is that, considering that the weights in BERTweet's layers are specifically adjusted to fit tweets' language style, using more data to fine-tune the model means only continue the initial training. It may be that lots of data may harm the learned weights of the model. Thus, we suggest that when fine-tuning Transformer-based models, such as BERT, RoBERTa, and BERTweet, samples of different sizes may be exploited instead of adopting a dataset with a massive number of instances.

Additionally, we present a comparison among all fine-tuned Transformer-based models against their original versions. Tables 6.13, 6.14, and 6.15 report this comparison in terms of the average rank position for BERT, RoBERTa, and BERTweet, respectively. We can see that the fine-tuned versions achieved meaningful predictive performances as compared to their original models, which indicates that fine-tuning strategies can boost classification performance in Twitter sentiment analysis. Moreover, from Tables 6.13 and 6.14, we note that the fine-tuned versions of BERT and RoBERTa benefited most from samples containing a large number of tweets. Conversely, as pointed out before, BERTweet achieved better overall performances by using smaller samples, as shown in Table 6.15.

Addressing research question RQ3, we could see that fine-tuning Transformer-based models improves the classification effectiveness in Twitter sentiment analysis. Nevertheless, using large sets of tweets does not guarantee better predictive performances, particularly for those models trained from scratch on tweets, such as BERTweet. We could observe that BERTweet benefited most from samples of tweets containing 50K tweets or less. Furthermore, regarding the classifiers, in general, MLP and LR seem to be good choices of classifiers to be employed after extracting the features from fine-tuned Transformer-based models.

Table 6.13: Average rank position results achieved for BERT model and its tuned models

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
BERT-250K	25.62	BERT-250K	25.62
BERT-5K	26.95	BERT-1.5M	26.45
BERT-1.5M	26.96	BERT-6.7M	26.69
BERT-500K	27.09	BERT-500K	26.70
BERT-6.7M	27.67	BERT-5K	27.69
BERT-0.5K	28.16	BERT-50K	28.36
BERT-50K	28.38	BERT-0.5K	28.40
BERT-1K	28.46	BERT-1K	28.95
BERT-10K	29.52	BERT (original)	29.50
BERT (original)	29.52	BERT-10K	29.68
BERT-25K	29.66	BERT-25K	29.95

Table 6.14: Average rank position results achieved for RoBERTa model and its tuned models

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
RoBERTa-50K	24.34	RoBERTa-50K	24.24
RoBERTa-500K	24.69	RoBERTa-1.5M	24.61
RoBERTa-1.5M	24.82	RoBERTa-500K	24.95
RoBERTa-5K	25.44	RoBERTa-5K	25.54
RoBERTa-250K	25.53	RoBERTa-250K	25.66
RoBERTa-25K	26.49	RoBERTa-25K	27.05
RoBERTa-10K	27.28	RoBERTa-10K	27.50
RoBERTa-1K	29.84	RoBERTa-1K	29.65
RoBERTa-0.5K	32.01	RoBERTa-0.5K	31.78
RoBERTa-6.7M	32.75	RoBERTa-6.7M	31.96
RoBERTa (original)	34.81	RoBERTa (original)	35.06

Table 6.15: Average rank position results achieved for BERTweet model and its tuned models

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
BERTweet-5K	20.74	BERTweet-5K	21.38
BERTweet-25K	22.62	BERTweet-25K	22.94
BERTweet-10K	22.85	BERTweet-10K	23.10
BERTweet-1K	23.73	BERTweet-1K	23.93
BERTweet-50K	25.25	BERTweet-50K	25.55
BERTweet-250K	26.67	BERTweet-250K	26.66
BERTweet-500K	30.31	BERTweet-500K	30.48
BERTweet-0.5K	31.70	BERTweet-0.5K	31.72
BERTweet (original)	33.80	BERTweet (original)	33.35
BERTweet-1.5M	34.80	BERTweet-1.5M	34.05
BERTweet-6.7M	35.53	BERTweet-6.7M	34.85

Chapter 7

Fine-tuning Transformer-based Models Using Sentiment Datasets

The experiments conducted in this chapter aim at answering the research question RQ4, stated as follows:

RQ4. Can Transformer-based autoencoder models benefit from a fine-tuning procedure with tweets from sentiment analysis datasets?

We address this research question by evaluating whether the sentiment classification of tweets benefits from fine-tuned language models using tweets from sentiment analysis datasets. For this purpose, we use the same collection of 22 benchmark datasets presented in Chapter 3 (Table 3.1). We perform this evaluation by assessing three distinct strategies to simulate three real-world scenarios. In addition, as done in Chapter 6, all experiments were performed three times using different seeds (12,34,56), with all the same hyperparameter and we report the average of the results.

The first fine-tuning strategy we investigate, referred to as *InData*, simulates the usage of a specific sentiment dataset itself as the new domain dataset to fine-tune a pre-trained language model. Precisely, each one of the 22 datasets is used once as the target dataset. For each of the 22 datasets, we use a 10-fold cross-validation procedure. In each of the ten executions, we use the tweets from nine folds as the source data (i.e., the training data) used to adjust a language model, which is then validated on the one remaining part of the data, referred to as the target dataset (i.e., the test data).

The second strategy, referred to as *LOO* (Leave One dataset Out), aims at simulating the situation where a collection of general sentiment datasets is available to fine-tune the language model. We use each dataset once as the target dataset while the tweets from

the remaining 21 datasets are combined to tune the language model. Although the target dataset contains sentiment labels for each tweet, these labels are not used in the fine-tuning process as we leverage the intermediate self-supervised masked language model task to fine-tune the network parameters.

The third and last strategy, referred to as *AllData*, is a combination of the two others. Specifically, as for strategy *InData*, for each assessed dataset (target dataset), and for each of the nine folds in the 10-fold cross-validation procedure, we combine the tweets from the nine folds (i.e., the training data of the target dataset) with the tweets from the remaining 21 datasets to fine-tune a language model. This last strategy evaluates the benefits of combining the tweets from a specific sentiment target dataset with a representative general sentiment dataset corpus in the fine-tuning process.

Table 7.1, 7.2, and 7.3 present the predictive performances achieved by fine-tuning each language model with strategies *InData*, *LOO*, and *AllData*, one at a time, by using the SVM classifier. As in previous chapters, for space constraints, we only report the detailed evaluation using the SVM classifier (refer to Appendix A for the detailed assessment of each classifier).

From Tables 7.1, 7.2, and 7.3, we can observe that BERT benefited most from strategy *InData*, which uses only the target dataset itself to adjust the language models. Conversely, fine-tuning RoBERTa and BERTweet models using strategies that combine tweets from distinct sentiment analysis corpora achieved the best results for most datasets. More clearly, *AllData*, which combines the tweets from the target dataset and tweets from a collection of sentiment datasets, achieved the best overall results with both RoBERTa and BERTweet. Also, regarding BERTweet, note that strategy *LOO* achieved comparable performances to *AllData*. It is also noteworthy that smaller datasets seem to have benefited most from fine-tuning RoBERTa and BERTweet by using strategy *LOO*. On the other hand, larger datasets achieved higher predictive performances when using strategy *AllData* to fine-tune RoBERTa and BERTweet. Table 7.4 shows a summary of the complete evaluation regarding all classifiers.

Regarding the overall results achieved for each dataset, Table 7.5 presents the best results. We can note that when fine-tuning the Transformer-based models with tweets from sentiment datasets, BERTweet outperformed BERT and RoBERTa for all datasets, except for datasets sarcasm (*sar*) and hobbit (*hob*). Interestingly, as mentioned before, while strategy *LOO* achieved the best results for smaller datasets, larger datasets seem to benefit from strategy *AllData*. Precisely, strategy *AllData* achieved the best overall

Table 7.1: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the SVM classifier and BERT

	Accuracy			F_1 -macro		
Dataset	BERT					
	AllData	LOO	InData	AllData	LOO	InData
iro	74.40	78.81	67.90	65.40	70.55	59.60
sar	71.0	70.18	64.10	68.50	68.58	60.20
ntu	85.00	82.74	88.10	84.70	82.39	87.80
S15	89.70	88.14	89.80	77.50	77.11	77.80
stm	88.80	90.25	89.90	88.70	90.24	89.80
per	84.40	85.66	82.00	82.20	83.49	80.00
hob	84.60	84.46	82.30	82.90	83.08	80.70
iph	82.70	83.09	83.00	80.80	81.07	81.40
mov	85.80	86.46	84.60	79.40	80.14	78.10
san	87.60	87.49	87.80	87.50	87.43	87.70
Nar	92.20	92.50	94.90	92.00	92.25	94.70
arc	89.10	88.42	89.90	88.90	88.20	89.80
S18	87.70	87.36	89.70	87.60	87.26	89.60
OMD	85.90	85.73	87.30	85.00	84.74	86.40
HCR	79.30	79.03	79.60	75.70	75.25	75.90
STS	91.70	90.71	93.50	90.50	89.35	92.60
SSt	84.70	84.71	87.50	84.30	84.39	87.20
Tar	85.70	86.24	86.90	85.70	86.23	86.90
Vad	89.90	90.16	91.50	88.50	88.84	90.30
S13	87.50	87.60	88.70	85.20	85.31	86.60
S17	91.80	91.56	92.90	91.30	91.04	92.40
S16	89.50	89.07	90.70	87.40	86.93	88.80
#wins	2	5	15	0	6	16
rank sums	49.0	49.0	34.0	51.0	48.0	33.0
position	2.5	2.5	1.0	3.0	2.0	1.0

Table 7.2: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the SVM classifier and RoBERTa

Dataset	RoBERTa					
	AllData	LOO	InData	AllData	LOO	InData
iro	46.70	46.67	46.70	31.00	31.00	31.00
sar	64.00	65.00	64.00	52.90	53.49	54.20
ntu	84.30	83.48	81.20	83.90	83.07	80.80
S15	87.20	86.31	86.20	74.90	74.28	72.40
stm	90.00	90.25	87.60	89.90	90.21	87.60
per	71.10	70.38	65.50	70.20	69.17	64.90
hob	72.80	73.94	71.30	71.90	73.10	70.50
iph	79.90	78.96	78.60	78.60	77.72	77.20
mov	81.50	79.87	72.10	75.30	74.01	66.50
san	87.50	87.17	85.50	87.40	87.04	85.30
Nar	93.10	92.50	92.10	92.90	92.35	91.90
arc	89.40	89.47	89.00	89.30	89.27	88.70
S18	88.40	88.54	88.00	88.30	88.44	87.80
OMD	85.60	84.58	85.70	84.50	83.60	84.70
HCR	76.90	76.04	78.10	73.20	72.59	74.20
STS	92.60	92.13	91.50	91.60	90.97	90.30
SSt	86.40	85.63	85.90	86.10	85.41	85.70
Tar	85.90	85.87	86.30	85.90	85.85	86.30
Vad	89.80	89.37	89.90	88.60	88.10	88.60
S13	86.60	86.41	86.10	84.60	84.41	83.90
S17	92.00	91.71	91.60	91.50	91.21	91.10
S16	89.50	89.71	89.30	87.60	87.75	87.40
#wins	12	6	5	13	4	5
rank sums	33.0	44.0	55.0	32.5	45.0	54.5
position	1.0	2.0	3.0	1.0	2.0	3.0

Table 7.3: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the SVM classifier and BERTweet

Dataset	BERTweet					
	AllData	LOO	InData	AllData	LOO	InData
iro	74.60	83.10	66.70	66.50	77.24	60.10
sar	68.60	67.50	61.80	65.50	64.32	56.40
ntu	92.10	93.54	90.10	91.80	93.33	89.80
S15	90.70	92.84	90.20	80.00	84.76	78.60
stm	92.70	92.75	90.50	92.60	92.73	90.50
per	86.10	86.55	82.50	84.10	84.48	80.70
hob	87.10	87.17	82.50	85.60	85.62	80.80
iph	85.10	83.48	83.30	83.50	81.79	81.80
mov	89.90	88.42	87.00	84.50	81.99	80.90
san	91.40	91.34	89.00	91.40	91.27	88.90
Nar	97.00	96.66	96.20	96.80	96.54	96.00
arc	91.40	90.57	90.70	91.30	90.40	90.50
S18	90.90	90.26	90.60	90.80	90.19	90.50
OMD	89.20	89.77	88.40	88.40	88.99	87.50
HCR	81.50	81.27	80.40	77.90	77.79	76.70
STS	95.20	94.99	94.70	94.50	94.21	93.90
SSt	89.10	88.51	88.80	88.90	88.25	88.50
Tar	87.70	87.63	87.30	87.70	87.62	87.30
Vad	92.50	92.73	92.30	91.40	91.70	91.20
S13	90.00	89.52	89.40	88.00	87.49	87.40
S17	93.60	93.59	93.20	93.10	93.17	92.80
S16	91.80	91.62	91.50	90.10	89.90	89.80
#wins	14	8	0	13	9	0
rank sums	30.0	39.0	63.0	31.0	39.0	62.0
position	1.0	2.0	3.0	1.0	2.0	3.0

Table 7.4: Overview of the results (number of wins, rank sum, and rank position, respectively) achieved by each classifier when fine-tuning the Transformer-Autoencoder models using strategies InData, LOO, and AllData in terms of accuracy and F_1 -macro

ACCURACY						
Strategy	LR	SVM	MLP	RF	XGB	Total
BERT						
AllData	1/51.0/2.0	2/49.0/2.5	3/43.5/2.0	5/42.0/2.0	2/47.5/2.0	13/233.0/10.5
LOO	1/56.0/3.0	5/49.0/2.5	3/56.5/3.0	6/49.0/3.0	3/53.0/3.0	18/263.5/14.5
InData	20/25.0/1.0	15/34.0/1.0	14/32.0/1.0	9/41.0/1.0	14/31.5/1.0	72/163.5/5.0
RoBERTa						
AllData	11/32.0/1.0	11/33.0/1.0	10/35.0/1.0	11/34.0/1.0	12/35.0/1.0	55/169.0/5.0
LOO	6/48.0/2.0	6/44.0/2.0	9/42.0/2.0	10/37.0/2.0	8/41.0/2.0	39/212.0/10.0
InData	3/52.0/3.0	4/55.0/3.0	2/55.0/3.0	1/61.0/3.0	2/56.0/3.0	12/279.0/15.0
BERTweet						
AllData	10/36.5/1.0	14/30.0/1.0	9/34.5/1.0	13/31.0/1.0	12/33.5/1.0	58/165.5/5.0
LOO	9/39.5/2.0	8/39.0/2.0	11/37.0/2.0	8/39.5/2.0	8/39.0/2.0	44/194.0/10.0
InData	2/56.0/3.0	0/63.0/3.0	1/60.5/3.0	1/61.5/3.0	1/59.5/3.0	5/300.5/15.0
F_1 -MACRO						
Strategy	LR	SVM	MLP	RF	XGB	Total
BERT						
AllData	1/51.0/2.0	0/51.0/3.0	3/45.0/2.0	6/41.0/2.0	3/48.5/2.0	13/236.5/11.0
LOO	2/56.0/3.0	6/48.0/2.0	4/55.0/3.0	5/51.0/3.0	3/52.0/3.0	20/262.0/14.0
InData	19/25.0/1.0	16/33.0/1.0	15/32.0/1.0	11/40.0/1.0	15/31.5/1.0	76/161.5/5.0
RoBERTa						
AllData	13/31.0/1.0	12/32.5/1.0	9/35.5/1.0	11/34.0/1.0	12/35.0/1.0	57/168.0/5.0
LOO	5/49.0/2.0	4/45.0/2.0	10/42.0/2.0	10/36.0/2.0	8/40.0/2.0	37/212.0/10.0
InData	4/52.0/3.0	4/54.5/3.0	2/54.5/3.0	1/62.0/3.0	2/57.0/3.0	13/280.0/15.0
BERTweet						
AllData	10/35.5/1.0	13/31.0/1.0	13/31.0/1.0	13/31.0/1.0	10/35.0/1.0	59/163.5/5.0
LOO	10/37.5/2.0	9/39.0/2.0	9/39.0/2.0	8/39.5/2.0	10/37.5/2.0	46/192.5/10.0
InData	1/59.0/3.0	0/62.0/3.0	0/62.0/3.0	1/61.5/3.0	0/59.5/3.0	2/304.0/15.0

Table 7.5: Best results achieved for each dataset by fine-tuning the Transformer-Autoencoder models using strategies InData, LOO, and AllData

Dataset	Accuracy	Classifier	Embedding	F_1 -macro	Classifier	Embedding
iro	83.33	LR	BERTweet-LOO	77.24	SVM	BERTweet-LOO
sar	80.66	MLP	RoBERTa-LOO	79.35	MLP	RoBERTa-LOO
ntu	93.89	LR	BERTweet-LOO	93.66	LR	BERTweet-LOO
S15	94.39	MLP	BERTweet-LOO	87.31	MLP	BERTweet-LOO
stm	92.75	SVM	BERTweet-LOO	92.73	SVM	BERTweet-LOO
per	88.52	MLP	BERTweet-LOO	86.14	LR	BERTweet-LOO
hob	89.50	MLP	RoBERTa-InData	88.00	MLP	RoBERTa-InData
iph	87.80	LR	BERTweet-InData	85.80	LR	BERTweet-AllData
mov	91.10	LR	BERTweet-AllData	85.40	LR	BERTweet-AllData
san	91.60	LR	BERTweet-AllData	91.50	LR	BERTweet-AllData
Nar	97.00	SVM	BERTweet-AllData	96.80	MLP	BERTweet-AllData
arc	92.10	MLP	BERTweet-AllData	91.90	MLP	BERTweet-AllData
S18	90.90	SVM	BERTweet-AllData	90.80	SVM	BERTweet-AllData
OMD	89.77	SVM	BERTweet-LOO	88.99	SVM	BERTweet-LOO
HCR	82.21	XGB	BERTweet-LOO	77.90	SVM	BERTweet-AllData
STS	95.20	SVM	BERTweet-AllData	94.50	SVM	BERTweet-AllData
SSt	89.10	SVM	BERTweet-AllData	88.90	SVM	BERTweet-AllData
Tar	87.70	SVM	BERTweet-AllData	87.70	SVM	BERTweet-AllData
Vad	93.14	LR	BERTweet-LOO	92.05	LR	BERTweet-LOO
S13	90.40	LR	BERTweet-InData	88.30	LR	BERTweet-InData
S17	93.60	SVM	BERTweet-AllData	93.17	SVM	BERTweet-LOO
S16	91.80	SVM	BERTweet-AllData	90.10	SVM	BERTweet-AllData

performances in ten out of the 22 datasets in terms of accuracy and in 11 out of the 22 datasets in terms of F_1 -macro. Strategy LOO achieved the best results in nine out of the 22 datasets for both accuracy and F_1 -macro. The better performance of the AllData strategy for larger target datasets indicates that the significant amount of information present in the target dataset is indispensable for the fine-tuning process, while the information present in smaller datasets seems not to contribute to the fine-tuning process, making the LOO strategy adequate for datasets with a limited amount of tweets.

Conversely, strategy InData did not achieve meaningful results. The inferior performance of the InData strategy in almost all datasets shows that, regardless of the size of the dataset, the use of external and more extensive data brings more information to the fine-tuning process, improving the final performance.

Next, we present an overall evaluation of combining all fine-tuned models and classifiers across the 22 datasets, in terms of the average rank position. Table 7.6 reports the top

Table 7.6: Top 10 average rank position results achieved for each combination Model-Classifer by evaluating Transformer-Autoencoder model

Model	Classifier	Accuracy avg. rank pos.	Model	Classifier	F_1 -macro avg. rank pos.
BERTweet-AllData	LR	4.86	BERTweet-AllData	LR	4.18
BERTweet-AllData	MLP	5.57	BERTweet-LOO	LR	5.34
BERTweet-LOO	MLP	6.07	BERTweet-AllData	MLP	5.45
BERTweet-LOO	LR	6.16	BERTweet-AllData	SVM	6.11
BERTweet-AllData	SVM	7.36	BERTweet-LOO	MLP	6.75
BERTweet-LOO	SVM	8.64	BERTweet-LOO	SVM	6.86
BERTweet-InData	LR	9.11	BERTweet-InData	LR	8.36
BERTweet-InData	MLP	11.27	RoBERTa-AllData	LR	11.89
RoBERTa-AllData	MLP	13.09	BERTweet-InData	MLP	11.95
RoBERTa-AllData	LR	13.48	BERTweet-InData	SVM	12.93

ten results among all 45 possible combinations (3 language models \times 3 fine-tuning strategies \times 5 classification algorithms). We can observe that the LR classifier trained with BERTweet embeddings fine-tuned via strategy AllData achieved the best overall predictive performances. Also, note that the fine-tuned BERTweet embeddings with strategies AllData and LOO, combined with LR, MLP, and SVM, appear at the top of the ranking (top six results). Another point worth highlighting is that BERTweet masters the top ten results, appearing in eight out of the ten positions in terms of accuracy and in nine out of the ten positions in terms of F_1 -macro.

Tables 7.7 and 7.8 show the results among all fine-tuned models and a summary of the results for each classifier, from best to worst, respectively, in terms of the average rank position. Once again, from Table 7.7, we can notice that all BERTweet fine-tuned models (InData, LOO, and AllData) were ranked in the top three results. Among the classifiers, as we can see in Table 7.8, MLP and LR achieved the best predictive performances and were ranked as the top two best classifiers. Conversely, RF was ranked as the worst classifier.

To evaluate the effectiveness of fine-tuning the Transformer-based models using tweets from sentiment datasets, we present a comparison among all fine-tuning strategies assessed in this study for each language model. Specifically, we compare the fine-tuned models presented in this chapter, by using strategies InData, LOO, and AllData, against the best fine-tuned models identified in Chapter 6, i.e., BERT-250K, RoBERTa-50K, and BERTweet-5K. Tables 7.9, 7.10, and 7.11 report these results in terms of the average rank position for BERT, RoBERTa, and BERTweet, respectively.

Table 7.7: Average position results achieved for each Embedding

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
BERTweet-AllData	12.99	BERTweet-AllData	13.51
BERTweet-LOO	14.04	BERTweet-LOO	14.46
BERTweet-InData	19.59	BERTweet-InData	19.95
RoBERTa-AllData	22.76	RoBERTa-AllData	22.70
RoBERTa-LOO	24.46	RoBERTa-LOO	24.00
BERT-InData	24.53	BERT-InData	24.32
RoBERTa-InData	27.90	RoBERTa-InData	27.70
BERT-AllData	30.12	BERT-AllData	29.85
BERT-LOO	30.61	BERT-LOO	30.51

Table 7.8: Average rank position results achieved for each Classifier

Classifier	Accuracy avg. rank pos.	Classifier	F_1 -macro avg. rank pos.
MLP	14.83	LR	14.41
LR	15.73	MLP	15.28
SVM	22.42	SVM	19.77
XGB	25.70	XGB	28.13
RF	36.32	RF	37.41

Table 7.9: Average rank position results achieved for BERT models

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
BERT-InData	8.04	BERT-InData	8.10
BERT-AllData	10.61	BERT-AllData	10.66
BERT-LOO	11.24	BERT-LOO	11.29
BERT-250K	12.12	BERT-250K	11.95

Table 7.10: Average rank position results achieved for RoBERTa models

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
RoBERTa-AllData	8.81	RoBERTa-AllData	8.85
RoBERTa-LOO	9.89	RoBERTa-LOO	9.84
RoBERTa-50K	11.61	RoBERTa-50K	11.52
RoBERTa-InData	11.68	RoBERTa-InData	11.80

Regarding BERT, as shown in Table 7.9, note that all fine-tuning strategies using tweets from sentiment datasets achieved better overall results than using the sample of 250K generic tweets. Moreover, strategy InData appears at the top of the ranking as the best fine-tuning strategy. It is worth mentioning that strategy InData uses only the tweets from the target dataset itself to adjust the language model. This means that the strategy InData used a number of tweets much smaller than the 250K tweets contained in the sample.

On the other hand, as we can see in Tables 7.10 and 7.11, strategy InData did not achieve meaningful results for RoBERTa and BERTweet models. Nevertheless, for these models, strategies AllData and LOO, which also use tweets from sentiment datasets, achieved rather comparable performances and were ranked as the top two best fine-tuning strategies.

To acknowledge the effectiveness of fine-tuning the Transformer-based models using tweets from sentiment datasets, we present an overall comparison among all fine-tuning strategies and all 47 models previously assessed in this study. Tables 7.12 and 7.13 present, respectively, the 10 best and the 10 worst overall combination of model and classifier, assessing the average rank position of all 280 (56 models and five classifier) model and classifier combinations. We note that BERTweet tuned with tweets from sentiment

Table 7.11: Average rank position results achieved for BERTweet models

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
BERTweet-AllData	9.00	BERTweet-AllData	9.11
BERTweet-LOO	9.76	BERTweet-LOO	9.76
BERTweet-5K	10.19	BERTweet-5K	10.20
BERTweet-InData	13.05	BERTweet-InData	12.93

datasets and combined with LR and MLP had the four best results, in terms of accuracy, and the two best results, in terms of F_1 -macro. These combinations were followed by BERTweet tuned with generic tweets. More specifically, combinations with the strategy AllData and LOO achieved better overall results. Independently of the language model, LR and MLP were the most frequent classifier in the top 10 results. Conversely, all the ten worst combinations are static representations combined with RF, which was unanimous in the worst model and classifiers combinations.

Assessing only the different kinds of embeddings, Tables 7.14 and 7.15 present, respectively, the best and the worst average rank position comparing all 56 representations (the nine models tuned with sentiment datasets and the 47 previous representations). This analysis confirms the good performance of fine-tuning the Transformer-based models using tweets from sentiment datasets. More specifically, the strategies AllData and LOO obtained the two best results. It is possible to notice that tuning BERTweet with generic tweets also brings performance improvement to BERTweet. Regarding the worst behaviors, presented in Table 7.15, it is possible to note that all the ten strategies are again static representations.

Lastly, regarding research question RQ4, we can highlight that fine-tuning Transformer-based models using tweets from sentiment datasets seems to boost classification performance in Twitter sentiment analysis. As a matter of fact, the strategies AllData and LOO exploited in this chapter, which use a collection of sentiment tweets to adjust a language model, achieved better overall results than using samples of unlabeled, or generic unlabeled, tweets. Although we do not use the labels of those tweets in the fine-tuning procedure, they may carry a lot of sentiment information as compared to the tweets from the Edinburgh corpus, which originated the samples of generic unlabeled tweets used in the experiments. Furthermore, BERTweet embeddings fine-tuned with strategy AllData seems to be very effective in determining the sentiment expressed in tweets, especially

Table 7.12: Top 10 average rank position results achieved for each combination Model-Classifier by evaluating all assessed model in this study

Model	Classifier	Accuracy avg. rank pos.	Model	Classifier	F_1 -macro avg. rank pos.
BERTweet-LOO	LR	17.30	BERTweet-LOO	LR	15.52
BERTweet-AllData	LR	18.16	BERTweet-AllData	LR	16.30
BERTweet-LOO	MLP	18.55	BERTweet-5K	LR	18.64
BERTweet-AllData	MLP	19.36	BERTweet-AllData	MLP	20.68
BERTweet-5K	LR	20.64	BERTweet-LOO	MLP	21.00
BERTweet-5K	MLP	23.14	BERTweet-25K	SVM	22.48
BERTweet-25K	MLP	24.30	BERTweet-10K	SVM	23.16
BERTweet-25K	LR	26.23	BERTweet-25K	LR	23.30
BERTweet-1K	MLP	26.91	BERTweet-10K	LR	23.86
BERTweet-50K	MLP	27.02	BERTweet-AllData	SVM	24.82

when used to train LR, MLP, and SVM classifiers.

Table 7.13: Tail 10 average rank position results achieved for each combination Model-Classifier by evaluating all assessed model in this study

Model	Classifier	Accuracy avg. rank pos.	Model	Classifier	F_1 -macro avg. rank pos.
EWE	RF	247.23	DeepMoji	RF	252.86
W2V-Araque	RF	249.75	BERTweet-static	LR	253.14
W2V-GN	RF	250.00	EWE	RF	256.48
GloVe-WP	RF	253.68	W2V-Araque	RF	259.80
fastText	RF	255.75	W2V-GN	RF	261.70
BERT-static	RF	257.32	GloVe-WP	RF	263.11
RoBERTa-static	RF	259.70	fastText	RF	266.95
BERT-static	LR	263.34	BERT-static	RF	267.75
BERTweet-static	RF	265.91	RoBERTa-static	RF	269.93
BERTweet-static	LR	274.43	BERTweet-static	RF	275.43

Table 7.14: Top 10 average rank position results achieved comparing all assessed models in this study

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
BERTweet-AllData	53.83	BERTweet-AllData	60.67
BERTweet-LOO	56.52	BERTweet-LOO	62.59
BERTweet-5K	60.56	BERTweet-5K	66.21
BERTweet-25K	65.38	BERTweet-25K	71.07
BERTweet-10K	65.51	BERTweet-10K	71.17
BERTweet-1K	68.59	BERTweet-1K	73.32
BERTweet-50K	72.31	BERTweet-50K	77.50
BERTweet-250K	78.13	BERTweet-250K	82.90
BERTweet-InData	83.80	BERTweet-InData	87.93
BERTweet-500K	86.10	BERTweet-500K	90.85

Table 7.15: Tail 10 average rank position results achieved comparing all assessed models in this study

Model	Accuracy avg. rank pos.	Model	F_1 -macro avg. rank pos.
SSWE	209.69	W2V-GN	204.38
GloVe-TWT	215.62	GloVe-TWT	207.22
DeepMoji	217.13	DeepMoji	208.08
EWE	217.46	EWE	208.43
TF-IDF	220.83	GloVe-WP	215.45
BERT-static	224.61	fastText	218.05
GloVe-WP	225.94	BERT-static	218.40
fastText	227.80	w2v-Araque	222.85
W2V-Araque	230.56	TF-IDF	224.34
BERTweet-static	244.21	BERTweet-static	237.01

Chapter 8

Conclusions and Future Works

In this work, we presented an extensive assessment of modern and classical word representations when used for the task of Twitter sentiment analysis. Specifically, we assessed the classification performance of 14 static representations, the most recent Transformer-based autoencoder models, including BERT, RoBERTa, and BERTweet, as well as different fine-tuning strategies of the language representation tasks in such models. All models were evaluated in the context of Twitter sentiment analysis using a rich set of 22 datasets and five classifiers from distinct natures. The main focus of this study was on identifying the most appropriate word representations for the sentiment analysis of English tweets.

Based on the results of the experiments performed in this study, we can highlight the following conclusions:

- Considering the static representations in limited resource scenario, we could note that Emo2Vec, w2v-Edin, and RoBERTa models seem to be well-suited representations for determining the sentiment expressed in tweets. The good performance achieved by Emo2Vec and w2v-Edin indicates that being trained from scratch with tweets can boost the classification performance of static representations when applied in Twitter sentiment analysis. Although RoBERTa was not trained from scratch with tweets, it is a Transformer-based autoencoder model, which holds state-of-the-art performance in several NLP tasks. Regarding the classifiers, we could see that SVM and MLP achieved the best overall performances, especially when used to train RoBERTa’s static embeddings.
- Regarding the Transformer-based models, we could observe that BERTweet is the most appropriate language model to be used in the sentiment classification of tweets. Specifically, the particular vocabulary tweets contain, combined with a language

model that was trained focused on learning their intrinsic structure, can effectively improve the performance of the Twitter sentiment analysis task. Considering the combination of language models and classifiers, we can point out that BERTweet achieved the best overall results when combined with LR and MLP. Furthermore, by comparing the Transformer-based models and the static representations, we could notice that the adaptation of the tokens' embeddings to the context they appear performed by the Transformer-based models benefits the sentiment classification task.

- When fine-tuning the Transformer-based models with a large set of English unlabeled tweets we could note that although it improves the classification performance, using as many tweets as possible does not necessarily means better results. In this context, we presented an extensive evaluation of sets of tweets with different sizes, varying from 0.5K to 1.5M. These results have shown that while BERT and RoBERTa achieved better predictive performances when tuned with sets of 250K and 50K tweets, respectively, BERTweet outperformed all fine-tuned models using only 5K tweets. This result indicates that models trained from scratch with tweets, such as BERTweet, needs less tweets to have its performance improved. Moreover, by comparing all fine-tuned models taking into account the classifiers, BERTweet combined with MLP, LR, and SVM achieved the best overall performances.
- Analyzing the fine-tuning of the language model based on Transformers autoencoders with sentiment analysis datasets, i.e., with tweets that express polarity, we can see that the tuned models' performance is better than when tuned with generic tweets. All fine-tuning strategies with sentiment analysis datasets performed better than the best-tuned models adjusted with generic tweets. We conclude then that it is worth fine-tuning a model based on Transformer autoencoders using a set of sentiment tweets. Among the fine-tuning strategies – using sentiment analysis tweets – explored in the study, it was possible to perceive that each Transformer model presented a better performance with different adjustment methods. The use of only the target dataset, for example, was a good option to be used with BERT. For RoBERTa and BERTweet, the combination of the target dataset with a set of tweets from other datasets presented a good strategy for fine-tuning the language model. In a general comparison, we noticed that BERTweet tuned with the union of the target dataset and the set of sentiment analysis tweets (BERTweet_AllData) performed better than the other adjusted models. Besides, we could observe that BERTweet_AllData presented a good performance when combined with LR and

MLP classifiers.

- After answering our research questions, we can briefly state that: (i) Transformer-based autoencoder models perform better than static representation, (ii) Transformer autoencoder models fine-tuned with English tweets behavior better than the respective original models and, finally, (iii) it is worth fine-tuning a language model originally trained with generic English tweets with tweets from sentiment analysis datasets. Considering all original and fine-tuned models, the best overall performance for the English tweets sentiment analysis task was achieved by the Transformer-Autoencoder model trained from scratch with generic tweets (BERTweet) when fine-tuned with tweets from a target sentiment dataset added by tweets from a large set of other sentiment datasets. This strategy was called BERTweet_AllData, which we consider a good suggestion for sentiment classification of English tweets, mainly when combined with MLP or LR classifiers.

For future work, we plan to investigate other methods for fine-tuning language models, mainly considering the polarity classification as the downstream tuning task. Transformer-Autoencoder pre-trained models, like BERT, RoBERTa and BERTweet, can have its weights adjusted looking for becoming more accurate in a specific task, like sentiment analysis. This adjustment is made by adding an extra classification layer in the top of the model and back-propagating the error in the final task through language models' weights. We intend then to compare the best results obtained in this study with the ones achieved by this specific-task category of fine-tuning. Some limitations and threats of the present work are the absence of statistical test concerning models performance assessment, a concentration on BERT-based architectures relative to the modern language models evaluated and the absence of classifier hyperparameters fine-tuning.

References

- [1] ADHIKARI, A., RAM, A., TANG, R., LIN, J. DoBERT: BERT for document classification. *CoRR abs/1904.08398* (2019).
- [2] AGRAWAL, A., AN, A., PAPAGELIS, M. Learning emotion-enriched word representations. In *Proceedings of the 27th International Conference on Computational Linguistics* (Santa Fe, New Mexico, USA, agosto de 2018), Association for Computational Linguistics, p. 950–961.
- [3] AJEET RAM PATHAK, BASANT AGARWAL, M. P. S. R. Application of Deep Learning Approaches for Sentiment Analysis. *Deep Learning-Based Approaches for Sentiment Analysis* (2020), 1–31.
- [4] AKKALYONCU YILMAZ, Z., WANG, S., YANG, W., ZHANG, H., LIN, J. Applying BERT to document retrieval with birch. 19–24.
- [5] ARAQUE, O., CORCUERA-PLATAS, I., SNCHEZ-RADA, J. F., IGLESIAS, C. A. Enhancing deep learning sentiment analysis with ensemble techniques in social applications. *Expert Syst. Appl.* 77, C (julho de 2017), 236–246.
- [6] BARBOSA, L., FENG, J. Robust sentiment detection on Twitter from biased and noisy data. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters* (2010), Association for Computational Linguistics, p. 36–44.
- [7] BENGIO, Y., SIMARD, P., FRASCONI, P. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks* 5, 2 (mar 1994), 157–166.
- [8] BRAVO-MARQUEZ, F., FRANK, E., MOHAMMAD, S. M., PFAHRINGER, B. Determining word-emotion associations from tweets by multi-label classification. In *2016 IEEE/WIC/ACM International Conference on Web Intelligence (WI)* (2016), p. 536–539.
- [9] CAMBRIA, E., PORIA, S., GELBUKH, A., THELWALL, M. Sentiment analysis is a big suitcase. *IEEE Intelligent Systems* 32, 6 (November 2017), 74–80.
- [10] CARVALHO, J. Exploiting Different Types of Features to Improve Classification Effectiveness in Twitter Sentiment Analysis. *indefinido* (2019).
- [11] CARVALHO, J., PLASTINO, A. On the evaluation and combination of state-of-the-art features in Twitter sentiment analysis. *Artif Intell Rev* (2020).
- [12] CHAYBOUTI, S., SAGHE, A., SHABOU, A. Efficientqa : a roberta based phrase-indexed question-answering system.

-
- [13] CHEN, L., WANG, W., NAGARAJAN, M., WANG, S., SHETH, A. P. Extracting diverse sentiment expressions with target-dependent polarity from twitter. In *ICWSM* (2012).
- [14] DEVLIN, J., CHANG, M.-W., LEE, K., TOUTANOVA, K. Bert: Pre-training of deep bidirectional transformers for language understanding, 2018.
- [15] DIAKOPOULOS, N. A., SHAMMA, D. A. Characterizing debate performance via aggregated twitter sentiment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2010), CHI '10, Association for Computing Machinery, p. 1195–1198.
- [16] DODGE, J., ILHARCO, G., SCHWARTZ, R., FARHADI, A., HAJISHIRZI, H., SMITH, N. Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping. *arXiv preprint arXiv:2002.06305* (2020).
- [17] DODGE, J., ILHARCO, G., SCHWARTZ, R., FARHADI, A., HAJISHIRZI, H., SMITH, N. A. Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping. *CoRR abs/2002.06305* (2020).
- [18] DONG, L., WEI, F., TAN, C., TANG, D., ZHOU, M., XU, K. Adaptive recursive neural network for target-dependent twitter sentiment classification. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)* (Baltimore, Maryland, June 2014), Association for Computational Linguistics, p. 49–54.
- [19] FAYYAD, U., PIATETSKY-SHAPIO, G., UTHURUSAMY, R. Summary from the kdd-03 panel - data mining: The next 10 years. *SIGKDD Explorations* 5 (01 2003), 191–196.
- [20] FELBO, B., MISLOVE, A., SØGAARD, A., RAHWAN, I., LEHMANN, S. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (Copenhagen, Denmark, setembro de 2017), Association for Computational Linguistics, p. 1615–1625.
- [21] G. SALTON, A. WONG, AND C. S. YANG. A Vector Space Model for Automatic Indexing. *Communications of the ACM* (1975), 613–620.
- [22] GAO, Z., FENG, A., SONG, X., WU, X. Target-dependent sentiment classification with bert. *IEEE Access* 7 (2019), 154290–154299.
- [23] GILBERT, C. H. E. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth International Conference on Weblogs and Social Media (ICWSM-14)*. Available at (20/04/16) [http://comp. social. gatech. edu/papers/icwsm14.vader.hutto.pdf](http://comp.social.gatech.edu/papers/icwsm14.vader.hutto.pdf) (2014).
- [24] GO, A., BHAYANI, R., HUANG, L. Twitter sentiment classification using distant supervision. *Processing* (2009), 1–6.

- [25] GONÇALVES, P., DALIP, D., REIS, J., MESSIAS, J., RIBEIRO, F., MELO, P., ARAÚJO, L., GONÇALVES, M., BENEVENUTO, F. Bazinga! caracterizando e detectando sarcasmo e ironia no twitter. In *Anais do IV Brazilian Workshop on Social Network Analysis and Mining* (Porto Alegre, RS, Brasil, 2015), SBC, p. .
- [26] GOODFELLOW, I., BENGIO, Y., COURVILLE, A. *Deep Learning*. The MIT Press, 2016.
- [27] HOWARD, J., RUDER, S. Universal language model fine-tuning for text classification, 2018.
- [28] LAN, Z., CHEN, M., GOODMAN, S., GIMPEL, K., SHARMA, P., SORICUT, R. Albert: A lite bert for self-supervised learning of language representations, 2020.
- [29] LE, H., VIAL, L., FREJ, J., SEGONNE, V., COAVOUX, M., LECOUTEUX, B., ALLAUZEN, A., CRABBÉ, B., BESACIER, L., SCHWAB, D. Flaubert: Unsupervised language model pre-training for french, 2020.
- [30] LIU, B. *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge University Press, 2020.
- [31] LIU, Y., OTT, M., GOYAL, N., DU, J., JOSHI, M., CHEN, D., LEVY, O., LEWIS, M., ZETTLEMOYER, L., STOYANOV, V. Roberta: A robustly optimized bert pre-training approach, 2019.
- [32] LOCHTER, J. V., ZANETTI, R. F., RELLER, D., ALMEIDA, T. A. Short text opinion detection using ensemble of classifiers and semantic indexing. *Expert Syst. Appl.* 62 (2016), 243–249.
- [33] MANNING, C. D., RAGHAVAN, P., SCHÜTZE, H. *Introduction to Information Retrieval*. Cambridge University Press, USA, 2008.
- [34] MIKOLOV, T., GRAVE, E., BOJANOWSKI, P., PUHRSCHE, C., JOULIN, A. Advances in pre-training distributed word representations, 2017.
- [35] MIKOLOV, T., SUTSKEVER, I., CHEN, K., CORRADO, G., DEAN, J. Distributed representations of words and phrases and their compositionality, 2013.
- [36] MOHAMMAD, S. M., BRAVO-MARQUEZ, F., SALAMEH, M., KIRITCHENKO, S. Semeval-2018 Task 1: Affect in tweets. In *Proceedings of International Workshop on Semantic Evaluation (SemEval-2018)* (New Orleans, LA, USA, 2018).
- [37] NAKOV, P., RITTER, A., ROSENTHAL, S., SEBASTIANI, F., STOYANOV, V. SemEval-2016 task 4: Sentiment analysis in Twitter. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)* (San Diego, California, junho de 2016), Association for Computational Linguistics, p. 1–18.
- [38] NARR, S., HÜLFENHAUS, M., ALBAYRAK, S. Language-independent twitter sentiment analysis. In *Language-Independent Twitter Sentiment Analysis* (2012).
- [39] NGUYEN, D. Q., VU, T., NGUYEN, A. T. Bertweet: A pre-trained language model for english tweets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* (2020), p. 9–14.

- [40] PAK, A., PAROUBEK, P. Twitter as a corpus for sentiment analysis and opinion mining. In *Proceedings of the 7th International Conference on Language Resources and Evaluation* (2010), p. 1320–1326.
- [41] PENNINGTON, J., SOCHER, R., MANNING, C. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Doha, Qatar, outubro de 2014), Association for Computational Linguistics, p. 1532–1543.
- [42] PETERS, M., NEUMANN, M., IYYER, M., GARDNER, M., CLARK, C., LEE, K., ZETTLEMOYER, L. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)* (2018), p. 2227–2237.
- [43] PETERS, M. E., NEUMANN, M., IYYER, M., GARDNER, M., CLARK, C., LEE, K., ZETTLEMOYER, L. Deep contextualized word representations. *CoRR abs/1802.05365* (2018).
- [44] PETROVIĆ, S., OSBORNE, M., LAVRENKO, V. The Edinburgh twitter corpus. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Linguistics in a World of Social Media* (Los Angeles, California, USA, junho de 2010), Association for Computational Linguistics, p. 25–26.
- [45] RADFORD, A., NARASIMHAN, K., SALIMANS, T., SUTSKEVER, I. Improving language understanding by generative pre-training.
- [46] ROSENTHAL, S., FARRA, N., NAKOV, P. SemEval-2017 task 4: Sentiment analysis in Twitter. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)* (Vancouver, Canada, agosto de 2017), Association for Computational Linguistics, p. 502–518.
- [47] SAIF, H. *Semantic Sentiment Analysis of Microblogs*. Tese de Doutorado, The Open University, June 2015.
- [48] SAIF, H., FERNANDEZ, M., HE, Y., ALANI, H. Evaluation datasets for twitter sentiment analysis: A survey and a new dataset, the sts-gold. In *Evaluation Datasets for Twitter Sentiment Analysis: A survey and a new dataset, the STS-Gold* (2013).
- [49] SPARCK JONES, K. *A Statistical Interpretation of Term Specificity and Its Application in Retrieval*. Taylor Graham Publishing, GBR, 1988, p. 132–142.
- [50] SPERIOSU, M., SUDAN, N., UPADHYAY, S., BALDRIDGE, J. Twitter polarity classification with label propagation over lexical links and the follower graph. In *Proceedings of the First Workshop on Unsupervised Learning in NLP* (USA, 2011), EMNLP ’11, Association for Computational Linguistics, p. 53–63.
- [51] TANG, D., WEI, F., YANG, N., ZHOU, M., LIU, T., QIN, B. Learning sentiment-specific word embedding for Twitter sentiment classification. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (Baltimore, Maryland, junho de 2014), Association for Computational Linguistics, p. 1555–1565.

- [52] THELWALL, M., BUCKLEY, K., PALTOGLOU, G. Sentiment strength detection for the social web. *J. Assoc. Inf. Sci. Technol.* 63, 1 (2012), 163–173.
- [53] TOMAS MIKOLOV, ILYA SUTSKEVER, KAI CHEN, GREG CORRADO, AND JEFFREY DEAN. Distributed representations of words and phrases and their compositionality. *NIPS* (2013), 3111–3119.
- [54] TOMAS MIKOLOV, ILYA SUTSKEVER, KAI CHEN, GREG CORRADO, AND JEFFREY DEAN. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics* (2017), 135–146.
- [55] TURNEY, P. D., PANTEL, P. From frequency to meaning: Vector space models of semantics. *Journal of artificial intelligence research* 37 (2010), 141–188.
- [56] TURNEY, P. D., PANTEL, P. From frequency to meaning: Vector space models of semantics. *CoRR abs/1003.1141* (2010).
- [57] VASWANI, A., SHAZEER, N., PARMAR, N., USZKOREIT, J., JONES, L., GOMEZ, A. N., KAISER, L., POLOSUKHIN, I. Attention is all you need, 2017.
- [58] XU, P., MADOTTO, A., WU, C.-S., PARK, J. H., FUNG, P. Emo2Vec: Learning generalized emotion representation by multi-task training. In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis* (Brussels, Belgium, outubro de 2018), Association for Computational Linguistics, p. 292–298.
- [59] ZHU, Y., KIROS, R., ZEMEL, R., SALAKHUTDINOV, R., URTASUN, R., TORRALBA, A., FIDLER, S. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books, 2015.

APPENDIX A – Extra Experimental Results

This appendix presents a complement to the experimental results reported in the previous chapters. Section A.1 presents experimental results complementary to those presented in Chapter 4. Section A.2 complements Chapter 5, Section A.3 refers to Chapter 6, and Section A.4 to Chapter 7.

A.1 Static Text Representation

Table A.1: Accuracies (%) achieved using the RF classifier

Dataset	w2v GN	GloVe WP	Fast Text	EWE	GloVe TWT	BERT static	TF- IDF	w2v Araque	w2v Edin	SSWE	Emo2 Vec	Deep Moji	RoB static	BTWT static
iro	69.29	66.43	66.19	63.33	67.86	69.52	66.19	62.86	67.86	72.62	72.38	64.76	67.86	64.52
sar	57.32	63.21	64.46	66.07	68.75	60.71	62.14	71.61	66.07	73.21	74.46	64.64	51.96	57.68
ntu	72.65	84.87	69.39	68.31	72.99	76.57	69.76	72.29	90.65	93.51	78.02	89.91	84.15	76.55
S15	85.36	85.67	85.36	85.36	85.67	85.36	89.1	85.36	85.67	85.98	85.98	85.36	85.36	85.36
stm	79.96	77.72	79.94	82.46	82.16	74.94	76.57	78.55	81.33	81.03	83.57	76.03	76.89	68.53
per	74.94	74.48	75.61	74.95	74.04	72.89	76.75	71.52	76.31	77.47	78.83	72.43	71.06	71.53
hob	83.13	80.26	79.11	87.16	82.94	80.85	92.72	84.87	81.42	81.24	86.59	84.47	79.5	80.64
iph	77.44	78.38	78.55	79.32	79.87	77.07	77.82	79.5	78.55	78.94	80.64	77.63	76.14	77.09
mov	82.0	82.0	82.0	82.0	82.18	82.18	82.89	82.0	82.36	83.41	87.17	82.53	82.0	82.0
san	79.9	75.98	78.83	78.75	78.99	74.03	79.17	77.93	79.49	81.86	80.55	79.74	73.94	73.87
Nar	78.97	75.47	79.38	78.16	82.8	72.54	78.49	80.44	83.21	88.67	88.35	83.05	74.58	70.99
arc	82.13	81.49	83.64	82.59	84.17	80.67	84.34	82.13	84.52	81.72	83.64	80.79	81.55	80.21
S18	76.81	74.02	77.3	76.71	77.41	72.94	78.59	72.08	77.67	80.1	79.93	75.31	74.07	68.86
OMD	78.44	77.07	74.97	75.39	75.55	76.7	76.29	73.19	78.18	76.34	76.71	73.66	76.18	76.28
HCR	74.16	73.38	73.95	74.58	75.1	74.32	78.46	74.11	74.58	74.95	74.69	74.11	74.21	73.27
STS	76.84	75.96	75.71	77.63	79.45	74.19	80.29	78.72	78.81	88.59	86.58	80.29	73.84	72.22
SSt	71.52	68.42	71.04	69.72	73.79	68.37	71.21	71.47	73.09	80.78	83.66	73.22	69.24	64.53
Tar	79.95	76.98	79.92	80.07	80.1	74.01	80.36	77.62	79.03	78.14	81.37	79.44	76.32	71.68
Vad	79.69	77.41	78.53	78.93	80.98	74.52	80.15	78.24	81.08	85.82	85.94	80.15	75.07	71.88
S13	76.22	74.99	75.35	75.33	77.59	74.6	77.59	75.54	76.68	83.23	87.46	76.91	74.62	73.85
S17	83.17	80.9	83.65	83.13	84.67	77.2	83.87	78.0	82.84	81.03	84.35	79.63	77.67	72.48
S16	80.54	79.76	80.05	80.48	81.96	76.56	81.5	77.77	80.84	80.75	82.98	78.48	76.56	73.94
#wins	1	0	0	0	1	0	3	0	1	6	10	0	0	0
rank sums	159.0	207.0	190.0	170.5	105.0	235.5	116.0	202.5	105.5	79.5	46.0	171.0	244.0	278.5
position	6.0	11.0	9.0	7.0	3.0	12.0	5.0	10.0	4.0	2.0	1.0	8.0	13.0	14.0

Table A.2: F_1 -macro scores (%) achieved by evaluating static representation using the RF classifier

Dataset	w2v GN	GloVe WP	Fast Text	EWE	GloVe TWT	BERT static	TF- IDF	w2v Araque	w2v Edin	SSWE	Emo2 Vec	Deep Moji	RoB static	BTWT static
iro	46.41	43.02	39.77	41.72	45.8	49.66	39.77	38.44	45.88	58.33	57.46	39.24	43.55	39.11
sar	55.42	59.85	60.6	64.08	65.79	55.44	61.34	68.91	64.05	71.65	72.87	61.88	47.6	51.59
ntu	70.52	83.96	65.88	64.28	70.21	74.73	66.47	69.28	90.26	93.29	77.35	89.37	82.96	74.07
S15	46.05	48.13	46.05	46.05	48.13	46.05	68.41	46.05	48.13	49.54	49.87	46.05	46.05	46.05
stm	79.9	77.58	79.77	82.42	82.02	74.82	76.1	78.49	81.11	81.01	83.51	75.9	76.79	68.32
per	55.83	57.89	57.99	59.07	56.55	48.74	70.06	48.16	60.96	66.79	70.75	51.34	44.26	44.52
hob	77.38	72.56	70.21	83.63	77.67	73.56	91.68	80.09	74.2	75.99	84.05	79.59	69.66	72.15
iph	67.42	69.35	68.24	70.11	70.78	66.53	67.49	69.61	69.64	72.87	74.96	68.44	64.27	66.23
mov	45.05	45.05	45.05	45.05	46.01	46.01	50.42	45.05	46.96	60.17	70.78	48.66	45.05	45.05
san	79.38	75.23	78.19	78.24	78.52	73.35	78.72	77.33	79.06	81.65	80.19	79.26	73.1	73.05
Nar	76.75	72.43	76.87	75.61	81.54	67.47	77.12	78.89	81.85	88.19	87.78	81.6	69.78	64.43
arc	80.94	80.12	82.48	81.55	83.25	79.34	83.82	80.71	83.52	80.69	82.72	79.01	80.37	78.82
S18	76.09	73.21	76.5	75.99	76.76	71.86	77.8	71.01	77.02	79.77	79.69	74.49	73.24	67.09
OMD	73.66	71.23	68.51	69.09	70.15	71.05	72.03	65.57	72.88	72.63	72.32	67.13	70.36	70.05
HCR	51.44	48.73	50.2	52.15	55.03	53.89	65.84	50.79	53.36	57.39	56.34	50.6	51.52	48.94
STS	64.53	63.21	61.79	66.08	70.38	58.24	73.01	68.82	69.21	86.14	83.25	71.65	56.87	52.8
SSt	68.21	64.4	67.36	66.25	71.64	64.04	69.48	68.11	71.01	80.07	82.99	70.67	64.88	58.86
Tar	79.94	76.97	79.91	80.06	80.09	73.98	80.34	77.6	79.02	78.09	81.33	79.4	76.3	71.63
Vad	70.25	65.06	67.25	68.26	73.08	58.05	73.23	67.7	73.31	82.54	82.69	71.59	59.68	50.69
S13	56.68	53.67	53.21	54.63	61.79	50.41	62.33	56.11	59.79	77.03	83.22	59.61	49.56	46.52
S17	80.86	77.91	81.37	80.74	82.72	72.22	82.2	73.53	80.62	78.74	82.76	75.89	73.43	64.49
S16	68.87	67.26	67.27	68.55	72.33	56.99	72.11	63.27	70.23	72.86	76.29	64.26	57.04	47.18
#wins	1	0	0	0	0	0	4	0	0	6	11	0	0	0
rank sums	161.5	208.0	200.0	165.5	109.5	234.0	106.5	204.5	105.0	70.0	42.0	173.5	245.5	284.5
position	6.0	11.0	9.0	7.0	5.0	12.0	4.0	10.0	3.0	2.0	1.0	8.0	13.0	14.0

Table A.3: Accuracies (%) achieved using the MLP classifier

Dataset	w2v GN	GloVe WP	Fast Text	EWE	GloVe TWT	BERT static	TF- IDF	w2v Araque	w2v Edin	SSWE	Emo2 Vec	Deep Moji	RoB static	BTWT static
iro	75.48	75.95	77.38	76.19	69.29	64.29	53.33	64.52	67.86	69.52	74.05	65.24	70.95	65.95
sar	70.36	60.36	58.93	64.82	77.32	65.0	51.96	70.36	67.5	83.04	61.96	57.5	59.11	53.39
ntu	80.2	88.84	77.02	73.73	76.98	86.72	75.89	79.17	95.3	92.46	79.46	91.76	88.81	90.95
S15	84.13	84.43	86.92	84.46	83.19	90.04	86.91	85.97	89.74	86.92	87.23	87.55	88.16	89.1
stm	80.5	80.78	80.49	81.88	81.34	85.76	80.76	79.92	82.99	79.1	81.05	76.34	84.1	76.62
per	83.83	80.63	83.39	82.0	81.77	79.27	80.19	79.73	79.95	75.64	79.5	79.73	79.49	77.43
hob	89.84	88.5	89.08	92.33	88.87	94.82	87.74	90.79	89.47	85.82	88.71	90.03	94.44	92.72
iph	81.02	80.82	79.68	81.57	84.39	81.37	78.93	80.64	78.57	78.55	81.78	81.95	81.94	82.14
mov	83.96	85.21	88.24	86.98	85.21	86.28	84.85	86.45	87.17	85.03	85.91	86.46	86.81	85.03
san	81.77	78.51	82.51	82.76	82.76	81.53	84.71	82.43	83.33	79.41	80.06	80.72	82.92	82.27
Nar	85.57	83.12	83.94	83.86	87.04	85.09	80.77	82.89	88.51	86.06	87.77	88.26	86.64	82.97
arc	84.75	84.4	84.81	85.45	86.26	85.97	85.51	83.76	87.14	79.81	83.12	84.23	87.6	85.57
S18	76.98	75.79	79.45	78.38	78.27	82.84	81.71	73.0	80.69	76.44	78.16	78.32	84.56	80.31
OMD	83.0	80.11	80.48	78.28	76.02	81.95	76.81	75.97	81.95	73.19	73.92	75.6	83.79	80.75
HCR	76.93	75.94	76.15	77.3	77.83	77.67	76.62	75.68	76.57	67.61	72.59	74.11	78.25	75.37
STS	85.16	83.68	86.82	86.73	86.92	87.02	84.61	86.28	88.1	87.86	85.25	88.35	87.46	85.79
SSt	77.06	72.13	76.37	76.23	76.98	77.11	71.43	73.7	79.12	75.67	82.13	77.63	78.29	73.75
Tar	81.48	78.34	80.99	80.53	79.58	80.33	78.51	80.04	80.39	74.47	79.67	79.87	81.57	79.95
Vad	87.18	86.53	85.56	85.89	85.75	87.37	82.55	85.37	88.18	84.72	84.46	86.73	89.2	85.41
S13	81.89	79.69	81.52	80.9	81.48	81.98	79.6	79.65	82.09	80.58	87.12	82.23	83.58	81.86
S17	87.25	85.68	87.43	86.42	86.66	87.13	85.71	82.24	85.57	77.52	83.19	83.55	87.65	86.92
S16	83.71	83.21	83.97	84.13	83.8	84.06	83.66	80.93	83.96	78.92	80.98	82.11	85.26	84.13
#wins	1	0	2	0	1	3	1	0	2	1	2	1	8	0
rank sums	151.5	212.5	148.5	142.0	152.0	117.5	225.0	221.0	100.5	234.0	192.0	167.5	74.0	172.0
position	6.0	11.0	5.0	4.0	7.0	3.0	13.0	12.0	2.0	14.0	10.0	8.0	1.0	9.0

Table A.4: F_1 -macro scores (%) achieved by evaluating static representation using the MLP classifier

Dataset	w2v GN	GloVe WP	Fast Text	EWE	GloVe TWT	BERT static	TF- IDF	w2v Araque	w2v Edin	SSWE	Emo2 Vec	Deep Moji	RoB static	BTWT static
iro	66.46	69.14	67.3	66.61	60.04	52.77	38.3	52.65	57.52	63.39	65.85	55.16	60.77	50.22
sar	69.56	57.48	56.76	63.38	76.49	59.85	49.12	68.59	64.74	82.58	61.3	51.48	56.76	49.76
ntu	79.77	88.39	76.33	72.38	76.14	86.45	74.92	78.58	95.19	92.26	78.95	91.51	88.52	90.76
S15	65.78	66.09	68.7	66.13	60.05	76.15	59.64	64.68	76.06	69.72	69.59	72.38	70.24	70.84
stm	80.4	80.69	80.38	81.8	81.31	85.73	80.67	79.81	82.91	78.97	80.95	76.18	84.05	76.48
per	80.23	75.31	79.88	77.39	77.42	73.81	73.57	74.74	75.33	70.19	74.49	74.65	74.45	71.14
hob	88.38	86.58	87.46	91.29	87.27	94.05	84.82	89.55	87.89	83.56	87.0	88.53	93.63	91.57
iph	77.51	76.75	76.16	77.32	81.15	77.63	73.61	76.88	74.66	73.92	77.76	78.26	78.2	78.46
mov	70.74	72.76	77.86	76.15	71.93	75.07	68.27	74.61	77.11	73.42	74.73	75.63	75.61	71.03
san	81.62	78.38	82.37	82.65	82.64	81.36	84.56	82.35	83.21	79.24	79.95	80.53	82.81	82.16
Nar	84.9	82.2	83.17	83.06	86.46	84.37	79.48	82.05	87.99	85.4	87.16	87.71	85.93	82.12
arc	84.33	83.94	84.39	85.05	85.9	85.6	85.04	83.32	86.8	79.28	82.65	83.74	87.29	85.22
S18	76.88	75.65	79.33	78.28	78.14	82.78	81.49	72.85	80.59	76.29	78.02	78.19	84.48	80.19
OMD	81.59	78.48	78.96	76.27	73.95	80.39	75.03	73.92	80.58	70.91	71.77	73.73	82.51	79.42
HCR	70.89	69.1	70.07	71.2	71.87	71.39	70.02	68.07	70.59	58.39	65.86	67.01	72.59	69.49
STS	82.37	80.75	84.48	84.32	84.38	84.72	81.23	83.85	86.04	85.85	82.34	86.27	85.2	83.05
SSt	76.36	71.2	75.52	75.4	76.15	76.36	70.23	72.86	78.47	74.95	81.63	76.87	77.55	72.83
Tar	81.47	78.33	80.98	80.52	79.57	80.3	78.5	80.04	80.38	74.44	79.65	79.85	81.56	79.94
Vad	84.85	83.85	82.92	83.22	83.14	85.14	78.86	82.72	86.16	82.15	81.66	84.39	87.32	82.82
S13	76.81	73.46	76.37	75.24	76.17	77.55	72.75	74.03	77.47	75.12	83.67	77.48	79.2	77.06
S17	86.35	84.7	86.53	85.51	85.73	86.22	84.65	80.91	84.57	75.82	82.09	82.4	86.74	85.99
S16	79.27	78.55	79.42	79.63	79.45	79.82	78.89	75.33	79.69	73.28	75.76	76.99	81.05	79.83
#wins	1	1	1	0	1	3	1	0	2	1	2	1	8	0
rank sums	146.5	213.0	149.5	144.0	152.0	112.5	243.0	222.0	95.0	228.0	191.0	164.0	75.5	174.0
position	5.0	11.0	6.0	4.0	7.0	3.0	14.0	12.0	2.0	13.0	10.0	8.0	1.0	9.0

Table A.5: Accuracies (%) achieved using the LR classifier

Dataset	w2v GN	GloVe WP	Fast Text	EWE	GloVe TWT	BERT static	TF- IDF	w2v Araque	w2v Edin	SSWE	Emo2 Vec	Deep Moji	RoB static	BTWT static
iro	70.0	70.71	66.9	66.9	69.05	60.95	57.86	65.71	59.76	71.43	78.81	61.43	62.62	52.14
sar	74.46	67.32	64.46	71.61	73.04	67.68	54.64	70.36	70.36	87.5	74.64	60.36	56.25	53.39
ntu	79.52	89.54	71.6	74.09	76.94	73.74	73.7	73.76	88.85	93.16	80.21	92.46	89.93	71.28
S15	81.32	81.32	74.81	85.08	77.59	84.72	90.35	79.74	82.55	76.03	80.4	84.43	86.6	81.3
stm	84.94	83.29	82.74	84.11	84.67	76.06	80.2	80.51	81.9	82.15	84.13	80.22	83.0	71.02
per	82.01	81.32	77.21	80.17	80.86	70.14	76.52	73.11	71.75	76.09	77.46	75.84	77.21	68.1
hob	84.47	88.32	76.82	91.56	87.54	80.82	91.37	88.5	75.09	80.07	86.61	88.31	90.79	79.3
iph	74.62	78.0	73.12	76.88	77.81	69.36	76.31	74.81	72.18	73.85	79.32	75.56	77.64	72.95
mov	79.51	80.57	77.72	81.64	81.29	75.24	84.31	78.62	72.21	82.53	82.35	77.9	82.53	69.72
san	80.14	79.0	76.63	80.63	80.79	70.19	81.86	78.19	77.69	80.88	80.15	79.73	80.47	72.63
Nar	84.19	84.35	82.64	84.43	87.37	77.59	80.44	83.29	83.05	89.0	88.26	87.29	85.33	74.91
arc	83.18	83.0	83.7	84.05	85.27	77.24	86.5	83.0	83.29	79.11	81.72	81.55	86.5	78.87
S18	80.31	78.75	80.1	80.53	82.04	75.53	82.3	75.68	78.48	81.23	81.23	77.94	84.89	75.52
OMD	81.58	80.48	79.07	78.33	77.38	73.55	78.59	75.81	76.55	75.34	76.28	75.02	81.53	75.92
HCR	72.43	71.9	68.92	73.0	74.26	65.2	78.09	68.86	66.77	65.09	67.55	68.34	74.36	65.77
STS	85.25	83.48	84.12	86.09	86.67	78.86	84.42	86.18	82.1	89.08	85.45	87.71	85.79	76.5
SSt	77.5	76.06	75.58	80.12	78.42	71.17	73.57	76.32	76.85	80.26	85.06	79.29	79.73	68.33
Tar	82.49	80.94	82.38	82.64	83.44	76.92	81.63	81.08	79.52	78.22	82.38	81.94	83.07	73.93
Vad	87.13	84.87	86.13	87.32	86.3	80.39	83.65	83.91	82.65	85.7	85.46	85.58	88.92	76.62
S13	80.36	78.78	78.42	79.67	80.24	72.77	78.73	77.25	76.06	81.75	87.83	79.99	81.11	71.13
S17	87.54	86.23	86.95	87.19	87.36	82.87	87.14	82.4	84.78	79.97	85.43	84.17	87.88	81.68
S16	83.69	82.43	83.22	83.46	83.79	79.54	84.18	78.57	80.3	79.19	82.28	80.89	84.61	76.19
#wins	3	0	0	1	1	0	5	0	0	4	4	0	5	0
rank sums	116.5	149.0	192.5	101.5	97.0	264.0	146.5	204.0	220.5	152.0	120.0	174.0	82.5	290.0
position	4.0	7.0	10.0	3.0	2.0	13.0	6.0	11.0	12.0	8.0	5.0	9.0	1.0	14.0

Table A.6: F_1 -macro scores (%) achieved by evaluating static representation using the LR classifier

Dataset	w2v GN	GloVe WP	Fast Text	EWE	GloVe TWT	BERT static	TF- IDF	w2v Araque	w2v Edin	SSWE	Emo2 Vec	Deep Moji	RoB static	BTWT static
iro	63.88	65.27	62.4	59.61	63.7	56.53	36.27	57.73	53.05	68.67	75.87	55.74	56.27	46.16
sar	74.02	66.17	62.88	70.84	71.25	66.16	50.32	68.0	68.45	87.19	73.32	57.75	54.23	52.23
ntu	79.2	89.21	71.0	73.22	76.14	73.46	72.77	73.25	88.71	92.99	79.87	92.27	89.57	70.89
S15	67.06	66.39	61.94	73.19	64.38	72.61	75.63	65.1	70.26	63.92	67.94	71.54	74.79	67.76
stm	84.89	83.21	82.62	84.08	84.65	75.95	80.1	80.33	81.71	82.12	84.08	80.13	82.92	70.79
per	79.75	78.36	74.63	77.56	77.71	66.57	73.01	70.55	69.58	73.02	75.21	73.51	75.24	64.82
hob	83.17	86.79	75.42	90.61	86.39	79.66	90.2	87.5	73.98	78.31	85.38	87.33	89.96	77.93
iph	72.53	76.05	71.69	74.98	75.59	68.06	72.07	72.84	70.8	71.27	77.33	73.85	75.55	71.42
mov	71.37	71.11	68.56	73.23	73.34	65.78	70.61	70.68	63.91	75.4	75.11	70.04	74.01	59.89
san	79.91	78.88	76.34	80.51	80.67	69.61	81.58	77.94	77.33	80.78	80.05	79.57	80.17	71.99
Nar	83.75	83.79	82.27	83.95	86.99	77.14	79.82	82.88	82.65	88.58	87.91	86.93	84.88	74.37
arc	82.84	82.6	83.32	83.75	84.94	76.66	86.11	82.62	82.78	78.62	81.23	80.98	86.15	78.27
S18	80.16	78.62	79.91	80.42	81.91	75.22	82.03	75.43	78.15	81.12	81.13	77.71	84.78	75.21
OMD	80.34	79.32	77.58	77.14	76.26	71.78	77.02	74.41	74.91	74.19	75.05	73.85	80.33	73.6
HCR	69.42	68.56	66.04	69.58	71.17	62.42	73.37	64.9	64.45	61.89	64.29	64.59	71.24	63.03
STS	83.57	81.63	82.17	84.38	85.04	76.3	81.73	84.47	79.85	87.7	83.74	86.1	84.05	74.02
SSt	77.16	75.68	75.25	79.79	78.05	70.94	72.99	75.94	76.66	79.91	84.77	78.98	79.36	68.08
Tar	82.49	80.92	82.37	82.63	83.43	76.9	81.61	81.06	79.51	78.21	82.36	81.9	83.06	73.89
Vad	85.52	83.02	84.47	85.71	84.55	78.28	81.18	82.1	81.02	84.0	83.74	83.92	87.45	74.57
S13	77.46	75.75	75.71	76.61	77.21	70.19	74.35	74.39	73.47	78.91	85.59	77.31	78.24	68.25
S17	86.92	85.54	86.3	86.56	86.7	81.94	86.32	81.39	83.98	78.99	84.69	83.27	87.22	80.63
S16	80.72	79.46	80.19	80.54	80.96	76.55	81.04	75.68	77.31	76.16	79.38	78.08	81.8	73.03
#wins	3	0	0	1	1	0	3	0	0	5	4	0	5	0
rank sums	113.0	152.0	189.0	102.5	96.0	261.0	166.0	200.0	221.0	152.0	117.5	169.0	82.0	289.0
position	4.0	6.5	10.0	3.0	2.0	13.0	8.0	11.0	12.0	6.5	5.0	9.0	1.0	14.0

Table A.7: Accuracies (%) achieved using the XGB classifier

Dataset	w2v GN	GloVe WP	Fast Text	EWE	GloVe TWT	BERT static	TF- IDF	w2v Araque	w2v Edin	SSWE	Emo2 Vec	Deep Moji	RoB static	BTWT static
iro	72.38	67.62	58.57	61.9	62.62	66.67	42.38	59.52	69.76	65.48	71.19	66.43	55.24	65.71
sar	63.57	63.21	66.25	70.54	71.79	58.93	51.96	58.93	57.5	73.21	70.36	56.25	46.25	49.29
ntu	75.22	82.37	68.73	71.24	75.19	76.59	62.61	75.2	89.92	92.08	76.22	87.79	82.75	78.4
S15	85.36	86.6	85.05	86.31	85.98	86.92	88.78	87.23	86.93	86.6	85.99	88.48	86.91	88.17
stm	80.49	75.49	81.33	77.15	79.95	74.11	69.09	79.91	81.34	81.05	83.56	76.05	78.0	67.41
per	77.21	79.27	78.82	79.5	78.12	73.58	73.57	76.08	77.91	77.0	79.05	78.35	79.95	69.7
hob	82.96	85.04	81.6	90.99	86.19	87.73	92.71	88.31	84.68	80.28	88.89	87.53	86.97	88.51
iph	79.13	81.38	81.0	79.5	81.94	77.82	73.3	80.64	82.12	78.74	81.59	80.83	78.96	76.32
mov	85.56	82.17	84.85	84.85	83.78	83.6	83.79	83.07	84.14	83.78	88.78	86.1	84.49	82.89
san	80.71	79.33	79.81	80.06	81.03	79.9	77.86	75.97	81.29	81.95	80.71	82.43	79.65	76.8
Nar	82.72	80.2	81.08	83.05	85.57	80.44	77.6	81.09	84.19	88.83	87.78	87.04	81.1	76.86
arc	84.34	82.71	85.62	84.17	85.86	82.71	83.99	83.88	87.14	82.77	84.46	82.95	84.75	82.48
S18	78.16	76.22	78.16	76.17	78.16	77.3	80.42	71.0	78.7	80.26	80.58	78.05	79.02	73.37
OMD	80.59	80.38	79.12	77.7	77.49	80.17	75.5	75.44	80.43	77.28	75.97	76.08	79.9	80.27
HCR	78.93	76.36	76.99	77.41	79.29	77.72	78.93	75.36	78.04	74.42	75.26	76.57	78.98	76.15
STS	83.97	81.81	84.07	83.88	83.97	81.61	81.61	84.96	83.38	88.89	86.93	86.18	81.32	78.56
SSt	76.58	72.7	74.97	73.96	76.01	72.87	70.12	74.71	76.32	79.86	83.92	76.89	73.92	68.98
Tar	80.33	79.98	80.93	80.24	81.22	78.31	79.98	79.15	80.18	78.28	81.86	79.64	78.91	76.15
Vad	85.7	83.13	85.51	84.03	85.6	81.7	81.63	83.53	85.87	86.15	86.8	85.37	84.08	78.93
S13	81.02	79.26	80.4	80.45	81.04	79.19	78.8	79.58	81.84	83.33	88.24	81.73	79.95	75.95
S17	87.32	85.05	87.16	85.58	87.06	85.28	85.19	82.07	85.33	81.24	85.03	83.52	85.33	80.73
S16	84.31	83.46	84.96	84.29	85.01	82.33	83.47	81.01	84.11	81.66	83.49	82.15	83.54	79.91
#wins	3	0	0	0	2	0	2	0	2	4	7	1	1	0
rank sums	125.5	195.5	151.5	153.5	123.0	207.5	215.5	214.5	106.5	153.0	98.5	144.0	167.5	254.0
position	4.0	10.0	6.0	8.0	3.0	11.0	13.0	12.0	2.0	7.0	1.0	5.0	9.0	14.0

Table A.8: F_1 -macro scores (%) achieved by evaluating static representation using the XGB classifier

Dataset	w2v GN	GloVe WP	Fast Text	EWE	GloVe TWT	BERT static	TF- IDF	w2v Araque	w2v Edin	SSWE	Emo2 Vec	Deep Moji	RoB static	BTWT static
iro	63.3	54.98	42.49	46.16	47.46	53.86	30.21	46.16	61.54	56.22	62.7	53.71	41.13	54.88
sar	61.87	61.94	65.06	69.41	70.85	53.87	49.32	56.77	56.1	72.17	68.52	55.59	43.19	44.94
ntu	74.12	81.55	66.91	69.66	73.9	75.17	60.93	73.57	89.58	91.87	75.41	87.35	82.05	77.55
S15	51.77	56.34	48.79	56.35	54.95	62.95	72.27	57.46	58.18	59.58	61.55	67.55	60.93	61.85
stm	80.43	75.24	81.25	77.05	79.83	73.94	68.8	79.87	81.18	80.97	83.47	75.98	77.95	66.91
per	68.33	71.55	71.29	72.1	69.99	60.54	66.27	66.93	69.67	69.02	73.31	70.46	72.51	55.75
hob	79.57	81.82	77.25	89.31	83.48	85.24	91.77	86.16	81.81	76.1	87.0	85.37	84.46	85.99
iph	73.42	75.1	75.54	73.28	77.1	70.45	65.9	74.83	77.84	73.31	77.2	75.39	71.69	68.78
mov	64.5	54.7	61.87	61.59	59.34	54.96	64.42	55.64	61.04	69.46	77.58	67.72	61.38	53.12
san	80.51	79.12	79.6	79.84	80.83	79.62	77.56	75.72	81.07	81.79	80.54	82.25	79.38	76.37
Nar	81.72	79.02	79.66	81.93	84.81	78.8	76.49	80.0	83.41	88.34	87.21	86.42	79.73	74.88
arc	83.75	82.06	85.08	83.7	85.38	82.02	83.56	83.24	86.65	82.0	83.82	82.16	84.25	81.82
S18	77.9	75.94	77.94	75.9	77.88	77.02	79.92	70.49	78.51	80.06	80.42	77.78	78.75	72.94
OMD	78.38	77.75	76.48	74.75	74.88	77.39	72.54	72.16	77.86	74.5	73.11	73.29	77.31	77.39
HCR	69.14	64.37	65.49	66.73	69.81	67.54	71.02	61.49	67.75	63.33	63.49	64.4	69.19	63.49
STS	79.78	77.0	79.72	79.74	80.07	75.87	76.91	81.37	79.07	86.79	84.4	82.98	75.01	70.87
SSt	75.52	71.24	73.8	72.62	74.82	71.15	68.54	73.53	75.22	79.26	83.38	75.9	72.42	67.16
Tar	80.32	79.97	80.92	80.24	81.21	78.29	79.97	79.13	80.17	78.26	81.83	79.62	78.9	76.12
Vad	82.16	78.77	81.87	79.96	82.17	76.04	75.91	79.41	82.42	83.53	84.22	82.11	79.45	71.12
S13	72.81	70.13	71.94	71.78	73.47	68.52	68.47	71.18	75.01	78.28	84.81	74.72	70.23	61.56
S17	86.31	83.75	86.09	84.37	85.98	83.81	83.79	80.09	84.07	79.46	83.82	81.92	83.94	78.51
S16	78.67	77.17	79.33	78.33	79.72	75.1	76.43	73.62	78.36	75.68	78.23	75.61	76.84	69.99
#wins	3	0	0	0	1	0	3	0	2	4	8	1	0	0
rank sums	129.0	195.5	153.0	160.5	123.0	217.5	211.5	216.5	110.0	135.0	81.5	143.0	177.0	257.0
position	4.0	10.0	7.0	8.0	3.0	13.0	11.0	12.0	2.0	5.0	1.0	6.0	9.0	14.0

A.2 Transformer-based Text Representation

Table A.9: Accuracies and F_1 -macro scores (%) achieved by evaluating Transformer-Autoencoder language models using the RF classifier

Dataset	Accuracy			F_1 -macro		
	RoBERTa	BERT	BERTweet	RoBERTa	BERT	BERTweet
iro	66.19	67.62	67.62	39.77	43.48	43.48
sar	66.07	67.32	64.64	61.2	64.83	62.33
ntu	83.1	83.06	84.88	82.09	81.77	83.58
S15	85.67	85.36	86.29	47.79	46.05	51.62
stm	86.62	83.02	88.31	86.58	82.97	88.26
per	76.08	77.23	76.09	59.57	61.8	60.49
hob	78.37	77.4	80.66	68.92	67.02	72.74
iph	80.06	80.25	82.13	71.21	71.96	75.55
mov	82.18	82.35	82.17	46.01	46.89	45.93
san	84.97	82.6	85.05	84.74	82.3	84.9
Nar	87.37	85.41	90.14	86.34	84.25	89.4
arc	88.36	86.09	88.36	87.9	85.51	87.93
S18	83.54	80.69	83.76	83.17	80.17	83.42
OMD	79.07	79.28	80.43	74.79	75.07	76.63
HCR	75.63	75.89	76.89	55.65	56.58	59.47
STS	82.89	81.71	82.4	76.1	74.36	75.04
SSt	80.08	78.68	79.12	78.63	77.15	77.45
Tar	81.83	82.35	82.72	81.82	82.34	82.71
Vad	82.41	82.55	81.08	75.51	75.67	72.66
S13	81.11	79.88	80.22	69.45	66.9	67.55
S17	87.52	86.54	87.84	86.2	85.1	86.57
S16	84.41	84.01	84.87	76.83	76.1	77.55
#wins	4	5	15	3	5	15
rank sums	46.5	52.5	33.0	48.0	52.5	31.5
position	2.0	3.0	1.0	2.0	3.0	1.0

Table A.10: Accuracies and F_1 -macro scores (%) achieved by evaluating Transformer-Autoencoder language models using the MLP classifier

Dataset	Accuracy			F_1 -macro		
	RoBERTa	BERT	BERTweet	RoBERTa	BERT	BERTweet
iro	70.71	80.48	73.33	63.36	70.87	64.96
sar	68.75	71.79	70.54	66.95	70.21	69.75
ntu	86.35	88.82	90.32	85.96	88.52	90.08
S15	88.48	91.89	92.23	70.62	78.77	82.42
stm	88.57	87.2	89.43	88.53	87.12	89.4
per	86.8	84.75	86.55	83.71	81.6	83.35
hob	87.74	85.06	86.0	85.73	82.9	83.65
iph	85.54	83.47	87.59	82.57	80.06	84.72
mov	89.47	88.41	87.69	81.11	78.5	78.11
san	88.88	87.58	91.17	88.78	87.5	91.11
Nar	91.2	91.12	95.6	90.8	90.73	95.43
arc	89.75	88.89	90.74	89.52	88.6	90.51
S18	87.41	85.85	88.65	87.34	85.77	88.57
OMD	85.15	84.84	87.15	83.98	83.73	86.12
HCR	80.76	79.56	81.08	75.53	74.03	76.14
STS	91.94	90.66	93.71	90.54	89.02	92.7
SSt	84.93	84.36	85.8	84.46	83.89	85.3
Tar	83.53	83.7	85.61	83.53	83.7	85.6
Vad	90.04	89.2	90.61	88.25	87.32	88.97
S13	87.44	87.09	89.54	84.0	83.67	86.87
S17	90.74	91.07	92.09	90.08	90.43	91.52
S16	89.75	88.53	90.96	86.99	85.43	88.5
#wins	3	2	17	3	2	17
rank sums	47.0	57.0	28.0	47.0	57.0	28.0
position	2.0	3.0	1.0	2.0	3.0	1.0

Table A.11: Accuracies and F_1 -macro scores (%) achieved by evaluating Transformer-Autoencoder language models using the LR classifier

Dataset	Accuracy			F_1 -macro		
	RoBERTa	BERT	BERTweet	RoBERTa	BERT	BERTweet
iro	65.95	80.48	72.14	57.55	73.08	65.1
sar	70.18	71.79	69.11	68.52	70.79	67.41
ntu	85.99	87.75	91.03	85.71	87.41	90.8
S15	86.93	90.66	92.23	69.44	78.66	83.33
stm	90.79	85.8	88.31	90.75	85.74	88.29
per	86.34	84.97	87.69	83.6	82.27	85.29
hob	87.93	86.01	87.16	86.29	84.27	85.4
iph	85.53	83.09	86.09	82.94	79.65	83.63
mov	88.93	87.7	88.95	81.2	78.63	82.12
san	89.05	87.01	91.01	88.95	86.93	90.95
Nar	91.11	91.04	95.11	90.73	90.67	94.95
arc	89.0	88.59	90.57	88.75	88.32	90.36
S18	87.04	83.76	88.54	86.96	83.65	88.47
OMD	84.37	84.15	87.15	83.34	83.09	86.26
HCR	80.45	79.24	80.5	76.42	74.75	76.77
STS	92.43	90.51	93.9	91.23	88.92	92.98
SSt	84.67	83.66	86.46	84.27	83.2	86.1
Tar	83.91	82.32	85.95	83.9	82.32	85.95
Vad	89.97	89.3	90.8	88.43	87.59	89.38
S13	87.89	86.78	89.61	85.15	83.7	87.37
S17	90.88	90.18	92.31	90.32	89.57	91.8
S16	89.39	88.63	91.03	87.08	86.06	89.05
#wins	2	2	18	2	2	18
rank sums	45.0	60.0	27.0	45.0	60.0	27.0
position	2.0	3.0	1.0	2.0	3.0	1.0

Table A.12: Accuracies and F_1 -macro scores (%) achieved by evaluating Transformer-Autoencoder language models using the XGB classifier

Dataset	Accuracy			F_1 -macro		
	RoBERTa	BERT	BERTweet	RoBERTa	BERT	BERTweet
iro	61.67	60.24	61.9	48.3	48.24	47.05
sar	64.64	62.14	68.93	63.08	60.37	67.47
ntu	83.08	82.72	85.62	82.54	82.24	85.04
S15	89.1	87.84	90.03	67.85	58.92	71.62
stm	86.6	82.46	81.02	86.57	82.43	80.97
per	81.09	80.87	82.0	73.83	74.16	75.37
hob	82.95	80.27	83.9	79.03	75.67	80.29
iph	83.09	80.82	83.09	77.99	75.27	78.88
mov	85.21	84.49	87.88	64.4	59.81	72.09
san	85.14	85.05	88.15	85.03	84.91	88.06
Nar	90.63	88.59	91.6	90.16	88.01	91.19
arc	89.47	87.31	90.63	89.18	86.93	90.4
S18	85.37	84.62	86.12	85.24	84.45	85.98
OMD	83.58	83.63	84.73	81.69	81.74	82.94
HCR	81.39	79.92	81.55	73.07	71.37	74.11
STS	89.92	87.86	90.51	87.73	85.04	88.36
SSt	83.62	82.79	82.83	82.99	82.06	82.08
Tar	83.36	83.33	84.51	83.34	83.32	84.5
Vad	87.63	87.61	88.01	84.85	84.75	85.13
S13	86.52	85.47	87.07	81.79	80.39	82.6
S17	89.85	89.93	90.53	89.06	89.13	89.78
S16	88.99	88.42	89.74	85.52	84.64	86.51
#wins	3	0	20	3	0	19
rank sums	43.5	63.0	25.5	44.0	61.0	27.0
position	2.0	3.0	1.0	2.0	3.0	1.0

A.3 Transformer-based Text Representation tuned With a Large Collection of English Tweets

Table A.13: Average classification accuracies (%) achieved by fine-tuning BERT with different samples of generic unlabeled tweets, using the RF classifier

Dataset	BERT									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	68.89	69.52	67.7	67.7	68.81	69.44	67.86	66.19	66.75	66.11
sar	68.33	65.53	72.14	68.33	70.24	75.42	67.8	70.36	68.93	65.24
ntu	82.1	82.58	80.07	80.7	77.57	79.61	82.36	88.84	88.24	89.67
S15	85.46	85.46	85.36	85.57	85.67	85.98	85.88	85.56	85.46	85.77
stm	83.66	85.71	84.12	84.04	84.22	83.12	84.97	85.62	84.95	83.57
per	76.69	77.3	77.83	78.59	77.83	78.3	78.51	76.46	74.87	74.42
hob	79.82	79.64	79.57	80.21	81.1	81.3	80.73	81.42	81.1	82.12
iph	82.14	82.51	82.25	82.57	82.2	81.5	81.26	80.81	81.38	82.19
mov	82.12	82.06	82.06	82.06	82.06	82.36	82.18	82.36	82.23	82.3
san	82.79	83.17	84.69	84.07	83.28	84.72	84.61	83.41	83.39	83.01
Nar	86.63	87.64	87.01	87.91	86.55	87.91	89.0	89.0	89.65	88.8
arc	86.56	86.77	86.77	87.21	86.59	87.25	88.01	87.99	87.58	88.69
S18	81.78	81.96	81.89	81.75	82.12	82.27	82.97	82.39	82.3	82.48
OMD	80.08	80.84	80.26	79.69	79.7	78.87	80.31	78.91	78.26	79.22
HCR	76.27	76.29	76.13	76.48	75.61	76.07	76.36	75.91	75.64	75.23
STS	82.56	82.73	83.02	82.66	82.81	83.27	84.7	84.96	84.88	84.6
SSt	80.76	80.4	80.56	80.05	79.67	79.54	80.63	79.35	79.87	79.56
Tar	82.45	82.6	83.1	82.35	82.22	82.22	82.47	81.81	81.92	81.53
Vad	82.97	83.14	83.65	83.09	83.01	83.65	84.32	83.95	84.29	84.33
S13	81.19	81.7	81.95	81.72	81.82	81.6	82.42	82.07	81.38	81.31
S17	86.6	87.12	87.02	87.1	87.38	87.08	87.27	86.38	85.12	85.07
S16	84.49	84.83	84.74	84.91	84.72	84.49	84.79	84.16	83.35	83.26
#wins	1	3	1	4	1	4	2	2	1	4
rank sums	153.0	114.0	120.5	119.0	134.0	113.5	73.5	114.0	131.5	137.0
position	10.0	3.5	6.0	5.0	8.0	2.0	1.0	3.5	7.0	9.0

Table A.14: Average F_1 -macro scores (%) achieved by fine-tuning BERT with different samples of unlabeled tweets, using the RF classifier

Dataset	BERT									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	46.05	46.81	46.33	44.26	45.46	47.31	43.55	39.77	41.03	41.54
sar	66.11	63.16	70.12	65.15	67.11	73.87	65.32	67.88	66.63	60.38
ntu	80.83	81.62	78.65	79.53	75.96	78.55	81.65	88.43	87.68	89.24
S15	46.63	46.63	46.05	47.21	47.79	49.54	48.96	47.21	46.63	48.48
stm	83.6	85.67	84.04	83.99	84.17	83.03	84.89	85.57	84.85	83.48
per	61.1	61.98	63.27	65.02	63.84	64.73	65.32	60.55	56.95	55.83
hob	71.4	71.1	71.37	72.28	74.3	74.42	73.21	75.01	74.18	76.11
iph	75.45	76.04	75.57	75.77	75.24	74.4	73.69	73.22	74.19	75.05
mov	45.66	45.67	45.34	45.34	45.37	46.94	45.98	46.94	46.27	46.87
san	82.51	82.94	84.46	83.84	83.04	84.5	84.37	83.2	83.11	82.69
Nar	85.59	86.72	86.02	86.98	85.5	87.15	88.32	88.28	88.97	87.99
arc	85.93	86.13	86.18	86.67	85.97	86.72	87.56	87.46	87.03	88.24
S18	81.28	81.51	81.49	81.34	81.7	81.82	82.61	81.93	81.84	82.04
OMD	76.24	77.24	76.43	75.81	75.69	74.74	76.52	74.46	73.5	74.79
HCR	57.72	57.87	57.67	58.59	56.1	57.35	58.83	57.08	56.32	55.11
STS	75.97	76.32	76.8	76.17	76.68	77.33	79.64	79.67	79.57	78.96
SSt	79.54	79.21	79.24	78.83	78.47	78.31	79.63	78.18	78.7	78.21
Tar	82.43	82.59	83.09	82.34	82.2	82.2	82.44	81.79	81.89	81.51
Vad	76.53	77.06	77.72	76.92	76.77	78.01	79.15	78.28	78.77	78.82
S13	70.02	71.53	71.73	71.24	71.5	71.69	73.23	72.2	70.71	70.88
S17	85.11	85.74	85.64	85.74	86.03	85.67	85.9	84.82	83.33	83.27
S16	76.96	77.8	77.55	77.84	77.57	77.21	77.63	76.38	74.74	74.61
#wins	0	3	1	1	1	5	6	2	1	3
rank sums	155.0	113.5	118.5	124.5	136.5	106.0	68.0	116.0	135.0	137.0
position	10.0	3.0	5.0	6.0	8.0	2.0	1.0	4.0	7.0	9.0

Table A.15: Average classification accuracies (%) achieved by fine-tuning BERT with different samples of generic unlabeled tweets, using the LR classifier

Dataset	BERT									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	74.05	70.71	70.95	66.67	75.24	72.14	76.43	81.9	75.0	76.35
sar	63.21	66.07	61.79	63.21	64.64	68.93	67.5	69.11	71.19	72.86
ntu	89.54	89.18	85.97	84.19	83.49	85.62	87.04	92.06	93.87	93.75
S15	91.89	91.59	90.96	88.47	89.08	90.97	88.14	90.68	88.16	88.37
stm	86.09	85.81	85.51	87.48	86.34	87.46	87.74	88.02	88.3	87.38
per	85.65	84.74	84.74	84.51	85.19	84.51	82.69	84.05	83.07	82.56
hob	86.97	86.78	89.27	87.73	88.11	86.4	88.51	88.31	87.61	88.95
iph	84.02	83.83	85.15	84.77	83.64	82.89	83.83	84.01	83.71	84.83
mov	86.64	88.06	88.23	88.05	87.34	87.16	86.98	86.99	88.53	89.37
san	87.01	86.03	86.19	86.27	86.68	87.09	88.06	86.6	87.31	88.18
Nar	90.38	90.54	90.71	91.04	90.87	90.87	92.75	93.32	93.97	93.29
arc	88.77	88.71	89.0	88.42	88.42	87.2	89.06	89.0	88.22	89.1
S18	85.21	83.76	84.56	83.7	84.62	83.59	84.99	85.48	85.91	85.87
OMD	83.42	84.15	83.52	84.0	84.63	83.21	83.95	83.95	84.07	84.94
HCR	78.82	77.67	79.66	77.93	78.56	78.77	77.56	77.78	78.33	77.51
STS	90.71	90.41	90.22	89.92	89.62	91.25	91.44	90.81	91.87	92.34
SSt	83.01	83.18	82.96	82.35	82.61	82.0	81.83	81.96	83.29	84.07
Tar	82.64	83.13	83.07	83.3	82.95	83.36	82.69	83.13	83.69	83.92
Vad	88.97	89.23	89.42	88.78	89.11	89.18	89.44	89.23	90.13	90.71
S13	86.96	86.71	87.6	86.3	87.0	87.07	86.93	87.6	87.46	87.79
S17	89.76	90.14	90.14	90.01	89.81	89.73	89.68	89.7	88.98	89.43
S16	88.83	88.56	88.68	88.71	88.38	87.8	88.56	87.84	87.96	88.17
#wins	3	1	4	0	0	0	0	1	4	10
rank sums	127.5	134.0	115.0	143.5	136.0	146.0	127.5	110.5	96.0	74.0
position	5.5	7.0	4.0	9.0	8.0	10.0	5.5	3.0	2.0	1.0

Table A.16: Average F_1 -macro scores (%) achieved by fine-tuning BERT with different samples of unlabeled tweets, using the LR classifier

BERT										
Dataset	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	65.37	60.93	59.85	57.44	64.84	64.25	67.52	75.87	68.01	67.32
sar	61.96	64.44	60.36	62.09	63.52	66.89	66.9	68.12	70.54	71.87
ntu	89.27	88.9	85.62	83.83	82.98	85.26	86.78	91.86	93.72	93.6
S15	80.13	79.03	78.24	73.86	75.57	78.95	75.2	79.71	75.72	74.85
stm	86.05	85.75	85.44	87.45	86.27	87.4	87.69	87.97	88.27	87.34
per	82.87	81.54	81.74	81.59	82.14	81.48	79.08	80.83	79.75	79.29
hob	85.15	84.77	87.84	85.99	86.54	84.49	86.88	86.75	85.9	87.49
iph	81.0	80.74	82.15	81.76	80.83	79.63	81.01	81.35	80.66	82.24
mov	77.9	79.57	80.68	80.17	78.38	78.93	78.65	77.88	81.4	82.7
san	86.93	85.94	86.09	86.17	86.61	87.01	88.0	86.52	87.24	88.12
Nar	89.99	90.15	90.35	90.66	90.51	90.51	92.46	93.05	93.73	93.04
arc	88.51	88.46	88.74	88.18	88.18	86.91	88.82	88.75	87.96	88.86
S18	85.13	83.66	84.47	83.61	84.53	83.48	84.89	85.39	85.83	85.79
OMD	82.41	83.07	82.47	82.97	83.63	82.23	82.94	82.98	83.08	84.0
HCR	74.51	73.02	74.95	73.39	74.0	74.5	73.21	73.11	74.04	73.36
STS	89.26	88.86	88.7	88.33	87.95	89.86	90.03	89.42	90.58	91.16
SSt	82.48	82.67	82.42	81.88	82.13	81.48	81.29	81.49	82.79	83.64
Tar	82.63	83.12	83.07	83.29	82.95	83.35	82.69	83.12	83.69	83.92
Vad	87.26	87.56	87.72	87.06	87.42	87.48	87.77	87.53	88.61	89.33
S13	83.99	83.73	84.71	83.1	84.12	84.18	83.93	84.72	84.72	85.17
S17	89.13	89.53	89.51	89.39	89.2	89.08	89.05	89.07	88.33	88.81
S16	86.37	86.04	86.18	86.23	85.83	85.13	86.09	85.2	85.41	85.7
#wins	3	1	2	0	0	0	0	1	4	11
rank sums	128.0	140.5	120.0	144.5	140.0	148.5	122.0	108.0	87.5	71.0
position	6.0	8.0	4.0	9.0	7.0	10.0	5.0	3.0	2.0	1.0

Table A.17: Average classification accuracies (%) achieved by fine-tuning BERT with different samples of generic unlabeled tweets, using the MLP classifier

Dataset	BERT									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	74.53	70.88	72.06	69.68	73.81	79.52	76.11	78.57	77.86	75.0
sar	69.88	68.45	68.45	67.02	64.64	73.57	68.39	72.68	71.67	73.75
ntu	89.3	88.59	85.61	85.74	83.6	86.59	87.64	92.06	93.27	93.52
S15	92.21	91.89	90.97	90.12	90.24	91.06	89.29	89.94	89.2	89.92
stm	85.9	86.45	86.53	86.36	86.89	86.81	88.21	87.28	88.76	86.92
per	85.8	86.18	84.59	84.66	84.81	85.35	83.37	83.37	83.22	83.76
hob	86.72	86.78	88.06	87.29	87.92	87.55	89.73	89.08	87.62	89.21
iph	84.96	84.21	84.53	84.97	84.02	83.33	84.59	84.07	84.4	85.7
mov	87.76	87.94	87.99	87.7	87.28	87.58	88.23	88.24	88.89	89.6
san	87.01	86.79	87.36	87.53	87.5	87.5	89.0	86.74	87.47	88.83
Nar	90.14	89.86	90.73	90.87	91.06	91.12	92.28	93.21	93.86	92.99
arc	88.59	88.77	89.37	89.14	89.29	88.2	89.21	89.31	88.67	88.75
S18	85.48	84.62	85.5	84.8	85.55	84.56	84.76	85.62	85.3	85.01
OMD	84.9	84.8	85.33	86.36	84.96	84.68	84.84	85.66	85.75	84.61
HCR	79.44	78.93	79.38	79.1	78.63	79.57	79.31	78.65	79.73	78.28
STS	91.22	91.07	90.61	90.97	90.52	91.54	91.97	91.71	91.92	92.08
SSt	83.55	83.54	83.47	83.09	83.3	82.6	83.14	83.49	83.94	84.0
Tar	83.45	84.03	84.1	84.43	83.68	84.5	84.34	83.31	83.22	83.53
Vad	89.54	89.5	89.58	89.13	89.78	89.86	90.57	89.78	90.13	90.41
S13	87.39	87.62	87.92	87.48	87.65	88.12	88.08	88.24	87.99	87.58
S17	90.78	90.71	90.85	90.63	90.5	89.98	90.04	89.55	88.81	88.86
S16	88.46	88.74	88.85	88.42	88.41	88.56	88.59	87.77	87.96	87.74
#wins	1	1	3	1	0	2	3	2	3	6
rank sums	129.0	139.5	114.5	142.0	141.0	121.5	101.5	108.0	107.0	106.0
position	7.0	8.0	5.0	10.0	9.0	6.0	1.0	4.0	3.0	2.0

Table A.18: Average F_1 -macro scores (%) achieved by fine-tuning BERT with different samples of unlabeled tweets, using the MLP classifier

BERT										
Dataset	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	64.35	59.96	60.1	59.76	62.98	72.36	66.84	70.86	70.15	66.15
sar	68.35	66.81	66.92	65.93	63.26	72.51	67.51	71.96	71.12	72.85
ntu	89.0	88.3	85.26	85.31	83.1	86.22	87.35	91.86	93.1	93.37
S15	80.21	78.89	76.36	75.15	76.32	78.0	75.22	76.34	74.53	74.66
stm	85.84	86.39	86.48	86.33	86.81	86.7	88.16	87.22	88.73	86.85
per	82.5	83.03	81.04	81.32	81.49	81.92	79.41	79.4	79.55	79.87
hob	84.59	84.6	86.12	85.51	86.22	85.58	88.08	87.44	85.7	87.52
iph	81.88	80.94	80.98	81.82	80.84	79.82	81.31	80.81	81.15	82.73
mov	77.55	77.84	78.31	77.39	76.62	77.95	79.03	78.51	80.25	81.18
san	86.92	86.71	87.28	87.44	87.41	87.43	88.94	86.66	87.39	88.77
Nar	89.71	89.42	90.33	90.48	90.69	90.73	91.94	92.92	93.61	92.69
arc	88.3	88.47	89.1	88.88	89.02	87.91	88.94	89.03	88.37	88.49
S18	85.38	84.51	85.4	84.68	85.45	84.45	84.65	85.5	85.2	84.91
OMD	83.62	83.58	84.11	85.23	83.72	83.46	83.66	84.52	84.67	83.51
HCR	73.71	73.11	73.87	72.96	72.82	73.91	73.48	72.43	74.49	72.61
STS	89.69	89.46	89.01	89.33	88.83	90.04	90.54	90.29	90.49	90.71
SSt	83.0	83.04	82.97	82.6	82.77	82.08	82.6	82.97	83.43	83.48
Tar	83.45	84.03	84.09	84.42	83.68	84.5	84.33	83.31	83.22	83.52
Vad	87.73	87.69	87.75	87.24	88.0	88.09	88.93	87.98	88.4	88.78
S13	84.13	84.4	84.73	84.22	84.47	85.06	84.91	85.2	84.9	84.39
S17	90.13	90.06	90.2	89.96	89.84	89.28	89.35	88.81	88.04	88.08
S16	85.41	85.8	85.86	85.32	85.3	85.5	85.56	84.37	84.69	84.46
#wins	1	1	3	1	0	2	3	2	3	6
rank sums	132.0	141.0	115.5	146.5	139.0	119.0	96.5	113.5	104.0	103.0
position	7.0	9.0	5.0	10.0	8.0	6.0	1.0	4.0	3.0	2.0

Table A.19: Average classification accuracies (%) achieved by fine-tuning BERT with different samples of generic unlabeled tweets, using the XGB classifier

	BERT									
Dataset	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	78.33	72.38	69.05	69.05	67.86	71.9	74.29	64.05	68.02	62.38
sar	57.32	50.54	63.04	58.93	66.07	68.93	60.36	57.86	69.4	59.52
ntu	83.11	84.48	80.95	81.31	81.26	81.28	83.82	90.29	87.53	89.92
S15	88.47	87.84	87.86	89.09	88.16	87.54	88.16	87.22	88.48	88.68
stm	84.96	83.56	84.69	85.24	81.32	83.85	84.69	83.27	82.92	86.35
per	79.28	83.38	82.23	82.46	83.15	82.01	82.46	79.95	78.44	78.82
hob	83.52	83.9	84.1	83.14	85.99	83.72	84.88	83.92	85.25	84.85
iph	82.89	84.2	83.46	83.08	85.53	83.63	81.95	84.4	82.82	80.51
mov	84.31	84.49	85.21	84.31	83.43	83.78	84.67	85.56	83.78	85.81
san	85.87	85.53	84.8	85.13	83.57	85.78	86.59	85.62	85.59	85.59
Nar	89.4	89.24	89.08	89.81	89.08	89.32	90.3	90.71	91.2	91.77
arc	88.94	88.13	88.13	88.07	88.83	88.19	88.77	89.81	89.06	89.04
S18	84.02	84.83	83.97	83.7	83.65	83.54	83.22	83.76	83.54	84.24
OMD	84.47	83.84	82.58	81.9	83.47	81.64	84.32	83.74	83.75	82.69
HCR	81.13	80.87	79.98	80.66	80.55	80.24	80.56	79.77	79.29	78.37
STS	87.27	87.95	87.37	88.15	87.46	88.1	89.43	89.48	89.12	90.05
SSst	83.57	82.92	82.74	82.44	81.52	82.74	83.53	82.53	83.24	81.84
Tar	83.18	83.96	84.22	84.19	83.01	83.91	83.65	83.33	83.61	83.72
Vad	87.82	87.46	87.82	88.25	87.65	87.94	87.96	88.08	88.35	88.38
S13	86.89	86.71	87.0	85.63	86.64	86.75	87.07	86.91	86.72	86.3
S17	89.92	89.6	89.76	89.3	89.49	89.55	89.52	88.7	88.7	88.64
S16	88.34	88.5	88.32	88.67	88.37	88.35	88.29	87.98	87.71	87.69
#wins	5	2	1	2	2	0	2	2	1	5
rank sums	108.0	113.5	128.0	125.5	146.0	131.5	97.5	119.5	122.0	118.5
position	2.0	3.0	8.0	7.0	10.0	9.0	1.0	5.0	6.0	4.0

Table A.20: Average F_1 -macro scores (%) achieved by fine-tuning BERT with different samples of unlabeled tweets, using the XGB classifier

Dataset	BERT									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	71.8	59.35	57.57	58.71	55.4	58.55	62.92	49.72	55.0	49.95
sar	56.06	49.58	59.97	58.0	64.49	68.15	57.05	56.54	66.52	57.29
ntu	82.27	83.73	80.3	80.78	80.5	80.67	83.24	90.02	86.99	89.62
S15	65.69	62.12	60.62	65.57	63.82	64.1	64.6	61.86	62.76	66.05
stm	84.92	83.48	84.66	85.18	81.3	83.72	84.61	83.17	82.8	86.29
per	71.38	77.86	76.33	75.95	76.83	75.75	76.14	72.98	71.0	71.03
hob	79.14	80.33	80.63	78.85	83.11	80.32	81.6	80.79	82.09	81.83
iph	78.71	80.53	79.09	78.76	81.91	79.35	76.93	79.86	77.63	74.78
mov	60.89	60.98	64.29	60.29	57.75	58.49	63.56	66.66	60.49	67.71
san	85.75	85.4	84.66	85.03	83.38	85.66	86.45	85.5	85.44	85.46
Nar	88.81	88.71	88.52	89.31	88.58	88.79	89.79	90.29	90.81	91.36
arc	88.63	87.78	87.81	87.74	88.51	87.85	88.48	89.52	88.76	88.76
S18	83.83	84.66	83.8	83.52	83.42	83.31	83.04	83.57	83.31	84.03
OMD	82.77	82.03	80.68	79.83	81.66	79.66	82.59	81.9	81.81	80.56
HCR	72.62	72.99	71.85	72.68	72.31	71.46	72.48	71.03	69.89	68.45
STS	84.29	85.24	84.42	85.58	84.63	85.35	87.2	87.12	86.67	87.81
SSt	82.87	82.27	82.04	81.74	80.77	82.1	82.94	81.86	82.62	81.13
Tar	83.18	83.96	84.22	84.19	83.0	83.9	83.64	83.32	83.59	83.71
Vad	85.11	84.66	85.11	85.59	84.7	85.31	85.3	85.47	85.83	85.85
S13	82.45	82.15	82.67	80.76	81.98	82.26	82.84	82.43	82.35	81.75
S17	89.12	88.76	88.98	88.47	88.66	88.69	88.71	87.77	87.73	87.68
S16	84.65	84.86	84.59	85.0	84.64	84.63	84.64	84.08	83.75	83.64
#wins	3	3	1	1	2	1	3	2	0	6
rank sums	105.5	112.0	129.5	126.0	146.5	131.5	95.5	120.0	128.0	115.5
position	2.0	3.0	8.0	6.0	10.0	9.0	1.0	5.0	7.0	4.0

Table A.21: Average classification accuracies (%) achieved by fine-tuning RoBERTa with different samples of generic unlabeled tweets, using the RF classifier

Dataset	RoBERTa									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	67.7	67.3	66.11	66.83	70.4	67.22	66.59	67.7	66.67	65.63
sar	64.58	63.81	70.78	67.08	65.95	65.65	68.45	64.23	66.07	61.31
ntu	82.6	83.8	81.32	81.17	81.4	86.35	87.67	91.96	91.11	87.87
S15	85.46	85.46	85.88	85.67	85.77	85.57	85.46	86.08	86.09	85.88
stm	86.64	86.62	87.94	88.21	89.24	89.79	88.96	88.03	86.73	85.9
per	75.92	76.47	77.83	77.22	76.99	77.45	76.32	76.47	77.22	75.17
hob	79.31	80.15	79.83	79.89	79.89	80.91	79.69	79.07	80.34	80.86
iph	81.25	81.13	81.94	81.5	81.95	82.26	81.95	81.32	80.95	79.63
mov	82.24	82.18	82.42	82.41	82.24	82.36	82.48	82.18	82.41	82.12
san	85.46	85.75	86.3	86.71	86.0	86.46	86.22	86.47	86.52	84.23
Nar	86.96	88.21	89.08	88.56	88.83	89.84	90.93	90.84	90.27	89.76
arc	88.09	87.79	88.14	88.22	88.44	89.06	88.88	89.52	89.06	88.11
S18	85.17	84.17	84.92	84.97	84.49	85.19	85.38	84.83	84.2	83.05
OMD	80.89	81.01	80.73	80.99	80.69	80.01	81.3	79.48	79.4	77.68
HCR	75.1	76.22	76.22	76.38	76.71	76.57	76.47	76.94	76.43	74.93
STS	82.98	83.84	84.91	83.73	84.68	85.82	86.56	86.09	86.2	84.09
SSst	80.73	81.74	82.5	81.58	81.8	81.64	83.2	83.18	83.01	80.37
Tar	82.59	82.59	83.86	82.79	83.38	82.94	83.2	82.79	83.34	81.35
Vad	82.9	83.1	84.0	83.33	83.57	85.7	86.6	85.98	85.76	84.23
S13	81.1	81.55	81.88	82.18	82.07	82.5	83.39	81.99	82.68	80.96
S17	87.46	87.08	88.4	88.07	88.26	88.52	88.14	88.13	87.6	85.6
S16	84.65	84.73	85.34	85.02	85.26	84.94	85.11	84.94	84.61	83.2
#wins	0	0	4	1	1	4	8	3	1	0
rank sums	171.5	165.0	104.0	120.0	104.5	85.0	77.5	100.5	98.5	183.5
position	9.0	8.0	5.0	7.0	6.0	2.0	1.0	4.0	3.0	10.0

Table A.22: Average F_1 -macro scores (%) achieved by fine-tuning RoBERTa with different samples of unlabeled tweets, using the RF classifier

RoBERTa										
Dataset	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	42.94	42.29	40.51	42.08	48.41	42.27	41.46	42.94	40.73	39.55
sar	60.06	60.18	67.56	64.17	62.26	62.9	65.85	60.86	62.89	57.94
ntu	81.69	82.97	80.47	80.24	80.44	85.77	87.25	91.67	90.72	87.2
S15	46.63	46.63	49.18	48.02	48.37	47.21	46.63	50.23	49.96	49.29
stm	86.6	86.58	87.89	88.19	89.22	89.77	88.9	87.99	86.7	85.82
per	59.21	59.93	64.06	62.69	62.4	62.44	60.38	60.56	62.58	56.82
hob	71.01	72.51	71.95	72.12	72.26	74.07	72.38	70.76	73.2	74.01
iph	73.53	73.47	74.83	74.53	74.9	75.68	75.33	73.34	73.48	70.95
mov	46.27	46.01	47.23	47.23	46.27	46.96	47.52	46.01	47.18	45.69
san	85.23	85.53	86.09	86.53	85.8	86.24	86.01	86.23	86.28	83.99
Nar	85.93	87.37	88.32	87.72	88.11	89.22	90.38	90.27	89.64	89.05
arc	87.6	87.28	87.66	87.77	88.01	88.67	88.5	89.14	88.63	87.59
S18	84.9	83.88	84.67	84.68	84.22	84.88	85.11	84.52	83.87	82.68
OMD	76.99	77.43	76.93	77.3	76.88	76.07	77.86	75.35	74.98	72.56
HCR	54.58	57.97	57.86	58.71	59.2	59.14	59.01	60.05	58.8	54.14
STS	76.25	77.68	79.48	77.63	79.14	80.95	82.32	81.43	81.58	78.05
SSt	79.41	80.53	81.39	80.42	80.75	80.64	82.31	82.29	82.06	79.13
Tar	82.58	82.58	83.84	82.78	83.38	82.92	83.18	82.77	83.31	81.33
Vad	76.5	76.61	78.21	77.18	77.63	81.05	82.44	81.45	81.09	78.61
S13	69.59	70.44	71.45	71.9	71.74	72.89	74.51	71.84	73.12	69.39
S17	86.06	85.61	87.17	86.83	87.03	87.34	86.88	86.85	86.29	83.87
S16	77.21	77.24	78.43	77.9	78.34	77.78	78.0	77.71	77.2	74.43
#wins	0	0	4	1	1	4	8	4	0	0
rank sums	170.5	165.0	105.5	118.5	108.5	83.0	71.0	102.0	101.0	185.0
position	9.0	8.0	5.0	7.0	6.0	2.0	1.0	4.0	3.0	10.0

Table A.23: Average classification accuracies (%) achieved by fine-tuning RoBERTa with different samples of generic unlabeled tweets, using the LR classifier

Dataset	RoBERTa									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	73.81	72.38	74.05	68.57	78.57	78.33	74.05	76.9	75.0	64.21
sar	67.5	70.36	71.79	71.79	74.64	73.04	71.79	74.46	67.5	63.15
ntu	87.79	87.42	88.52	85.98	87.79	89.95	90.3	92.8	92.79	92.2
S15	85.39	85.06	87.25	87.54	88.48	89.42	87.84	88.78	89.75	88.68
stm	88.57	89.13	90.81	89.67	90.25	90.82	89.42	87.75	89.42	87.38
per	84.29	86.79	86.79	87.71	83.61	85.66	84.52	85.9	87.26	85.34
hob	88.12	89.09	86.98	88.89	87.35	89.46	87.54	87.16	86.78	86.21
iph	86.28	86.84	86.47	84.97	86.86	85.15	85.71	85.71	85.53	82.83
mov	89.47	91.07	88.4	89.48	88.94	86.8	88.58	86.81	89.3	86.93
san	89.62	88.89	90.6	89.46	89.05	89.46	88.55	89.87	90.11	88.89
Nar	93.48	93.23	94.21	93.31	93.72	94.54	93.97	92.83	94.54	94.27
arc	90.57	90.1	90.63	89.64	88.88	89.7	90.05	90.34	89.87	89.95
S18	87.58	88.76	87.79	87.31	87.95	87.74	88.17	87.84	87.04	86.84
OMD	85.78	84.78	86.31	85.99	85.52	85.94	85.36	84.84	83.95	85.08
HCR	80.24	79.45	79.24	79.03	79.4	79.19	78.46	78.56	79.92	79.61
STS	93.02	93.31	93.9	92.82	93.51	92.77	93.51	94.3	93.71	92.81
SSst	84.75	86.19	85.54	85.1	86.15	86.28	85.32	86.46	85.1	84.75
Tar	84.02	83.53	84.57	84.51	84.08	84.31	84.42	84.37	85.09	84.38
Vad	90.99	91.21	91.23	91.3	91.37	91.71	91.13	91.83	91.94	91.24
S13	88.94	88.24	89.01	88.88	89.04	89.13	88.99	88.56	88.63	88.78
S17	91.41	91.44	91.6	91.82	91.67	91.89	91.79	90.91	91.35	89.89
S16	89.74	89.69	90.05	90.02	90.26	89.74	89.94	89.99	89.91	89.21
#wins	1	2	3	1	4	5	0	3	4	0
rank sums	140.0	137.0	92.0	125.0	103.0	100.5	124.0	113.5	109.0	166.0
position	9.0	8.0	1.0	7.0	3.0	2.0	6.0	5.0	4.0	10.0

Table A.24: Average F_1 -macro scores (%) achieved by fine-tuning RoBERTa with different samples of unlabeled tweets, using the LR classifier

Dataset	RoBERTa									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	66.95	67.21	67.2	59.58	74.6	71.86	67.04	71.31	68.34	56.94
sar	66.7	68.85	71.17	70.99	73.96	72.0	70.81	73.11	66.1	62.01
ntu	87.54	87.1	88.3	85.68	87.62	89.79	90.1	92.6	92.65	92.02
S15	68.91	68.25	71.89	73.63	75.69	77.69	73.93	77.11	78.8	76.42
stm	88.55	89.12	90.76	89.62	90.23	90.78	89.36	87.7	89.41	87.32
per	81.02	84.02	84.01	85.29	80.69	82.91	81.33	83.45	84.93	82.41
hob	86.65	87.67	85.15	87.41	85.57	87.96	85.91	85.51	84.98	84.44
iph	83.85	84.66	84.15	82.55	84.65	82.58	83.04	83.1	82.87	79.71
mov	82.27	85.1	79.87	82.85	81.43	78.01	79.86	77.48	81.0	77.63
san	89.55	88.82	90.54	89.4	88.96	89.4	88.48	89.8	90.05	88.82
Nar	93.22	92.95	93.99	93.03	93.46	94.32	93.73	92.55	94.32	94.06
arc	90.37	89.89	90.42	89.39	88.63	89.47	89.85	90.12	89.65	89.74
S18	87.51	88.7	87.71	87.23	87.89	87.67	88.11	87.77	86.96	86.77
OMD	84.79	83.8	85.43	85.07	84.56	84.94	84.4	83.93	82.94	84.1
HCR	76.07	75.14	74.72	74.65	74.99	74.67	74.0	74.26	75.96	75.62
STS	91.88	92.21	92.9	91.71	92.49	91.65	92.51	93.38	92.67	91.66
SSt	84.38	85.82	85.16	84.69	85.78	85.89	84.96	86.08	84.73	84.34
Tar	84.01	83.52	84.56	84.51	84.07	84.3	84.41	84.36	85.08	84.37
Vad	89.53	89.85	89.86	89.94	90.03	90.42	89.73	90.51	90.67	89.86
S13	86.38	85.62	86.5	86.34	86.5	86.69	86.49	85.93	86.06	86.3
S17	90.87	90.92	91.09	91.31	91.15	91.39	91.29	90.35	90.81	89.29
S16	87.49	87.44	87.81	87.8	88.08	87.48	87.72	87.77	87.68	86.87
#wins	1	3	3	1	3	5	0	2	5	0
rank sums	140.0	132.5	90.0	125.5	103.5	99.0	128.0	115.0	108.5	168.0
position	9.0	8.0	1.0	6.0	3.0	2.0	7.0	5.0	4.0	10.0

Table A.25: Average classification accuracies (%) achieved by fine-tuning RoBERTa with different samples of generic unlabeled tweets, using the MLP classifier

Dataset	RoBERTa									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	70.24	71.59	73.01	77.3	74.6	76.67	73.01	74.68	76.51	65.71
sar	67.98	67.02	74.16	70.84	76.07	73.51	67.68	73.99	69.41	64.58
ntu	88.6	87.3	86.83	87.9	87.67	90.19	90.66	92.21	92.67	92.8
S15	87.23	87.45	89.0	89.1	89.32	90.05	88.9	90.75	91.28	90.14
stm	88.48	88.49	88.75	88.94	89.97	90.16	88.21	88.02	89.14	87.66
per	84.96	86.64	85.19	86.71	84.52	85.73	83.31	84.23	86.66	84.52
hob	88.76	89.15	87.43	87.94	87.16	87.87	86.59	88.12	87.74	87.04
iph	87.9	87.1	87.54	85.15	88.16	84.96	86.27	87.46	86.96	83.89
mov	88.64	91.97	88.47	88.18	88.71	88.17	87.87	88.34	88.58	88.23
san	89.3	89.0	90.06	89.62	89.05	90.14	88.85	90.11	90.38	89.19
Nar	92.91	92.85	94.19	93.39	93.45	94.08	93.75	92.94	94.18	94.24
arc	89.89	89.19	89.93	90.24	89.09	90.14	90.51	90.16	90.34	89.68
S18	87.09	88.17	87.7	87.5	87.95	88.31	88.58	87.61	87.36	86.77
OMD	84.91	85.17	85.99	85.45	85.63	85.4	86.41	85.45	84.94	84.86
HCR	79.99	80.29	80.35	79.45	80.08	80.83	79.02	80.87	80.33	79.12
STS	91.84	92.85	93.44	92.36	93.23	92.89	93.15	93.9	93.49	92.98
SSst	85.06	86.16	85.36	85.1	85.9	86.6	85.39	86.24	85.52	84.24
Tar	83.96	84.05	84.12	84.7	84.59	84.74	84.44	84.67	84.74	83.96
Vad	90.43	91.01	91.5	91.07	91.28	91.89	91.35	91.4	92.08	90.89
S13	88.18	88.23	88.85	88.75	88.81	89.0	88.68	88.58	88.56	88.33
S17	91.13	91.35	91.78	91.88	91.47	91.7	91.77	91.29	91.18	90.18
S16	89.49	89.99	90.19	90.31	90.12	90.01	90.27	90.0	89.97	89.15
#wins	0	2	0	4	2	4	3	2	4	2
rank sums	165.5	143.0	99.5	111.5	107.5	85.5	131.5	99.5	91.5	175.0
position	9.0	8.0	3.5	6.0	5.0	1.0	7.0	3.5	2.0	10.0

Table A.26: Average F_1 -macro scores (%) achieved by fine-tuning RoBERTa with different samples of unlabeled tweets, using the MLP classifier

RoBERTa										
Dataset	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	61.23	63.53	63.54	68.53	64.24	68.09	64.29	63.35	68.07	54.57
sar	66.38	63.82	73.37	70.2	75.45	72.9	66.1	72.69	68.01	63.81
ntu	88.26	87.0	86.5	87.61	87.42	89.99	90.45	92.01	92.51	92.62
S15	68.2	69.48	72.79	73.63	74.01	75.23	73.49	77.97	78.76	76.05
stm	88.44	88.47	88.69	88.88	89.93	90.11	88.14	87.97	89.11	87.6
per	81.15	83.42	81.6	83.74	80.96	82.46	79.7	81.02	83.76	80.63
hob	87.04	87.47	85.43	86.13	85.05	85.76	84.73	86.3	85.82	85.03
iph	85.26	84.53	85.09	82.29	85.85	82.08	83.29	84.71	84.22	80.53
mov	78.97	85.29	78.93	78.63	78.98	78.21	77.33	78.31	78.94	78.18
san	89.22	88.93	89.98	89.55	88.97	90.08	88.78	90.04	90.32	89.11
Nar	92.59	92.53	93.94	93.1	93.17	93.83	93.49	92.64	93.94	94.01
arc	89.65	88.91	89.69	90.0	88.8	89.91	90.29	89.93	90.11	89.44
S18	87.02	88.09	87.62	87.42	87.87	88.25	88.52	87.53	87.28	86.68
OMD	83.78	84.1	84.95	84.35	84.51	84.24	85.34	84.4	83.82	83.68
HCR	75.14	74.76	74.78	73.41	74.43	75.82	73.68	75.35	75.25	72.92
STS	90.43	91.58	92.3	91.1	92.06	91.67	91.97	92.82	92.33	91.76
SSt	84.63	85.76	84.92	84.66	85.47	86.18	84.98	85.81	85.09	83.76
Tar	83.95	84.04	84.11	84.7	84.58	84.74	84.42	84.67	84.74	83.96
Vad	88.73	89.43	90.03	89.55	89.77	90.5	89.85	89.9	90.65	89.28
S13	85.06	85.14	85.93	85.8	85.84	86.21	85.71	85.54	85.58	85.28
S17	90.5	90.74	91.2	91.31	90.86	91.12	91.18	90.68	90.56	89.49
S16	86.7	87.34	87.64	87.75	87.52	87.32	87.65	87.35	87.17	86.27
#wins	0	2	0	3	2	5	3	1	5	2
rank sums	163.0	142.0	102.5	112.0	109.0	85.5	126.0	104.0	87.0	179.0
position	9.0	8.0	3.0	6.0	5.0	1.0	7.0	4.0	2.0	10.0

Table A.27: Average classification accuracies (%) achieved by fine-tuning RoBERTa with different samples of generic unlabeled tweets, using the XGB classifier

Dataset	RoBERTa									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	69.29	72.62	64.29	70.71	72.62	73.81	58.33	67.38	68.1	67.94
sar	55.89	58.93	61.96	74.46	66.25	56.43	70.36	50.71	54.64	68.51
ntu	84.15	83.81	76.61	82.37	80.58	84.55	89.22	88.88	92.06	88.02
S15	88.48	88.48	90.97	88.78	87.85	90.34	89.1	89.73	87.55	88.68
stm	86.36	86.05	85.25	85.79	87.48	86.63	86.63	84.4	83.85	83.76
per	80.17	81.32	82.0	81.32	82.68	80.62	82.7	81.79	81.12	82.24
hob	81.01	83.33	81.6	81.42	83.54	83.34	84.09	83.92	81.98	85.05
iph	81.96	84.4	84.78	85.71	84.77	83.65	80.65	84.2	83.08	83.26
mov	85.03	86.99	87.87	85.91	86.99	84.84	86.63	87.88	86.45	86.69
san	86.68	87.33	86.93	88.23	87.09	87.9	86.35	87.25	87.83	86.46
Nar	90.06	91.44	92.42	91.03	90.87	92.26	91.77	92.09	91.61	91.8
arc	89.81	89.0	89.58	89.35	90.1	89.81	90.45	90.16	89.23	89.14
S18	85.58	87.2	85.8	86.39	86.98	86.5	86.82	85.8	84.62	84.76
OMD	84.94	84.79	85.99	85.26	85.36	84.62	85.36	83.58	82.74	83.72
HCR	80.03	82.18	81.08	80.4	80.5	81.45	80.5	81.13	80.19	79.75
STS	89.68	90.76	90.76	90.46	90.81	91.45	91.49	91.05	90.86	90.66
SSt	84.01	83.88	85.23	84.4	84.93	84.75	85.02	85.06	85.37	83.0
Tar	83.85	84.14	85.15	84.48	85.0	84.51	83.79	83.5	84.8	83.28
Vad	88.44	88.27	89.28	89.16	88.85	89.39	90.16	90.18	90.06	89.13
S13	86.64	87.12	87.67	87.83	87.76	87.62	88.24	87.57	87.71	86.82
S17	90.47	90.72	90.8	90.88	90.42	90.91	90.91	90.74	90.09	88.56
S16	88.84	89.37	89.62	89.51	89.78	89.55	89.55	89.4	89.67	88.31
#wins	0	2	4	3	2	2	5	2	2	1
rank sums	173.0	130.5	100.0	118.5	98.0	97.0	88.5	109.5	136.0	159.0
position	10.0	7.0	4.0	6.0	3.0	2.0	1.0	5.0	8.0	9.0

Table A.28: Average F_1 -macro scores (%) achieved by fine-tuning RoBERTa with different samples of unlabeled tweets, using the XGB classifier

Dataset	RoBERTa									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	54.77	61.69	47.99	58.14	57.59	67.22	42.02	54.39	53.42	54.93
sar	53.42	56.62	59.94	73.89	64.55	54.6	68.13	49.36	50.89	67.67
ntu	83.63	83.39	75.86	81.64	79.9	84.13	88.92	88.63	91.87	87.65
S15	67.46	63.75	74.11	66.24	62.03	73.59	67.08	69.26	61.61	66.92
stm	86.32	86.0	85.21	85.77	87.43	86.61	86.58	84.32	83.77	83.63
per	72.78	73.45	75.6	75.1	76.31	74.48	77.13	74.92	74.8	76.17
hob	76.79	79.35	77.42	76.73	80.1	79.58	80.82	80.22	78.26	82.06
iph	77.2	79.94	80.96	82.27	81.08	79.67	75.66	79.93	78.21	78.86
mov	64.06	70.45	73.38	66.84	68.1	66.22	69.41	72.07	68.51	68.91
san	86.56	87.22	86.8	88.14	86.99	87.8	86.24	87.12	87.73	86.33
Nar	89.57	91.0	92.06	90.61	90.4	91.92	91.4	91.76	91.21	91.42
arc	89.53	88.65	89.29	89.06	89.84	89.56	90.22	89.88	88.94	88.84
S18	85.46	87.06	85.68	86.27	86.86	86.33	86.7	85.68	84.43	84.61
OMD	83.23	83.07	84.4	83.69	83.75	82.83	83.73	81.68	80.76	81.73
HCR	71.23	74.87	72.89	72.13	72.09	73.59	72.17	73.06	71.71	70.58
STS	87.5	88.8	88.7	88.46	88.84	89.61	89.71	89.14	88.95	88.6
SSt	83.38	83.23	84.71	83.78	84.34	84.21	84.43	84.5	84.8	82.36
Tar	83.84	84.13	85.14	84.47	85.0	84.51	83.78	83.49	84.79	83.27
Vad	85.85	85.72	86.89	86.75	86.44	87.19	88.04	88.13	87.98	86.75
S13	82.03	82.75	83.57	83.82	83.63	83.73	84.38	83.54	83.73	82.31
S17	89.72	89.98	90.07	90.17	89.69	90.2	90.21	90.01	89.31	87.6
S16	85.29	85.99	86.44	86.27	86.59	86.3	86.31	86.06	86.4	84.46
#wins	0	2	5	3	2	1	5	1	2	1
rank sums	171.0	132.0	100.5	118.5	105.0	92.5	87.0	110.5	136.5	156.5
position	10.0	7.0	3.0	6.0	4.0	2.0	1.0	5.0	8.0	9.0

Table A.29: Average classification accuracies (%) achieved by fine-tuning BERTweet with different samples of generic unlabeled tweets, using the RF classifier

Dataset	BERTweet									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	66.19	66.67	67.14	66.19	68.1	66.19	67.14	66.67	65.71	65.71
sar	68.45	67.56	69.88	70.42	72.38	67.02	69.46	66.67	68.57	69.88
ntu	84.64	86.7	87.79	87.3	87.66	88.26	88.63	89.58	88.02	87.18
S15	86.4	86.4	86.71	86.6	86.81	86.19	85.88	86.09	85.67	85.36
stm	87.37	88.3	89.98	89.14	88.22	88.95	88.39	89.32	89.42	87.18
per	79.12	79.95	81.17	80.56	80.71	78.59	79.34	77.44	77.45	74.72
hob	79.5	80.98	81.24	81.17	80.14	80.71	80.34	81.62	81.56	80.28
iph	82.21	83.08	83.33	83.21	83.52	83.77	83.03	82.96	80.02	81.34
mov	82.29	82.77	82.89	82.89	82.77	82.65	82.77	82.65	82.12	82.12
san	86.84	88.43	88.97	89.13	88.54	88.81	88.21	87.15	85.81	85.46
Nar	90.52	91.91	92.77	92.42	92.28	91.33	91.91	91.19	89.97	89.41
arc	89.33	90.07	90.26	90.04	90.63	89.87	90.32	90.03	89.13	89.54
S18	85.42	86.34	87.32	86.7	87.09	85.51	85.6	84.62	83.97	83.22
OMD	80.41	82.18	82.42	81.76	82.06	82.57	81.29	80.87	79.91	79.56
HCR	76.05	76.78	77.34	77.74	77.81	77.1	76.52	76.45	75.37	74.95
STS	84.38	86.78	86.89	86.95	86.61	86.09	85.51	86.22	85.24	84.24
SSst	80.59	82.61	83.86	82.65	82.36	81.96	82.07	81.85	80.98	80.79
Tar	83.46	83.95	84.0	83.89	83.9	83.47	83.11	83.36	82.14	81.03
Vad	82.3	84.5	85.01	84.43	84.97	84.5	84.34	84.36	83.97	83.6
S13	81.37	82.37	83.0	82.66	82.41	81.38	81.7	81.63	81.1	80.51
S17	88.76	89.49	89.8	89.77	89.46	88.82	88.83	88.04	86.29	86.36
S16	85.08	85.66	86.35	85.91	86.02	85.1	85.24	84.7	83.91	83.41
#wins	0	0	11	3	5	2	0	2	0	0
rank sums	174.5	99.0	42.5	72.5	71.0	117.0	117.0	138.0	176.0	202.5
position	8.0	4.0	1.0	3.0	2.0	5.5	5.5	7.0	9.0	10.0

Table A.30: Average F_1 -macro scores (%) achieved by fine-tuning BERTweet with different samples of unlabeled tweets, using the RF classifier

Dataset	BERTweet									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	39.77	41.51	41.68	40.55	43.28	40.55	41.68	40.73	39.59	39.59
sar	66.13	66.1	67.19	66.67	69.79	64.23	66.92	61.76	65.65	66.39
ntu	83.74	85.86	87.08	86.75	86.97	87.55	87.81	89.05	87.43	86.45
S15	51.39	51.82	53.47	53.2	54.26	50.77	49.18	50.19	47.9	46.05
stm	87.35	88.27	89.96	89.09	88.17	88.91	88.33	89.29	89.37	87.12
per	67.18	69.57	71.54	70.5	70.95	66.74	67.67	63.21	63.49	56.65
hob	71.22	74.03	74.23	74.13	72.82	73.42	72.93	74.63	74.5	73.0
iph	75.59	76.75	77.22	76.69	77.37	77.72	76.43	76.78	71.69	73.92
mov	46.59	48.89	49.66	49.37	48.98	48.38	49.01	48.31	45.69	45.69
san	86.66	88.28	88.85	89.01	88.41	88.7	88.07	86.99	85.58	85.24
Nar	89.81	91.36	92.35	91.95	91.79	90.72	91.38	90.64	89.27	88.62
arc	88.9	89.71	89.92	89.72	90.31	89.55	90.0	89.7	88.69	89.1
S18	85.13	86.1	87.1	86.46	86.87	85.23	85.33	84.3	83.6	82.85
OMD	76.56	78.88	79.29	78.61	78.76	79.36	77.72	77.06	75.85	75.16
HCR	57.46	59.43	61.11	62.03	62.36	60.45	58.81	58.71	55.23	54.09
STS	78.43	82.26	82.6	82.56	82.17	81.31	80.33	81.44	80.04	78.34
SSt	79.26	81.45	82.85	81.53	81.23	80.74	80.85	80.69	79.87	79.56
Tar	83.45	83.94	83.99	83.87	83.89	83.47	83.09	83.34	82.12	81.02
Vad	74.93	78.75	79.67	78.65	79.58	78.75	78.51	78.56	77.97	77.22
S13	70.31	72.52	73.74	73.1	72.59	70.41	71.0	71.0	69.81	68.3
S17	87.61	88.44	88.83	88.82	88.43	87.71	87.74	86.87	84.82	84.89
S16	78.17	79.09	80.35	79.59	79.74	78.12	78.28	77.44	75.89	74.9
#wins	0	0	12	1	5	2	0	2	0	0
rank sums	174.0	99.5	39.5	76.5	71.0	117.0	116.0	135.5	178.0	203.0
position	8.0	4.0	1.0	3.0	2.0	6.0	5.0	7.0	9.0	10.0

Table A.31: Average classification accuracies (%) achieved by fine-tuning BERTweet with different samples of generic unlabeled tweets, using the LR classifier

Dataset	BERT									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	67.62	67.62	75.24	73.81	70.24	72.14	72.14	66.19	67.14	64.52
sar	67.68	66.25	70.36	68.93	73.04	73.21	73.04	67.5	67.98	73.21
ntu	88.89	90.67	91.39	92.45	91.73	91.38	93.15	92.8	92.67	92.45
S15	91.9	93.47	93.14	91.9	93.47	90.64	88.79	90.66	90.96	88.18
stm	90.54	91.64	91.65	89.14	90.53	91.09	90.81	91.92	91.27	92.2
per	89.05	88.83	88.6	89.51	88.37	87.46	86.55	86.34	83.97	83.83
hob	89.25	89.45	89.26	89.45	89.27	88.11	87.35	88.1	88.37	87.73
iph	87.6	86.84	87.22	87.59	86.65	87.22	86.47	87.41	85.21	86.84
mov	89.49	90.73	90.37	90.2	90.19	92.87	90.2	89.84	89.6	89.3
san	90.44	91.74	91.17	91.83	90.93	91.75	91.34	90.36	89.4	90.77
Nar	96.01	96.98	96.82	96.57	96.49	96.49	96.74	95.92	95.79	94.7
arc	90.92	91.04	91.56	91.21	92.08	91.38	91.27	90.86	91.27	91.97
S18	89.4	89.56	89.19	88.92	88.92	87.63	88.38	86.82	86.73	86.5
OMD	86.83	88.3	88.56	86.99	87.83	88.04	87.83	87.41	86.95	86.04
HCR	81.45	81.18	81.55	81.02	81.6	81.18	81.97	80.03	77.93	77.88
STS	94.69	95.08	94.94	94.4	95.08	95.18	94.98	94.3	93.86	93.76
SSst	86.89	88.34	88.38	87.94	88.07	87.29	87.42	87.37	86.51	87.29
Tar	85.55	86.53	85.84	85.98	85.64	85.72	85.43	85.09	85.39	84.8
Vad	91.21	92.8	92.61	92.59	92.37	92.16	92.73	91.99	92.19	92.4
S13	89.81	90.36	90.7	90.36	90.29	89.56	90.04	88.67	89.39	88.81
S17	92.6	93.21	92.89	93.0	93.08	92.66	92.41	91.63	91.16	91.13
S16	90.82	91.34	91.63	91.47	91.58	91.29	90.99	90.53	90.38	89.73
#wins	1	7	5	3	2	3	2	0	0	2
rank sums	143.0	78.5	69.5	92.0	95.0	110.0	114.5	162.0	172.5	173.0
position	7.0	2.0	1.0	3.0	4.0	5.0	6.0	8.0	9.0	10.0

Table A.32: Average F_1 -macro scores (%) achieved by fine-tuning BERTweet with different samples of unlabeled tweets, using the LR classifier

Dataset	BERTweet									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	61.52	60.15	67.29	67.48	63.67	65.71	66.28	59.57	60.79	55.21
sar	66.75	65.08	69.42	67.88	72.02	72.03	72.0	66.32	67.32	72.08
ntu	88.63	90.5	91.2	92.29	91.53	91.14	92.95	92.6	92.51	92.23
S15	82.86	86.05	84.98	83.7	85.7	80.63	76.14	80.2	80.17	76.85
stm	90.51	91.61	91.63	89.11	90.48	91.04	90.79	91.91	91.24	92.18
per	87.07	86.71	86.51	87.53	86.01	84.98	84.05	83.86	81.25	80.9
hob	87.8	87.94	87.65	87.83	87.87	86.73	85.59	86.52	86.81	86.18
iph	85.4	84.58	84.91	85.36	84.34	85.06	84.17	85.31	82.83	84.63
mov	82.7	84.93	84.35	84.33	83.89	88.27	83.79	83.27	83.31	82.35
san	90.38	91.68	91.12	91.77	90.86	91.7	91.27	90.29	89.32	90.7
Nar	95.85	96.87	96.7	96.44	96.35	96.35	96.61	95.76	95.63	94.51
arc	90.73	90.85	91.38	91.04	91.92	91.2	91.09	90.67	91.08	91.8
S18	89.34	89.5	89.11	88.84	88.84	87.56	88.3	86.73	86.65	86.41
OMD	86.0	87.48	87.76	86.09	87.0	87.23	87.0	86.56	86.07	85.06
HCR	77.77	77.36	78.02	77.26	77.83	77.19	78.18	76.05	73.64	73.63
STS	93.85	94.3	94.12	93.55	94.3	94.41	94.21	93.42	92.9	92.79
SSt	86.54	88.02	88.05	87.59	87.72	86.94	87.1	87.04	86.18	86.97
Tar	85.54	86.53	85.83	85.98	85.63	85.72	85.43	85.08	85.38	84.79
Vad	89.81	91.64	91.43	91.4	91.15	90.89	91.56	90.72	90.97	91.24
S13	87.56	88.18	88.59	88.2	88.1	87.26	87.81	86.22	87.1	86.4
S17	92.13	92.76	92.43	92.56	92.64	92.19	91.91	91.1	90.61	90.58
S16	88.81	89.43	89.75	89.57	89.68	89.35	89.01	88.47	88.3	87.54
#wins	1	7	4	3	1	2	2	0	0	2
rank sums	142.0	79.5	71.0	89.5	94.0	110.5	116.5	163.0	173.0	171.0
position	7.0	2.0	1.0	3.0	4.0	5.0	6.0	8.0	10.0	9.0

Table A.33: Average classification accuracies (%) achieved by fine-tuning BERTweet with different samples of generic unlabeled tweets, using the MLP classifier

Dataset	BERTweet									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	73.73	72.06	75.64	72.54	75.4	82.3	75.4	69.53	67.14	65.08
sar	67.2	71.96	72.32	70.89	72.09	72.85	69.41	66.61	68.45	72.32
ntu	89.0	91.86	91.99	91.84	92.34	90.77	92.44	93.28	92.69	93.4
S15	92.12	94.18	92.01	92.64	93.15	91.58	91.07	92.21	92.1	90.03
stm	89.69	91.27	90.72	89.41	89.78	90.16	89.69	90.81	90.43	91.73
per	86.91	87.69	87.85	88.9	87.45	87.38	87.31	87.01	83.9	84.07
hob	87.98	89.32	88.62	87.98	88.69	89.83	86.83	87.91	88.43	88.06
iph	87.08	86.35	87.35	86.52	87.16	87.16	86.67	87.72	85.71	87.9
mov	90.25	90.97	90.07	91.38	90.55	93.29	91.32	89.96	90.02	91.03
san	90.25	91.12	91.61	91.58	91.55	91.75	91.01	90.28	89.32	90.96
Nar	95.95	97.04	96.66	96.52	95.92	96.25	96.63	95.95	95.68	94.84
arc	90.94	91.13	91.46	91.19	91.92	91.0	90.59	90.52	91.21	91.64
S18	89.01	89.04	89.13	89.08	88.76	87.22	88.27	87.11	86.84	86.43
OMD	86.17	86.97	88.29	86.04	87.56	87.01	87.51	87.06	86.34	86.53
HCR	80.9	81.53	81.91	81.62	81.57	81.41	80.85	81.01	78.81	78.7
STS	94.33	94.76	94.92	94.09	94.86	95.38	94.49	94.46	94.2	93.1
SSst	86.26	87.23	87.51	87.87	87.57	87.81	87.65	87.04	86.12	86.59
Tar	85.62	86.54	85.92	86.24	86.23	86.23	85.95	85.63	84.78	84.57
Vad	90.9	92.3	92.32	92.45	92.57	92.37	92.31	91.52	91.78	91.98
S13	89.29	90.14	90.12	90.39	90.3	89.91	90.08	89.36	89.42	88.85
S17	92.06	92.65	92.82	92.83	92.85	92.51	92.23	91.4	90.73	90.69
S16	90.53	91.02	91.16	91.44	91.31	91.13	90.68	90.53	90.3	89.87
#wins	0	3	3	4	3	6	0	0	0	3
rank sums	167.0	95.0	75.5	95.5	78.5	91.0	128.0	150.0	174.0	155.5
position	9.0	4.0	1.0	5.0	2.0	3.0	6.0	7.0	10.0	8.0

Table A.34: Average F_1 -macro scores (%) achieved by fine-tuning BERTweet with different samples of unlabeled tweets, using the MLP classifier

Dataset	BERTweet									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	65.37	63.53	66.75	62.64	67.86	73.99	67.83	59.79	57.83	55.25
sar	66.56	71.03	71.34	69.94	71.17	71.48	68.11	63.92	67.17	70.95
ntu	88.7	91.65	91.81	91.66	92.11	90.54	92.17	93.09	92.49	93.21
S15	81.41	86.22	81.45	83.07	84.09	80.54	79.26	80.56	80.26	75.31
stm	89.65	91.23	90.69	89.37	89.75	90.11	89.66	90.79	90.4	91.71
per	83.83	84.87	84.88	86.38	84.47	84.31	84.27	84.03	79.88	80.62
hob	85.96	87.59	86.83	86.17	86.92	88.3	84.65	86.07	86.54	86.25
iph	84.06	83.56	84.84	83.75	84.64	84.55	84.06	85.42	82.88	85.37
mov	81.79	83.64	82.29	84.63	82.98	87.66	84.27	82.06	82.37	83.27
san	90.19	91.06	91.55	91.51	91.49	91.7	90.94	90.21	89.24	90.89
Nar	95.78	96.91	96.52	96.38	95.74	96.08	96.48	95.78	95.5	94.62
arc	90.74	90.94	91.26	90.99	91.74	90.79	90.37	90.28	91.0	91.42
S18	88.95	88.97	89.06	89.01	88.67	87.14	88.19	87.01	86.76	86.33
OMD	85.17	86.03	87.32	85.09	86.62	86.0	86.57	86.06	85.27	85.45
HCR	76.53	76.79	77.33	76.92	76.5	76.17	76.14	75.98	73.23	73.44
STS	93.36	93.87	94.05	93.07	93.99	94.59	93.57	93.51	93.22	91.87
SSt	85.83	86.82	87.11	87.49	87.17	87.42	87.27	86.63	85.69	86.17
Tar	85.61	86.54	85.91	86.24	86.23	86.22	85.95	85.62	84.77	84.56
Vad	89.32	90.96	90.98	91.12	91.26	91.02	90.93	90.05	90.35	90.57
S13	86.53	87.59	87.47	87.82	87.74	87.27	87.46	86.59	86.62	85.89
S17	91.51	92.12	92.31	92.32	92.35	91.96	91.68	90.78	90.06	90.02
S16	88.07	88.72	88.85	89.18	89.03	88.71	88.19	88.01	87.74	87.1
#wins	0	3	3	4	3	6	0	1	0	2
rank sums	169.0	89.0	74.0	95.0	77.0	94.0	126.5	151.5	175.0	159.0
position	9.0	3.0	1.0	5.0	2.0	4.0	6.0	7.0	10.0	8.0

Table A.35: Average classification accuracies (%) achieved by fine-tuning BERTweet with different samples of generic unlabeled tweets, using the XGB classifier

Dataset	BERTweet									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	67.38	73.33	68.57	64.52	63.1	65.71	69.52	76.19	64.21	61.67
sar	53.21	56.07	60.71	55.18	63.39	64.82	56.43	56.43	57.74	59.29
ntu	83.85	89.55	85.62	87.78	85.97	87.82	86.38	88.54	88.02	87.12
S15	90.04	91.89	90.64	89.11	90.33	90.02	88.14	88.15	87.54	86.92
stm	85.25	86.07	86.91	85.79	85.8	88.3	85.8	87.74	87.47	86.36
per	86.32	85.64	83.14	86.34	84.06	83.37	82.68	84.08	83.59	81.33
hob	81.81	87.16	84.48	85.26	84.1	84.67	83.15	84.88	84.35	81.98
iph	85.72	85.71	87.21	84.78	84.59	84.78	84.97	85.73	83.27	85.15
mov	87.88	88.06	88.59	88.77	88.24	90.38	87.17	88.06	88.17	87.53
san	88.72	89.87	90.61	90.93	88.97	88.56	88.8	88.4	88.32	86.84
Nar	91.85	94.62	93.81	93.48	93.07	93.15	93.23	93.15	92.64	92.01
arc	89.82	90.86	91.79	90.98	91.03	90.28	90.57	90.63	89.6	90.86
S18	87.36	88.22	88.38	87.68	88.11	87.14	87.47	85.48	85.23	84.56
OMD	84.68	85.0	86.99	85.78	85.31	84.78	86.52	85.26	84.49	84.47
HCR	81.13	80.76	81.66	81.81	80.77	81.5	81.45	81.39	79.44	78.88
STS	91.25	91.84	92.38	92.28	92.03	91.64	91.35	91.64	90.89	90.9
SSt	85.32	86.37	86.94	87.2	85.1	85.06	84.4	85.71	84.18	84.71
Tar	85.15	85.55	85.81	85.32	85.35	84.31	84.22	84.66	83.68	82.81
Vad	88.47	90.16	90.37	90.11	89.61	89.54	90.06	89.73	89.78	89.13
S13	87.85	88.53	88.99	88.15	87.85	87.62	87.71	87.6	86.98	86.23
S17	91.43	91.92	92.37	92.36	92.3	91.52	91.41	90.8	89.59	89.74
S16	90.21	90.85	91.11	90.85	90.75	90.28	90.37	89.55	89.45	88.8
#wins	0	4	10	4	0	3	0	1	0	0
rank sums	152.5	76.5	56.0	81.0	112.0	119.5	138.0	117.0	171.0	186.5
position	8.0	2.0	1.0	3.0	4.0	6.0	7.0	5.0	9.0	10.0

Table A.36: Average F_1 -macro scores (%) achieved by fine-tuning BERTweet with different samples of unlabeled tweets, using the XGB classifier

Dataset	BERTweet									
	0.5K	1K	5K	10K	25K	50K	250K	500K	1.5M	6.7M
iro	54.36	61.9	53.1	49.09	51.4	53.28	62.24	66.75	53.0	42.85
sar	49.87	53.96	59.51	53.78	62.33	62.79	54.39	55.69	56.94	55.57
ntu	83.16	89.22	85.29	87.38	85.11	87.31	85.77	88.22	87.64	86.64
S15	71.56	77.95	73.8	69.48	73.47	70.08	63.77	63.21	61.07	58.37
stm	85.2	85.99	86.81	85.72	85.73	88.27	85.71	87.71	87.42	86.28
per	82.43	81.37	77.92	81.86	78.86	77.46	77.25	78.55	77.76	74.77
hob	77.56	84.41	81.24	81.87	80.0	81.5	79.35	80.79	80.85	78.1
iph	81.95	82.07	84.28	80.85	80.71	80.71	81.08	82.04	78.76	81.09
mov	71.91	73.18	76.03	75.34	73.91	78.82	70.84	72.74	73.3	70.2
san	88.62	89.79	90.55	90.87	88.88	88.46	88.7	88.29	88.19	86.71
Nar	91.39	94.35	93.48	93.12	92.72	92.78	92.88	92.83	92.26	91.6
arc	89.52	90.66	91.59	90.76	90.81	90.04	90.35	90.4	89.33	90.62
S18	87.25	88.09	88.26	87.55	87.99	87.03	87.35	85.32	85.07	84.41
OMD	83.01	83.39	85.56	84.26	83.75	83.14	85.0	83.56	82.72	82.69
HCR	73.08	72.69	74.28	74.31	73.16	73.97	74.0	73.57	70.16	69.09
STS	89.3	90.18	90.85	90.69	90.47	89.82	89.35	89.98	88.91	88.98
SSt	84.74	85.88	86.41	86.7	84.54	84.48	83.79	85.15	83.59	84.18
Tar	85.13	85.54	85.8	85.31	85.34	84.3	84.21	84.64	83.67	82.8
Vad	85.68	87.98	88.26	87.9	87.32	87.14	87.89	87.45	87.55	86.75
S13	83.74	84.73	85.43	84.33	83.94	83.48	83.49	83.42	82.51	81.36
S17	90.76	91.27	91.79	91.77	91.71	90.88	90.74	90.09	88.76	88.9
S16	87.16	88.07	88.43	88.06	87.94	87.26	87.4	86.3	86.11	85.17
#wins	1	4	10	3	0	3	0	1	0	0
rank sums	153.0	74.0	53.0	83.0	108.5	122.5	136.0	121.0	169.0	190.0
position	8.0	2.0	1.0	3.0	4.0	6.0	7.0	5.0	9.0	10.0

A.4 Transformer-based Text Representation Tuned With Sentiment Datasets

Table A.37: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the LR classifier and BERT

	Accuracy			F_1 -macro		
Dataset	BERT					
	AllData	LOO	InData	AllData	LOO	InData
iro	76.6	75.48	72.4	68.8	65.16	65.6
sar	67.7	67.5	69.2	65.8	66.05	67.8
ntu	85.5	84.55	89.1	85.2	84.11	88.8
S15	90.0	90.02	90.9	78.5	79.78	78.9
stm	87.2	85.79	89.9	87.1	85.75	89.8
per	86.2	86.79	87.5	83.7	84.21	85.2
hob	87.8	87.92	88.2	86.2	86.17	86.7
iph	84.3	84.77	86.7	81.8	81.91	84.4
mov	88.0	89.48	88.5	80.0	82.8	81.0
san	87.6	85.87	88.9	87.5	85.77	88.8
Nar	92.0	92.1	94.6	91.6	91.75	94.4
arc	89.6	88.82	90.8	89.4	88.56	90.6
S18	86.6	85.53	88.5	86.5	85.43	88.4
OMD	84.8	84.05	86.7	83.8	82.95	85.9
HCR	78.8	77.67	79.9	74.4	73.32	76.0
STS	91.4	90.71	93.4	90.0	89.21	92.3
SSt	82.9	82.53	86.8	82.4	82.05	86.4
Tar	82.8	83.3	84.7	82.8	83.29	84.7
Vad	89.4	89.54	91.4	87.8	87.92	90.1
S13	88.0	87.67	89.5	85.2	84.9	87.1
S17	91.0	90.36	92.1	90.4	89.75	91.6
S16	89.1	89.07	90.6	86.7	86.64	88.5
#wins	1	1	20	1	2	19
rank sums	51.0	56.0	25.0	51.0	56.0	25.0
position	2.0	3.0	1.0	2.0	3.0	1.0

Table A.38: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the RF classifier and BERT

	Accuracy			F_1 -macro		
Dataset	BERT					
	AllData	LOO	InData	AllData	LOO	InData
iro	69.8	70.4	66.7	48.0	48.93	41.8
sar	70.1	73.99	64.6	68.1	72.47	61.3
ntu	82.0	78.66	83.9	81.1	77.63	82.9
S15	86.0	85.77	85.9	49.6	48.37	49.1
stm	82.6	86.54	85.9	82.6	86.5	85.8
per	80.2	79.88	77.5	69.3	68.28	63.6
hob	79.8	80.46	79.2	72.7	73.83	71.2
iph	82.6	82.95	82.8	76.1	76.47	76.6
mov	82.5	82.24	82.5	47.6	46.33	47.8
san	84.8	84.12	85.5	84.5	83.89	85.4
Nar	88.9	88.97	89.5	88.3	88.24	88.7
arc	88.5	87.49	88.9	88.0	86.95	88.5
S18	84.6	83.4	85.3	84.2	83.01	85.0
OMD	80.5	79.7	81.5	76.8	75.6	78.0
HCR	77.6	77.39	76.6	61.8	61.31	58.9
STS	85.7	85.4	85.2	80.9	80.59	80.0
SSt	80.3	80.24	82.0	79.3	79.22	80.9
Tar	82.5	82.28	83.4	82.5	82.27	83.4
Vad	84.3	84.72	84.7	79.1	79.91	79.2
S13	82.9	82.74	82.0	74.2	73.97	72.0
S17	88.5	88.36	89.1	87.4	87.15	88.1
S16	85.7	85.21	85.7	79.3	78.66	79.2
#wins	7	6	11	6	5	11
rank sums	42.0	49.0	41.0	41.0	51.0	40.0
position	2.0	3.0	1.0	2.0	3.0	1.0

Table A.39: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the MLP classifier and BERT

	Accuracy			F_1 -macro		
Dataset	BERT					
	AllData	LOO	InData	AllData	LOO	InData
iro	76.0	73.89	74.5	67.4	61.34	66.8
sar	67.7	71.78	69.7	66.1	70.27	68.6
ntu	85.1	84.55	88.7	84.8	84.1	88.5
S15	91.2	91.17	91.1	78.5	79.76	77.5
stm	87.6	87.83	89.0	87.5	87.76	89.0
per	87.0	87.63	86.6	84.1	84.95	83.5
hob	88.4	87.48	87.6	86.8	85.6	85.7
iph	85.0	84.71	87.3	82.2	81.42	84.7
mov	88.8	90.02	88.7	79.8	81.91	80.1
san	88.5	87.99	88.5	88.4	87.92	87.9
Nar	91.9	91.82	94.4	91.6	91.47	94.1
arc	90.4	89.82	90.9	90.2	89.57	90.6
S18	86.7	86.46	88.7	86.6	86.37	88.6
OMD	85.9	86.24	86.3	84.7	85.14	85.2
HCR	80.1	79.91	80.1	74.8	74.56	75.1
STS	91.5	91.28	93.2	90.0	89.74	92.0
SSt	84.0	83.31	86.3	83.5	82.8	85.9
Tar	84.1	84.1	84.8	84.1	84.09	84.8
Vad	90.2	89.96	91.4	88.5	88.24	89.8
S13	88.4	88.1	89.5	85.3	85.02	86.8
S17	91.4	91.19	92.2	90.8	90.57	91.6
S16	89.3	89.19	90.7	86.5	86.26	88.2
#wins	5	3	16	3	4	15
rank sums	43.5	56.5	32.0	45.0	55.0	32.0
position	2.0	3.0	1.0	2.0	3.0	1.0

Table A.40: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the XGB classifier and BERT

	Accuracy			F_1 -macro		
Dataset	BERT					
	AllData	LOO	InData	AllData	LOO	InData
iro	65.5	75.48	69.0	52.9	64.16	56.5
sar	65.1	68.75	57.7	62.9	67.34	54.8
ntu	81.5	80.97	83.2	80.7	80.39	82.6
S15	89.1	88.48	89.1	65.3	64.26	67.5
stm	83.2	85.24	84.6	83.1	85.19	84.5
per	83.1	83.38	84.2	77.3	77.54	78.4
hob	84.3	83.53	82.9	80.9	79.89	79.2
iph	84.5	84.39	84.5	80.9	80.25	80.7
mov	85.7	85.75	86.9	65.8	66.22	70.0
san	86.6	84.8	86.6	86.5	84.64	86.5
Nar	89.6	89.48	91.9	89.1	88.93	91.5
arc	89.5	87.78	90.5	89.2	87.43	90.2
S18	85.4	85.37	86.7	85.2	85.17	86.5
OMD	84.2	83.9	85.3	82.4	82.05	83.7
HCR	81.8	81.6	80.8	74.5	74.12	72.8
STS	89.3	89.48	90.8	87.0	87.19	88.8
SSt	82.8	82.22	85.0	82.1	81.48	84.4
Tar	83.5	83.36	84.8	83.5	83.35	84.8
Vad	88.1	88.61	89.4	85.6	86.19	87.1
S13	87.0	87.23	87.8	82.8	83.09	83.8
S17	90.8	90.78	91.5	90.0	90.05	90.8
S16	89.0	88.87	90.2	85.6	85.42	87.3
#wins	5	3	17	4	3	16
rank sums	47.5	53.0	31.5	48.5	52.0	31.5
position	2.0	3.0	1.0	2.0	3.0	1.0

Table A.41: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the LR classifier and RoBERTa

	Accuracy			F_1 -macro		
Dataset	RoBERTa					
	AllData	LOO	InData	AllData	LOO	InData
iro	73.8	74.52	70.8	66.6	67.99	63.5
sar	70.0	75.89	66.8	67.9	75.06	65.0
ntu	90.3	90.32	88.4	90.1	90.08	88.2
S15	89.5	89.73	87.6	77.1	79.24	73.8
stm	90.6	91.1	90.1	90.6	91.05	90.0
per	86.9	85.65	85.6	84.4	82.98	82.8
hob	88.6	88.12	88.6	87.2	86.5	87.3
iph	86.5	86.28	85.9	84.2	84.12	83.5
mov	90.4	89.83	90.2	83.7	82.84	83.6
san	90.3	89.79	90.2	90.2	89.73	90.1
Nar	94.3	93.89	93.7	94.1	93.65	93.4
arc	90.4	89.75	90.8	90.2	89.51	90.6
S18	88.6	88.49	88.0	88.5	88.4	87.9
OMD	87.6	85.94	87.8	86.8	85.02	86.9
HCR	80.0	79.08	80.0	75.8	74.81	75.7
STS	93.8	93.31	93.5	92.8	92.21	92.4
SSt	86.6	86.24	86.3	86.2	85.85	85.9
Tar	84.9	85.15	84.7	84.9	85.14	84.6
Vad	91.4	91.13	91.2	90.1	89.74	89.8
S13	90.0	89.29	89.1	87.6	86.83	86.6
S17	92.3	92.25	91.8	91.8	91.76	91.3
S16	90.6	90.44	90.7	88.5	88.31	88.6
#wins	13	6	5	13	5	4
rank sums	32.0	48.0	52.0	31.0	49.0	52.0
position	1.0	2.0	3.0	1.0	2.0	3.0

Table A.42: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the RF classifier and RoBERTa

	Accuracy			F_1 -macro		
Dataset	RoBERTa					
	AllData	LOO	InData	AllData	LOO	InData
iro	66.5	67.14	66.3	43.2	43.99	41.9
sar	71.8	78.33	63.6	68.6	76.11	59.5
ntu	86.9	86.94	83.3	86.4	86.46	82.5
S15	87.0	87.23	85.8	55.6	56.54	48.6
stm	88.8	89.61	88.6	88.7	89.58	88.5
per	79.3	77.6	76.9	67.6	63.75	61.8
hob	79.3	80.72	80.3	72.0	73.82	73.2
iph	83.1	82.71	82.3	77.0	76.58	76.2
mov	82.6	82.18	82.2	48.2	46.01	46.0
san	88.4	88.15	86.6	88.3	88.0	86.4
Nar	91.4	90.68	89.7	90.9	90.15	89.0
arc	90.3	89.99	89.7	89.9	89.64	89.3
S18	86.3	86.57	85.3	86.1	86.36	85.0
OMD	82.9	81.56	83.4	79.7	78.1	80.4
HCR	77.5	77.74	76.8	61.7	62.31	59.5
STS	86.9	85.43	85.9	82.8	80.38	81.1
SSt	84.0	83.16	82.4	83.2	82.23	81.4
Tar	83.5	83.63	82.8	83.5	83.62	82.8
Vad	86.1	85.92	85.3	81.7	81.53	80.5
S13	83.1	82.79	82.3	74.1	73.54	72.5
S17	89.6	89.49	88.9	88.6	88.45	87.8
S16	86.5	86.85	86.3	80.6	81.23	80.4
#wins	11	10	1	11	10	1
rank sums	34.0	37.0	61.0	34.0	36.0	62.0
position	1.0	2.0	3.0	1.0	2.0	3.0

Table A.43: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the MLP classifier and RoBERTa

	Accuracy			F_1 -macro		
Dataset	RoBERTa					
	AllData	LOO	InData	AllData	LOO	InData
iro	75.3	77.38	68.1	67.3	69.95	57.2
sar	71.0	80.66	65.8	69.4	79.35	63.8
ntu	89.6	90.2	88.9	89.3	89.95	88.6
S15	89.8	89.51	89.3	74.9	75.43	74.2
stm	90.4	89.7	89.0	90.4	89.65	88.9
per	86.2	84.59	85.0	82.9	81.1	81.6
hob	88.6	87.55	89.5	86.9	85.26	88.0
iph	86.3	87.72	86.5	83.6	85.29	83.7
mov	91.0	89.71	90.7	83.7	81.35	82.9
san	91.1	89.86	89.8	91.1	89.79	89.8
Nar	93.9	94.13	93.3	93.6	93.88	93.0
arc	90.5	89.77	90.4	90.2	89.51	90.2
S18	88.6	88.81	87.7	88.5	88.74	87.6
OMD	86.8	85.78	87.1	85.8	84.67	86.1
HCR	80.4	80.31	80.4	75.7	75.32	75.5
STS	93.3	93.43	92.7	92.2	92.31	91.4
SSt	87.0	86.21	86.1	86.6	85.78	85.7
Tar	84.9	85.23	84.8	84.9	85.23	84.7
Vad	91.6	90.96	91.2	90.2	89.39	89.7
S13	89.7	89.4	89.1	87.0	86.65	86.2
S17	92.4	92.31	91.9	91.9	91.76	91.3
S16	90.6	90.61	90.6	88.1	88.13	88.0
#wins	11	9	3	10	10	3
rank sums	35.0	42.0	55.0	35.5	42.0	54.5
position	1.0	2.0	3.0	1.0	2.0	3.0

Table A.44: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the XGB classifier and RoBERTa

	Accuracy			F_1 -macro		
Dataset	RoBERTa					
	AllData	LOO	InData	AllData	LOO	InData
iro	70.1	61.19	62.5	60.1	48.44	49.4
sar	62.5	77.32	63.5	59.9	76.75	61.6
ntu	85.4	84.91	82.6	84.9	84.35	81.9
S15	91.0	89.73	88.1	73.5	70.33	65.2
stm	86.3	88.6	87.8	86.2	88.55	87.8
per	81.8	83.37	82.5	75.5	77.72	76.1
hob	83.1	80.46	83.6	79.4	76.06	80.1
iph	85.3	85.52	84.0	81.7	81.9	80.2
mov	88.6	86.28	86.5	75.6	70.66	67.5
san	89.8	90.03	88.2	89.8	89.96	88.1
Nar	92.9	91.85	91.6	92.5	91.5	91.2
arc	90.8	90.92	89.6	90.6	90.7	89.3
S18	87.6	88.22	86.8	87.5	88.11	86.6
OMD	86.7	86.31	87.3	85.2	84.79	85.9
HCR	82.0	81.87	81.6	74.7	74.52	73.6
STS	91.6	91.4	91.1	89.9	89.52	89.2
SSt	86.2	85.58	84.5	85.8	85.05	83.9
Tar	84.8	84.28	84.7	84.8	84.27	84.7
Vad	90.1	89.61	89.6	88.1	87.53	87.4
S13	88.3	88.05	87.5	84.5	84.17	83.4
S17	91.8	91.7	91.4	91.2	91.05	90.7
S16	90.3	90.36	90.1	87.3	87.47	87.1
#wins	12	8	2	12	8	2
rank sums	35.0	41.0	56.0	35.0	40.0	57.0
position	1.0	2.0	3.0	1.0	2.0	3.0

Table A.45: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the LR classifier and BERTweet

	Accuracy			F_1 -macro		
Dataset	BERTweet					
	AllData	LOO	InData	AllData	LOO	InData
iro	75.9	83.33	72.8	69.1	76.47	66.6
sar	69.1	76.07	65.5	67.9	74.75	63.6
ntu	92.0	93.89	89.4	91.7	93.66	89.2
S15	92.3	92.52	92.1	83.3	84.76	81.7
stm	91.4	92.47	91.5	91.3	92.45	91.4
per	88.0	88.38	88.1	85.8	86.14	85.6
hob	89.1	89.47	88.2	87.5	87.96	86.7
iph	87.8	85.9	87.8	85.8	83.66	85.7
mov	91.1	89.3	89.3	85.4	82.63	82.4
san	91.6	90.85	90.1	91.5	90.76	90.0
Nar	96.9	96.82	96.3	96.7	96.7	96.1
arc	91.9	91.62	91.6	91.7	91.44	91.4
S18	90.0	88.49	89.9	90.0	88.42	89.8
OMD	88.6	88.98	88.3	87.8	88.22	87.5
HCR	81.2	81.04	80.4	77.4	77.22	76.6
STS	95.1	94.99	94.8	94.3	94.18	94.0
SSt	88.3	88.55	88.6	87.9	88.25	88.2
Tar	86.3	86.07	85.6	86.2	86.06	85.6
Vad	92.7	93.14	92.6	91.6	92.05	91.4
S13	90.3	89.54	90.4	88.1	87.23	88.3
S17	93.3	93.13	92.9	92.9	92.68	92.5
S16	91.7	91.46	91.6	89.8	89.56	89.7
#wins	11	9	3	11	11	1
rank sums	36.5	39.5	56.0	35.5	37.5	59.0
position	1.0	2.0	3.0	1.0	2.0	3.0

Table A.46: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the RF classifier and BERTweet

	Accuracy			F_1 -macro		
Dataset	BERTweet					
	AllData	LOO	InData	AllData	LOO	InData
iro	68.1	66.67	66.7	44.9	40.73	41.8
sar	69.3	68.16	65.1	66.8	65.41	61.5
ntu	88.9	88.14	85.2	88.3	87.4	84.3
S15	87.5	87.33	86.1	58.0	56.77	50.2
stm	90.6	90.99	87.3	90.6	90.97	87.2
per	80.6	81.77	77.9	70.4	72.54	64.7
hob	81.0	81.74	79.7	74.3	74.54	72.2
iph	83.1	82.14	83.5	77.1	74.96	77.7
mov	82.8	82.65	82.6	48.9	48.36	48.1
san	88.7	89.52	86.9	88.6	89.4	86.8
Nar	93.0	92.58	91.0	92.6	92.1	90.4
arc	91.4	90.61	90.0	91.1	90.29	89.6
S18	87.1	86.57	86.3	86.8	86.34	86.0
OMD	83.4	83.38	82.2	80.6	80.4	79.0
HCR	77.2	77.86	76.5	60.8	62.44	58.8
STS	87.4	87.76	86.1	83.3	83.81	81.3
SSt	83.8	83.18	82.4	82.8	82.16	81.2
Tar	83.9	84.49	83.6	83.9	84.48	83.6
Vad	85.5	85.49	84.5	80.6	80.48	78.8
S13	82.9	83.1	82.1	73.7	74.27	71.8
S17	89.9	89.6	89.6	88.9	88.6	88.6
S16	86.3	85.95	86.0	80.4	79.53	79.7
#wins	13	8	1	13	8	1
rank sums	31.0	39.5	61.5	31.0	39.5	61.5
position	1.0	2.0	3.0	1.0	2.0	3.0

Table A.47: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the MLP classifier and BERTweet

	Accuracy			F_1 -macro		
Dataset	BERTweet					
	AllData	LOO	InData	AllData	LOO	InData
iro	76.0	80.87	74.4	68.0	72.95	67.2
sar	71.6	76.07	68.3	70.3	74.87	67.2
ntu	91.8	92.58	89.2	91.6	92.36	88.9
S15	93.3	94.39	91.7	83.9	87.31	79.2
stm	91.4	91.0	90.4	91.4	90.96	90.4
per	87.5	88.52	86.8	84.6	86.05	83.4
hob	89.1	89.14	87.4	87.4	87.3	85.4
iph	87.5	86.53	87.6	85.1	83.87	85.0
mov	91.0	90.85	89.8	84.1	83.69	82.0
san	91.6	90.96	90.1	91.5	90.87	89.5
Nar	96.9	96.09	96.0	96.8	95.92	95.8
arc	92.1	91.64	91.5	91.9	91.45	91.3
S18	89.9	88.29	89.5	89.8	88.21	89.4
OMD	88.1	88.07	87.4	87.2	87.16	86.4
HCR	81.1	81.75	80.2	76.0	76.99	75.3
STS	94.8	94.77	94.4	94.0	93.86	93.5
SSt	87.9	87.93	87.6	87.5	87.55	87.2
Tar	85.9	86.0	85.5	85.8	86.0	85.5
Vad	92.8	93.0	92.2	91.5	91.8	90.8
S13	90.4	89.71	90.4	88.0	87.04	87.9
S17	93.2	93.21	92.7	92.8	92.73	92.2
S16	91.4	91.25	91.3	89.2	88.95	89.1
#wins	10	11	2	13	9	0
rank sums	34.5	37.0	60.5	31.0	39.0	62.0
position	1.0	2.0	3.0	1.0	2.0	3.0

Table A.48: Accuracies and F_1 -macro scores (%) achieved by evaluating InData, LOO, and AllData fine-tuning strategies using the XGB classifier and BERTweet

	Accuracy			F_1 -macro		
Dataset	BERTweet					
	AllData	LOO	InData	AllData	LOO	InData
iro	70.4	61.9	65.9	58.1	45.24	51.6
sar	68.3	65.0	60.5	66.1	62.72	57.1
ntu	88.0	86.72	86.0	87.4	86.13	85.4
S15	91.0	90.96	90.5	75.1	76.98	72.4
stm	88.7	91.08	84.0	88.6	91.05	84.0
per	84.0	85.19	85.6	78.8	81.01	80.9
hob	86.0	85.45	82.8	83.2	81.95	79.0
iph	86.4	87.59	84.8	82.8	84.22	80.8
mov	89.7	88.05	87.8	77.5	72.95	73.1
san	89.9	89.87	87.7	89.9	89.79	87.6
Nar	94.1	92.83	93.3	93.8	92.45	93.0
arc	91.7	91.85	91.3	91.5	91.65	91.0
S18	88.2	88.06	87.3	88.1	87.94	87.1
OMD	86.7	86.31	86.0	85.3	84.85	84.5
HCR	82.0	82.21	81.0	74.7	75.26	73.0
STS	92.7	92.33	91.5	91.2	90.81	89.8
SSt	85.7	86.24	85.9	85.2	85.66	85.3
Tar	85.6	85.92	85.2	85.6	85.92	85.2
Vad	90.5	90.54	89.9	88.4	88.42	87.6
S13	88.3	88.6	88.0	84.5	85.06	84.0
S17	92.2	92.12	91.8	91.5	91.5	91.2
S16	91.0	90.94	91.0	88.3	88.18	88.3
#wins	13	8	2	12	11	1
rank sums	33.5	39.0	59.5	35.0	37.5	59.5
position	1.0	2.0	3.0	1.0	2.0	3.0