

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Progressive Sentiment Analysis for Code-Switched Text Data

Permalink

<https://escholarship.org/uc/item/5gj9344w>

Author

Ranjan, Sudhanshu

Publication Date

2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Progressive Sentiment Analysis for Code-Switched Text Data

A Thesis submitted in partial satisfaction of the
requirements for the degree of Master of Science

in

Computer Science

by

Sudhanshu Ranjan

Committee in charge:

Professor Jingbo Shang, Chair
Professor Julian John Mcauley
Professor Ndapandula Nakashole

2022

Copyright

Sudhanshu Ranjan, 2022

All rights reserved.

The Thesis of Sudhanshu Ranjan is approved and is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2022

TABLE OF CONTENTS

Thesis Approval Page	iii
Table of Contents	iv
List of Figures	v
List of Tables	vi
Acknowledgements	vii
Abstract of the Thesis	ix
Chapter 1 Introduction	1
1.1 Preliminaries	3
1.2 Related work	4
1.3 Broader impact	4
Chapter 2 Methodology	5
2.1 Source Dataset Pretraining	5
2.2 Bucket Creation	6
2.3 Progressive Training	6
2.4 Class ratio based instance selection	8
Chapter 3 Experiments	9
3.1 Datasets	9
3.2 Model training	10
3.3 Evaluation	11
3.4 Baselines	11
3.5 Performance comparison	12
3.6 Performance comparison across buckets	14
3.7 Probing task : Out-Of-Distribution (OOD) detection	14
3.8 Comparison with other multilingual models	17
3.9 Hyper-parameter sensitivity analysis	17
3.10 Qualitative analysis	18
Chapter 4 Conclusion and Future work	21
Bibliography	22

LIST OF FIGURES

Figure 1.1.	An example of code-switched text, its transliterated and the translation versions. A code-switched text generally contains phrases from multiple languages in a single sentence. The text in blue are words in Hindi that have been written in the Latin script.	1
Figure 1.2.	S is the source labelled dataset. Unlabelled code-switched dataset T is divided into buckets. As we move from bucket B_1 to B_3 , samples are dominated by low-resource language. The progressive training strategy includes samples from B_1 to B_i in i -th iteration.	2
Figure 3.1.	Distribution of fraction of English words in a sample vs number of samples for the Spanish-English, Hindi-English and Tamil-English datasets (left-to-right). For Hindi-English, we can observe two spikes in the graph. For the other two datasets, we observe the progression to be more gradual.	11
Figure 3.2.	Model performance comparison across buckets. B_1 contains English dominated instances and B_2 contains low resource language dominated instances. Values are reported across 5 runs. Points on the same vertical line are from the same run i.e. both models were initialized with same initial weights. . .	14
Figure 3.3.	Figure showing bucketwise OOD for different datasets using the model m_{pt} . In most cases, samples in B_2 are more OOD compared to samples in B_1 across different values of α	16
Figure 3.4.	Figure showing model wise OOD for bucket B_2 . We compare two models, m_{pt} and m_1 . We observe that samples in B_2 are more OOD for model m_{pt} compared to m_1	17
Figure 3.5.	F1 score using different values of parameter δ for the three datasets (Spanish-English, Hindi-English and Tamil-English). The plots represent macro, micro and weighted F1 score (left-to-right).	18

LIST OF TABLES

Table 3.1.	Statistics of datasets. SemEval2017 is a supervised English dataset and the rest are code-switched datasets.	10
Table 3.2.	Performance on Spanish-English, Hindi-English & Tamil-English datasets. For Tamil-English dataset, - <i>Ratio</i> method increases F1 score of majority class by $\sim 2\%$ but F1 score of negative class drops by $\sim 9\%$. Thus, we observe an increase in weighted F1 score but a decrease in macro F1 score.	13
Table 3.3.	Performance using multiple multilingual models. First three rows denote performance without using progressive training and the last row denotes the performance when the model with best performance is used with progressive training.	18
Table 3.4.	Model performance on the three datasets for different number of buckets (k) using mBERT.	18
Table 3.5.	Instances where model m_{pt} was incorrect but model m_1 learnt the correct label.	19
Table 3.6.	Instances where model m_1 was incorrect but model m_2 learnt the correct label.	20

ACKNOWLEDGEMENTS

I would like to acknowledge Jingbo for his immense support and guidance. His insights have been extremely helpful throughout the research in setting the problem statement, coming up with approaches and conducting experiments. This theses wouldn't have been possible without him. I would also like to thank Dheeraj for being a wonderful collaborator. It's hard to imagine getting through the uncertainties of research without his support. I would like to thank my family and friends who have been supportive throughout my Masters.

This thesis in full, is a reprint of the material as it would appear in Ranjan, Sudhanshu; Mekala, Dheeraj; Shang, Jingbo. "Progressive Sentiment Analysis for Code-Switched Text Data" in Findings of Empirical Methods in Natural Language Processing (Findings of EMNLP), 2022. The dissertation/thesis author was the primary investigator and author of this paper.

To my mom & dad for their love and support.

ABSTRACT OF THE THESIS

Progressive Sentiment Analysis for Code-Switched Text Data

by

Sudhanshu Ranjan

Master of Science in Computer Science

University of California San Diego, 2022

Professor Jingbo Shang, Chair

Multilingual transformer language models have recently attracted much attention from researchers and are used in cross-lingual transfer learning for many NLP tasks such as text classification and named entity recognition. However, similar methods for transfer learning from monolingual text to code-switched text have not been extensively explored mainly due to the following challenges: (1) Code-switched corpus, unlike monolingual corpus, consists of more than one language and existing methods can't be applied efficiently, (2) Code-switched corpus is usually made of resource-rich and low-resource languages and upon using multilingual pre-trained language models, the final model might bias towards resource-rich language. In this work, we focus on code-switched sentiment analysis where we have a labelled resource-rich

language dataset and unlabelled code-switched data. We propose a framework that takes the distinction between resource-rich and low-resource language into account. Instead of training on the entire code-switched corpus at once, we create buckets based on the fraction of words in the resource-rich language and progressively train from resource-rich language dominated samples to low-resource language dominated samples. Extensive experiments across multiple language pairs demonstrate that progressive training helps low-resource language dominated samples.

Chapter 1

Introduction

Code-switching is the phenomena where the speaker alternates between two or more languages in a conversation. The lack of annotated data and diverse combinations of languages with which this phenomenon can be observed, makes it difficult to progress in NLP tasks on code-switched data. And also, the prevalence of different languages is different, making annotations expensive and difficult.

Intuitively, multilingual language models like mBERT [1] can be used for code-switched text since a single model learns multilingual representations. Although the idea seems straightforward, there are multiple issues. Firstly, mBERT performs differently on different languages depending on their script, prevalence and predominance. mBERT performs well in medium-resource to high-resource languages, but is outperformed by non-contextual subword embeddings in a low-resource setting [2]. Moreover, the performance is highly dependent on the script [3]. Secondly, pre-trained language models have only seen monolingual sentences during the unsupervised pretraining, however code-switched text contains phrases from both the languages in a

Original	fixing mein saja hone ka gift
Transliterated	fixing मे सजा होने का gift
Translated	gift of punishment for fixing

Figure 1.1. An example of code-switched text, its transliterated and the translation versions. A code-switched text generally contains phrases from multiple languages in a single sentence. The text in blue are words in Hindi that have been written in the Latin script.

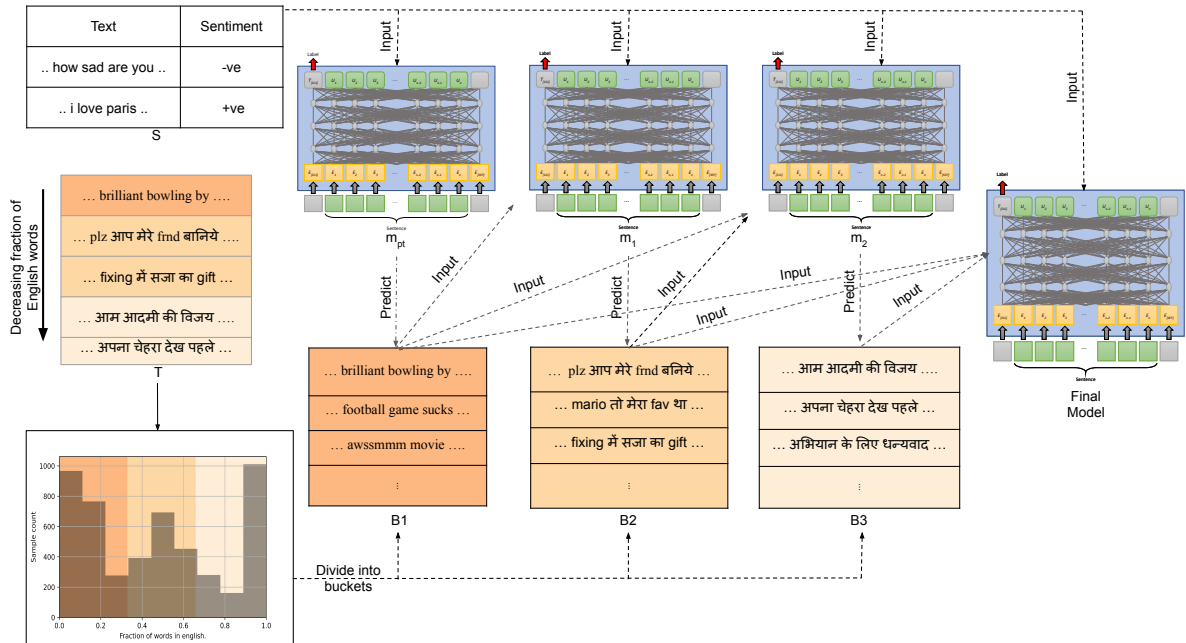


Figure 1.2. S is the source labelled dataset. Unlabelled code-switched dataset T is divided into buckets. As we move from bucket B_1 to B_3 , samples are dominated by low-resource language. The progressive training strategy includes samples from B_1 to B_i in i -th iteration.

single sentence as shown in Figure 1.1, thus making it an entirely new scenario for the language models. Thirdly, there is difference in the languages based on the amount of unsupervised corpus that is used during pretraining. For e.g., mBERT is trained on the wikipedia corpus. English has ~ 6.3 million articles, whereas Hindi and Tamil have only $\sim 140K$ articles each. This may lead to under-representation of low-resource languages in the final model. Further, English has been extensively studied by NLP community over the years, making the supervised data and tools more easily accessible. Thus, the model would be able to easily learn patterns present in the resource-rich language segments and motivating us to attempt transfer learning from English supervised datasets to code-switched datasets.

The main idea behind the approach can be summarised as follows: *When doing zero shot transfer learning from a resource-rich language (LangA) to code switched language (say LangA-LangB, where LangB is a low-resource language compared to LangA), the model is more likely to be wrong when the instances are dominated by LangB. Thus, instead of self-training on the entire corpus at once, we propose to progressively move from LangA-dominated instances*

to *LangB-dominated instances while transfer learning*. As illustrated in Figure 1.2, model trained on the annotated resource-rich language dataset is used to generate pseudo-labels for code-switched data. Progressive training uses the resource-rich language dataset and (unlabelled) resource-rich language dominated code-switched samples together to generate better quality pseudo-labels for (unlabelled) low-resource language dominated code-switched samples. Lastly, annotated resource-rich language dataset and pseudo-labelled code-switched data are then used together for training which increases the performance of the final model.

Our key contributions are summarized as:

- We propose a simple, novel training strategy that demonstrates superior performance. Since our hypothesis is based on the pretraining phase of the multilingual language models, it can be combined with any transfer learning method.
- We conduct experiments across multiple language-pair datasets, showing efficiency of our proposed method.
- We create probing experiments that verify our hypothesis.

1.1 Preliminaries

Our problem is a sentiment analysis problem where we have a labelled resource-rich language dataset and unlabelled code-switched data. From here onwards, we refer the labelled resource-rich language dataset as the source dataset and the unlabelled code-switched dataset as target dataset. Since code-switching often occurs in language pairs that include English, we refer to English as the resource-rich language. The source dataset, say S , is in English and has the text-label pairs $\{(x_{s_1}, y_{s_1}), (x_{s_2}, y_{s_2}), \dots, (x_{s_m}, y_{s_m})\}$ and the target dataset, say T , is in code-switched form and has texts $\{x_{cs_1}, x_{cs_2}, \dots, x_{cs_n}\}$, where m is significantly greater than n . The objective is to learn a sentiment classifier to detect sentiment of code-switched data by leveraging labelled source dataset and unlabelled target dataset.

1.2 Related work

Multiple tasks like Language Identification, Named Entity Recognition, Part-of-Speech, Sentiment Analysis, Question Answering and NLI have been studied in the code-switched setting. For sentiment analysis, [4] showed that multilingual approaches can outperform pipelines of monolingual models on code-switched data. [5] use CNN based network for the same. [6] use hierarchical meta embeddings to combine multilingual word, character and sub-word embeddings for the NER task. [7] augment morphological clues to language models and uses them for transfer learning from English to code-switched data with labels. [8] uses translation API to create synthetic code-switched text from English datasets and use this for transfer learning from English to code-switched text without labels in the code-switched case. [9] use synthetically generated code-switched data to enhance zero-shot cross-lingual transfer learning. Recently, [10] released the GLUECoS benchmark to study the performance of multiple models for code-switched tasks across two language pairs En-Es and En-Hi. The benchmark contains 6 tasks, 11 datasets and has 8 models for every task. Multilingual transformers fine tuned with masked-language-model objective on code-switched data can outperform generic multilingual transformers. Results from [10] show that sentiment analysis, question answering and NLI are significantly harder than tasks like NER, POS and LID. In this work, we focus on the sentiment analysis task in the absence of labeled code-switched data using multilingual transformer models, while taking into account the distinction between resource-rich and low-resource languages.

1.3 Broader impact

The broader impact of our work is to inspire methods to take into account the biases present in the current methodologies for solving problems related to code-switching and bridge the gap between performance in resource-rich language datasets and code-switched datasets.

Our code is publicly available at : <https://github.com/s1998/progressiveTrainCodeSwitch>

Chapter 2

Methodology

Our methodology can be broken down into three main steps: (1) Source dataset pretraining, which uses the resource-rich language labelled source dataset S for training a text classifier. This classifier is used to generate pseudo-labels for the target dataset T . (2) Bucket creation, which divides the unlabelled data T into buckets based on the fraction of words from resource-rich language. Some buckets would contain samples that are more resource-rich language dominated while others contain samples dominated by low-resource language. (3) Progressive training, where we initially train using S and the samples dominated by resource-rich language and gradually include the low-resource language dominated instances while training. For rest of the thesis, pretraining refers to step 1 and training refers to the training in step 3. And, we also use class ratio based instance selection to prevent the model getting biased towards majority label.

2.1 Source Dataset Pretraining

Resource-rich languages have abundant resources which includes labeled data. Intuitively, sentences in T that are similar to positive sentiment sentences in S would also be having positive sentiment (and same for the negative sentiment). Therefore, we can treat the predictions made on T by multilingual model trained on S as their respective pseudo-labels. This would assign noisy pseudo-labels to unlabeled dataset T . The source dataset pretraining step is a text classification

task. Let the model obtained after pretraining on dataset S be called m_{pt} . This model is used to generate the initial pseudo-labels and to select the instances to be used for progressive training.

2.2 Bucket Creation

Since progressive training aims to gradually progress from training on resource-rich language dominated samples to low-resource language dominated samples, we divide the dataset T into buckets based on fraction of words in resource-rich language. This creates buckets that have more resource-rich language dominated instances and also buckets that have more low-resource language dominated instances as well. In Figure 1.2, we can observe that the instances in the leftmost bucket are dominated by the English, whereas the instances in the rightmost bucket are dominated by Hindi. More specifically, we define:

$$f_{eng}(x_i) = \frac{n_{eng}(x_i)}{n_{words}(x_i)}$$

where $n_{eng}(x_i)$ and $n_{words}(x_i)$ denotes the number of English words and total number of words in the text x_i . Then, we sort the texts in dataset T in decreasing order of $f_{eng}(x_i)$ and create k buckets (B_1, \dots, B_k) with equal number of texts in each bucket. Thus, bucket B_1 contains the instances mostly dominated by English language and as we move towards buckets with higher index, instances would be dominated by the low-resource language.

2.3 Progressive Training

As the model m_{pt} is obtained by fine-tuning on a resource-rich language dataset S , it is more likely to perform better on resource-rich language dominated instances. Therefore, we choose to start progressive training from resource-rich language dominated samples. However, note that the pseudo-labels generated for dataset T are noisy, thus we sample high confident resource-rich language dominated samples to obtain better quality pseudo-labels for the rest of the instances.

Algorithm 1: Pseudocode for our progressive training framework.

Input: Source dataset S , target dataset T
Parameter: Selection fraction δ , number of buckets k
Output: Predictions on target dataset T
Model $m_{bb} \in (mBERT, MuRIL, IndicBERT)$
Model $m_{pt} \leftarrow m_{bb}$ trained on dataset S
// Bucketing step
 $T' \leftarrow ((f_{eng}(x_i), x_i, m_{pt}(x_i)) \text{ for } x_i \text{ in } T)$
 $T' \leftarrow reverse_sorted(T')$
 $(B_1, B_2, \dots, B_k) \leftarrow$ divide T' into k equal buckets
// Class ratio based instance selection
 $X_{st} = \phi$
for $i = 0, 1$ **do**
 $X_{class_i} \leftarrow$ Samples in T' predicted to be in class i using m_{pt}
 $X_{st} = X_{st} \cup (\delta \text{ most confident samples in } X_{class_i})$
// Progressive training step
Model $m_0 \leftarrow m_{pt}$
for $i = 1$ **to** k **do**
 $X_{st_i} \leftarrow X_{st} \cap B_i$
 $T_i \leftarrow \cup_{r=1}^{r=i} ((x, m_{r-1}(x)) \text{ for } x \text{ in } X_{st_r})$
 Model $m_i \leftarrow m_{bb}$ trained on $S \cup T_i$
Return $m_k(T)$

Firstly, we use m_{pt} to obtain all the high confidence samples from dataset T to be used for progressive training and their respective pseudo-labels. Among the samples to be used for progressive training, we select the samples from B_1 and use them along with S to train a second classifier which is further used to generate pseudo-labels for the rest of the samples to be used for progressive training. Then we select samples from B_2 and use them along with samples from previous iterations (i.e. samples selected from B_1 and S) to get a third classifier. We continue this process until we reach the last bucket and use the model obtained at the last iteration to make the final predictions.

More formally, we use m_{pt} to select the most confident δ fraction of samples from the dataset T , considering probability as the proxy for the confidence. Let X_{st} denote the δ fraction of samples with the highest probability of the majority class to be used for progressive training. Let $X_{st_i} = X_{st} \cap B_i$, where X_{st_i} is the subset of samples from bucket B_i that would be used for the progressive training. To train across k buckets, we use k iterations. Let m_j denote the model

obtained after training for iteration j and m_0 refers to model m_{pt} . Iteration j is trained using texts $((\cup_{i=1}^j X_{sti}) \cup S)$. The true labels for texts in S are available and for texts X_{sti} , predictions obtained using model m_{i-1} are considered as their respective labels. The model obtained at the last iteration i.e. m_k is used for final predictions.

2.4 Class ratio based instance selection

Datasets frequently have a significant amount of class imbalance. Therefore, when selecting the samples for progressive training, we often end up selecting a very small amount or no samples from the minority class which leads to very poor performance. Hence, instead of selecting δ fraction of samples from the entire dataset T , we select δ fraction of samples per class. Specifically, let X_+ and X_- denote the set of samples for which the pseudo-labels are positive and negative sentiment respectively. For progressive training, we choose δ fraction of most confident samples from X_+ and δ fraction of most confident samples from X_- .

The pseudo-code for algorithm is in Algorithm 1.

Chapter 3

Experiments

We describe the details relevant to the experiments in this section and also elaborate on the probing tasks.

3.1 Datasets

For source dataset pretraining, we use the English Twitter dataset from SemEval 2017 Task 4 [11]. We use three code-switched datasets for our experiments : Hindi-English [12], Spanish-English [13], and Tamil-English [14]. Hindi-English, Spanish-English are collected from Twitter and the Tamil-English is collected from YouTube comments. The statistics of the dataset can be found in Table 3.1. Two out of the three datasets have a class imbalance, the maximum being in the case of Tamil-English where the positive class is $\sim 5x$ of the negative class. We upsample the minority class to create a balanced dataset.

Most of the sentences in the datasets are written in the Roman script. The words in Hindi and Tamil are converted into the Devanagari script. We use the processed dataset provided by GLUECoS for Hindi-English and Spanish-English datasets. The processed version of Hindi-English dataset has the Hindi words written in Devanagari script. Since we deal with low-resource languages for which tools might not be well developed, we use heuristics to detect the words in English. For Spanish-English and Tamil-English datasets, we use the Brown corpus

from NLTK ¹ to detect the English words in the sentence. Words that are not present in the corpus are considered of another language. For the Tamil-English dataset, we use the AI4Bharat Transliteration python library ² to get the transliterations.

Table 3.1. Statistics of datasets. SemEval2017 is a supervised English dataset and the rest are code-switched datasets.

Dataset	Total	Positive	Negative
SemEval2017	27608	19799	7809
Spanish-English	914	489	425
Hindi-English	6190	3589	2601
Tamil-English	10097	8484	1613

3.2 Model training

In all the experiments we use multilingual-bert-base-cased (mBERT) classifier. The supervised English dataset has a 80-20 train-validation split. Following [15], δ is set to 0.5. We observe that in most datasets, the number of spikes in the distribution plot of $f_{eng}(x_i)$ is either 1 or 2. For example, we observe there are only two spikes for the Hindi-English dataset in Figure 3.1. Therefore, we set $k=2$. We train the classifier in both pre-training and training phase for 4 epochs. During pretraining with supervised English dataset, we choose the best weights using the validation set. While training, we use the weights obtained after fourth epoch. For final evaluation, we use the model obtained from the last iteration. The batch size is 64, sequence length is 128 and learning rate is 5e-5. These hyperparameters are same for all the models used during pre-training and training. Every iteration takes approximately $\sim 1-2$ seconds and ~ 12 GB of memory on a GPU.

In the rest of the thesis, we refer to the model pretrained on the resource-rich language source dataset as model m_{pt} , the model trained on source dataset along with bucket B_1 as m_1 ,

¹<https://www.nltk.org/>

²<https://pypi.org/project/ai4bharat-transliteration/>




Figure 3.1. Distribution of fraction of English words in a sample vs number of samples for the Spanish-English, Hindi-English and Tamil-English datasets (left-to-right). For Hindi-English, we can observe two spikes in the graph. For the other two datasets, we observe the progression to be more gradual.

and the model trained on source dataset along with the buckets B_1 and B_2 as m_2 . m_2 is used for final predictions.

3.3 Evaluation

As the datasets are significantly skewed between the two classes, we choose to report micro, macro and weighted f1 scores as done in [16]. For code-switched datasets, we use all the samples without labels during the self-training. The final score is obtained using the predictions made by model m_2 on all the samples and their true labels. For each dataset, we run the experiment with 5 seeds and report the mean and standard deviation.

3.4 Baselines

We consider four baselines as described below:

- Deep Embedding for Clustering (**DEC**) [17] has been used in WeSTClass [18] for self-training using unlabeled documents after pretraining on generated pseudo documents. We adapt DEC similarly to our setting, by pretraining on S and self-training using DEC objective only on T .
- No Progressive Training (**No-PT**) initially trains the model on the source dataset S . As

done in [15], it selects δ fraction of the code-switched data with pseudo-labels and trains a classifier on selected samples and the source dataset S without any progressive training.

- Unsupervised Self-Training (**Unsup-ST**) [19] starts with a pretrained sentiment analysis model and then self-trains using code-switched dataset. We use the default version which doesn't require human annotations. We use the model m_{pt} to initiate the self-training for fair comparison.
- Zero shot (**ZS**) [20] denotes the zero shot performance when the model is pretrained on monolingual resource-rich language dataset S .

We also compare with two ablation versions of our method, denoted by - *Source* and - *Ratio*. - *Source* uses only the code-switched dataset with its corresponding pseudo-labels without the source dataset S for training. - *Ratio* chooses the most confident samples for training without taking the class ratio into account.

We also report the performance in the supervised setting, denoted by *Supervised*. For each dataset, we train the model only on dataset T but use true labels to do the same. This is the possible upper bound.

3.5 Performance comparison

The results for all the three datasets are reported in Table 3.2. In almost all the cases, we observe a performance improvement using our method as compared to the baselines, maximum improvement being upto $\sim 1.2\%$ in the case of Spanish-English.

The comparison between ZS and our method shows the necessity of target code-switched dataset and the comparison between No-PT and our method shows that progressive training has a positive impact. In most cases, the final performance is within $\sim 10\%$ of the supervised setting. We believe our improvements are significant since the baselines are close to the supervised model in terms of the performance and yet our progressive training strategy makes a significant

Table 3.2. Performance on Spanish-English, Hindi-English & Tamil-English datasets. For Tamil-English dataset, - *Ratio* method increases F1 score of majority class by $\sim 2\%$ but F1 score of negative class drops by $\sim 9\%$. Thus, we observe an increase in weighted F1 score but a decrease in macro F1 score.

Methods	Spanish-English			Hindi-English			Tamil-English		
	Macro-F1	Micro-F1	Weighted-F1	Macro-F1	Micro-F1	Weighted-F1	Macro-F1	Micro-F1	Weighted-F1
Supervised	71.1	71.5	71.6	74.7	74.9	75.1	62.7	84.3	82.8
DEC	75.7 \pm 0.4	75.3 \pm 0.3	75.5 \pm 0.5	66.0 \pm 0.5	67.9 \pm 0.7	67.1 \pm 0.4	51.7 \pm 0.9	77.8 \pm 0.8	75.5 \pm 0.6
No-PT	76.2 \pm 0.6	76.3 \pm 0.6	76.3 \pm 0.6	67.1 \pm 0.8	69.2 \pm 0.3	68.4 \pm 0.6	53.3 \pm 0.4	78.6 \pm 0.6	76.7 \pm 0.2
ZS	76.1 \pm 0.7	76.2 \pm 0.7	76.0 \pm 0.7	63.8 \pm 0.9	67.4 \pm 0.3	65.6 \pm 0.7	52.7 \pm 0.6	79.1 \pm 0.1	76.3 \pm 0.2
Unsup-ST	73.9 \pm 1.7	74.3 \pm 1.6	74.1 \pm 1.7	66.6 \pm 1.5	66.7 \pm 1.4	66.8 \pm 1.3	49.8 \pm 2.1	69.2 \pm 5.3	69.8 \pm 8.4
Ours	77.4\pm0.8	77.5\pm0.8	77.5\pm0.8	67.8\pm0.8	69.9\pm0.5	69.1\pm0.7	53.1\pm0.4	80.5 \pm 0.2	77.5 \pm 0.1
- Source	76.5 \pm 0.7	76.7 \pm 0.6	76.6 \pm 0.7	67.2 \pm 0.5	68.7 \pm 0.4	68.3 \pm 0.4	52.7 \pm 0.7	79.2 \pm 1.0	76.8 \pm 0.6
- Ratio	77.1 \pm 0.7	77.3 \pm 0.6	77.2 \pm 0.6	64.7 \pm 0.5	68.4 \pm 0.2	66.5 \pm 0.4	49.9 \pm 0.3	83.3\pm0.1	77.7\pm0.1

improvement. We perform statistical significance test between our method and other baselines. In all the cases, we observe the p-value to be less than 0.001. The progressively trained model for Spanish-English does better than its corresponding supervised setting, outperforming it by $\sim 6\%$. We hypothesize, this is because of having a large number of instances in the source dataset S , the progressively trained model has access to more information and successfully leveraged it to improve the performance on target code-switched dataset. The comparison between our method and its ablated version - *Source* demonstrates the importance of source dataset while training the classifier. We can note that our proposed method is efficiently transferring the relevant information from the source dataset to the code-switched dataset, thereby improving the performance. On comparing our method with - *Ratio*, we observe that using class ratio based instance selection improves the performance in two out of three cases. For the Tamil-English dataset, we observe that the weighted & micro F1 score are higher for - *Ratio* method but the macro F1 score is poor. This is because the F1 score of the positive class increases by $\sim 2\%$ but F1 score of negative class drops by $\sim 9\%$ when using - *Ratio* method instead of ours. Since the dataset is skewed in the favor of the positive class, this lead to a higher weighted and micro F1 score.

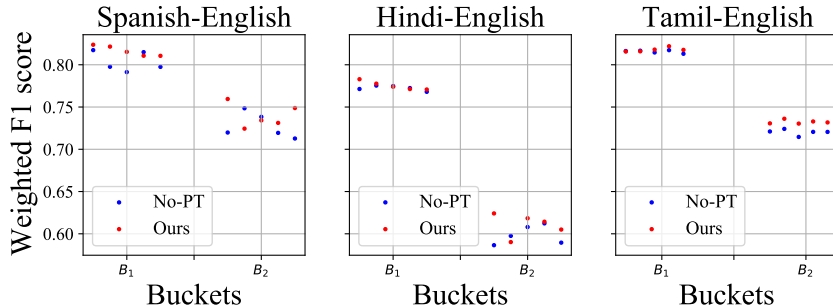


Figure 3.2. Model performance comparison across buckets. B_1 contains English dominated instances and B_2 contains low resource language dominated instances. Values are reported across 5 runs. Points on the same vertical line are from the same run i.e. both models were initialized with same initial weights.

3.6 Performance comparison across buckets

In Figure 3.2, we plot the performance obtained by No-PT and our method on both buckets. Since our method aims at improving the performance of low-resource language dominated instances, we expect our model m_2 to perform better on bucket B_2 and we observe the same. As shown in Figure 3.2, in most of the cases, our method performs better than the baseline on bucket B_2 . For bucket B_1 , we observe a minor improvement in the case of Spanish-English, whereas it stays similar for other datasets. Detailed qualitative analysis is present in 3.10.

3.7 Probing task : Out-Of-Distribution (OOD) detection

As previously mentioned, our proposed framework is based on two main hypotheses: (1.) A transformer model trained on resource-rich language dataset is more likely to be correct/robust on resource-rich language dominated samples compared to the low-resource language dominated samples, (2.) The models obtained using the progressive training framework is more likely to be correct/robust on the low-resource dominated samples compared to the models self-trained on the entire code-switched corpus at once. To confirm our hypotheses, we perform a probing task where we compute the fraction of the samples that are OOD. More specifically, we ask two questions: a) Is the fraction of OOD samples same for both the buckets for model m_{pt} ? b) Is

there a change in OOD fraction for bucket B_2 if we use model m_1 instead of model m_{pt} ? The first question helps in verifying the first part of the hypothesis and the second question helps in verifying the second part of the hypothesis.

Since the source dataset S is in English and the target dataset T is code-switched, the entire dataset T might be considered as out-of-distribution. However, transformer models are considered robust and can generalise well to OOD data [21]. Determining if a sample is OOD is difficult until we know more about the difference in the datasets. However, model probability can be used as a proxy. We use the method based on model’s softmax probability output similar to [22] to do OOD detection. Higher the probability of the predicted class, more is the confidence of the model, thus less likely the sample is out of distribution.

For a given model trained on a dataset, a threshold p_α is determined using the development set (or the unseen set of samples) to detect OOD samples. p_α is the probability value such that only α fraction of samples from the development set (or the unseen set of samples) have probability of the predicted class less than p_α . For example, if $\alpha = 10\%$, 90% of samples in the development set have probability of predicted class greater than p_α . If a new sample from another dataset (or bucket) has probability of predicted class less than p_α , we would consider it to be OOD. Using p_α , we can determine the fraction of samples from the new set that are OOD. Since, there is no method to know the exact value of α to be used, we report OOD using three values of α : 0.01, 0.05 and 0.10. For model m_{pt} , we use the development split from the dataset S to determine the value of p_α , and for model m_1 , we use the set of samples from bucket B_1 that are not used in self-training (i.e. $B_1 - X_{st_1}$) to determine p_α . Based on the value of α , we conduct two experiments and answer our two questions.

Is the fraction of OOD samples same for both the buckets for model m_{pt} ? In the first experiment, we consider the model trained on the source dataset and try to find the fraction of OOD samples in both the buckets. Since the first bucket contains more resource-rich language dominated samples, we expect a lesser fraction of samples to be out-of-distribution compared to the second bucket. In Figure 3.3, we plot the bucketwise OOD for different datasets. Values are

reported across 5 runs, points on the same vertical line are from the same run i.e. once a model m_{pt} has been trained, the same model is used to evaluate the fraction of OOD data in both the buckets. We observe that lesser fraction of samples from first bucket are OOD in all the datasets except Spanish-English. This shows that instances dominated by resource-rich language are less likely to be out-of-distribution for the classifier trained on S compared to instances dominated by low-resource language, thus providing empirical evidence in support of the first part of our hypothesis. We also report the zero-shot performance in Appendix ?? i.e. we use the classifier trained on source dataset for inference with no training on target dataset. For all the datasets, we observe zero-shot model performs better on B_1 compared to B_2 , thus further bolstering our hypothesis.




Figure 3.3. Figure showing bucketwise OOD for different datasets using the model m_{pt} . In most cases, samples in B_2 are more OOD compared to samples in B_1 across different values of α .

Is there a change in OOD fraction for bucket B_2 if we use model m_1 instead of model m_{pt} ? In the second experiment, we compare the fraction of OOD data in bucket B_2 for the models m_{pt} and m_1 . We report the values for 5 runs in Figure 3.4. Points on the same vertical line are from the same run i.e. both m_{pt} and m_1 were initialized with same initial weights. We observe a lesser fractions of samples in bucket B2 are OOD for model m_1 compared to model m_{pt} . This is expected since the model m_1 has seen samples with low-resource language words while training, thus providing empirical evidence in the support of our proposed training strategy. Although, the samples from B_2 would still have noisy labels, we expect them to be more accurate when predicted by m_1 than m_{pt} .




Figure 3.4. Figure showing model wise OOD for bucket B_2 . We compare two models, m_{pt} and m_1 . We observe that samples in B_2 are more OOD for model m_{pt} compared to m_1 .

3.8 Comparison with other multilingual models

Recently, multiple multilingual transformer models focusing on Indian languages have been proposed. We experiment with MuRIL [23] and IndicBERT [24]. Firstly, we obtain the performance of three language models: mBERT, MuRIL, and IndicBERT without progressive training on all datasets and we use progressive training on top of the best performing model corresponding to each dataset and verify whether it further improves the performance. The F1 scores are reported in Table 3.3. We observe that performance either increases or stays very competitive in all the cases, thus showing our method is capable of improving performance even when used with the best multilingual model for the task.

3.9 Hyper-parameter sensitivity analysis

There are two hyper-parameters in our experiments: the number of buckets (k) and the ratio of samples selected for self-training (δ). We vary k from 2 to 4 to study the effect of the number of buckets on the performance and the F1-scores are reported in Table 3.4. Our method is fairly robust to the values of k . For almost all values of k , our method does better than the baselines. As mentioned earlier, the number of spikes in the distribution plot of f_{eng} is 1 or 2 for all the datasets. In presence of more number of spikes, higher value of k is recommended.

For studying the effect of hyper-parameter δ , we plot macro, micro, and weighted F1 scores across multiple values of δ in Figure 3.5. With low δ , there wouldn't be enough sentences

Table 3.3. Performance using multiple multilingual models. First three rows denote performance without using progressive training and the last row denotes the performance when the model with best performance is used with progressive training.

Model	Spanish-English			Hindi-English			Tamil-English		
	Macro-F1	Micro-F1	Weighted-F1	Macro-F1	Micro-F1	Weighted-F1	Macro-F1	Micro-F1	Weighted-F1
mBERT	76.2±0.6	76.3±0.6	76.3±0.6	67.1±0.8	69.2±0.3	68.4±0.6	53.3±0.4	78.6±0.6	76.7±0.2
MuRIL	-	-	-	77.0±0.4	77.7±0.3	77.7±0.4	54.2±0.2	64.2±0.4	68.8±0.3
IndicBERT	-	-	-	73.5±0.5	74.5±0.3	74.3±0.4	54.6±0.1	68.0±0.6	71.3±0.4
Ours + Best	77.4±0.8	77.5±0.8	77.5±0.8	77.0±0.4	77.6±0.4	77.6±0.4	53.1±0.4	80.5±0.2	77.5±0.1

Table 3.4. Model performance on the three datasets for different number of buckets (k) using mBERT.

Buckets	Spanish-English			Hindi-English			Tamil-English		
	Macro-F1	Micro-F1	Weighted-F1	Macro-F1	Micro-F1	Weighted-F1	Macro-F1	Micro-F1	Weighted-F1
k=2	77.4±0.8	77.5±0.8	77.5±0.8	67.8±0.8	69.9±0.5	69.1±0.7	53.1±0.4	80.5±0.2	77.5±0.1
k=3	76.1±0.7	76.2±0.7	76.2±0.7	67.2±0.6	69.4±0.5	68.6±0.6	53.1±0.2	79.9±0.4	77.2±0.1
k=4	76.6±1.6	76.7±1.7	76.7±1.6	67.6±0.4	69.7±0.3	68.9±0.4	52.9±0.4	80.6±0.7	77.5±0.3

for self-training to help whereas with high δ , the samples would be too noisy. Thus, a value in the middle i.e. 0.4-0.6 should be reasonable choice.




Figure 3.5. F1 score using different values of parameter δ for the three datasets (Spanish-English, Hindi-English and Tamil-English). The plots represent macro, micro and weighted F1 score (left-to-right).

3.10 Qualitative analysis

As discussed previously, on the low-resource language dominated bucket, our model is correct more often than the *No-PT* baseline. We focus on samples from bucket B_2 for qualitative analysis. For the sample, "fixing me saja hone ka gift", the Hindi word "saja" refers to punishment which is negative in sentiment whereas the word "gift" is positive in

sentiment. Thus, the contextual information in the Hindi combined with that of the English is necessary to make correct prediction. For the sample ”*Mera bharat mahan, padhega India tabhi badhega India*”, the model has to identify Hindi words ”mahan” & ”badhega” to make the correct predictions.

We also do the qualitative analysis by looking at predictions of samples between successive iterations. In Table 3.5, we randomly choose samples which are predicted incorrectly by model m_{pt} but are predicted correctly by model m_1 . Out of 8 samples, 6 samples had sentiment specifically present in the Hindi words. In Table 3.6, we randomly choose samples which are predicted incorrectly by model m_1 but are predicted correctly by model m_2 . Out of 8 samples, 4 samples had sentiment specifically present in the Hindi words and 2 samples required understanding both the Hindi and English words simultaneously. The blue highlighted words are relevant to determining the sentiment of the sentence.

Table 3.5. Instances where model m_{pt} was incorrect but model m_1 learnt the correct label.

Bucket	Text	Label
0	yes we teachers really need sar dard ki goli aftr taking class	negative
0	they are promising moon right now to get the cm post . . . waade aise hone chahiye jo janta k welfare k liye ho . . . free wahi baantna chahta hai jo desperate ho kisi tarah ek baar bas kursi mil jaye .	negative
0	oooo ! grandfather bas ab nahi kitna natak karoge	negative
0	i agree with this cartoon . bahut achcha doston .	positive
0	seriously maza boht ata tha , , , mere pass mast collection hota tha	positive
1	school mein milne wale laddoo kitni khushi dete they :d	positive
1	guddu itni paas se tv dekhega toh aankhe button ho jayengi !	negative
1	arun lal ki commentary yaad aa gayi . . usse zyada manhoos koi nahi .	negative

Table 3.6. Instances where model m_1 was incorrect but model m_2 learnt the correct label.

Bucket	Text	Label
0	Osm party salman bhai apka gift	positive
0	mat dekho brothers and sisters , dekha nahi jayega	negative
0	Ap late ho ap ne apne cometment pori nai ki	negative
0	looks like 'kaho na pyaar hai' phase ended for modi #aapsweep #aapkidilli #delhidecides	negative
1	comment krne se jyada sabke padhne me maja aata h	positive
1	guddu tumhara hi school aisa hoga yaar . mera school n mere teacher to ache h	positive
1	sir kya msg krtte kartte mar jau	negative
1	Hum bhuke mur rahe hai sir	negative

Chapter 4

Conclusion and Future work

In this work, we propose progressive training framework that takes distinction between low-resource and resource-rich language into account while doing zero-shot transfer learning for code-switched sentiment analysis. We show that our framework improves performance across multiple datasets. Further, we also create probing tasks to provide empirical evidence in support of our hypothesis. In future, we want to extend the framework to other tasks like question-answering and natural language inference.

Limitations. A key potential limitation of the current framework is that if the number of samples in buckets are very disproportionate, the progressive learning might not result in significant improvement.

Ethical consideration. This work proposes a progressive training framework to transfer knowledge from resource-rich language data to low-resource code-switched data. We work on sentiment classification task which is a standard NLP problem. Based on our experiments, we don't see any major ethical concerns with our work.

Bibliography

- [1] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” 2019.
- [2] B. Heinzerling and M. Strube, “Sequence tagging with contextual and non-contextual subword representations: A multilingual evaluation,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, (Florence, Italy), pp. 273–291, Association for Computational Linguistics, July 2019.
- [3] T. Pires, E. Schlinger, and D. Garrette, “How multilingual is multilingual bert?,” 2019.
- [4] D. Vilares, M. A. Alonso, and C. Gómez-Rodríguez, “Sentiment analysis on monolingual, multilingual and code-switching Twitter corpora,” in *Proceedings of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, (Lisboa, Portugal), pp. 2–8, Association for Computational Linguistics, Sept. 2015.
- [5] Y. K. Lal, V. Kumar, M. Dhar, M. Shrivastava, and P. Koehn, “De-mixing sentiment from code-mixed text,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, (Florence, Italy), pp. 371–377, Association for Computational Linguistics, July 2019.
- [6] G. I. Winata, Z. Lin, J. Shin, Z. Liu, and P. Fung, “Hierarchical meta-embeddings for code-switching named entity recognition,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, (Hong Kong, China), pp. 3541–3547, Association for Computational Linguistics, Nov. 2019.
- [7] G. Aguilar and T. Solorio, “From English to code-switching: Transfer learning with strong morphological clues,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, (Online), pp. 8033–8044, Association for Computational Linguistics, July 2020.
- [8] B. Samanta, N. Ganguly, and S. Chakrabarti, “Improved sentiment detection via label transfer from monolingual to synthetic code-switched text,” *arXiv preprint arXiv:1906.05725*, 2019.
- [9] L. Qin, M. Ni, Y. Zhang, and W. Che, “Cosda-ml: Multi-lingual code-switching data augmentation for zero-shot cross-lingual nlp,” *ArXiv*, vol. abs/2006.06402, 2020.

- [10] S. Khanuja, S. Dandapat, A. Srinivasan, S. Sitaram, and M. Choudhury, “GLUECoS: An evaluation benchmark for code-switched NLP,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, (Online), pp. 3575–3585, Association for Computational Linguistics, July 2020.
- [11] S. Rosenthal, N. Farra, and P. Nakov, “SemEval-2017 task 4: Sentiment analysis in Twitter,” in *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, (Vancouver, Canada), pp. 502–518, Association for Computational Linguistics, Aug. 2017.
- [12] B. G. Patra, D. Das, and A. Das, “Sentiment analysis of code-mixed indian languages: An overview of sail_code-mixed shared task @icon-2017,” 2018.
- [13] D. Vilares, M. A. Alonso, and C. Gómez-Rodríguez, “EN-ES-CS: An English-Spanish code-switching Twitter corpus for multilingual sentiment analysis,” in *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, (Portorož, Slovenia), pp. 4149–4153, European Language Resources Association (ELRA), May 2016.
- [14] B. R. Chakravarthi, R. Priyadharshini, V. Muralidaran, N. Jose, S. Suryawanshi, E. Sherly, and J. P. McCrae, “Dravidiancodemix: Sentiment analysis and offensive language identification dataset for dravidian languages in code-mixed text,” *Language Resources and Evaluation*, 2016.
- [15] Z. Wang, D. Mekala, and J. Shang, “X-class: Text classification with extremely weak supervision,” in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, (Online), pp. 3043–3053, Association for Computational Linguistics, June 2021.
- [16] D. Mekala and J. Shang, “Contextualized weak supervision for text classification,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, (Online), pp. 323–333, Association for Computational Linguistics, July 2020.
- [17] J. Xie, R. Girshick, and A. Farhadi, “Unsupervised deep embedding for clustering analysis,” 2016.
- [18] Y. Meng, J. Shen, C. Zhang, and J. Han, “Weakly-supervised neural text classification,” *CoRR*, vol. abs/1809.01478, 2018.
- [19] A. Gupta, S. Menghani, S. K. Rallabandi, and A. W. Black, “Unsupervised self-training for sentiment analysis of code-switched data,” in *Proceedings of the Fifth Workshop on Computational Approaches to Linguistic Code-Switching*, (Online), pp. 103–112, Association for Computational Linguistics, June 2021.
- [20] T. Pires, E. Schlinger, and D. Garrette, “How multilingual is multilingual bert?,” *CoRR*, vol. abs/1906.01502, 2019.

- [21] D. Hendrycks, X. Liu, E. Wallace, A. Dziedzic, R. Krishnan, and D. Song, “Pretrained transformers improve out-of-distribution robustness,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, (Online), pp. 2744–2751, Association for Computational Linguistics, July 2020.
- [22] D. Hendrycks and K. Gimpel, “A baseline for detecting misclassified and out-of-distribution examples in neural networks,” 2018.
- [23] S. Khanuja, D. Bansal, S. Mehtani, S. Khosla, A. Dey, B. Gopalan, D. K. Margam, P. Aggarwal, R. T. Nagipogu, S. Dave, S. Gupta, S. C. B. Gali, V. Subramanian, and P. P. Talukdar, “Muril: Multilingual representations for indian languages,” *CoRR*, vol. abs/2103.10730, 2021.
- [24] D. Kakwani, A. Kunchukuttan, S. Golla, G. N.C., A. Bhattacharyya, M. M. Khapra, and P. Kumar, “IndicNLPSuite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for Indian languages,” in *Findings of the Association for Computational Linguistics: EMNLP 2020*, (Online), pp. 4948–4961, Association for Computational Linguistics, Nov. 2020.