

Two-stage Pipeline for Multilingual Dialect Detection

Ankit Vaidya and Aditya Kane

Pune Institute of Computer Technology, Pune
{ankitvaidya1905, adityakane1}@gmail.com

Abstract

Dialect Identification is a crucial task for localizing various Large Language Models. This paper outlines our approach to the VarDial 2023 DSL-TL shared task. Here we have to identify three or two dialects from three languages each which results in a 9-way classification for Track-1 and 6-way classification for Track-2 respectively. Our proposed approach consists of a two-stage system and outperforms other participants' systems and previous works in this domain. We achieve a score of 58.54% for Track-1 and 85.61% for Track-2. Our codebase is available publicly¹.

1 Introduction

Language has been the primary mode of communication for humans since the pre-historic ages. Studies have explored the evolution of language and outlined mathematical models that govern the intricacies of natural language (Nowak and Krakauer, 1999; Hauser et al., 2014). Inevitably, as humans established civilization in various parts of the world, this language was modified by, and for the group of people occupied by that particular geographical region. This gave rise to multiple national dialects of the same language.

The VarDial workshop (Aepli et al., 2023) (co-located with EACL 2023) explores various dialects and variations of the same language. We participated in the Discriminating Between Similar Languages - True Labels (DSL-TL) shared task. In this task, the participants were provided with data from three languages, with each language having three varieties. This shared task consisted of two tracks – Track-1 featuring nine-way classification and Track-2 featuring six-way classification. The second track included two particular national dialects of each language (eg. American English and British English), and the first track had one general

¹https://github.com/ankit-vaidya19/EACL_VarDial2023




Figure 1: Class distribution of dialects

label (English) in addition to the aforementioned two.

We ranked 1st in both of the tracks. Moreover, we beat the next best submission by a margin of 4.5% in the first task and 5.6% in the second task. We were the only team to surpass the organizer baseline scores. We present our winning solution in this paper. We used an end-to-end deep learning pipeline which consisted of a language identification model and three language-specific models, one for each language. We converged upon the best combination by doing an elaborate analysis of various models available. Furthermore, in this work we also analyze the performance of the pipeline as a whole and also provide an ablation study. Lastly, we provide some future directions in this area of research.

2 Related Work

The present literature encompasses various aspects of dialect identification. We study this from three perspectives: large language models, language identification and dialect classification problems.

2.1 Large Language Models

The success of transformers and BERT (Devlin et al., 2019) based models was inevitable since the initial boom of the transformer (Vaswani et al.,




Figure 2: System diagram for dialect classification. The LID classifies the input into one of 3 languages. The sample is then further classified into dialects by language specific models.

2017) model. In recent years, many other architectures like RoBERTa (Liu et al., 2019) and ELECTRA (Clark et al., 2020) have further pushed the state-of-the-art in this domain. Moreover, autoregressive models like GPT (Radford and Narasimhan, 2018) and GPT-2 (Radford et al., 2019) have also shown their prowess. Multilingual versions of RoBERTa, namely XLM-RoBERTa (Conneau et al., 2020) are also available. Lastly, language specific models like Spanish BERT (la Rosa y Eduardo G. Ponferrada y Manu Romero y Paulo Villegas y Pablo González de Prado Salas y María Grandury, 2022) and Portuguese BERT (Souza et al., 2020) are available as well. Our winning solution makes use of these large language models trained on specific languages.

2.2 Language Identification Models

Many multilingual language identification models have been developed in order to classify the language of the input sentence beforehand. Even though the initial works used n -gram models and generative mixture models (Lui et al., 2014; Baldwin and Lui, 2010; Al-Badrashiny and Diab, 2016; Bernier-Colborne et al., 2021; Bestgen, 2021) or even conditional random fields (Al-Badrashiny and Diab, 2016) and other classical machine learning methods like naive bayes (Jauhiainen et al., 2020, 2022, 2021), modern methods have shifted to the use of deep learning for language identification (Mathur et al., 2017; Thara and Poornachandran, 2021; Romero et al., 2021; Bernier-Colborne et al.,

2022). Recent works have focused on deep learning based methods, where handling codemixed data is a big challenge in the domain. For our experiments, we use a version of XLM-RoBERTa finetuned on a language identification dataset². This model has near-perfect test accuracy of 99.6%.

2.3 Dialect Classification

Dialect classification has been previously solved using statistical methods like Gaussian Mixture Models and Frame Selection Decoding or Support Vector Machines (SVM) (Lei and Hansen, 2011; Tillmann et al., 2014). It has been explored relatively sparsely, mostly in the case for local languages (Hegde et al., 2020). Deep learning approaches have been explored in previous editions of the VarDial workshop shared tasks (Rebeja and Cristea, 2020) and otherwise (Lin et al., 2020). Dialect classification was also explored previously as a part of other shared tasks (Khered et al., 2022). We want to stress that given the multilingual nature of the dataset, using the present methods directly was not an option. In our work, although we take inspiration from the previous works, we propose a novel system that surpasses the performance of the previous systems by a large margin.

3 Data

The dataset (Zampieri et al., 2023) contained a total of 11,610 sentences belonging to 3 languages:

²This model is available [here](#) and dataset is available [here](#)

Model-EN	Model-ES	Model-PT	Validation F1	Test F1
RoBERTa _{base}	Spanish RoBERTa _{base}	XLM-RoBERTa _{base}	60.74%	-
RoBERTa _{base}	Spanish BERT _{base}	XLM-RoBERTa _{base}	60.05%	-
BERT _{base}	Spanish RoBERTa _{base}	XLM-RoBERTa _{base}	59.08%	-
BERT _{base}	Spanish BERT _{base}	Portuguese BERT _{base}	60.17%	-
BERT _{base}	Spanish BERT _{base}	XLM-RoBERTa _{base}	58.40%	-
BERT_{base}	Spanish RoBERTa_{base}	Portuguese BERT_{base}	60.85%	58.54%
RoBERTa _{base}	Spanish RoBERTa _{base}	Portuguese BERT _{base}	62.51%	58.09%
RoBERTa _{base}	Spanish BERT _{base}	Portuguese BERT _{base}	61.83%	57.03%

Table 1: Our complete results for Track-1 using the two-stage dialect detection pipeline. Model-* denotes the language of the models used for the experiments.

Model-EN	Model-ES	Model-PT	Validation F1	Test F1
RoBERTa _{base}	XLM-RoBERTa _{base}	Portuguese BERT _{base}	80.84%	-
RoBERTa _{base}	XLM-RoBERTa _{base}	XLM-RoBERTa _{base}	79.16%	-
BERT _{base}	Spanish RoBERTa _{base}	XLM-RoBERTa _{base}	82.22%	-
BERT _{base}	XLM-RoBERTa _{base}	Portuguese BERT _{base}	80.55%	-
BERT _{base}	XLM-RoBERTa _{base}	XLM-RoBERTa _{base}	78.87%	-
RoBERTa_{base}	Spanish RoBERTa_{base}	Portuguese BERT_{base}	84.19%	85.61%
BERT _{base}	Spanish RoBERTa _{base}	Portuguese BERT _{base}	83.90%	85.11%
RoBERTa _{base}	Spanish RoBERTa _{base}	XLM-RoBERTa _{base}	82.51%	83.68%

Table 2: Our complete results for Track-2 using the two-stage dialect detection pipeline. Model-* denotes the language of the models used for the experiments.

English (EN), Spanish (ES), Portuguese (PT) and each language had 3 corresponding varieties. The varieties for English were: American English ($EN - US$), British English ($EN - GB$) and Common English Instances (EN). Similarly varieties corresponding to Spanish and Portuguese were: European/Peninsular Spanish ($ES - ES$), Argentine Spanish ($ES - AR$), Common Spanish Instances (ES) and European Portuguese ($PT - PT$), Brazilian Portuguese ($PT - BR$), Common Portuguese Instances (PT). These were divided into a training set containing 8,745 sentences and the validation set containing 2,865 sentences. The system was evaluated on a separate testing set containing 1,290 sentences. This dataset has acute class imbalance. We observed that the class PT-BR had the most number of samples (2,724) and the class EN had the least number of samples (349), and thus the imbalance ratio was almost 1:8. We have illustrated the data distribution in Figure 1. We tried to mitigate this imbalance using over-sampling and weighted sampling methods. However, the improved data sampling method did not affect the performance.

4 System Description

This was a problem of multi-class classification having 9 classes for Track-1 and 6 classes for Track-2. The samples were belonging to 3 languages having 3 varieties each, so the classification pipeline was made in 2 stages. The Language Identification (LID) model which is the first stage classifies the sentence into 3 languages: English (EN), Spanish (ES) and Portuguese (PT). The LID is a pretrained XLM-RoBERTa that is fine-tuned for the task of language identification. It is able to classify the input sentence into 20 languages. We classify and separate the samples according to their language. The samples corresponding to the specific languages are then fed into the language specific models for dialect identification. For dialect identification we have used models like BERT and RoBERTa with a linear layer connected to the pooler output of the models. Then fine-tuning is done on the models for dialect identification using the samples corresponding to the specific languages. For the task of dialect identification we experimented with several pretrained models

Lg	Model	Train F1	Val F1
EN	RoBERTa	79.74%	71.34%
EN	BERT	80.71%	70.19%
EN	ELECTRA	65.02%	66.60%
EN	XLM-RoBERTa	71.64%	66.12%
EN	GPT-2	56.78%	49.74%
ES	Spanish RoBERTa	74.36%	62.96%
ES	XLM-RoBERTa	59.46%	61.58%
ES	Spanish BERT	67.40%	60.76%
ES	Spanish GPT-2	34.33%	46.11%
PT	Portuguese BERT	67.63%	55.15%
PT	XLM-RoBERTa	64.33%	48.46%
PT	Portuguese ELECTRA	62.11%	46.34%
PT	Portuguese GPT-2	38.52%	34.19%

Table 3: Performance on Track-1 validation dataset of individual models used in the two-stage pipeline. "Lg" stands for language of the model used.

Lg	Model	Train F1	Val F1
EN	RoBERTa	91.70%	88.75%
EN	BERT	94.24%	88.32%
EN	XLM-RoBERTa	87.61%	84.68%
ES	Spanish RoBERTa	96.05%	87.05%
ES	XLM-RoBERTa	89.25%	80.29%
PT	Portuguese BERT	89.49%	79.21%
PT	XLM-RoBERTa	81.61%	75.91%

Table 4: Performance on Track-2 validation dataset of individual models used in the two-stage pipeline. "Lg" stands for language of the model used.

like XLM-RoBERTa, BERT, ELECTRA, GPT-2 and RoBERTa. All models were fine-tuned for 20 epochs with a learning rate of $1e-6$ and weight decay $1e-6$ with a batch size of 8. The best performing model checkpoint was chosen according to the epoch-wise validation macro-F1 score.

5 Experiments and Results

5.1 Experiments Using Large Language Models

For the task of Dialect Identification we have tried various language specific models like XLM-RoBERTa, BERT, ELECTRA, RoBERTa and GPT-2. The base variant of all these models were used and all the models were used through the HuggingFace (Wolf et al., 2020) library. The pooler output

Lg	Model	Adapted F1	F.T. F1
EN	RoBERTa	85.20%	88.75%
EN	BERT	83.21%	88.32%
EN	XLM-RoBERTa	81.21%	84.68%
ES	Spanish RoBERTa	78.45%	87.05%
ES	XLM-RoBERTa	66.89%	80.29%
PT	Portuguese BERT	72.17%	79.21%
PT	XLM-RoBERTa	71.89%	75.91%

Table 5: Comparative results of two-way classification using the finetuned (F.T.) predictions and predictions adapted from three-way classification models.

of these models was passed through a linear layer and the models were fine-tuned. First, we experimented with different models for Track-1. All the models were trained for 20 epochs with learning rate $1e-6$, weight decay $1e-6$ and a batch size of 8. We used XLM-RoBERTa as the baseline for all 3 languages. The best performing models for the English language were RoBERTa and BERT whereas GPT-2 was the worst performing. Similarly the language specific versions of RoBERTa and BERT performed well for the Spanish and Portuguese respectively. Overall the worst performing model was GPT-2 across all 3 languages. The validation F1 scores are present in Table 3. The two best-performing models for every language were chosen for Track-2. The same procedure as specified above was used and the F1 scores are present in Table 4. The train and validation F1 scores for 2-class classification are higher for all models as compared to the F1 score of the same models for 3-class classification. This was mainly due to the poor representation and accuracy of classification of the third class. We observed symptoms of overfitting in all models after 12-15 epochs and the best validation F1 score was obtained in the range of 4-8 epochs.

5.2 LID Experiments

The pipeline for dialect identification is divided into two parts as the sentences in the dataset belong to different languages. The stages are described in Section 4. The XLM-RoBERTa we have used for language classification has a test accuracy of 99.6% meaning it correctly classifies all input sentences and hence, can be considered as a perfect classifier. For the final pipeline we experimented using the two best performing models for each language in

Track-1 and Track-2. For both the tracks we experimented with all 8 (2^3) possible combinations of models and calculated the validation F1 score for the combined validation dataset which had sentences belonging to all languages. The validation scores for Track-1 and Track-2 are shown in Table 1 and Table 2 respectively. For both the tracks, the three pipelines with the best validation F1 scores were chosen for submission.

5.3 Using 3-way Classifier as a 2-way Classifier

In Track-1, participants are expected to train a classifier which classifies amongst 9 classes, and in Track-2, participants are expected to train a classifier which classifies amongst 6 classes. These 6 classes are a proper subset of the 9 classes from Track-1. Thus, an intuitive baseline for Track-2 is to use the model finetuned for Track-1, whilst considering only the relevant classes for the latter task. The classes *EN*, *ES* and *PT*, i.e. the classes without any national dialect associated with them are not included in Track-2 as compared to Track-1. Thus, we calculate the predictions for the Track-2 validation dataset using the models for Track-1 and exclude the metrics for Track-1 specific classes to get the metrics for this "adapted" 2-way classification. We show the results of this experiment in Table 5 and observe that, as expected, the adapted 2-way classification performs worse compared to the explicitly finetuned variant.

5.4 Results for Track-1 and Track-2

We now present our experiments and their performance for both tracks. Our experiments for Track-1 are described in Table 1 and our experiments for Track-2 are described in Table 2. The participants were allowed three submissions for evaluation on the test set, so we submitted predictions using the three systems which performed the best on the validation set. As mentioned in Section 5.2, we performed 2^3 , i.e. a total of 8 experiments using the two best models for each language. We observed that RoBERTa_{base} on English, Spanish BERT_{base} on Spanish and Portuguese BERT_{base} performed the best on the testing set for Track-1. The same combination, with RoBERTa_{base} for English, worked best for Track-2. All of our submissions were the top submissions for each track, which surpassed the next best competitors by a margin of 4.5% and 5.6% for Track-1 and Track-2 respectively.

5.5 Ablation of Best Submissions

We hereby make some observations of our submissions and other experiments. To assist this, we plot the confusion matrices of our best submissions for Track-1 and Track-2 in Figures 3 and 4 respectively. Note that these confusion matrices have their rows (i.e. true labels axes) normalized according to the number of samples in the class. Here are observations from our experiments:

- 1. BERT-based models outperform other models across all languages:** We observe that BERT-based models outperform ELECTRA-based and GPT-2-based models, as shown in Table 3. We speculate this is because of the inherent architecture of BERT, which combines semantic learning with knowledge retention. This combination of traits is particularly useful for this task.
- 2. Common labels perform the worst across all languages:** We observe that the common labels *EN*, *ES* and *PT* perform the worst, both in the individual as well as the two-stage setup. We hypothesize this is because of the absence of dialect specific words, or words that are specific to the geographical origin of the national dialect (for example, "Yankees" for EN-US and "Oxford" for EN-GB).
- 3. English models work better than models of other languages:** It can be noted from Figures 4 and 3 that the English models have the best performance across all classes. This can be attributed to two reasons: absence of national dialect specific words and lesser pre-training data in the case of Portuguese.
- 4. British English is most correctly classified class:** We can observe that the Spanish or Portuguese models make equal number of mistakes in the case of either national dialect, in the case of Track-2 (see Figure 4). However, in the case of English, the label *EN* – *GB* is correctly classified for more than 95% of the cases. We speculate this is because British English involves slightly distinctive grammar and semantics, which help the model separate it from other classes.
- 5. The proposed 2-step method is scalable for multiple language dialect classification:** We can strongly assert that the novel 2-step deep




Figure 3: Confusion matrix of 9-way classification. Note that rows are normalized according to the number of samples in that class.




Figure 4: Confusion matrix of 6-way classification. Note that rows are normalized according to the number of samples in that class.

learning method for multilingual dialect classification is a scalable method for the task due to two specific reasons: firstly, the multilingual models (like XLM-RoBERTa) might not have the vocabulary as well as the learning capabilities to learn the minute differences between individual dialects. Secondly, this system can be quickly expanded for a new language by simply adding a language specific dialect classifier, provided the language identification model supports that particular language.

6 Conclusion

In this paper we propose a two-stage classification pipeline for dialect identification for multilingual

corpora. We conduct thorough ablations on this setup and provide valuable insights. We foresee multiple future directions for this work. The first is to expand this work to many languages and dialects. Secondly, it is a worthwhile research direction to distill this multi-model setup into a single model with multiple prediction heads.

Limitations

The obvious limitation of this system is the excessive memory consumption due to the usage of language specific models. For low resource languages this system is difficult to train and scale. We hope that these problems will be addressed by researchers in future works.

References

- Noëmi Aepli, Çağrı Çöltekin, Rob van der Goot, Tommi Jauhiainen, Mourhaf Kazzaz, Nikola Ljubešić, Kai North, Barbara Plank, Yves Scherrer, and Marcos Zampieri. 2023. Findings of the VarDial evaluation campaign 2023. In *Proceedings of the Tenth Workshop on NLP for Similar Languages, Varieties and Dialects*, Dubrovnik, Croatia. Association for Computational Linguistics.
- Mohamed Al-Badrashiny and Mona Diab. 2016. [LILI: A simple language independent approach for language identification](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1211–1219, Osaka, Japan. The COLING 2016 Organizing Committee.
- Timothy Baldwin and Marco Lui. 2010. [Language identification: The long and the short of the matter](#). In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 229–237, Los Angeles, California. Association for Computational Linguistics.
- Gabriel Bernier-Colborne, Serge Leger, and Cyril Goutte. 2021. [N-gram and neural models for uralic language identification: NRC at VarDial 2021](#). In *Proceedings of the Eighth Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 128–134, Kiyv, Ukraine. Association for Computational Linguistics.
- Gabriel Bernier-Colborne, Serge Leger, and Cyril Goutte. 2022. [Transfer learning improves French cross-domain dialect identification: NRC @ VarDial 2022](#). In *Proceedings of the Ninth Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 109–118, Gyeongju, Republic of Korea. Association for Computational Linguistics.

- Yves Bestgen. 2021. [Optimizing a supervised classifier for a difficult language identification problem](#). In *Proceedings of the Eighth Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 96–101, Kiyv, Ukraine. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: pre-training text encoders as discriminators rather than generators](#). *CoRR*, abs/2003.10555.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Marc D. Hauser, Charles Yang, Robert C. Berwick, Ian Tattersall, Michael J. Ryan, Jeffrey Watumull, Noam Chomsky, and Richard C. Lewontin. 2014. [The mystery of language evolution](#). *Frontiers in Psychology*, 5.
- Pradyoth Hegde, Nagaratna B. Chittaragi, Siva Krishna P. Mothukuri, and Shashidhar G. Koolagudi. 2020. [Kannada dialect classification using cnn](#). In *Mining Intelligence and Knowledge Exploration*, pages 254–259, Cham. Springer International Publishing.
- Tommi Jauhiainen, Heidi Jauhiainen, and Krister Lindén. 2021. [Naive Bayes-based experiments in Romanian dialect identification](#). In *Proceedings of the Eighth Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 76–83, Kiyv, Ukraine. Association for Computational Linguistics.
- Tommi Jauhiainen, Heidi Jauhiainen, and Krister Lindén. 2022. [Italian language and dialect identification and regional French variety detection using adaptive naive Bayes](#). In *Proceedings of the Ninth Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 119–129, Gyeongju, Republic of Korea. Association for Computational Linguistics.
- Tommi Jauhiainen, Heidi Jauhiainen, Niko Partanen, and Krister Lindén. 2020. [Uralic language identification \(ULI\) 2020 shared task dataset and the wanca 2017 corpora](#). In *Proceedings of the 7th Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 173–185, Barcelona, Spain (Online). International Committee on Computational Linguistics (ICCL).
- Abdullah Salem Khered, Ingy Yasser Hassan Abdou Abdelhalim, and Riza Batista-Navarro. 2022. [Building an ensemble of transformer models for Arabic dialect classification and sentiment analysis](#). In *Proceedings of the The Seventh Arabic Natural Language Processing Workshop (WANLP)*, pages 479–484, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Javier De la Rosa y Eduardo G. Ponferrada y Manu Romero y Paulo Villegas y Pablo González de Prado Salas y María Grandury. 2022. [Bertin: Efficient pre-training of a spanish language model using perplexity sampling](#). *Procesamiento del Lenguaje Natural*, 68(0):13–23.
- Yun Lei and John H. L. Hansen. 2011. [Dialect classification via text-independent training and testing for arabic, spanish, and chinese](#). *IEEE Transactions on Audio, Speech, and Language Processing*, 19(1):85–96.
- Wanqiu Lin, Maulik Madhavi, Rohan Kumar Das, and Haizhou Li. 2020. [Transformer-based arabic dialect identification](#).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *ArXiv*, abs/1907.11692.
- Marco Lui, Jey Han Lau, and Timothy Baldwin. 2014. [Automatic detection and language identification of multilingual documents](#). *Transactions of the Association for Computational Linguistics*, 2:27–40.
- Priyank Mathur, Arkajyoti Misra, and Emrah Budur. 2017. [LIDE: language identification from text documents](#). *CoRR*, abs/1701.03682.
- Martin A. Nowak and David C. Krakauer. 1999. [The evolution of language](#). *Proceedings of the National Academy of Sciences*, 96(14):8028–8033.
- Alec Radford and Karthik Narasimhan. 2018. [Improving language understanding by generative pre-training](#).
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Petru Rebeja and Dan Cristea. 2020. [A dual-encoding system for dialect classification](#). In *Proceedings of the 7th Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 212–219, Barcelona, Spain (Online). International Committee on Computational Linguistics (ICCL).
- David Romero, Luis Fernando D’Haro, and Christian Salamea. 2021. [Exploring Transformer-based Language Recognition using Phonotactic Information](#). In *Proc. IberSPEECH 2021*, pages 250–254.

- Fábio Souza, Rodrigo Nogueira, and Roberto Lotufo. 2020. BERTimbau: pretrained BERT models for Brazilian Portuguese. In *9th Brazilian Conference on Intelligent Systems, BRACIS, Rio Grande do Sul, Brazil, October 20-23 (to appear)*.
- S. Thara and Prabaharan Poornachandran. 2021. Transformer based language identification for malayalam-english code-mixed text. *IEEE Access*, 9:118837–118850.
- Christoph Tillmann, Saab Mansour, and Yaser Al-Onaizan. 2014. Improved sentence-level Arabic dialect classification. In *Proceedings of the First Workshop on Applying NLP Tools to Similar Languages, Varieties and Dialects*, pages 110–119, Dublin, Ireland. Association for Computational Linguistics and Dublin City University.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Marcos Zampieri, Kai North, Tommi Jauhiainen, Mariano Felice, Neha Kumari, Nishant Nair, and Yash Bangera. 2023. Language variety identification with true labels. *arXiv preprint arXiv:2303.01490*.