

Task-Specific Pre-Training and Cross Lingual Transfer for Code-Switched Data

Akshat Gupta, Sai Krishna Rallabandi, Alan Black

Carnegie Mellon University

akshatgu@andrew.cmu.edu, {srallaba, awb}@cs.cmu.edu

Abstract

Using task-specific pre-training and leveraging cross-lingual transfer are two of the most popular ways to handle code-switched data. In this paper, we aim to compare the effects of both for the task of sentiment analysis. We work with two Dravidian Code-Switched languages - Tamil-English and Malayalam-English and four different BERT based models. We compare the effects of task-specific pre-training and cross-lingual transfer and find that task-specific pre-training results in superior zero-shot and supervised performance when compared to performance achieved by leveraging cross-lingual transfer from multilingual BERT models.

1 Introduction

Code-Switching is a common phenomenon which occurs in many bilingual and multilingual communities around the world. It is characterized by the usage of more than one language in a single utterance (Sitaram et al., 2019; Jose et al., 2020; Priyadharshini et al., 2020). India is one such place with many communities using different Code-Switched languages, with Hinglish (Code-Switched Hindi-English) being the most popular one. Dravidian languages like Tamil, Kannada, Malayalam and Telugu are also usually code-mixed with English. For over 2600 years, recorded Tamil literature has been documented. The 500 BCE Tamil script found at the site of Porunthal is situated 12 km southwest of Palani, Tamil Nadu. Tamil was the first to be listed as a classical language of India, one of 22 scheduled languages in the Constitution of India, and is one of the world’s longest-surviving classical languages. These code-mixed languages are commonly used to interact with social media, which is why it is essential to build systems that are able to handle Code-Switched data (Chakravarthi et al., 2021a; Chakravarthi and

Muralidaran, 2021; Chakravarthi et al., 2021b; Suryawanshi and Chakravarthi, 2021).

Sentiment analysis poses the task of inferring opinion and emotions of a text query (usually social media comments or Tweets) as a classification problem, where each query is classified to have a positive, negative or neutral sentiment (Thavareesan and Mahesan, 2019, 2020a,b; Nanli et al., 2012). It has various applications like understanding the sentiment of different tweets, facebook and youtube comments, understanding product reviews etc (Puranik et al., 2021; Hegde et al., 2021; Yaraswini et al., 2021; Ghanghor et al., 2021b,a). In times of the pandemic, where the entire world is living online, it has become an even more important tool. Robust systems for sentiment analysis already exist for high resource languages like English (Barbieri et al., 2020), yet progress needs to be made for lower resource and code-switched languages.

One of the major bottlenecks in dealing with Code-Switched languages is the lack of availability of annotated datasets in the Code-Mixed languages. To alleviate this problem for Dravidian languages, datasets have been released for Sentiment Analysis in Tamil-English (Chakravarthi et al., 2020b) and Malayalam-English (Chakravarthi et al., 2020a). Various shared tasks (Patwa et al., 2020) (Chakravarthi et al., 2020c) have also been accompanied by the release of these datasets to advance research in this domain.

Models built on top of contextualized word embeddings have recently received huge amount of success and are used in most state of the art models. Systems built on top of BERT (Devlin et al., 2018) and its multilingual variants mBERT and XLM-RoBERTa (Conneau et al., 2019) have been the top ranking systems in both the above competitions.

In this paper, we train four different BERT based models for sentiment analysis for two different code-switched Dravidian languages. The main con-

tributions of this paper are:

- Comparing the effects of Task-specific pre-training and cross-lingual transfer for code-switched sentiment analysis. In our experiments, we find that the performance with task-specific pre-training on English BERT models is consistently superior when compared to multilingual BERT models.
- We present baseline results for the Malayalam-English (Chakravarthi et al., 2020a) and Tamil-English (Chakravarthi et al., 2020b) dataset for a three-class sentiment classification problem, classifying each sentence into positive, negative and neutral sentiments. Our results can be used as baselines for future work. Previous work (Chakravarthi et al., 2020c) on these datasets treated the problem as a five-class classification problem.

2 Related work

Various datasets for Code-Switched Sentiment Analysis have been released in the last few years, some of which have also been accompanied by shared tasks (Patra et al., 2018) (Patwa et al., 2020) (Chakravarthi et al., 2020c) in the respective languages. The shared task released with (Chakravarthi et al., 2020c) focused on Sentiment Analysis for Tamil-English and Malayalam-English datasets. The best performing systems for both these tasks were built on top of BERT variants. (Chakravarthi et al., 2020b) (Chakravarthi et al., 2020a) have provided baseline results for sentiment analysis with a dataset of Youtube comments in Tamil-English and Malayalam-English datasets respectively, using various classification algorithms including Support Vector Machines, Decision Trees, K-Nearest Neighbours, BERT based models etc. In our paper, we use BERT, mBERT (Devlin et al., 2018), XLM-RoBERTa (Conneau et al., 2019) and a RoBERTa based sentiment classification model (Barbieri et al., 2020) for sentiment classification.

Sentiment classification is usually done by classifying a query into one of three sentiments - positive, negative and neutral. In this work, we perform a three-class classification to be able to leverage the power of the TweetEval sentiment classifier (Barbieri et al., 2020), which was trained on a dataset of English Tweets. The TweetEval model is a monolingual model trained on an out-of-

Language	Positive	Negative	Neutral
Tam-Eng	10,559	2,037	850
Mal-Eng	2,811	738	1,903
Hinglish	6,616	5,892	7,492

Table 1: Dataset statistics for Tamil-English, Malayalam-English and Hinglish (Sentimix) dataset. These numbers for the depict the entire dataset which is then divided into train, development and test sets by respective authors.

domain dataset for our task (the Tamil-English and Malayalam-English datasets are made from scrapping Youtube comments). We also use mBERT and XLM-RoBERTa (Conneau et al., 2019) models for classification, which are trained on more than 100 languages and are thus able to leverage the power of cross-lingual transfer for sentiment classification. Previous works (Jayanthi and Gupta, 2021) has also shown that mBERT and XLM-RoBERTa based models achieve state of the art performance when dealing with code-switched Dravidian languages.

3 Dataset

In this paper, we primarily test our models on Tamil-English (Chakravarthi et al., 2020b) and Malayalam English (Chakravarthi et al., 2020a). The dataset was collected by scrapping Youtube comments from Tamil and Malayalam movies. All the sentences in the dataset are in the latin script. We also use the Sentimix Hinglish dataset (Patwa et al., 2020) to leverage cross-lingual transfer from Hinglish. The datasets statistics are summarized in Table 1. The numbers shown are for the entire dataset which was then split into train, development and test sets by the respective authors.

4 Models

We train sentiment analysis models based on top of four BERT variants:

- **BERT** (Devlin et al., 2018) : The original BERT model was trained using Masked Language Modelling (MLM) and Next Sentence Prediction (NSP) objectives on English Wikipedia and BookCorpus (Zhu et al., 2015) and has approximately 110M parameters. We use the uncased-base implementation from the Hugging Face library for our work.
- **mBERT** (Devlin et al., 2018): This is a multi-lingual BERT model trained on 104 languages

and has approximately 179M parameters. We again use the uncased-base model for our work.

- **XLM-RoBERTa** (Conneau et al., 2019): This is a multilingual RoBERTa model trained on 100 languages and has approximately 270M parameters. The RoBERTa models were an incremental improvement over BERT models with optimized hyperparameters. The most significant change was the removal of the NSP objective used while training BERT. The XLM-RoBERTa model is trained large multilingual corpus of 2.5TB of webcrawled data. We use the uncased-base XLM-RoBERTa model.
- **TweetEval, a RoBERTa based sentiment classifier**: The paper by (Barbieri et al., 2020) is a benchmark of Tweet classification models trained on top of RoBERTa. We use its sentiment classification model, which is referred to as the TweetEval model in this paper. The sentiment classification model was trained on a dataset of 60M English Tweets. The underlying RoBERTa model was trained on English data and has 160M parameters.

We use the Hugging Face library implementation of these models. We expect mBERT and XLM-RoBERTa based models to leverage cross-lingual transfer from a large pool of languages it is trained on. The TweetEval model was trained on a dataset of English Tweets for the task of sentiment analysis, but is still out-of domain for Youtube comments datasets. It is important to note that all our chosen models either have a different domain, language or task on which they were initially trained and hence are not directly suitable for the task of code-switched sentiment analysis of Youtube comments.

4.1 Metrics

We evaluate our results based on weighted average scores of precision, recall and F1. When calculating the weighted average, the precision, recall and F1 scores are calculated for each class and a weighted average is taken based on the number of samples in each class. This metric is apt as the datasets used are unbalanced. We use the same sklearn implementation¹ of the weighted av-

¹https://scikit-learn.org/stable/modules/generated/sklearn.metrics.classification_report.html

Model	Precision	Recall	F1
Baseline	0.76	0.80	0.78
BERT	0.74	0.78	0.75
mBERT	0.75	0.75	0.75
XLM-RoBERTa	0.76	0.76	0.76
TweetEval	0.77	0.77	0.77

Table 2: Monolingual Results for Malayalam-English. All scores are weighted average scores.

Model	Precision	Recall	F1
Baseline	0.66	0.79	0.66
BERT	0.74	0.74	0.74
mBERT	0.75	0.77	0.76
XLM-RoBERTa	0.75	0.78	0.76
TweetEval	0.76	0.79	0.76

Table 3: Monolingual Results for Tamil-English. All scores are weighted average scores.

erage metric as used in (Chakravarthi et al., 2020b) (Chakravarthi et al., 2020a). All the numbers shown in the paper are weighted average scores.

5 Experiments

In this paper, we aim to understand the effects of task-specific pre-training and cross-lingual transfer in improving performance of BERT based models on code-switched datasets. *Task-Specific Pretraining* refers pre-training a model on the same task for which a larger dataset is available and fine-tuning the so-obtained model on the target dataset, which is usually much smaller. In this case, we use the TweetEval model trained for the task of sentiment analysis on a large English corpus, which we fine-tune on the Dravidian code-mixed datasets. *Cross-Lingual Transfer* is commonly referred to the phenomenon of leveraging features and contextual information learnt from a different set of languages for a task with a new target language. Leveraging cross-lingual transfer is one of the main reasons behind training multilingual BERT models, where we expect multilingual BERT models to perform better on an unknown language when compared to an English BERT model.

5.1 Monolingual Results

We first present Monolingual Sentiment Classification Results for Tamil-English and Malayalam-English datasets, as shown in Table 2 and Table 3 respectively. The baseline model for Malayalam-English is based on mBERT while the baseline

Train Lang	Test Lang: Tamil			Test Lang: Malayalam			Test Lang: Hinglish		
	mBERT	xlm-r	TE	mBERT	xlm-r	TE	mBERT	xlm-r	TE
<i>English</i>	-	-	0.197	-	-	0.345	-	-	.506
<i>Hinglish</i>	0.389	0.393	0.403	0.423	0.413	0.473	-	-	-
<i>Tamil</i>	-	-	-	0.389	0.523	0.427	0.379	0.321	0.432
<i>Malayalam</i>	0.562	0.504	0.621	-	-	-	0.376	0.321	0.538

Table 4: Zero-shot prediction results for different models trained on the *Train Language*. Here we only report the weighted average F1 scores. We use '-' to represent cells that are Not Applicable to cross-lingual transfer. TE stands for the model built on top of the TweetEval sentiment classification model. xlm-r refers to the XLM-RoBERTa model.

model for Tamil-English was a Random Forest classifier as presented in (Chakravarthi et al., 2020c). The Baseline results were trained for a five-way classification problem, hence the weighted average scores have been re-weighted so that they correspond to a three-way classification problem. We also train an English BERT model to act as a baseline for understanding the effects of task-specific pre-training and cross lingual transfer.

We can see that the TweetEval Model improves on the Baseline results. The improvement is quite significant for Tamil-English. The unusually high improvement in the Tamil-English dataset can be due to the fact that the two models are trained for two different problems. The Baseline results are for a five-way classification problem where the F1 scores are re-weighted to only include three classes - positive, negative and neutral. Our models are trained specifically for a three-class classification problem. Our results can provide baselines for future work on sentiment analysis for the Dravidian datasets, which is usually studied as a three-class classification problem. The TweetEval model also performs the best out of all the models tested. We see that the multilingual BERT models perform better than the English BERT models, although the improvement is not very drastic. XLM-RoBERTa consistently performs better than mBERT.

The pre-trained mBERT and XLM-RoBERTa models are trained on multiple languages and in multiple scripts. We expect these models to leverage cross-lingual transfer and perform better than the English BERT model in the code-switched domain. Although we do see improvement in performance by multilingual BERT models over the English BERT models, the improvements are not very drastic. On the other hand, we see a consistently larger improvement due to task-specific pre-training when fine-tuning the TweetEval model

on the Dravidian language datasets. We hypothesize two possible reasons for cross-lingual transfer being less effective for the Dravidian datasets. The first is that even though the multilingual BERT models were trained on multiple languages, they were trained in languages in their original scripts. The datasets we are considering contain Malayalam and Tamil in Romanized form. Due to this, the multilingual BERT models do not have representations and contexts of Malayalam and Tamil tokens in Romanized form. Thus the multilingual BERT model has to learn the representation of these new tokens just as an English BERT model would. The second reason is that although the multilingual BERT models were trained on multiple languages, they were not trained on a code-switched dataset.

5.2 Cross-Lingual Transfer

In this section, we look at the zero-shot transfer between the different sets of languages for the above used models for sentiment analysis. The results are shown in Table 4. The first column of Table 4 refers to the language the given models were trained on. For Hindi-English (Hinglish), we used the Hinglish Sentimix dataset (Patwa et al., 2020).

We first look at the zero shot transfer of English language for Tamil-English, Malayalam-English and Hinglish datasets (referring to the entire first row of Table 4). This is same as looking at the zero-shot performance of the TweetEval model since it was trained on an English corpus. We see that an English sentiment analysis model has the best zero-shot performance for the Hinglish dataset. We hypothesize that this is because the TweetEval model is an in-domain model for the Hinglish dataset (Patwa et al., 2020) as the Hinglish dataset is also a Twitter corpus. The Malayalam and Tamil datasets are out of domain since they are made from Youtube comments.

We can also see that the best zero-shot transfer

Model	Precision	Recall	F1
XLM-RoBERTa	0.76	0.76	0.76
TweetEval	0.77	0.77	0.77
TweetEval + Hinglish	0.78	0.78	0.78

Table 5: Comparison between performance of XLM-RoBERTa, TweetEval and TweetEval model pretrained on Hinglish data for the Malayalam-English dataset.

Model	Precision	Recall	F1
XLM-RoBERTa	0.75	0.78	0.76
TweetEval	0.76	0.79	0.76
TweetEval + Hinglish	0.76	0.79	0.77

Table 6: Comparison between performance of XLM-RoBERTa, TweetEval and TweetEval model pretrained on Hinglish data for the Tamil-English dataset.

results are obtained when the TweetEval model is fine-tuned on a linguistically closer language than English for each of the three datasets. For example, the best zero-shot results on the Tamil dataset is achieved when we fine-tune the TweetEval model on Malayalam. Similarly, the best zero-shot results for Tamil are achieved when we fine-tune the TweetEval model on Hinglish. Finally, the best zero-shot results for the Hinglish dataset are achieved when we fine-tune the TweetEval model on Malayalam. These results show that task-specific pre-training is more effective for zero-shot performance and hint at the superiority of task-specific pre-training over cross-lingual transfer.

In the next experiment, we try to leverage cross-lingual transfer along with task-specific pre-training. To do this, we first fine-tune the TweetEval model on the Hinglish dataset. We expect this fine-tuned model to begin to learn code-switching and recognizing new tokens in Hinglish which are not a part of its vocabulary. To make sure that we don't let the model overfit on the Hinglish dataset, we only fine-tune the model on the Hinglish dataset for 1 epoch. Then we fine-tune this model on the target datasets - Tamil-English and Malayalam-English. The results for this experiment are shown in Table 5 and Table 6. Though the combined models show improvements for both datasets, the improvements are not statistically significant. Rigorous experiments to explore this idea will be part of our future work.

6 Conclusion

In this paper we present various experiments to compare the effects of task-specific pre-training and cross-lingual transfer on performance of sentiment classification models for code-switched data. To do so, we check the performance of four BERT models on two code-switched languages - Malayalam-English and Tamil English. We find that task-specific pre-training is superior to cross-lingual transfer for our chosen code-switched datasets.

The results presented in this paper for four different BERT models can be used as baselines for future work on sentiment analysis for the chosen datasets. We also present first results for the TweetEval (Barbieri et al., 2020) sentiment classification model for code-switched data.

References

- Francesco Barbieri, Jose Camacho-Collados, Leonardo Neves, and Luis Espinosa-Anke. 2020. Tweeteval: Unified benchmark and comparative evaluation for tweet classification. *arXiv preprint arXiv:2010.12421*.
- Bharathi Raja Chakravarthi, Navya Jose, Shardul Suryawanshi, Elizabeth Sherly, and John Philip McCrae. 2020a. A sentiment analysis dataset for code-mixed Malayalam-English. In *Proceedings of the 1st Joint Workshop on Spoken Language Technologies for Under-resourced languages (SLTU) and Collaboration and Computing for Under-Resourced Languages (CCURL)*, pages 177–184, Marseille, France. European Language Resources association.
- Bharathi Raja Chakravarthi and Vigneshwaran Muralidaran. 2021. Findings of the shared task on Hope Speech Detection for Equality, Diversity, and Inclusion. In *Proceedings of the First Workshop on Language Technology for Equality, Diversity and Inclusion*. Association for Computational Linguistics.
- Bharathi Raja Chakravarthi, Vigneshwaran Muralidaran, Ruba Priyadarshini, and John Philip McCrae. 2020b. Corpus creation for sentiment analysis in code-mixed Tamil-English text. In *Proceedings of the 1st Joint Workshop on Spoken Language Technologies for Under-resourced languages (SLTU) and Collaboration and Computing for Under-Resourced Languages (CCURL)*, pages 202–210, Marseille, France. European Language Resources association.
- Bharathi Raja Chakravarthi, Ruba Priyadarshini, Shubhanker Banerjee, Richard Saldhana, John Philip McCrae, Anand Kumar M, Parameswari Krishnamurthy, and Melvin Johnson. 2021a. Findings of the shared task on Machine Translation

- in Dravidian languages. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.
- Bharathi Raja Chakravarthi, Ruba Priyadharshini, Navya Jose, Anand Kumar M, Thomas Mandl, Prasanna Kumar Kumaresan, Rahul Ponnusamy, Hariharan V, Elizabeth Sherly, and John Philip McCrae. 2021b. Findings of the shared task on Offensive Language Identification in Tamil, Malayalam, and Kannada. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.
- Bharathi Raja Chakravarthi, Ruba Priyadharshini, Vigneshwaran Muralidaran, Shardul Suryawanshi, Navya Jose, Elizabeth Sherly, and John P. McCrae. 2020c. [Overview of the Track on Sentiment Analysis for Dravidian Languages in Code-Mixed Text](#). In *Forum for Information Retrieval Evaluation, FIRE 2020*, page 21–24, New York, NY, USA. Association for Computing Machinery.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Nikhil Kumar Ghanghor, Parameswari Krishnamurthy, Sajeetha Thavareesan, Ruba Priyadharshini, and Bharathi Raja Chakravarthi. 2021a. IITK@DravidianLangTech-EACL2021: Offensive Language Identification and Meme Classification in Tamil, Malayalam and Kannada. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*, Online. Association for Computational Linguistics.
- Nikhil Kumar Ghanghor, Rahul Ponnusamy, Prasanna Kumar Kumaresan, Ruba Priyadharshini, Sajeetha Thavareesan, and Bharathi Raja Chakravarthi. 2021b. IITK@LT-EDI-EACL2021: Hope Speech Detection for Equality, Diversity, and Inclusion in Tamil, Malayalam and English. In *Proceedings of the First Workshop on Language Technology for Equality, Diversity and Inclusion*, Online.
- Siddhanth U Hegde, Adeep Hande, Ruba Priyadharshini, Sajeetha Thavareesan, and Bharathi Raja Chakravarthi. 2021. UVCE-IITTT@DravidianLangTech-EACL2021: Tamil Troll Meme Classification: You need to Pay more Attention. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.
- Sai Muralidhar Jayanthi and Akshat Gupta. 2021. Sj.aj@dravidianlangtech-eacl2021: Task-adaptive pre-training of multilingual bert models for offensive language identification. *arXiv preprint arXiv:2102.01051*.
- Navya Jose, Bharathi Raja Chakravarthi, Shardul Suryawanshi, Elizabeth Sherly, and John P. McCrae. 2020. [A Survey of Current Datasets for Code-Switching Research](#). In *2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS)*, pages 136–141.
- Zhu Nanli, Zou Ping, Li Weiguang, and Cheng Meng. 2012. Sentiment analysis: A literature review. In *2012 International Symposium on Management of Technology (ISMOT)*, pages 572–576. IEEE.
- Braja Gopal Patra, Dipankar Das, and Amitava Das. 2018. Sentiment analysis of code-mixed indian languages: An overview of sail_code-mixed shared task@ icon-2017. *arXiv preprint arXiv:1803.06745*.
- Parth Patwa, Gustavo Aguilar, Sudipta Kar, Suraj Pandey, Srinivas PYKL, Björn Gambäck, Tanmoy Chakraborty, Tamar Solorio, and Amitava Das. 2020. Semeval-2020 task 9: Overview of sentiment analysis of code-mixed tweets. *arXiv e-prints*, pages arXiv–2008.
- Ruba Priyadharshini, Bharathi Raja Chakravarthi, Mani Vegupatti, and John P. McCrae. 2020. [Named Entity Recognition for Code-Mixed Indian Corpus using Meta Embedding](#). In *2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS)*, pages 68–72.
- Karthik Puranik, Adeep Hande, Ruba Priyadharshini, Sajeetha Thavareesan, and Bharathi Raja Chakravarthi. 2021. IITTT@LT-EDI-EACL2021-Hope Speech Detection: There is always hope in Transformers. In *Proceedings of the First Workshop on Language Technology for Equality, Diversity and Inclusion*. Association for Computational Linguistics.
- Sunayana Sitaram, Khyathi Raghavi Chandu, Sai Krishna Rallabandi, and Alan W Black. 2019. A survey of code-switched speech and language processing. *arXiv preprint arXiv:1904.00784*.
- Shardul Suryawanshi and Bharathi Raja Chakravarthi. 2021. Findings of the shared task on Troll Meme Classification in Tamil. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.
- Sajeetha Thavareesan and Sinnathamby Mahesan. 2019. [Sentiment Analysis in Tamil Texts: A Study on Machine Learning Techniques and Feature Representation](#). In *2019 14th Conference on Industrial and Information Systems (ICIIS)*, pages 320–325.

- Sajeetha Thavareesan and Sinnathamby Mahesan. 2020a. [Sentiment Lexicon Expansion using Word2vec and fastText for Sentiment Prediction in Tamil texts](#). In *2020 Moratuwa Engineering Research Conference (MERCCon)*, pages 272–276.
- Sajeetha Thavareesan and Sinnathamby Mahesan. 2020b. [Word embedding-based Part of Speech tagging in Tamil texts](#). In *2020 IEEE 15th International Conference on Industrial and Information Systems (ICIIS)*, pages 478–482.
- Konthala Ysaswini, Karthik Puranik, Adeep Hande, Ruba Priyadharshini, Sajeetha Thavareesan, and Bharathi Raja Chakravarthi. 2021. [IIITT@DravidianLangTech-EACL2021: Transfer Learning for Offensive Language Detection in Dravidian Languages](#). In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *Proceedings of the IEEE international conference on computer vision*, pages 19–27.