Findings of the Shared Task on Troll Meme Classification in Tamil

Shardul Suryawanshi and **Bharathi Raja Chakravarthi**
Insight SFI Research Centre for Data Analytics,
Data Science Institute,
National University of Ireland Galway
Galway, Ireland

Abstract

The internet has facilitated its user-base with a platform to communicate and express their views without any censorship. On the other hand, this freedom of expression or free speech can be abused by its user or a troll to demean an individual or a group. Demeaning people based on their gender, sexual orientation, religious believes or any other characteristics –trolling– could cause significant distress in the online community. Hence, the content posted by a troll needs to be identified and dealt with before causing any more damage. Amongst all the forms of troll content, memes are most prevalent due to their popularity and ability to propagate across cultures. A troll uses a meme to demean, attack or offend its targetted audience. In this shared task, we provide a resource (TamilMemes) that could be used to train a system capable of identifying a troll meme in the Tamil language. In our TamilMemes dataset, each meme has been categorized into either a "troll" or a "not_troll" class. Along with the meme images, we also provided the Latin transcripted text from memes. We received ten system submissions from the participants, which were evaluated using the weighted average F1-score. The system with the weighted average F1-score of 0.55 secured the first rank.

1   Introduction

We have seen a rise in the usage of memes on the internet. Memes could come in many forms and languages, but we emphasize on the image with text (IWT) memes (Du et al., 2020) in the Tamil language. IWT memes are inherently multimodal as their meaning is not understood when just the image or the text considered in isolation; hence, both image and text should be considered. These memes can propagate and mutate through cultures like the selfish gene (Dawkins, 2016). Propagation and mutation of memes ensure widespread on the internet; hence they could be weaponized by trolls (Mandl et al., 2020). Trolls are individuals or users on the internet who tend to attack or offend other individuals or groups in a demeaning manner (Tomaiuolo et al., 2020).

A meme used by troll –troll meme– could come in any languages, and since languages are representative of cultures, a meme in one language could represent a culture. In this shared task, we provide the TamilMemes dataset (Suryawanshi et al., 2020) that consists of 2,969 memes to the participants. Additionally, we also provided the text (caption) associated with each meme. The text has been transcripted in Latin by a native Tamil speaker, which was later evaluated and corrected by the expert.

Tamil (ISO 639-3: tam) is an official language in Tamil Nadu, Puducherry, Sri Lanka, Singapore and a recognised minority language in Malaysia and South Africa with a 75 million speaker base (Chakravarthi et al., 2018; Thavareesan and Mahesan, 2019, 2020a,b). Tamil inscription in pottery from Kodumanal and Porunthal [1] [2] was the oldest inscription in India dating to 580 BCE then Asoka inscription in Prakrit, Greek and Aramaic dating to 260 BCE (Mahadevan, 2002; Rajakumar and Bharathi, 2012). In Sanskrit, the oldest known inscriptions are from the 1st century CE, such as Dhana's Ayodhya Inscription and Ghosundi-Hathibada. The first book printed in India in an Indian language was Tampiran Vanakkam in Tamil on 1578 CE, a 16-page translation of the Portuguese "Doctrina Christam" (Balachandran, 2005). In the modern Tamil script, there are 12 vowels, 18 consonants and one special character, the āytam (Chakravarthi, 2020b). The vowels and consonants combine in order to form 216 compound characters, giving a total of 247 characters.

---

[1]Tamil inscriptions wiki
[2]Iron Age - Early Historic Transition in South Indian Appraisal

However, Tamil is often transcribed in Latin Script by a multilingual speaker on social media platforms (Chakravarthi, 2020a). This behaviour could be attributed to the ease or comfort in scripting Tamil in Latin script (Chakravarthi et al., 2019; Hande et al., 2020). Also, these multilingual speakers tend to switch to another language such as English. This phenomenon, also known as code-mixing or code-switching is commonly observed amongst multilingual speakers (Jose et al., 2020; Priyadharshini et al., 2020). A great deal of work and resources (Chakravarthi et al., 2020a,c) has been created for Dravidian languages but the multimodal aspect of code-mixed content in the form of meme remains unexplored.

In recent years, we have seen a rise in awareness of tackling multimodal offensive (Sharma et al., 2020) or hate (Kiela et al., 2020) content in memes in the English language, but other under-resourced languages remain unexplored. Our "Troll Meme Classification in Tamil" task aims to promote research in multimodal troll meme classification in under-resourced Tamil language. Moreover, this task provides a unique opportunity to study the effect of code-switching or code-mixing in the English transcript of Tamil (Chakravarthi et al., 2020b) in association with the image from the meme.

2 Task Description

The goal of the "Troll Meme Classification in Tamil" shared task was to classify if a given meme is a "troll" or "not-troll" based on the image and text associated with the meme in the Tamil language. The text from the meme is written in either the Tamil grammar and English lexicon or English grammar and Tamil lexicon. However, for consistency, we transcripted the text in Latin.

Troll meme is a meme, which consists of offensive text and non-offensive images, offensive images with non-offensive text, sarcastically offensive text with non-offensive images, or sarcastic images with offensive text to provoke, distract and has digressive or off-topic content with intend to demean or offend particular people, group or race, otherwise, a not-troll meme (Suryawanshi et al., 2020). Figure 1 shows examples of a troll and not-troll meme from the TamilMeme dataset. Example 1 is a troll meme targeted towards the potato chip brand called "Lays". In this example, an image is harmless with just a picture of the potato chips

Split	Troll	Not-troll	Total
Train	1,282	1,018	2,300
Test	395	272	667
Total	1,677	1,290	2,967

Table 1: A class distribution in the training and test set of the TamilMemes dataset

packet, but the translation of the text is *"If you buy one packet air, then 5 chips free"* which is offensive for the brand. The translation of Example 2 would be *"Sorry my friend (girl)"*. This example does not contain any provoking or offensive image or text and hence, it is a not-troll meme.

Previously, we developed this TamilMemes dataset and treated the task of identifying a troll meme as an image classification problem. Since the text associated with the meme acts as a context of the image, we enhanced our TamilMemes dataset by providing the text as a separate modality for the shared task. We expected our participant to approach the task in a multimodal way.

(a) An example on troll meme from TamilMemes dataset

(b) An example on not-troll meme from TamilMemes dataset

Figure 1: Examples for a troll and not-troll meme from the TamilMemes dataset

Rank	Team	Precision	Recall	F1-score
1	Codewithzichao (Li, 2021)	0.57	0.60	0.55
2	IIITK (Ghanghor et al., 2021)	0.56	0.59	0.54
3	NLP@CUET (Hossain et al., 2021)	0.55	0.58	0.52
4	SSNCSE_NLP (Silvia A and B, 2021)	0.58	0.60	0.50
5	Simon_work (Que et al., 2021)	0.53	0.58	0.49
6	TrollMeta (J and HS, 2021)	0.45	0.41	0.48
7	UVCE-IIITT (Hegde et al., 2021)	0.60	0.60	0.46
8	cean	0.53	0.57	0.43
9	HUB (Huang and Bai, 2021)	0.50	0.54	0.40
10	iiit_dwd (Mishra and Saumya, 2021)	0.52	0.59	0.30

Table 2: The rank list with detailed report on Precision (P), Recall (R), F1-score (F1) of submissions for the "Troll Meme Classification in Tamil" task

3 Evaluation

Table 1 shows the class distribution in the training and test set for the TamilMemes dataset. We provided a training set of 2,500 memes (with the Latin scripted text) to the participants. Later, we evaluated all the systems on the held-out or test set of 667 memes. We considered a weighted average version of the F1-score as a primary evaluation metric by taking class imbalance into the account. The weighted average F1-score[3] is calculated by averaging the support-weighted mean F1 score per-class.

4 Methodology

In system submissions, we saw a variety of methodologies used by the participants which were as simple as Logistic Regression (Wright, 1995) and as complex as BERT (Devlin et al., 2018). By keeping multimodality in mind, most of the participants implemented a method that could leverage both text and image. Mostly, the image part has been processed using Convolutional Neural Network (CNN) such as ResNet152 (He et al., 2016) or the custom residual networks while the text has been processed RNN (LSTM (Hochreiter and Schmidhuber, 1997)), transformer (BERT, ROBERTA (Liu et al., 2019)). However, in one exceptional case, one participant has used a transformer for processing both modalities.

- Li (2021) utilized a transformer-based approach that leverages the pre-trained BERT and ResNet152 model to derive the text and image features. The novelty of the work

comes from the implementation of multi-modal attention which considers the whole text caption in the context of the image by mapping both image and text features in the same semantic space. Their model achieved a 0.55 weighted average F1 score and ranked first in the shared task.

- Ghanghor et al. (2021) utilized a transformer-based approach to identify the troll and not-troll Tamil meme. The state of the art text classifier (BERT) and image (CNN) classifiers were used to extract useful attention features. Their system submission for Tamil troll meme classification achieved a 0.54 weighted average F1 score.

- Hossain et al. (2021) utilized a systematic text-based, image-based and multimodal approach to identify a troll meme. The paper laid out a detailed investigation into the problem presented in the shared task by capturing visual and textual features using CNN, VGG16 (Simonyan and Zisserman, 2014), Inception (Szegedy et al., 2017), m-BERT, XLM-RoBERTa, XLNet (Yang et al., 2019). Multimodal features were extracted by combining image (CNN, ResNet50, Inception) and text (BiLSTM) (Zhou et al., 2016) features using an early fusion technique. But the results showed that the text-based approach with XLNet achieved the highest weighted F1-score of 0.58, and enabled their system implementation to secure 3rd rank.

- Silvia A and B (2021) used traditional NLP approaches such as Multilayer perceptron (MLP (Haykin and Network, 2004)), Random forest

---

[3]Weighted average F-1 score is calculated with the help of the sklearn classification report utility

(RF) (Breiman, 2001) and K-nearest neighbour (KNN) (Cunningham and Delany, 2020) on the text features derived from tf-idf, count vector and mBERT embeddings. The paper opts for a text-based approach over the multimodal approach and achieved a weighted average F1-score of 0.50.

- Que et al. (2021) used a text-based approach to deal with the multimodal issue of identifying troll memes. Here the author trains XLM-Roberta in combination with the custom CNN on the text modality and achieves a weighted average F1-score of 0.49.

- J and HS (2021) utilized ResNet-50 to identify if a given meme is a troll or not based on the visual features. However, no text modality has been used as a feature while training the system. Their system achieved a weighted average F1-score of 0.48 and ranked sixth.

- Hegde et al. (2021) utilized a transformer-based approach to identify the troll and not-troll Tamil meme. The state of the art text (mBERT) and image (Vision transformer) (Dosovitskiy et al., 2020) classifiers were used to extract useful attention features. They achieve a weighted average F1-score of 0.46 and ranked seventh.

- Huang and Bai (2021) used a BiGRU (Dey and Salem, 2017) and custom CNN to capture text and image features, which later are concatenated to form a multimodal representation of a meme. They ranked ninth with a weighted average F1-score of 0.40.

- Mishra and Saumya (2021) used a hybrid approach that combines text and image features using CNN and BiLSTM. Their proposed model obtained 10th rank in the shared task and reported a weighted F1-score of 0.30.

5 Results and Discussion

We have received 10 submissions for the shared task. Table 2 shows that the maximum weighted average F1-score could only reach 0.55, which is less. While the submissions claimed that the evaluation metric exceeded the validation set, the poor performance on the test set point towards overfitting. Overfitting is commonly seen when complex models with large numbers of parameters are trained on small datasets.

In this shared task, we provided 2,967 samples which might be less for the data-hungry models such as BERT and it's predecessor such as ROBERTa, XLM-RoBERTa. Moreover, due to the high dimensionality presented by the images, the traditional machine learning models such as logistic regression, Naive Bayes are not able to generalize. The poor performance of the traditional machine learning model shows that they are rather too simple to learn useful features from the high dimensional data presented in the form of image and text. Furthermore, the Tamil text (could be with or without code-mixed) from the meme is transcribed in the Latin script leaves little room for using pre-trained models (e.g. multilingual BERT) which are primarily trained in English or zero-shot transferred to other languages written in their native script. Hence, such multilingual models need to be finetuned on the more code-mixed, transliterated data. In the case of images, the visual features have been captured using a CNN pre-trained on the imagenet weights.

In the previous study, we observed that the visual features captured by such CNNs tend to be less useful for identifying a troll Tamil meme. The same effect has been observed in the system submitted to the shared task. Overall, we need a robust approach that takes into account all the aspects of this multimodal classification problem –lack of data, high dimensionality, low resource language processing and code-mixing.

6 Conclusion

In the "Troll Meme Classification in Tamil" shared task, we presented a unique multimodal classification problem by providing a newly improved TamilMeme dataset which now has Tamil text (in the form of transcribed Latin text) from memes. This task was not only a multimodal classification problem but also posed challenges such as natural language processing of low-resourced language, transcribed code-mixed text. The submissions received from the participants showed multiple ways to approach the problem and we hope that their contribution along with our enriched dataset will kindle the research in this less explored area.

References

R Balachandran. 2005. Pioneers of Tamil Literature: Transition to Modernity. *Indian Literature*, 49(2 (226):179–184.

Leo Breiman. 2001. Random forests. *Machine learning*, 45(1):5–32.

Bharathi Raja Chakravarthi. 2020a. HopeEDI: A multilingual hope speech detection dataset for equality, diversity, and inclusion. In *Proceedings of the Third Workshop on Computational Modeling of People's Opinions, Personality, and Emotion's in Social Media*, pages 41–53, Barcelona, Spain (Online). Association for Computational Linguistics.

Bharathi Raja Chakravarthi. 2020b. *Leveraging orthographic information to improve machine translation of under-resourced languages*. Ph.D. thesis, NUI Galway.

Bharathi Raja Chakravarthi, Mihael Arcan, and John P. McCrae. 2018. Improving wordnets for under-resourced languages using machine translation. In *Proceedings of the 9th Global Wordnet Conference*, pages 77–86, Nanyang Technological University (NTU), Singapore. Global Wordnet Association.

Bharathi Raja Chakravarthi, Mihael Arcan, and John P. McCrae. 2019. WordNet gloss translation for under-resourced languages using multilingual neural machine translation. In *Proceedings of the Second Workshop on Multilingualism at the Intersection of Knowledge Bases and Machine Translation*, pages 1–7, Dublin, Ireland. European Association for Machine Translation.

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.

Bharathi Raja Chakravarthi, M Anand Kumar, John Philip McCrae, Premjith B, Soman KP, and Thomas Mandl. 2020b. Overview of the track on HASOC-Offensive Language Identification-DravidianCodeMix. In *Working Notes of the Forum for Information Retrieval Evaluation (FIRE 2020). CEUR Workshop Proceedings. In: CEUR-WS. org, Hyderabad, India*.

Bharathi Raja Chakravarthi, Vigneshwaran Muralidaran, Ruba Priyadharshini, and John Philip McCrae. 2020c. 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.

Padraig Cunningham and Sarah Jane Delany. 2020. k-nearest neighbour classifiers–. *arXiv preprint arXiv:2004.04523*.

Richard Dawkins. 2016. *The selfish gene*. Oxford university press.

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*.

R. Dey and F. M. Salem. 2017. Gate-variants of gated recurrent unit (gru) neural networks. In *2017 IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS)*, pages 1597–1600.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.

Yuhao Du, Muhammad Aamir Masood, and Kenneth Joseph. 2020. Understanding visual memes: An empirical analysis of text superimposed on memes shared on twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, pages 153–164.

Nikhil Ghanghor, Bharathi Raja Chakravarthi, Ruba Priyadharshini, Sajeetha Thavareesan, and Parameswari Krishnamurthy. 2021. IIITK@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*. Association for Computational Linguistics.

Adeep Hande, Ruba Priyadharshini, and Bharathi Raja Chakravarthi. 2020. KanCMD: Kannada CodeMixed dataset for sentiment analysis and offensive language detection. In *Proceedings of the Third Workshop on Computational Modeling of People's Opinions, Personality, and Emotion's in Social Media*, pages 54–63, Barcelona, Spain (Online). Association for Computational Linguistics.

Simon Haykin and N Network. 2004. A comprehensive foundation. *Neural networks*, 2(2004):41.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.

Siddhanth U Hegde, Adeep Hande, Ruba Priyadharshini, Sajeetha Thavareesan, and Bharathi Raja Chakravarthi. 2021. Uvce-iiitt@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.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.

Eftekhar Hossain, Omar Sharif, and Mohammed Moshiul Hoque. 2021. NLP-CUET@DravidianLangtech-EACL2021: Investigating Visual and Textual Features to Identify Trolls from Multimodal Social Media Memes. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Bo Huang and Yang Bai. 2021. HUB@DravidianLangTech-EACL2021: Meme Classification for Tamil Text-Image Fusion. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Manoj Balaji J and Chinmaya HS. 2021. Trollmeta@DravidianLangtech-EACL2021: Meme classification using deep learning. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

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.

Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020. The hateful memes challenge: Detecting hate speech in multimodal memes. *arXiv preprint arXiv:2005.04790*.

Zichao Li. 2021. Codewithzichao@DravidianLangtech-EACL2021: Exploring Multimodal Transformers for Meme Classification in Tamil Language. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

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 preprint arXiv:1907.11692*.

Iravatham Mahadevan. 2002. Aryan or dravidian or neither? a study of recent attempts to decipher the indus script (1995-2000). *Electronic Journal of Vedic Studies*, 8(1):1–19.

Thomas Mandl, Sandip Modha, Anand Kumar M, and Bharathi Raja Chakravarthi. 2020. Overview of the HASOC Track at FIRE 2020: Hate Speech and Offensive Language Identification in Tamil, Malayalam, Hindi, English and German. In *Forum for Information Retrieval Evaluation*, FIRE 2020, page 29–32, New York, NY, USA. Association for Computing Machinery.

Ankit Kumar Mishra and Sunil Saumya. 2021. IIIT_DWD@DravidianLangtech-EACL2021: Identifying Tamil troll meme using a hybrid deep learning approach. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

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.

Qinyu Que, Yuanchi Qu, and Suidong Qu. 2021. Simon@DravidianLangtech-EACL2021: Meme Classification for Tamil with BERT. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

S Rajakumar and V Subbiah Bharathi. 2012. Eighth century tamil consonants recognition from stone inscriptions. In *2012 International Conference on Recent Trends in Information Technology*, pages 40–43. IEEE.

Chhavi Sharma, William Paka, Scott, Deepesh Bhageria, Amitava Das, Soujanya Poria, Tanmoy Chakraborty, and Björn Gambäck. 2020. Task Report: Memotion Analysis 1.0 @SemEval 2020: The Visuo-Lingual Metaphor! In *Proceedings of the 14th International Workshop on Semantic Evaluation (SemEval-2020)*, Barcelona, Spain. Association for Computational Linguistics.

Agnusimmaculate Silvia A and Bharathi B. 2021. SSNCSE NLP@DravidianLangtech-EACL2021: Meme classification for Tamil using machine learning approach. In *Proceedings of the First Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

Shardul Suryawanshi, Bharathi Raja Chakravarthi, Pranav Verma, Mihael Arcan, John Philip McCrae, and Paul Buitelaar. 2020. A dataset for troll classification of TamilMemes. In *Proceedings of the WILDRE5– 5th Workshop on Indian Language Data: Resources and Evaluation*, pages 7–13, Marseille, France. European Language Resources Association (ELRA).

Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander Alemi. 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.

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 (MERCon)*, 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.

Michele Tomaiuolo, Gianfranco Lombardo, Monica Mordonini, Stefano Cagnoni, and Agostino Poggi. 2020. A survey on troll detection. *Future Internet*, 12(2):31.

Raymond E Wright. 1995. Logistic regression.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*.

Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. Attention-based bidirectional long short-term memory networks for relation classification. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 2: Short papers)*, pages 207–212.