Overview of Abusive Comment Detection in Tamil - ACL 2022

Ruba priyadharshini[1], **Bharathi Raja Chakravarthi**[2], **Subalalitha Chinnaudayar Navaneethakrishnan** [3], **Thenmozhi Durairaj** [4], **Malliga Subramanian**[5] **Kogilavani Shanmugavadivel** [5], **Siddhanth U Hegde** [6], **Prasanna Kumar Kumaresan** [7]
[1]Madurai Kamaraj University,Tamil Nadu, India, [2]National University of Ireland Galway,
[3]SRM Institute Of Science And Technology, Tamil Nadu, India,
[4]SSN College of Engineering, Tamil Nadu, India,
[5]Kongu Engineering College, Tamil Nadu, India,
[7]University Visvesvaraya College of Engineering, Karnataka, India,
[6]Indian Institute of Information Technology and Management, Kerala, India.
bharathi.raja@insight-centre.org

Abstract

The social media is one of the significant digital platforms that create a huge impact in peoples of all levels. The comments posted on social media is powerful enough to even change the political and business scenarios in very few hours. They also tend to attack a particular individual or a group of individuals. This shared task aims at detecting the abusive comments involving, Homophobia, Misandry, Counter-speech, Misogyny, Xenophobia, Transphobic. The hope speech is also identified. A dataset collected from social media tagged with the above said categories in Tamil and Tamil-English code-mixed languages are given to the participants. The participants used different machine learning and deep learning algorithms. This paper presents the overview of this task comprising the dataset details and results of the participants.

1 Introduction

Their distribution of digital information has increased to a greater extent. The importance of the Online Social Networks (OSNs) has grown significantly in recent years, and they have become a go-to source for acquiring news, information, and entertainment (Halevy et al., 2022; Priyadharshini et al., 2021; Kumaresan et al., 2021). However, despite many positive impacts of employing OSNs, a growing body of evidence indicates that there is an ever-increasing number of malevolent actors who are exploiting these networks to spread poison and cause harm to other individuals (Chakravarthi, 2020; Chakravarthi and Muralidaran, 2021). The term "Hate Speech" (HS) refers to any form of communication that is abusive, insulting, intimidating, and/or incites violence or discrimination and that disparages an individual or a vulnerable group on the basis of characteristics such as ethnicity, gender, sexual orientation, or religious affiliation (Whillock and Slayden, 1995; Sampath et al., 2022; Ravikiran et al., 2022; Chakravarthi et al., 2022; Bharathi et al., 2022; Priyadharshini et al., 2022). Because of this diversity in thematic foci, we refer to them as themes. Examples of topics include misogyny, sexism, racism, transphobia, homophobia, and xenophobia (Chakravarthi et al., 2020, 2021; Ghanghor et al., 2021a,b; Yasaswini et al., 2021). The abusive comments targeting people have a huge impact on them psychologically(Wiegand et al., 2021). This task lays a foundation on how these comments can be detected for Dravidian language Tamil. Tamil is a Dravidian classical language used by the Tamil people of South Asia. Tamil is an official language of Tamil Nadu, Sri Lanka, Singapore, and the Union Territory of Puducherry in India. Significant minority speak Tamil in the four other South Indian states of Kerala, Karnataka, Andhra Pradesh, and Telangana, as well as the Union Territory of the Andaman and Nicobar Islands (Subalalitha, 2019; Srinivasan and Subalalitha, 2019; Narasimhan et al., 2018). It is also spoken by the Tamil diaspora, which may be found in Malaysia, Myanmar, South Africa, the United Kingdom, the United States, Canada, Australia, and Mauritius. Tamil is also the native language of Sri Lankan Moors (Sakuntharaj and Mahesan, 2021, 2017, 2016; Thavareesan and Mahesan, 2019, 2020a,b, 2021). Tamil, one of the 22 scheduled languages in the Indian Constitution, was the first to be designated as a classical language of India. Tamil is one of the world's longest-surviving classical languages. The earliest epigraphic documents discovered on rock edicts and "hero stones" date from the 6th century BC. Tamil has the oldest ancient non-Sanskritic Indian literature of any Indian language (Anita and Subalalitha, 2019b,a; Subalalitha and Poovammal, 2018). Since the comments posted online contain mixture of languages that are familiar with the users that are posting the comments, the task also considers detecting

Proceedings of the Second Workshop on Speech and Language Technologies for Dravidian Languages, pages 292 - 298

the comments from the Tamil-English code mixed language.

The goal of this task is to identify whether a given comment contains abusive comment. A comment/post within the corpus may contain more than one sentence but the average sentence length of the corpora is 1. The annotations in the corpus are made at a comment/post level. The participants were provided development, training and test dataset in Tamil and Tamil-English languages. The dataset is tagged using various classes namely, Homophobia, Misandry, Counter-speech, Misogyny, Xenophobia, Transphobic and Hope Speech. To the best of our knowledge, this is the first shared task on abusive detection in Tamil at this fine-grained level. 11 teams participated for detecting abusive comments in Tamil language and Tamil-English language tasks.

2 Task Description

The task is primarily a comment/post-level classification task. Given a YouTube comment, the systems submitted by the participants should classify it abusive categories. The participants were provided with development, training and test dataset in Tamil and Tamil-English. The dataset is tagged using various classes namely, Homophobia, Misandry, Counter-speech, Misogyny, Xenophobia, Transphobic and hope speech. 10 teams participated for detecting abusive comment in Tamil language and 11 teams participated for the Tamil-English language.

3 Data Description

The Tamil language training data contains 2240 comments, the validation set contains 560 comments, and the test data set includes 699 comments. The Tamil-English language test data set contains 5943 comments, the validation set contains 1486 comments and the 1857 test comments. The distribution of the seven categories in the whole dataset is shown in Table 1.

4 Participant's methodology

4.1 Pre-processing strategies

The participants have predominantly used "transliteration" as one of the pre-processing strategies. The Tamil-English code-mixed texts necessitate this approach. Apart from transliteration, removal of punctuation, stop words have also been used. Class balancing of the data has also been attempted as the distribution of the class labels in the given training dataset.

4.2 Participant's Systems

Term Frequency- Inverse Document Frequency (TF- IDF) and BERT embeddings have been used to extract and represent the features in the feature extraction phase. The participants have used a wide variety of machine learning algorithms, deep learning models, and transformers. Logistic Regression, Linear Support Vector Machines, Gradient Boost classifier, and K neighbor classifier have been used as machine learning algorithms. Ensemble models attempted composed of a mixture of these machine learning models. Multi-layered perceptron, Recurrent Neural Networks (RNN), Vanilla LSTM (Schuster and Paliwal, 1997) were opted as deep learning models. On the transformers front, mBERT(Devlin et al., 2018), MuRIL BERT (Khanuja et al., 2021), XLM RoBERTa (Liu et al., 2019), and ULMFit (Howard and Ruder, 2018) models have been opted. The MuRIL BERT models have shown the best performance compared to the other models. This is primarily because it is trained exclusively for Indian languages. The ranking of the teams for both of the language tasks is shown in Tables 2 and 3. The ranking is given based on their f1 score and how intense their system is, which counts their pre-processing techniques and the number of models used to prove their performance.

5 Error Analysis of the Systems

The participants have used the standard metrics such as Weighted Precision, Weighted Recall, and Weighted F-score to evaluate the performance of their systems. The equations of these metrics are given below.

$$Precision = \frac{TP}{TP + FP} \quad (1)$$

$$Recall = \frac{TP}{TP + FP} \quad (2)$$

where, TP= Number of True Positives and FP= Number of false Positives

$$F - Score = 2 * \frac{Precision * Recall}{Precision + Recall} \quad (3)$$

Comment category	Count in the datasets
None of the above	5011
Misandry	1276
Counter-speech	497
Xenophobia	392
Misogyny	336
Hope Speech	299
Homophobia	207
Transphobic	163

Table 1: Distribution of Comment Categories in the dataset

TeamName	Precision	Recall	F1-Score	Rank
CEN-Tamil(S N et al., 2022)	0.380	0.290	0.320	1
COMBATANT	0.290	0.330	0.300	2
DE-ABUSE(Palanikrmar et al., 2022)	0.330	0.29	0.290	3
DLRG(Diraphe et al., 2022)	0.340	0.260	0.270	4
TROPER	0.400	0.230	0.250	5
abusive-checker	0.140	0.140	0.140	6
Optimize_Prime(Patankar et al., 2022)	0.130	0.130	0.130	7
GJG	0.130	0.140	0.130	8
umuteam	0.130	0.130	0.130	9
MUCIC	0.120	0.130	0.120	10
BpHigh(Pahwa, 2022)	0.180	0.120	0.060	11
SSNCSE_NLP(Varsha and Bharathi, 2022)	0.130	0.140	0.090	12

Table 2: Rank list based on weighted average F1-score along with other evaluation metrics (Precision and Recall) for Tamil Language

TeamName	Precision	Recall	F1-Score	Rank
abusive-checker	0.460	0.380	0.410	1
GJG	0.370	0.340	0.350	2
umuteam	0.350	0.370	0.350	3
pandas(G L et al., 2022)	0.330	0.370	0.340	4
Optimize_Prime(Patankar et al., 2022)	0.310	0.380	0.320	5
MUCIC	0.400	0.280	0.290	6
CEN-Tamil(S N et al., 2022)	0.300	0.230	0.250	7
SSNCSE_NLP(Varsha and Bharathi, 2022)	0.260	0.240	0.250	8
IIITDWD	0.380	0.170	0.180	9
DLRG(Diraphe et al., 2022)	0.180	0.150	0.140	10
BpHigh(Pahwa, 2022)	0.140	0.160	0.100	11

Table 3: Rank list based on weighted average F1-score along with other evaluation metrics (Precision and Recall) for Tamil-English Language

$$P_{weighted} = \sum_{i=1}^{L}(Precisionofi \times Weightofi) \quad (4)$$

, where i is the test sample size.

$$R_{weighted} = \sum_{i=1}^{L}(Recallofi \times Weighti) \quad (5)$$

$$F{-}Score_{weighted} = \sum_{i=1}^{L}(F{-}Scoreofi{\times}Weighti) \quad (6)$$

The participants have also used accuracy, Macro-Precision, Macro-Recall, and Macro-F-scores to evaluate the system. It can be observed that the highest F-score achieved by the systems is 0.41. This is primarily due to the inability of the techniques to handle the errors observed consistently in all the systems during the classification. The various scenarios of errors are explained below.

Scenario 1: The systems fail to classify the sentences whenever the sentences do not contain even a single Tamil word. In other words, the sentences contain only the English transliterated words. For example, the comment, "World health enda ilukara ara kora nayae, " is classified as "Xenophobia" by all the systems while the actual label is "None of the above. The comment is actually against a xenophobic person. On the other comment, "sornam lakshmi mudiyathu mooditu" is classified as "Misandry" by all the systems while the actual class is "Misogyny." The name "sornam lakshhmi " refers to a woman but none of the systems labeled this right.

Scenario 2: The comments contain spelling mistakes and could not be handled during the pre-processing step. For example, This is classified

சாதி வெறி காட்டுமிராண்டி நாய்ங்க எவன்

as "None of the above " by all the systems while it is supposed to be "Misandry." This is due to the spelling mistake in the comment. The word

"எவன்" should have been "இவன்"

Scenario 3: The pre-processing strategies have had a harmful effect on the text and have resulted

"திருட வந்த பயல்களின் மதம் பின்பற்றும் நீ இந்தியனல்ல" is changed to "பயல மதம யனல"

in spelling mistakes. For example, the text, This has lead to the misclassification.

Scenario 4: Certain comments were too short and had references that were not captured by the systems. For example, the comment give below is

"ரெட் shirt நாகரீகம் தெரியாதவன்"

supposed to be classified as "Misandry." It is instead classified as "None of the above." Apart from these scenarios, the systems could never classify incomplete comments and double entendre comments correctly. Specific comments had hyperlinks that had the main content, which was missed by the systems.

6 Conclusion

This shared task aims at detecting the categories of abusive comments that are posted on social media. This kind of analysis would quantify the negativity that is spread in the society, which in turn should be turned into positivity either by enacting laws to enforce restrictions on posting comments on social media. This has been the motivation behind hosting this shared task which has attempted to aggregate the comments from social media in two languages, namely, Tamil and in code mixed language containing Tamil and English scripts. These comments were trained by various machine learning, deep learning, and transfer learning models. 11 teams participated in Tamil and Tamil-English languages tasks. 7 categories of abusive categories were tagged in the collected comments. The ranking of the teams was done based on the performance shown by the systems that were used by the participants and the in-depth analysis done by them. It was observed that the transformer models showed better performance when compared to that of the rest of the systems.

References

R Anita and CN Subalalitha. 2019a. An approach to cluster Tamil literatures using discourse connectives. In *2019 IEEE 1st International Conference on Energy, Systems and Information Processing (ICESIP)*, pages 1–4. IEEE.

R Anita and CN Subalalitha. 2019b. Building discourse parser for Thirukkural. In *Proceedings of the 16th*

International Conference on Natural Language Processing, pages 18–25.

B Bharathi, Bharathi Raja Chakravarthi, Subalalitha Chinnaudayar Navaneethakrishnan, N Sripriya, Arunaggiri Pandian, and Swetha Valli. 2022. Findings of the shared task on Speech Recognition for Vulnerable Individuals in Tamil. In *Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion*. Association for Computational Linguistics.

Bharathi Raja Chakravarthi. 2020. 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 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*, pages 61–72, Kyiv. Association for Computational Linguistics.

Bharathi Raja Chakravarthi, Ruba Priyadharshini, Thenmozhi Durairaj, John Phillip McCrae, Paul Buitaleer, Prasanna Kumar Kumaresan, and Rahul Ponnusamy. 2022. Findings of the shared task on Homophobia Transphobia Detection in Social Media Comments. In *Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion*. Association for Computational Linguistics.

Bharathi Raja Chakravarthi, Ruba Priyadharshini, Vigneshwaran Muralidaran, Shardul Suryawanshi, Navya Jose, Elizabeth Sherly, and John P McCrae. 2020. Overview of the track on sentiment analysis for Dravidian languages in code-mixed text. In *Forum for Information Retrieval Evaluation*, pages 21–24.

Bharathi Raja Chakravarthi, Ruba Priyadharshini, Rahul Ponnusamy, Prasanna Kumar Kumaresan, Kayalvizhi Sampath, Durairaj Thenmozhi, Sathiyaraj Thangasamy, Rajendran Nallathambi, and John Phillip McCrae. 2021. Dataset for identification of homophobia and transophobia in multilingual YouTube comments. *arXiv preprint arXiv:2109.00227*.

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*.

Ankita Diraphe, Ratnavel Rajalakshmi, and Antonette Shibani. 2022. Dlrg@dravidianlangtech-acl2022: Abusive comment detection intamil using multilingual transformer models. In *Proceedings of the Second Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Gayathri G L, Krithika S, Divyasri K, Thenmozhi Durairaj, and B. Bharathi. 2022. Pandas@tamilnlp-acl2022: Abusive comment detection in tamil code-mixed data using custom embeddings with labse. In *Proceedings of the Second Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Nikhil Ghanghor, Parameswari Krishnamurthy, Sajeetha Thavareesan, Ruba Priyadharshini, and Bharathi Raja Chakravarthi. 2021a. 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*, pages 222–229, Kyiv. Association for Computational Linguistics.

Nikhil Ghanghor, Rahul Ponnusamy, Prasanna Kumar Kumaresan, Ruba Priyadharshini, Sajeetha Thavareesan, and Bharathi Raja Chakravarthi. 2021b. IIITK@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*, pages 197–203, Kyiv. Association for Computational Linguistics.

Alon Halevy, Cristian Canton-Ferrer, Hao Ma, Umut Ozertem, Patrick Pantel, Marzieh Saeidi, Fabrizio Silvestri, and Ves Stoyanov. 2022. Preserving integrity in online social networks. *Communications of the ACM*, 65(2):92–98.

Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*.

Simran Khanuja, Diksha Bansal, Sarvesh Mehtani, Savya Khosla, Atreyee Dey, Balaji Gopalan, Dilip Kumar Margam, Pooja Aggarwal, Rajiv Teja Nagipogu, Shachi Dave, et al. 2021. Muril: Multilingual representations for indian languages. *arXiv preprint arXiv:2103.10730*.

Prasanna Kumar Kumaresan, Ratnasingam Sakuntharaj, Sajeetha Thavareesan, Subalalitha Navaneethakrishnan, Anand Kumar Madasamy, Bharathi Raja Chakravarthi, and John P McCrae. 2021. Findings of shared task on offensive language identification in tamil and malayalam. In *Forum for Information Retrieval Evaluation*, pages 16–18.

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*.

Anitha Narasimhan, Aarthy Anandan, Madhan Karky, and CN Subalalitha. 2018. Porul: Option generation and selection and scoring algorithms for a tamil flash card game. *International Journal of Cognitive and Language Sciences*, 12(2):225–228.

Bhavish Pahwa. 2022. Bphigh@tamilnlp-acl2022: Augmentation strategies for indic transformer-based abusive comment detection in tamil. In *Proceedings of the Second Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Vasanth Palanikrmar, Sean Benhur, Adeep Hande, and Bharathi Raja Chakravarthi. 2022. Deabuse@tamilnlp-acl 2022: Transliteration as data augmentation for abuse detection in tamil. In *Proceedings of the Second Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Shantanu Patankar, Omkar Gokhale, Onkar Litake, Aditya Mandke, and Dipali Kandam. 2022. Optimize_prime@tamilnlp-acl2022: Abusive comment detection in tamil. In *Proceedings of the Second Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Ruba Priyadharshini, Bharathi Raja Chakravarthi, Subalalitha Chinnaudayar Navaneethakrishnan, Thenmozhi Durairaj, Malliga Subramanian, Kogilavani Shanmugavadivel, Siddhanth U Hegde, and Prasanna Kumar Kumaresan. 2022. Findings of the shared task on Abusive Comment Detection in Tamil. In *Proceedings of the Second Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Ruba Priyadharshini, Bharathi Raja Chakravarthi, Sajeetha Thavareesan, Dhivya Chinnappa, Durairaj Thenmozhi, and Rahul Ponnusamy. 2021. Overview of the DravidianCodeMix 2021 shared task on sentiment detection in Tamil, Malayalam, and Kannada. In *Forum for Information Retrieval Evaluation*, pages 4–6.

Manikandan Ravikiran, Bharathi Raja Chakravarthi, Anand Kumar Madasamy, Sangeetha Sivanesan, Ratnavel Rajalakshmi, Sajeetha Thavareesan, Rahul Ponnusamy, and Shankar Mahadevan. 2022. Findings of the shared task on Offensive Span Identification in code-mixed Tamil-English comments. In *Proceedings of the Second Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Prasanth S N, R Aswin Raj, Adhithan P, Premjith B, and Soman K P. 2022. Cen-tamil@dravidianlangtech-acl2022: Abusive comment detection in tamil using tf-idf and random kitchen sink algorithm. In *Proceedings of the Second Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Ratnasingam Sakuntharaj and Sinnathamby Mahesan. 2016. A novel hybrid approach to detect and correct spelling in Tamil text. In *2016 IEEE International Conference on Information and Automation for Sustainability (ICIAfS)*, pages 1–6.

Ratnasingam Sakuntharaj and Sinnathamby Mahesan. 2017. Use of a novel hash-table for speeding-up suggestions for misspelt Tamil words. In *2017 IEEE International Conference on Industrial and Information Systems (ICIIS)*, pages 1–5.

Ratnasingam Sakuntharaj and Sinnathamby Mahesan. 2021. Missing word detection and correction based on context of Tamil sentences using n-grams. In *2021 10th International Conference on Information and Automation for Sustainability (ICIAfS)*, pages 42–47.

Anbukkarasi Sampath, Thenmozhi Durairaj, Bharathi Raja Chakravarthi, Ruba Priyadharshini, Subalalitha Chinnaudayar Navaneethakrishnan, Kogilavani Shanmugavadivel, Sajeetha Thavareesan, Sathiyaraj Thangasamy, Parameswari Krishnamurthy, Adeep Hande, Sean Benhur, Kishor Kumar Ponnusamy, and Santhiya Pandiyan. 2022. Findings of the shared task on Emotion Analysis in Tamil. In *Proceedings of the Second Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Mike Schuster and Kuldip K Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681.

R Srinivasan and CN Subalalitha. 2019. Automated named entity recognition from tamil documents. In *2019 IEEE 1st International Conference on Energy, Systems and Information Processing (ICESIP)*, pages 1–5. IEEE.

C. N. Subalalitha. 2019. Information extraction framework for Kurunthogai. *Sādhanā*, 44(7):156.

CN Subalalitha and E Poovammal. 2018. Automatic bilingual dictionary construction for Tirukural. *Applied Artificial Intelligence*, 32(6):558–567.

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.

Sajeetha Thavareesan and Sinnathamby Mahesan. 2021. Sentiment analysis in Tamil texts using k-means and k-nearest neighbour. In *2021 10th International Conference on Information and Automation for Sustainability (ICIAfS)*, pages 48–53.

Josephine Varsha and B. Bharathi. 2022. Ssncse nlp@tamilnlp-acl2022: Transformer based approach for detection of abusive comment for tamil language. In *Proceedings of the Second Workshop on Speech and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Rita Kirk Whillock and David Slayden. 1995. *Hate speech*. ERIC.

Michael Wiegand, Josef Ruppenhofer, and Elisabeth Eder. 2021. Implicitly abusive language–what does it actually look like and why are we not getting there? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 576–587. Association for Computational Linguistics.

Konthala Yasaswini, 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*, pages 187–194, Kyiv. Association for Computational Linguistics.