

PANDAS@TamilNLP-ACL2022: Abusive Comment Detection in Tamil Code-Mixed Data Using Custom Embeddings with LaBSE

Gayathri G L

SSN College of Engineering

gayathri2010090@ssn.edu.in

Krithika S

SSN College of Engineering

krithika2010039@ssn.edu.in

Divyasri K

SSN College of Engineering

divyasri2011037@ssn.edu.in

Durairaj Thenmozhi

SSN College of Engineering

theni_d@ssn.edu.in

B. Bharathi

SSN College of Engineering

bharathib@ssn.edu.in

Abstract

Abusive language has lately been prevalent in comments on various social media platforms. The increasing hostility observed on the internet calls for the creation of a system that can identify and flag such acerbic content, to prevent conflict and mental distress. This task becomes more challenging when low-resource languages like Tamil, as well as the often-observed Tamil-English code-mixed text, are involved. The approach used in this paper for the classification model includes different methods of feature extraction and the use of traditional classifiers. We propose a novel method of combining language-agnostic sentence embeddings with the TF-IDF vector representation that uses a curated corpus of words as vocabulary, to create a custom embedding, which is then passed to an SVM classifier. Our experimentation yielded an accuracy of 52% and a macro F1-score of 0.54.

1 Introduction

In recent times, with rapid digitisation, people are increasingly using social media and various other forums available online for interpersonal communication (Riehm et al., 2020). However, these platforms also come with their own share of drawbacks, such as the propagation of fake news (Waszak et al., 2018) and cyberbullying (Whittaker and Kowalski, 2015), to list a few.

Comments that are found to be offensive and often degrading, that may be targeted at an individual or a community as a whole, are categorised as abusive comments. These comments often have negative effects on the mental well-being of people (O'Reilly et al., 2018), with an apparent relation between the time spent on social media and increasing levels of depression (Karim et al., 2020). There is a pressing need for moderation on these websites, which motivates the creation of a system that will be able to classify abusive comments into one of many categories. It could also be useful in

identifying and filtering out vitriolic content.

A major challenge faced with this task is that most of the data available contains a mixture of languages (Lin et al., 2021), with people often transliterating from their native language into English, thus posing a hurdle, as most resources available for the task of Abusive language detection are pre-trained on English text.

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. Tamil is one of the world's longest-surviving classical languages. Malayalam is Tamil's closest significant cousin; the two began splitting during the 9th century AD (Anita and Subalalitha, 2019b,a; Subalalitha and Poovammal, 2018; Subalalitha, 2019; Srinivasan and Subalalitha, 2019; Narasimhan et al., 2018; Sakuntharaj and Mahesan, 2021, 2017, 2016; Thavareesan and Mahesan, 2019, 2020a,b, 2021).

The Task A of Abusive Comment Detection in Tamil-ACL 2022 (Priyadharshini et al.) involves classification of purely Tamil text, whereas task B deals with the classification of code-mixed Tamil English text into 8 categories as listed in Table 1. Our approach for Task B was to create embeddings for each data record and then pass them to the various classifiers. Three types of embeddings were employed - a multilingual BERT that produces language-agnostic embeddings, TF-IDF vectorizer and a combination of both.

The remainder of this paper is organised as follows. Section 2 is dedicated to related works obtained from the literature survey. Section 3 proceeds to describe the dataset used. Section 4 covers the details of the preprocessing steps, outlines the feature extraction process and describes the model employed for this task. Section 5 summarises the results and Section 6 concludes the paper.

Category	Definition	Example	Train	Dev
None-of-the-above	Does not belong in any of the other categories	Bala kumar wat ur asking.? 1st olunga kealviya kealunga.	3715	917
Misandry	These are comments indicating contempt against men.	Poda H cha naaye	830	218
Counter-Speech	It is a way of undermining a harsh remark by giving alternate narratives of the story.	Manickam Anbu ammaavai pathi pesurathu sari kidaiyaathu.	348	95
Xenophobia	These are comments that involve hatred towards people of a different culture/ country.	kudisekiram tamilnadu china controll poidum...	297	70
Hope-Speech	These contain sentences that include phrases indicative of hope and other such positive emotions.	DMKJambu Lingaa OK manaviyai mathippom. Malai pola valgaiyil uyarvom.	213	53
Misogyny	These are hateful statements against women.	Gh Wb u pondatti pundaila en Pola Vidava... ungomma punda naaruthu	211	50
Homophobia	Statements with a negative connotation, targeted towards homosexuality.	Nee Naam gay sax panalam	172	43
Transphobia	Referring to those hateful comments having a prejudice against transgender people.	Pitchakara Moothavinga Train la Ukkara uda maatithunga Parathesinga thuu. Ithungaluku daily azha vendiyatha iruku	157	40

Table 1: Description of the dataset

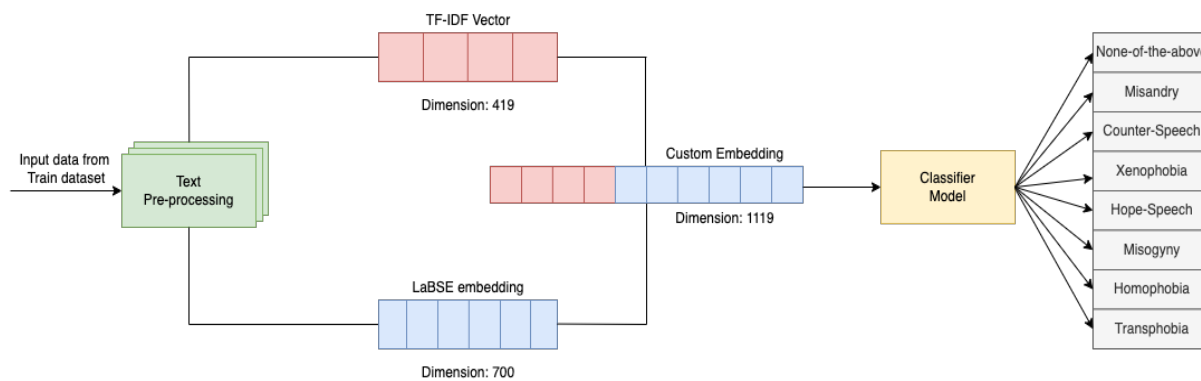


Figure 1: Diagram of the model

2 Related Works

Identification and classification of offensive tasks in a fast and effective manner is very important in the moderation of online platforms (Priyadharshini et al., 2021; Kumaresan et al., 2021; Chakravarthi, 2020; Chakravarthi and Muralidaran, 2021; Sampath et al., 2022; Ravikiran et al., 2022; Chakravarthi et al., 2022; Bharathi et al., 2022). We explored various models to achieve the same.

Ravishankar et al. (Ravishankar and Raghunathan, 2017) proposed three different approaches to classify tamil tweets based on syntactic patterns. These include Tweet weight model, TF-IDF and Domain-Specific Tags (DST), and used Tamil Dictionary (Agarathi). The authors collected tweets from 100 movies which amounted upto 7000 tweets. They proposed three other feature extraction models which include TF-IDF, adjective rules, negation rules, and adjective rules which could be passed into classifiers.

Alison P. Ribeiro et al. (Ribeiro and Silva, 2019) presented their model to identify hate speech against women and immigrants. It consisted of pre-trained word embeddings using FastText and GloVe which they passed through a CNN network.

Younes Samih et al. (Modha et al., 2021) modelled an architecture with Support Vector Machine (SVM) and Deep Neural Networks (DNNs) for task to identify Hate Speech and offensive content. They experimented four different approaches and combined them into an ensemble. They used FastText for the first one, FFNN architecture with four hidden layers for the second one and for the third one they created pretrained word embeddings using Mazajak method which was then passed into a CNN layer and a BiLSTM layer. Their next approach was using BERT. They combined these to

create an ensemble which performed well for the given dataset.

Anna Glazkova et al. (Glazkova et al., 2021) for the HASOC 2021 task which focused on detecting offensive, profane and hate content in tweets in six languages. They proposed various models which include pretrained BERT, RoBERTa and LaBSE. Though the performance of the models were similar for the English datasets, LaBSE outperformed the others for Hindi and Marathi datasets.

Shervin et al (Malmasi and Zampieri, 2017) used a corpus of 14.5k English tweets and modelled an approach to classify them as hate speech and non-hate speech. Their model uses character n-grams, word n-grams and word skip-grams for feature extraction which was passed onto a linear SVM classifier.

Burnap et al. in their paper (Burnap and Williams, 2014) wrote about their model - they used unigram, bigram feature extraction techniques and POS (Parts of Speech tagging), they also used the Stanford Lexical Parser, along with a context-free lexical parsing model, to extract typed dependencies within the tweet text. This was further passed to classifiers like Bayesian Logistic Regression, Random Forest Decision Trees and Support Vector Machines.

Aswathi Saravananaraj et al. proposed an approach for the automatic identification of cyberbullying words and rumours. They modelled a Naive Bayes and a Random Forest approach which obtained a greater accuracy than pre-existing models.

From the literature survey performed, it is inferred that an approach involving feature extraction using TF-IDF delivers good results and that transformer models like LaBSE work the best for Indian language datasets, with a particularly high accuracy for the Tamil language. The SVM classifier works

well for Hate/Abusive language recognition.

Although various innovative models have been experimented on in the studies discussed above, a model involving TF-IDF feature extraction, LaBSE and SVM, will be a novel approach to this task.

3 Dataset

The dataset used for the study is made up of comments made by subscribers in Tamil language, in relation to a video available on the streaming platform YouTube. The text contents were retrieved and stored following manual annotation. The text statements were organized into 8 different classes: transphobia, counter-speech, misandry, homophobia, hope-speech, xenophobia, misogyny and none-of-the-above, depending on the sentiment reflected through them. The dataset has a highly disproportionate share of expressions being brought under 'none-of-the-above' - 62% from the train dataset and 61% from the development dataset. The data distribution of the train and development datasets is depicted in Table 1. The average number of sentences in a comment is 1.42 with the maximum number being 20 and the minimum number of sentences per data point being 1. The average word count for each row is 12.092.

4 Methodology

The proposed methodology for this task involves extracting lexical and sentence features from the data and applying classifier models, such as SVM, MLP and K neighbours classifier, to them. This is illustrated in Figure 1.

4.1 Preprocessing

Tamil, a Dravidian Language predominantly spoken in the South Asian region, consists of an intricate script consisting of 12 vowels and 18 consonants that can be combined in various ways to give 216 compound characters. The sentences are arranged in the order of Subject Object Verb, and use postpositions. The main difficulty during preprocessing of the code mixed data is the use of different spellings for the same word while typing Tamil sentences in English.

Before mining the text for valuable information, the raw unstructured textual data is stripped of the noise it contains, in the form of punctuations and stop words, to produce meaningful features that might prove instrumental in classifying the samples into the eleven classes available.

1. Text normalisation: Words with different capitalisation may be considered to be different words and, to prevent that from happening, the dataset was standardised by the conversion of the text to lowercase.
2. Removal of punctuations: Since the model involves creating a corpus of the most frequently occurring words in every category, punctuations are removed. The list of punctuations from the string library was used in this process.
3. Removal of extra unwanted characters: The dataset contained a significant number of lines containing noise including emojis and iOS flags which have been filtered out, using RegEx.
4. Removal of stop words: Stop words are words in a language that are used in abundance as a part of the grammatical structure but do not necessarily add to the meaning of the sentence as a whole. These involve propositions, pronouns and articles among others. To achieve this, a curated list of Tamil-English stop words has been created and used. It includes words such as “r” (are), “ur” (your), “nee” (translates to you in Tamil) and “intha” (translates to this in Tamil).
5. Encoding: In the dataset, the data is classified into categories with textual labels. To ensure that the machine learning model is able to understand the data it is being fed, a label encoder is used on the target variable.

4.2 Feature extraction

The training dataset was first preprocessed as described in the above sub-section. Following this, embeddings of the textual data were created as outlined below.

4.2.1 Statistical feature extraction utilising TF-IDF

TF-IDF, standing for term frequency-inverse document frequency, is a method of quantifying a sentence, based on the words it contains. Each row is vectorized using a technique in which every word is essentially given a score that is indicative of its importance in the overall document.

In our implementation, a vocabulary list was first created by extracting the top 100 of the most frequently used words of each category. To ensure

Feature	Classifier	Accuracy	Precision	Recall	F1-score
TF-IDF+LaBSE	SVM	0.74	0.44	0.49	0.70
	MLP	0.67	0.44	0.45	0.49
	Random Forest	0.68	0.34	0.39	0.5
	Gradient Boosting Classifier	0.69	0.55	0.40	0.53
TF-IDF	SVM	0.71	0.28	0.31	0.75
	MLP	0.70	0.41	0.45	0.52
	K Neighbours Classifier	0.66	0.27	0.31	0.44
LaBSE	SVM	0.71	0.32	0.38	0.67
	MLP	0.66	0.38	0.40	0.43

Table 2: Macro-averaged Performance scores of the models deployed

that stop words were not included in the list of most frequent words, they were removed before vectorisation. A TF-IDF vectoriser was then initialised using this custom vocabulary, so that the resultant vector depends upon only the words that are statistically more likely to be found in an abusive comment.

4.2.2 LaBSE feature extraction

Language-Agnostic BERT Sentence Embedding, also known as LaBSE, is the state-of-the-art model in sentence embedding, and works by encoding sentences into a shared embedding space, where similar sentences lie closer to each other (Feng et al., 2020).

LaBSE proves to be a good fit for this task since it is language agnostic and is proven to work better than other previously existing sentence embedders like Doc2Vec and SentenceBERT with regard to languages with low resources (Firmiano and Da Silva, 2021) (Zhu et al., 2021). It is trained on bilingual low-resource sentences as well and works in a way so that it maximises the compatibility between the source sentence and its translation and minimises it with the other samples.

The pre-processed training data is made ready to be classified by encoding it using LaBSE, which creates an embedding with a dimension of 700, for each sentence. The model includes 630 dense layers and 1 sigmoid layer. For this task, the laBSE model was used with the default parameters. The learning rate was set to 0.001.

4.2.3 Custom embeddings

The imbalance of data was first tackled by selecting a number of random data points such that the real life variance is still retained, but the disparity between the number of samples of each type was reduced. In this run, both LaBSE and TF-IDF en-

codings were used so that the advantages of both these embedding methods could be harnessed in one model.

Individual embeddings of each type were initially created using the methods as described in the above two runs. They were then appended to each other to obtain a custom embedding.

4.3 Classifier Models

Following this, simple ML models such as SVM, MLP, and K Neighbours Classifier were used to classify the embeddings obtained. They are explained as follows;

SVM, also known as Support Vector Machine works by choosing the best hyperplane such that the data classes are segregated better (Mathur and Foody, 2008). RBF kernel has been used in the SVM classifier to optimise the results.

Multilayer Perceptron, abbreviated to MLP, is a feedforward deep learning network consisting of an input layer and output layer that are completely connected to each other by paths. Hidden Layer size of 200 has been used for the optimisation of this particular model, along with rectified linear unit as the activation function.

K neighbours classifier uses the closest K neighbours and identifies the most common class found among them. This label is then assigned to the data point that is to be classified. Here, the classifier uses a value of 3 for the number of neighbours.

Scikit-learn library (Pedregosa et al., 2011) in Python was used to successfully deploy these models and the results are tabulated in Table 2.

5 Results and Analysis

5.1 Performance metrics

This task is evaluated on the Macro averages of Precision, Recall and F1-score, which computes the

Team	Accuracy
abusive-checker	0.650
GJG_TamilEnglish_deBERTa	0.600
UMUteam	0.590
Pandas	0.520
Optimize_Prime_Tamil_English_Run2	0.450

Table 3: Results of the shared task

metric individually for each class and then averages the values, so that each class gets equal priority.

In classification, precision refers to the probability that a correct classification has been done. It is the ratio of correctly classified points to the total number of points that have been predicted to be of that class.

$$Precision = \frac{TP}{TP + FP}$$

Recall, on the other hand, gives an idea of the number of classifications of a type that are rightly performed. It is the ratio of the correctly classified points of a particular class to the sum of the correctly and incorrectly classified point of the same class.

$$Recall = \frac{TP}{TP + FN}$$

F1-score is the weighted average of Precision and recall, and is used mostly when a balance of both these metrics is required or when a large class imbalance is encountered.

$$F1 - score = 2 * \frac{Precision * Recall}{Precision + Recall}$$

5.2 Results

The development dataset was used for the unbiased evaluation of the performance of the models that were fit on the training dataset. For this task, the performance metrics used for analysis were accuracy and macro averages, which include precision, recall and F1-score.

For each type of embedding used, SVM was found to be the best classifier and performs better with Radial Basis Function kernel than with Linear kernel, with accuracy scores of 0.70 and 0.71 respectively. From the table, we can also come to the conclusion that an SVM classifier with the custom embedding gives the best performance, with an accuracy of 0.74, outperforming the models employ-

ing either LaBSE or TF-IDF, each with an accuracy of only 0.71.

This run secured the 4th rank in Task B which used Tamil English data as is shown in Table 3. The model performed on the test set with a macro F1-Score of 0.34, a precision score of 0.33 and a recall score of 0.37.

6 Conclusion

This paper discusses our approach for the DravidianTechLang ACL 2022 shared task, which aims to identify and classify abusive content in Tamil-English code-mixed text collected from social media. This research contributes to this task by analysing a set of classification models for identifying various types of abusive comments. We have used a combination of embeddings using TF-IDF and LaBSE with the SVM classifier. Our results showed that using this pre-trained multilingual model along with the SVM classifier yielded better results for the code-mixed data.

This model gave us a macro F-1 score of 0.49 and an accuracy of 0.74 on the development dataset, and an accuracy of 0.520 and a weighted F1-score of 0.54 on the test dataset. In the future, we would like to improve our results by using better preprocessing techniques, which may be achieved by acknowledging and utilising the significance of relevant special characters and emoticons in the given text, instead of removing them altogether.

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.
- Peter Burnap and Matthew Leighton Williams. 2014. Hate speech, machine classification and statistical modelling of information flows on twitter: Interpretation and communication for policy decision making.
- 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 Priyadarshini, 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.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. [Language-agnostic BERT sentence embedding](#). *CoRR*, abs/2007.01852.
- Alan Firmiano and Ticiano L Coelho Da Silva. 2021. Identifying duplicate police reports. In *2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 244–247. IEEE.
- Anna Glazkova, Michael Kadantsev, and Maksim Glazkov. 2021. Fine-tuning of pre-trained transformers for hate, offensive, and profane content detection in english and marathi. *arXiv preprint arXiv:2110.12687*.
- Fazida Karim, Azeezat A Oyewande, Lamis F Abdalla, Reem Chaudhry Ehsanullah, and Safeera Khan. 2020. Social media use and its connection to mental health: a systematic review. *Cureus*, 12(6).
- 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.
- Xiaotian Lin, Nankai Lin, Kanoksak Wattanachote, Shengyi Jiang, and Lianxi Wang. 2021. Multilingual text classification for dravidian languages. *arXiv preprint arXiv:2112.01705*.
- Shervin Malmasi and Marcos Zampieri. 2017. Detecting hate speech in social media. *arXiv preprint arXiv:1712.06427*.
- Ajay Mathur and Giles M Foody. 2008. Multiclass and binary svm classification: Implications for training and classification users. *IEEE Geoscience and remote sensing letters*, 5(2):241–245.
- Sandip Modha, Thomas Mandl, Gautam Kishore Shahi, Hiren Madhu, Shrey Satapara, Tharindu Ranasinghe, and Marcos Zampieri. 2021. Overview of the hasoc subtrack at fire 2021: Hate speech and offensive content identification in english and indo-aryan languages and conversational hate speech. In *Forum for Information Retrieval Evaluation*, pages 1–3.
- Anitha Narasimhan, Aarthi 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.
- Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive language detection in online user content. In *Proceedings of the 25th international conference on world wide web*, pages 145–153.
- Michelle O’Reilly, Nisha Dogra, Natasha Whiteman, Jason Hughes, Seyda Eruyar, and Paul Reilly. 2018. Is social media bad for mental health and wellbeing? exploring the perspectives of adolescents. *Clinical child psychology and psychiatry*, 23(4):601–613.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Ruba Priyadarshini, Bharathi Raja Chakravarthi, Subalalitha Chinnaudayar Navaneethakrishnan, Thenmozhi Durairaj, Malliga Subramanian, Kogilavani Shanmugavadivel, Siddhanth U Hegde, and booktitle = Kumaresan, Prasanna Kumar”. Findings of the shared task on Abusive Comment Detection in Tamil.
- Ruba Priyadarshini, 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.
- Nadana Ravishankar and Shriram Raghunathan. 2017. Corpus based sentiment classification of tamil movie tweets using syntactic patterns.
- Alison Ribeiro and Nádia Silva. 2019. Inf-hateval at semeval-2019 task 5: Convolutional neural networks for hate speech detection against women and immigrants on twitter. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 420–425.
- Kira E Riehm, Calliope Hologue, Luther G Kalb, Daniel Bennett, Arie Kapteyn, Qin Jiang, Cindy B Veldhuis, Renee M Johnson, M Daniele Fallin, Frauke Kreuter, et al. 2020. Associations between media exposure and mental distress among us adults at the beginning of the covid-19 pandemic. *American Journal of Preventive Medicine*, 59(5):630–638.
- 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.
- 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.
- Przemyslaw M Waszak, Wioleta Kasprzycka-Waszak, and Alicja Kubanek. 2018. The spread of medical fake news in social media—the pilot quantitative study. *Health policy and technology*, 7(2):115–118.
- Elizabeth Whittaker and Robin M Kowalski. 2015. Cyberbullying via social media. *Journal of school violence*, 14(1):11–29.
- Jie Zhu, Braja G Patra, and Ashraf Yaseen. 2021. Recommender system of scholarly papers using public datasets. In *AMIA Annual Symposium Proceedings*, volume 2021, page 672. American Medical Informatics Association.