Journal of Computing & Biomedical Informatics
ISSN: 2710 - 1606

Volume 04 Issue 02
2023

Research Article
https://doi.org/10.56979/402/2023

Islamophobia Content Detection Using Natural Language Processing

Abdul Jaleel[1], Mehmoon Anwar[1], Farooq Ali[1], Raza Mukhtar[1] and Muhammad Farooq[2]

[1]Department of Computer Science, University of Engineering and Technology, Taxila, 47080, Pakistan
[2]Department of Information Technology, University of the Punjab, Lahore Pakistan
*Corresponding Author: Farooq Ali. Email: farooq.ali@uettaxila.edu.pk

Received: October 29, 2022 **Accepted:** February 23, 2023 **Published:** March 29, 2023.

Abstract: With growing hate and discrimination based on caste and race, Islamophobia is one of the major and most populated phenomena nowadays. Islamophobia, in general refers to irrational antagonism, fear, or hate of Islam, Muslims, and Islamic culture, as well as actual discrimination against these groups or people within them. Islamophobia has been steadily rising over the last five years, and this trend has continued during the last year or so. A little decline occurred at the start of the year, mid-year, and towards the end of 2021, indicating that the tendency is very variable over the months, but the general trend is growing. Furthermore, in terms of magnitude, Europe deserves special attention, followed by Asia and North America. There are various researches targeting the above-mentioned issue but it always feels that this particular domain needs to be addressed more and proper systems or filters should be created to avoid any gender, race or culture-based discrimination at least from the social space used by millions. This research focuses on identifying Islamophobic content over social media and twitter in particular. In this article, the domain of islamophobia is explored on the social platform especially Twitter. To our knowledge, this is one of the very few studies that addresses Islamophobia using such advanced algorithms (BERT). Our objective is to find a model that can appropriately categorize Islamophobic tweets collected from Twitter. Initially, we constructed a dataset by extracting tweets through the use of specific keywords. Next, we classified the extracted data as either hateful or non-hateful based on whether or not it displayed Islamophobia. Subsequently, the dataset underwent pre-processing to decrease any extraneous information, such as punctuation, stop words, empty entries, and duplicates. Following that, two models, LSTM and BERT were implemented on the dataset and LSTM yielded an accuracy of 93.3 percent and BERT yielded 97.1 percent accuracy.

Keywords: BERT; Islamophobia; LSTM; Twitter.

1. Introduction

The bias against specific human groups can be traced back to the early days of our society. However, the value of social media and websites as sources of information has increased over time, especially with the advent of smartphones, laptops, tablets, and other portable devices that allow us to access the internet from anywhere. Communication, information retrieval, payments, shopping, and transportation are among the various services that have been significantly impacted by technological advancements. In the past, accessing these services required us to rely on our friends, colleagues, or relatives, but now we depend on digitalized mediums to stay connected with them. However, it is becoming increasingly difficult to shield ourselves from the bad influence of these channels and protect ourselves from potential threats.

Social media utilization to have dialogues has its negatives too and one of the most prominent one is Islamophobia or in other words online hatred directed towards Islam, Muslims, and its culture. Religious hatred is the outcome of violence done merely to preserve Truth is subjective and can be influenced by biased language, which can affect communication among believers. With the advent of communication technologies, information can be transmitted quickly without proper validation, leading to the rise of Islamophobia. As the world's largest religion, Islamophobia can have political and security implications. Social media platforms, especially Twitter, have become a massive source of positive and negative comments. Islamophobia is one of the most frequently mentioned negative statements, with various forms of expression, such as rallies, laws prohibiting Islamic emblems, and negative ideas on social media. Analyzing the classification of discussion material and community patterns that lead to statements concerning Islamophobia on social media is intriguing. Sentiment analysis, which determines the representation of emotion through text, can help identify positive or negative sentiment. Data on social media discussions with hate speech hashtags organized by the city can assist in analyzing society's social model and character. The intersection of cutting-edge internet-based technology and society's engagement in its development is an area of interest several unrealized fields, notably in the domain of ethics, have had outstanding societal impact.

Previous studies have examined the problem at hand in their own way, but as we have discussed, most of them have focused on discrimination in general. In contrast, our research study is specifically concerned with online Islamophobia. Our paper deals with the issue of Islamophobic behavior on the Twitter platform, which is significant for a couple of reasons. First, there is very little research that has been conducted on the detection of Islamophobia on social media. Second, there is no benchmark dataset available, but dataset designed by us, and analysis technique can assist digital mediums networks in constructing a better filtering process to eliminate hateful content. This, in turn, can provide a secure cyberspace where digital media users are protected from victimization and attackers. This recognition is crucial as the first step in reducing the societal attitude associated with hate speech in community disputes. It is anticipated that a community would form to participate in the forums specified to better understand the true teachings of religion and to reduce the racial hatred community on social networks.

This research aims to find suitable Islamophobic content classifier using respective mining and classification algorithms. The key objectives of this work are as follows:

- To design an Islamophobic based hatred dataset as a benchmark.
- To find an efficient algorithm for classification and emotion mining.
- To evaluate the experimental results using performance metrics.

2. Related Work

Cyberbullying is a significant issue faced by individuals on social media platforms, and one form of cyberbullying is Islamophobia, which has affected many Muslims in the last decade due to the increased availability of the internet. Islamophobia is most harmful bad effects of the digital media development platforms, with obvious evidence showing that religion-based hate can cause severe and long-term issues in individuals. These negative effects include suicide, depression, anxiety, self-harm, negative emotions, and psychosomatic symptoms. The 14th OIC report on Islamophobia revealed a rising trend of hate towards Muslims, with the majority of it originating from Europe, particularly France and the UK.

Many researchers have worked on the automatic detection of bullying or specifically Islamophobic content. Consequently, there is an undeniable need for a detection method that can identify Islamophobic content over the internet. The majority of previous studies have focused on using sentiment analysis to get criminal patterns from social media or internet-based applications, studies focus on single factors such as harassment, personal data breaches, extremist ties, etc. Most of the studies have focused on bullying in general, with very little research conducted on Islamophobic-based hatred. Our approach differs from previous studies in several ways, including the dataset used and algorithms implemented. Therefore, we provide a brief summary of the previous studies and their findings.

Fachrul Kurniawan did a study with the goal of discovering the issue of sentiment analysis connected to Islamophobia in social media, specifically Twitter. To identify the data, they employed machine learning algorithms such as support vector machine (SVM) as well as long-short term memory (LSTM). The drop duplication operation generated 4339 from the preceding 10997 in the pre-processing step, and the result language identified 31 languages. They achieved 73.797% accuracy for LSTM and 60.22% accuracy for SVM using Polynomial kernel [2]. Bertie Vidgen and Taha Yasseri create an automated software tool that identifies between non-Islamophobic, mild Islamophobic, and strong Islamophobic text based on in-depth conceptual study. It obtained 77.6% accuracy and 83% balanced accuracy. The technology will allow for future quantitative study on the causes, spread, prevalence, and consequences of Islamophobic hate speech on social media [3]. In another study, Buğra AYAN used emotion analysis to search Twitter for tweets that were Islamophobic or not. Precision, recall, and F1 estimates were produced using models trained with linear ridge models. The system was trained using 80% of the data set, which included 162,000 tweets, then tested with the remaining 20%. The accuracy rate of the system was also evaluated after applying word-based TF-IDF weighting to the data set, Ridge Regression, and the Naive Bayes model. The accuracy gained in Tweet categorization for these models was 96.9% in Ridge Regression in the F1 criteria, whereas it was 95.4% in Naive Bayes. Furthermore, Ridge Period was shorter than the Naive Bayes Classifier as training time [4]. E Ferrara and WQ Wang introduced a machine learning system that detects the researchers conducted a study on identifying extremist users and predicting their content adoption and interaction reciprocity on social media platforms. They used a combination of metadata, network, and temporal factors and analyzed. The study involved three forecasting tasks: (i) identifying extremist users, (ii) approximating the adoption of extremist content by regular users, and (iii) predicting users' response to contacts initiated by extremists. The researchers evaluated the forecasting, up to 72% AUC for interaction mutuality estimating in different forecasting scenarios [5].

In 2021, Almutairi and Al-Hagery presented a model for detecting and classifying the data related to cyberbullying of Arabic language on Twitter in the context was classified using PMI and SVM, with SVM achieving an F1-score of 82% and PMI achieving 50% [6]. Another study by Bouma et al. used anomaly detection to detect crime on Twitter, specifically focusing on Turkish Kurdish data and the 4daagse corpus. The findings showed that the system could effectively evaluate messages and detect changes in sentiment [7]. Van Hee et al. programmed cyberbullying recognition is explored in content of digital media is studied by using linear SVM and study of type of classification like binary. The data is used from ASKfm to create Dutch and English corpora and found that AUC grades were more challenging to data disparity than recollection, precision or F-scores [8]. Some researchers have used deep learning algorithms for bullying classification, and one study found that a model combining the combination of GloVe840 word embedding method and BLSTM yielded the most favorable outcome, with precision, F1 score and accuracy of, 96.60%, 94.20%, and 92.60% respectively, as reported in [9]. Moreover, another approach known as CNN-CB surpassed conventional content-based cyberbullying detection by attaining 95% accuracy through a CNN-based model and word embeddings [10].

3. Methodology

The methodology delivers a strong overview of the research process, as depicted in Figure 1. The first phase involves Twitter data obtained using its API. This data is then constructed along with necessary outcomes and features to classify them into Islamophobic and non-Islamophobic categories.

Finally, we draw conclusions and discuss the results obtained from the experiments, highlighting the strengths and limitations of the proposed approach. It is important to note that during the pre-processing phase, we applied various techniques such as tokenization, stop-word removal, stemming, and lemmatization to make the text data more structured and meaningful. Additionally, we performed data balancing to ensure that both Islamophobic and non-Islamophobic tweets were represented equally in the training and testing datasets.

During the classification phase, we used LSTM and BERT, which are two popular deep learning algorithms for natural language processing tasks. LSTM is a recurrent neural network that can learn long-term dependencies in sequential data, while BERT is a transformer-based model that can capture context-dependent word representations. To evaluate this approach, we used various evaluation metrics such as accuracy, precision, recall, and F1 score. We also performed a statistical significance test to determine if the differences in performance between the two algorithms were significant.

Overall, the proposed methodology provides a comprehensive and systematic approach to detect and classify Islamophobic content on social media, which can help social media platforms to identify and filter out harmful content and create a safer online environment for all users.

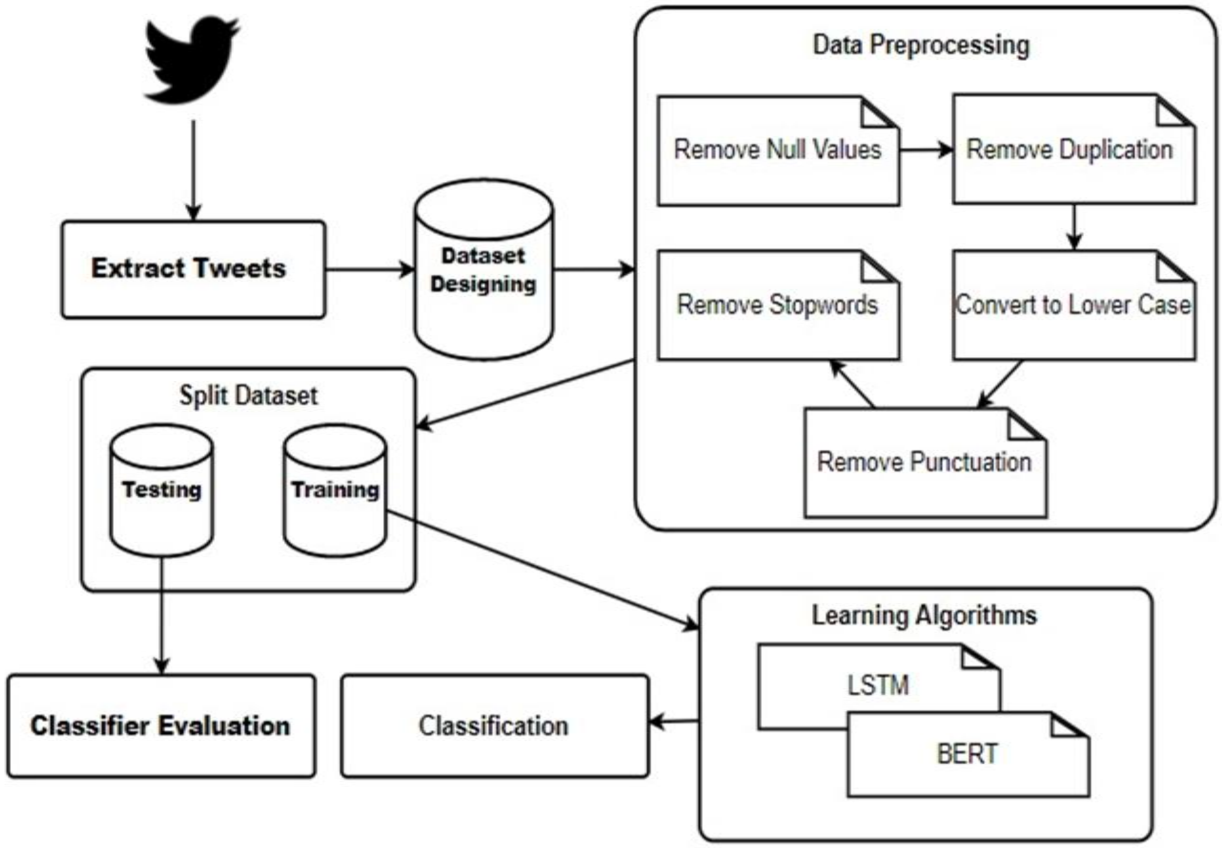

Figure 1. Proposed Methodology of Work

3.1. Dataset Designing

Acquiring data is the initial and most crucial phase of any investigation. The creation and formulation of our dataset presented a challenge due to the chosen language and limited available content. Through Twitter's API, we were able to extract 55,000 tweets using specific keywords to avoid any privacy violations, while abiding by Twitter's network policies. Accessing the Twitter API requires an authorization key which permits individuals to gather data. Our data collection process involved designing a code block that prompts for input keywords and number of tweets, which returned data based on our filters. Following the data extraction, the tweets were categorized as Islamophobic or non-Islamophobic. Islamophobic tweets were labeled as true and non-Islamophobic tweets were labeled as false. To ensure accurate labeling, Muslims (primarily friends and colleagues) verified the dataset by voting true or false on each tweet. Only tweets receiving two or more votes out of three were marked as a specific category, while mixed or balanced voting tweets were disregarded. In the end, we constructed a final dataset containing 5,000 tweets, with 2,500 Islamophobic and 2,500 non-Islamophobic tweets.

3.2. Data Labeling and Collection

Data collection is a key footstep in NLP. Our research intended towards gather info from digital media mainly Twitter thorough tweets associated to Islamophobia. To approach Twitter's details, I had to gain API route, which gave endorsement to retrieve the data for exploration direction. An API, short for Application Programming Interface, acts as an channel between two applications, facilitating the correspondence of call between them. You can use the Twitter API to know and write Twitter data, such as tweets, profiles, and follower data, and admittance various tweets about a specific topic or region. After acquiring the tweet data, we designed the datasets by as well as the basic characteristics that each dataset should have. To recognize between Islamophobic and non-Islamophobic tweets, the dataset was flagged by Muslim colleagues and fellows. A statement is flagged as Islamophobic statement only if it had an odd and higher number of votes. For instance, a tweet was considered Islamophobic if it received 2 out of 3 votes. In case of any conflicts, the data in question was removed from the dataset.

3.3. Data Preprocessing

Preprocessing is an important step in natural language processing that seeks to enhance text performance of the classifier through data cleaning. A broad range of methods are used to make datasets more efficient for classification and extraction of features. The goal is to convert the main features in data which can be create appropriate prediction simulation while also improving unsupervised learning training efficiency. We conduct pre - processing steps such as transforming the text to uncapitalized, expelling punctuation, stop-words and so on. These phases were taken to eliminate noise from the data and enable the algorithms to learn more efficiently, resulting in more accurate and trustworthy outcomes during the algorithm testing phase. Pre-processing is a critical aspect of any NLP work is an essential step in Natural Language Processing as it helps to improve the efficiency and accuracy of the learning algorith-

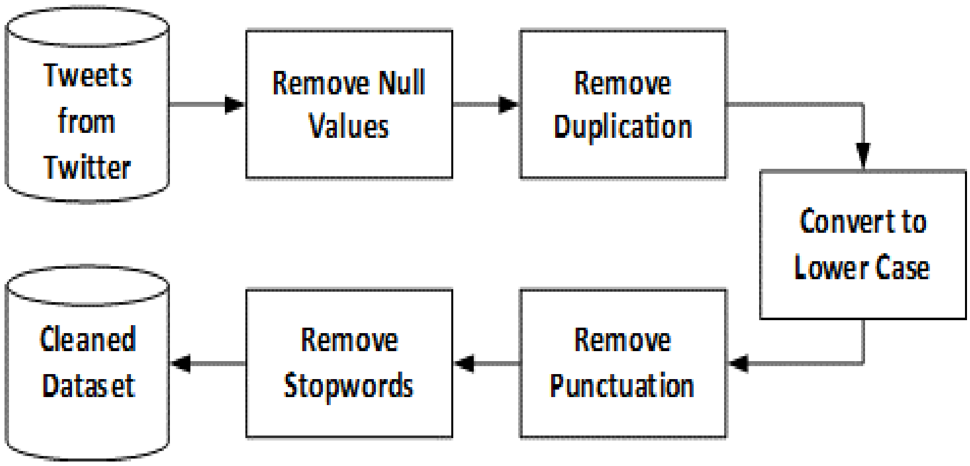

Figure 2. Flow of Data Preprocessing

ms used for text classification. The pre-processing steps include techniques such as lowercasing, stop word removal, and punctuation removal. These steps are necessary to remove noise from the dataset and enable the algorithms to learn more effectively. The sequence of pre-processing steps used in our research is illustrated in Figure 2.

3.4. Remove Null Values

In this phase, the benchmark developed dataset is analyzed and removed any null values present. This was done to ensure that the learning and classification algorithms could operate effectively without encountering any errors due to null values. Many algorithms cannot handle invalid values, which may inaccuracies during the model learning and model testing process. Therefore, removing null values from the dataset is crucial to ensure the accuracy and reliability of the results obtained from these algorithms.

3.5. Removing Duplication

During the pre-processing step, it is essential to remove all duplicated items from the dataset to make each entry of dataset unique. If the dataset has a high proportion of repeated entries, it can result in a bias towards those items during categorization, which can lead to inaccurate results. Therefore, removing duplicates can help to ensure that the model is trained and tested on a more diverse and representative dataset, which can improve the overall accuracy of the classification model.

3.6. Text Conversion to Lowercase

During this pre-processing step, we converted the data in the form of text data to small letter using a simple line of code in Python. This helps the machine to process the text in a more consistent and logical manner. However, we only applied this conversion to the features that were necessary for training and testing the learning algorithms. This is because there may be certain cases where the text's original case might have some significance and should not be converted to lowercase.

3.7. Removing Stopwords

During the pre-processing step, removing stop words is essential to enhance the understanding of the sentence's meaning. Stop words are commonly used words such as "the," "a," "an," "and," "or," and "in" that have little to no semantic meaning. Removing stop words improves the performance of natural language processing algorithms by leaving only relevant words behind. This process is also used by search engines like Google to improve the speed and accuracy of search results.

4. Results and Discussion

Next, we will delve into the outcomes of the appropriate machine learning algorithms used in our study. We evaluate the results using performance matrices of the learning algorithms and compare the outcomes of two algorithms applied on our dataset. To facilitate comparison, we present the results in a table and a graph.

To measure the performance of algorithms, we use performance metrics that include various measures such as Accuracy, Recall, F1 Score, and Precision. We use these metrics to analyze the performance and determine the best algorithm for our work. We discuss each of these measures based on the actual and prediction-based metrics presented in Table 1.

Table 1. Evaluation Metrics

Evaluation Metrics		**Actual**	
		Positive	**Negative**
Prediction	Positive	True Positive (TP)	False Positive (FP)
	Negative	False Negative (FN)	True Negative (TN)

The accuracy metric evaluates the proportion of correct predictions made by a classifier. However, it may not be a reliable measure when dealing with imbalanced datasets as it can result in biased predictions towards the category that appears more frequently, leading to incorrect predictions for the other categories.

$$\boldsymbol{Accuracy = \frac{(TP + TN)}{(TP + FP + FN + TN)}} \tag{1}$$

The terms TP (True Positive), TN (True Negative), FN (False Negative), and FP (False Positive) are used to calculate the values for Precision, Recall, and F1-Score. These measures are more reliable than accuracy in evaluating classifier performance, particularly when the dataset is imbalanced. The formulas for these metrics are provided below.

$$\boldsymbol{Precision = \frac{(TP)}{(TP + FP)}} \tag{2}$$

$$\boldsymbol{Recall = \frac{TP}{(TP + FN)}} \tag{3}$$

$$\boldsymbol{F1\ Score = 2 * \frac{Precision * Recall}{Precision + Recall}} \tag{4}$$

After implementing the selected algorithms on our dataset for classification of Islamophobic hatred data points, the finding of classification algorithms was calculated using performance matrices discussed above. The LSTM model takes in a sequence of words as input and uses the memory cell to remember important information from the previous words in the sequence. This allows the LSTM to understand context and meaning in the text, which is important for text classification. LSTM achieved an accuracy score of 93.3 percent. It performed well on our dataset and outperformed many of the previous research's accuracy scores. On the other hand, BERT (Bidirectional Encoder Representations from Transformers) is a transformer-based pre-trained model that is designed to understand the context of a sentence by looking at the words both before and after a given word**.** LSTM (Long Short-Term Memory) is a type of recurrent neural network that can process sequential data by using a memory cell to remember information from previous time steps.

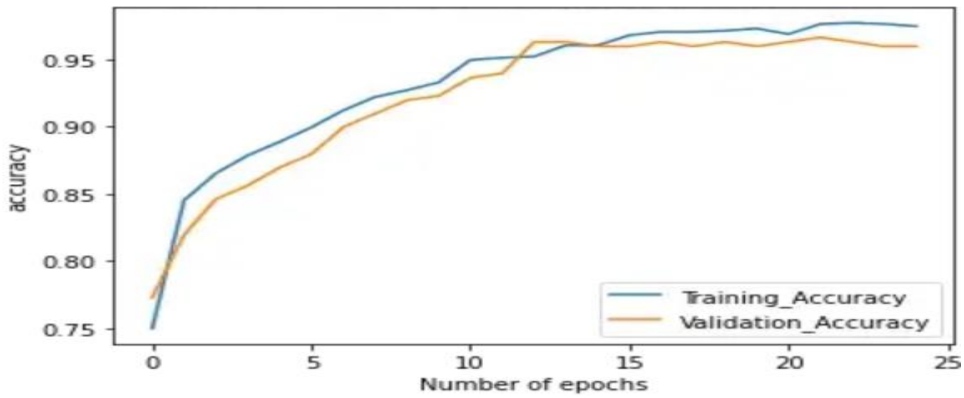

Figure 3. Comparison training and validation accuracy wrt epochs

As in the Figure 3 the reason shown that we achieved the accuracy during the training and validation is almost same which is reason to achieve these best results for BERT and it is increases with number of epochs.

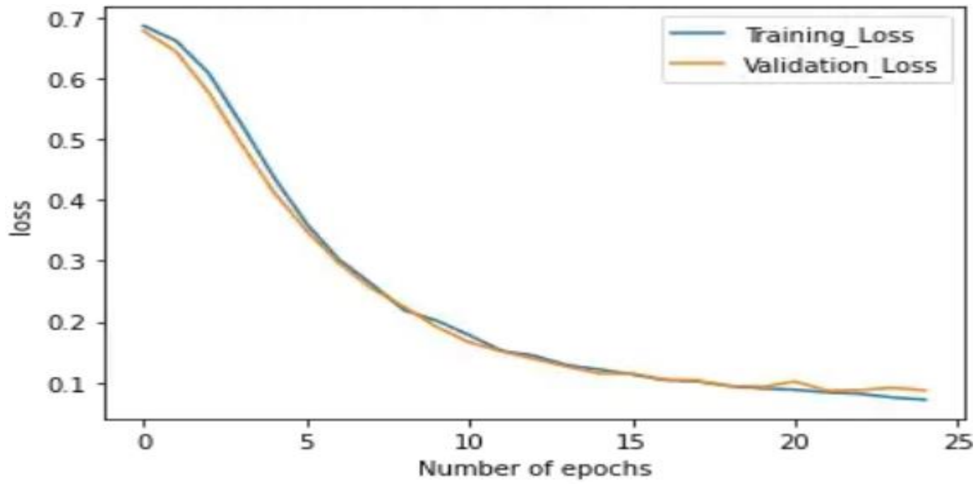

Figure 4. Comparison training and validation loss wrt epochs

Additionally, the graph of loss shown in Figure 4 shows that the loss also decreases when the number of epochs increases minimum difference.

BERT uses a technique called pre-training, where the model is trained on a large corpus of text before being fine-tuned on a specific task, such as text classification. BERT achieved an outstanding accuracy score of 97.1 %, outperforming LSTM. We can see in Table 2 that LSTM accuracy is less than BERT because the loss of during training and validation increases which drops its validation accuracy over training accuracy.

Table 2. Classifiers' Accuracy

Classification Algorithm	Accuracy
LSTM	93.4
BERT	97.1

5. Results Comparison

Now, in Fig. 3, both accuracies are compared of algorithms (LSTM and BERT) to represent the best fit algorithm for Islamophobic content classification by both proposed models comparison. BERT is a transformer-based architecture that has been shown to outperform LSTM (Long Short-Term Memory) in text classification tasks. This is because BERT can consider the context of a word in both the left and right sides of the sentence, whereas LSTM only considers the context from the past. Additionally, BERT uses attention mechanisms which allow the model to focus on specific parts of the input when making a prediction, whereas LSTM does not have this capability. This allows BERT to better understand the meaning of a sentence and make more accurate predictions. Overall, BERT has been found to be more effective than LSTM in text classification tasks because of its ability to consider context from both the left and right sides of the sentence and its use of attention mechanisms.

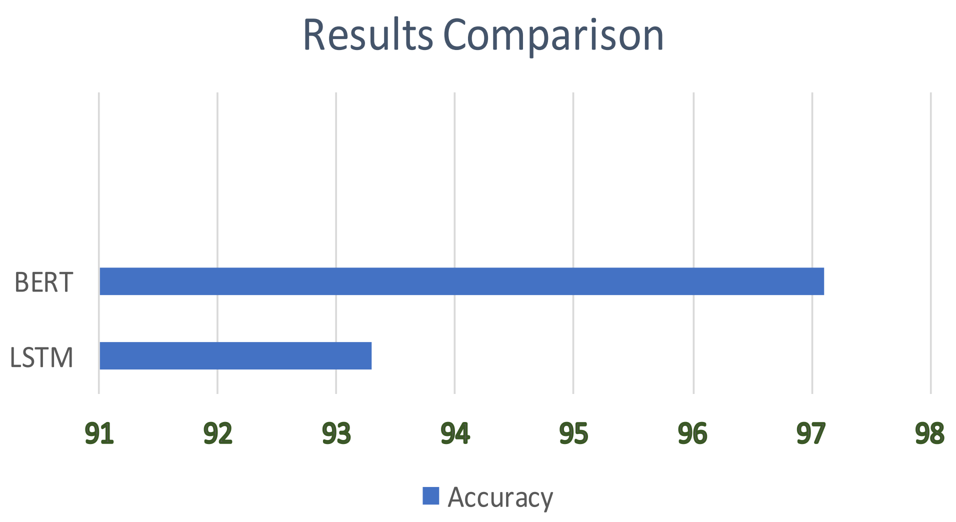

Figure 5. Comparison of Both Algorithms

When we compare both results, we found that BERT outperforms LSTM with an accuracy of 97.1 percent.

6. Conclusion and Future Work

In this study we have developed a benchmark dataset to address the need for an advanced word filtering mechanism for identifying Islamophobic content in tweets and posts on social media. Further, it is evident by the above graphical representation that BERT has outperformed LSTM for the designed dataset, achieving an accuracy of 97.1 percent.

To enhance our research in the future, we can augment our constructed dataset with more data to achieve more precise and dependable outcomes with heightened accuracy, various techniques such as stemming and lemmatization can be employed for preprocessing, which can help extract more relevant features using different strategies for feature extraction. There are numerous future paths in the realm of identifying Islamophobic content, with most research studies currently focused on English. However, social media platforms are used in several regional and international languages, presenting unique opportunities for further research. Moreover, improving the features derived from social media sites can enhance the classification models for identifying Islamophobia. Our proposed methodology can aid in designing an effective filter mechanism for detecting Islamophobia on different social media platforms. Considering that Europe has over 300 million active social media users, our study can serve as a crucial filter mechanism for preventing cyberbullying (Islamophobia) in cyberspace. We believe that our model can efficiently assist in the automatic detection of Islamophobic content on any social media platform.

Acknowledgments: I'd like to extend my gratefulness to everyone at UET Taxila, including my friends and family, who have played an instrumental role in making this work possible. Their support has contributed significantly to my research experience, which I will cherish for a lifetime.

References

1. 14th Annual Report on Islamophobia March 2022.
2. Kurniawan, Fachrul, and Aji Prasetya Wibawa Badruddin. "Identification of Islamophobia Sentiment Analysis on Twitter Using Text Mining Language Detection." Journal of Positive School Psychology (2022): 8286-8294.
3. Vidgen, Bertie, and Taha Yasseri. "Detecting weak and strong Islamophobic hate speech on social media." Journal of Information Technology & Politics 17.1 (2020): 66-78.
4. Ayan, Buğra, B. Kuyumcu, and B. Ciylan. "Detection of Islamophobic Tweets on Twitter Using Sentiment Analysis." Gazi University Journal of Science Part C 7.2 (2019): 495-502.
5. Ferrara, Emilio, et al. "Predicting online extremism, content adopters, and interaction reciprocity." International conference on social informatics. Springer, Cham, 2016.
6. Almutairi, Amjad Rasmi, and Muhammad Abdullah Al-Hagery. "Cyberbullying Detection by Sentiment Analysis of Tweets' Contents Written in Arabic in Saudi Arabia Society." International Journal of Computer Science & Network Security vol. 21, no.3, pp. 112-119, 2021.
7. Bouma, Henri, et al. "On the early detection of threats in the real world based on open-source information on the internet." Information Technologies and Security. 2012.
8. Van Hee, Cynthia, et al. "Automatic detection of cyberbullying in social media text." PloS one vol. 13, no.10, pp. e0203794, 2018.
9. Mouheb, Djedjiga, et al. "Detection of arabic cyberbullying on social networks using machine learning." 2019 IEEE/ACS 16th International Conference on Computer Systems and Applications (AICCSA). IEEE, 2019.
10. Bharti, Shubham, et al. "Cyberbullying detection from tweets using deep learning." Kybernetes 2021.
11. Green, Todd H. The fear of Islam: An introduction to Islamophobia in the West. Fortress press, 2019.
12. Najib, Kawtar, and Peter Hopkins. "Where does Islamophobia take place and who is involved? Reflections from Paris and London." Social & Cultural Geography 21.4 (2020): 458-478.
13. Uenal, Fatih, et al. "The nature of Islamophobia: A test of a tripartite view in five countries." Personality and Social Psychology Bulletin 47.2 (2021): 275-292.
14. Iner, Derya, et al. Islamophobia in Australia-II (2016-2017). Charles Sturt University, 2019.
15. R. M. Silva, R. L. Santos, T. A. Almeida, and T. A. Pardo, "Towards automatically filtering fake news in Portuguese," Expert Systems with Applications, vol. 146, p. 113199, 2020.
16. F. H. Khan, U. Qamar, and S. Bashir, "A semi-supervised approach to sentiment analysis using revised sentiment strength based on SentiWordNet," Knowledge and information Systems, vol. 51, pp. 851-872, 2017.
17. L. Weilin and G. K. Hoon, "Personalization of trending tweets using like-dislike category Model," Procedia Computer Science, vol. 60, pp. 236-245, 2015.
18. K. Gao, H. Xu, and J. Wang, "Emotion classification based on structured information," in 2014 International conference on multisensor fusion and information integration for intelligent systems (MFI), 2014, pp. 1-6.
19. Y. Jia, Z. Chen, and S. Yu, "Reader emotion classification of news headlines," in 2009 International Conference on Natural Language Processing and Knowledge Engineering, 2009, pp. 1-6.
20. F. Liangzu, L. Ruifan, and Z. Yanquan, "Extracting key sentiment sentences from internet news via multiple source features," in 2014 4th IEEE International Conference on Network Infrastructure and Digital Content, 2014, pp. 126-130.
21. P. Bahad, P. Saxena, and R. Kamal, "Fake news detection using bi-directional LSTM-recurrent neural network," Procedia Computer Science, vol. 165, pp. 74-82, 2019.
22. C. J. Rameshbhai and J. Paulose, "Opinion mining on newspaper headlines using SVM and NLP," International Journal of Electrical and Computer Engineering (IJECE), vol. 9, pp. 2152-2163, 2019.
23. A. Samuels and J. Mcgonical, "News Sentiment Analysis," arXiv preprint arXiv:2007.02238, 2020.
24. D. B. Bracewell, F. Ren, and S. Kuriowa, "Multilingual single document keyword extraction for information retrieval," in 2005 International Conference on Natural Language Processing and Knowledge Engineering, 2005, pp. 517-522.
25. Y. Choi and H. Lee, "Data properties and the performance of sentiment classification for electronic commerce applications," Information Systems Frontiers, vol. 19, pp. 993-1012, 2017.
26. J. Kamps, M. Marx, R. J. Mokken, and M. De Rijke, "Using WordNet to measure semantic orientations of adjectives," in LREC, 2004, pp. 1115-1118.
27. C. Hutto and E. Gilbert, "Vader: A parsimonious rule-based model for sentiment analysis of social media text," in Proceedings of the International AAAI Conference on Web and Social Media, 2014.
28. A. Mulahuwaish, K. Gyorick, K. Z. Ghafoor, H. S. Maghdid, and D. B. Rawat, "Efficient classification model of web news documents using machine learning algorithms for accurate information," Computers & Security, vol. 98, p. 102006, 2020.
29. G. Gadek and P. Guélorget, "An interpretable model to measure fakeness and emotion in news," Procedia Computer Science, vol. 176, pp. 78-87, 2020.
30. D. Rambaccussing and A. Kwiatkowski, "Forecasting with news sentiment: Evidence with UK newspapers," International Journal of Forecasting, vol. 36, pp. 1501-1516, 2020.