

# Explainable Identification of Hate Speech towards Islam using Graph Neural Networks

Azmine Toushik Wasi

Shahjalal University of Science and Technology, Bangladesh  
azmine32@student.sust.edu

## Abstract

Islamophobic language on online platforms fosters intolerance, making detection and elimination crucial for promoting harmony. Traditional hate speech detection models rely on NLP techniques like tokenization, part-of-speech tagging, and encoder-decoder models. However, Graph Neural Networks (GNNs), with their ability to utilize relationships between data points, offer more effective detection and greater explainability. In this work, we represent speeches as nodes and connect them with edges based on their context and similarity to develop the graph. This study introduces a novel paradigm using GNNs to identify and explain hate speech towards Islam. Our model leverages GNNs to understand the context and patterns of hate speech by connecting texts via pretrained NLP-generated word embeddings, achieving state-of-the-art performance and enhancing detection accuracy while providing valuable explanations. This highlights the potential of GNNs in combating online hate speech and fostering a safer, more inclusive online environment.

**Disclaimer:** *This manuscript may contain examples of hateful or offensive language, as it discusses hate speech in the context of detection and analysis. These instances are included strictly for research purposes and do not reflect the authors' views. Reader discretion is advised.*

## 1 Introduction

Detecting and eliminating hate speech on social media platforms is of utmost importance for the promotion of harmony and tranquillity in society (Rawat et al., 2024; Kovács et al., 2021; Davidson et al., 2017). The escalating presence of hate speech specifically targeting Islam or Muslim communities on online discussion platforms is a growing concern (Ayuningtias et al., 2022). This form of hate speech not only fosters an environment of

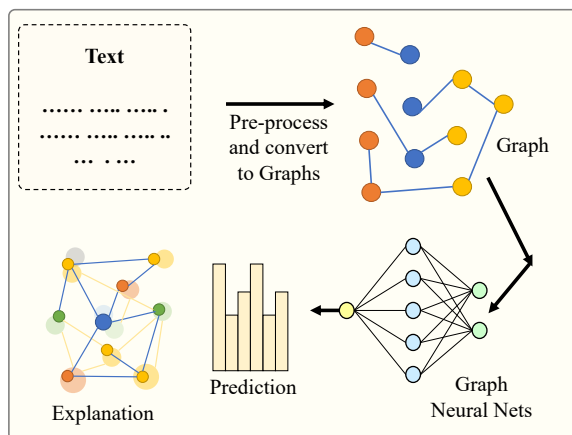


Figure 1: Our approach of Hate Speech towards Islam using GNNs

intolerance and hostility but can also have severe psychological impacts on individuals and communities, leading to real-world violence and discrimination (Saha et al., 2019).

To address this issue, researchers have increasingly turned to advanced technologies; using text-processing approaches in AI. Natural Language Processing (NLP) techniques are frequently employed for hate speech detection, with some offering severity assessment of hate speech (Mozaferi et al., 2020; Schmidt and Wiegand, 2017; Cao et al., 2020). These methods utilize sophisticated algorithms to analyse vast amounts of textual data, identifying patterns and features indicative of hate speech. For instance, deep learning models, like recurrent neural networks (RNNs), can learn complex representations of text data, enabling them to detect subtle and context-dependent instances of hate speech (Schmidt and Wiegand, 2017). Modern NLP techniques, on the other hand, can enhance these models by providing richer linguistic insights. Tokenization, part-of-speech tagging, and named entity recognition are just a few NLP techniques that help in breaking down and understanding the

text’s structure and meaning. Moreover, the integration of latest NLP model and transformers, like BERT (Devlin et al., 2019) and GPT (Radford et al., 2019; Brown et al., 2020), has significantly improved the ability of models to understand context (Ferrando et al., 2023), sarcasm (A et al., 2021), and implicit hate speech (Cao et al., 2020; Mozafari et al., 2020), which are often challenging to detect. Another interesting approach is to use human-centric perspectives of AI using some benchmark dataset (Wasi et al., 2024a,b).

Researchers have tried to employ GNNs in hate speech classification (De la Peña Sarracén and Rosso, 2022; Hebert et al., 2022; Bölücü and Canbay, 2021), but still needs more focus on this area. Despite their potential, GNNs have not been actively employed for the purpose of interpretable identification of hate speech, particularly in Islamic contexts. Islamophobic<sup>1</sup> content often exhibits close word choices and hate speakers from the same community, which GNNs can leverage to reveal and explain patterns, alongside impressive classification scores.

In this study, we introduce a novel approach employing graph neural networks for the identification and explication of hate speech directed at Islam (XG-HSI), as demonstrated in Figure 2. We pre-process the dataset to focus on Islamic contexts, utilize pretrained NLP models for word embeddings, establish connections between texts, and employ a series of graph encoders for hate speech target identification, which achieves state-of-the-art performance.

## 2 Background

Graph Neural Networks (GNNs) are powerful neural networks designed for processing non-Euclidean<sup>2</sup> data organized in complex, interconnected graphs (Scardapane, 2024; Waikhom and Patgiri, 2021). Using their ability to utilize relations between different data points (Xu et al., 2019), GNNs have shown tremendous promise in

<sup>1</sup>In this work, the terms "hate speech towards Islam" and "Islamophobic hate speech" is used interchangeably.

<sup>2</sup>In this context, non-Euclidean data refers to data that lacks a regular grid structure, like graphs where nodes (data points) are connected by edges (relationships). Unlike Euclidean data such as images or sequences, which follow a structured grid or linear pattern, non-Euclidean data is irregular and interconnected, such as a social network where users (nodes) have varying numbers of connections (edges). GNNs are designed to process this complex graph structure, capturing relationships and patterns that traditional neural networks cannot handle effectively (Scardapane, 2024).

text classification and detection tasks (Lu et al., 2020; Zhang et al., 2020; Pal et al., 2020). GNNs have the ability to enhance hate speech detection on social media by modeling complex relationships between users and content, capturing contextual information from interactions. They propagate information across the network, identifying coordinated and evolving hate speech patterns. We also present a case study in Section 5 to illustrate how incorporating related information enhances the process. Recent graph-based approaches uses LLM to develop knowledge graphs BanglaAutoKG (Wasi et al., 2024c), World Knowledge in Implicit Hate Speech Detection (Lin, 2022), and HRGraph (Wasi, 2024).

In this work, we adopted a general bag of words-based approach to create graphs, without LLMs. By integrating with pretrained NLP models, GNNs leverage contextual word embeddings to better understand the subtleties of hate speech. This combined approach improves the accuracy, context-awareness, and adaptability of detection systems, making them more effective in identifying hate speech directed at Islam and potentially generalizing to other targeted groups.

## 3 Methodology

### 3.1 Notations

Let a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, X)$ , where  $\mathcal{V}$  represents nodes,  $\mathcal{E}$  denotes edges. We also define  $N$  and  $M$  as the numbers of nodes and edges, respectively. Each node  $v$  is associated with a feature  $x_i \in \mathbb{R}^F$ , and the node feature matrix for the entire graph is denoted as  $X \in \mathbb{R}^{N \times F}$ , where  $F$  represents the feature vector length. In our approach, each content<sup>3</sup> denotes a node, contextual similarity between two nodes is denoted by an edge and word embeddings are node features of the graph. The task involves a node classification task to detect hate speech and Islamophobic content.

### 3.2 Data Pre-Processing

Initially, the dataset was filtered to focus on hate speech targeting Islam. Next, pretrained NLP models is applied to the text to obtain word embeddings  $X$  as node features for all nodes  $\mathcal{V}$ . Edges  $\mathcal{E}$  are determined using cosine similarity between embeddings with a threshold of 0.725. Subsequently, GNN is applied for the classification task.

<sup>3</sup>Each content denotes the full post, that was labelled as hate speech or not.

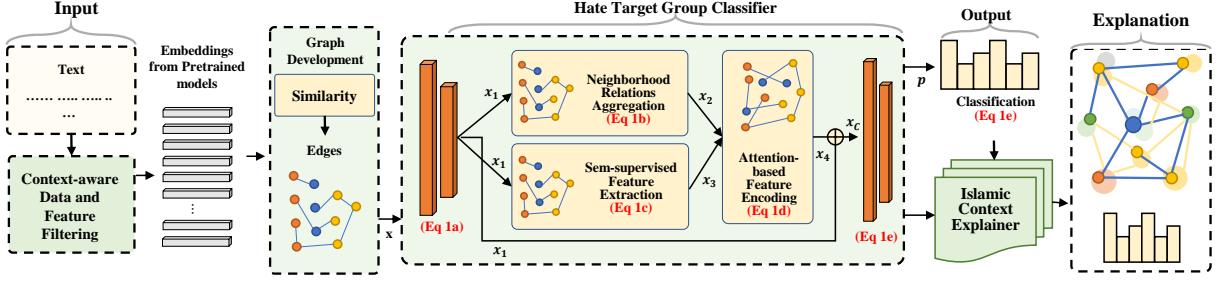


Figure 2: Our framework for Explainable Identification of Hate Speech towards Islam using GNNs.

### 3.3 Graph Encoder

After data pre-processing, every data point<sup>4</sup>  $x \in X$  undergoes a series of transformations to get output  $p$ . First, it is processed by a linear layer producing  $x_1$  (Equation 1).

$$x_1 = Wx + b \quad (1)$$

Subsequently,  $x_1$  is passed into two initial graph encoders to aggregate neighborhood information, feature extraction, and yield  $x_2, x_3$  utilizing  $\mathcal{G}$  and concatenated to  $x_{23}$  (Equation 2,3, 4). Here in Equation 2, we aggregate features from a node’s local neighborhood, to learn different characteristics (Hamilton et al., 2017). In Equation 3 and 4, we use a semi-supervised learning on graph-structured data, employing an efficient variant of convolutional neural networks that operate directly on graphs (Kipf and Welling, 2017).

$$x_2 = W_1 x_1 + W_2 \cdot \text{mean}_{j \in \mathcal{N}(i)} x_{1j} \quad (2)$$

$$x_3 = W_1 x_{1i} + W_2 \sum x_{1j} \quad (3)$$

$$x_{23} = \text{concat}(x_2, x_3); \quad (4)$$

Here,  $\mathcal{N}$  is the set of neighbouring nodes. Following this,  $x_{23}$  is passed through another graph layer employing attention-based feature extraction, utilizing masked self-attentional layers to implicitly assign different weights to nodes in a neighbourhood (Veličković et al., 2018), producing  $x_4$  (Equation 5 and 6).

$$x_4 = \alpha_{i,i} \Theta x_{23_i} + \sum \alpha_{i,j} \Theta x_{23_j} \quad (5)$$

$$\alpha = \frac{\exp(\text{LeakyReLU}(a^\top [\Theta x_i \parallel \Theta x_{23_j}]))}{\sum \exp(\text{LeakyReLU}(a^\top [\Theta x_{23_i} \parallel \Theta x_{23_k}]))} \quad (6)$$

<sup>4</sup>the BERT-embeddings, denoting the features of each sentence.

Here,  $\theta$  refers to trainable model weights.  $\alpha$  is the attention value, calculated by the equation mentioned.

Finally,  $x_4$  is passed through a final linear layer to obtain logits  $p_l$ , which are then subjected to a softmax operation to derive probabilities  $p$  (Equation 7 and 8).

$$x_c = \text{concat}(x_1, x_4); p_l = Wx_c + b \quad (7)$$

$$p = \text{softmax}(p_l) \quad (8)$$

An illustration of the network is presented in Figure 2.

### 3.4 Loss Function

Cross Entropy loss (Mao et al., 2023) is designed to minimize the difference between the predicted probabilities and true values, as follows:

$$l_{CE} = -\frac{1}{n} \sum_{i=1}^n (p'_i \log \sigma(p_i) + (1 - p'_i) \log (1 - \sigma(p_i))) \quad (9)$$

### 3.5 Graph Explanation

GNNExplainer (Ying et al., 2019) is used to derive explanations from the graph encoder network for interpreting the results and find underlying relations and causation. It works by taking a trained GNN model and its predictions as input, and returns explanations in the form of compact subgraph structures and subsets of influential node features. This model-agnostic approach can explain predictions of any GNN-based model on various graph-based machine learning tasks, including node classification, link prediction, and graph classification. GNNExplainer formulates explanations as rich subgraphs of the input graph, maximizing mutual information with the GNN’s predictions. It achieves this by employing a mean field variational approximation to learn real-valued graph masks that select

Table 1: Experimental Results ( $\uparrow$ )

Model	Accuracy	Macro F1
CNN-GRU	0.628	0.604
BiRNN	0.591	0.578
BiRNN-HateXplain	0.612	0.621
BERT	0.692	0.671
BERT-HateXplain	0.693	0.681
<b>XG-HSI-BiRNN (Ours)</b>	<b>0.742</b>	<b>0.737</b>
<b>XG-HSI-BERT (Ours)</b>	<b>0.751</b>	<b>0.747</b>

important subgraphs and feature masks that highlight crucial node features. Through this process, GNNExplainer offers insights into the underlying reasoning of GNN predictions, enhancing model interpretability and facilitating error analysis.

## 4 Experiments

### 4.1 Experimental Setup

**Dataset.** We use HateXplain (Mathew et al., 2020), a benchmark hate speech dataset designed for addressing bias and interpretability. The dataset has hate speech targets labelled. We use this labelling to collect only Muslim-focused sentences and created a subset to work on this project. We have used a 6:2:2 train, validation and test split in our work.

**Baselines.** The baseline models are: CNN-GRU, BiRNN (Schuster and Paliwal, 1997), BiRNN-HateXplain (Mathew et al., 2020), BERT (Devlin et al., 2019), BERT-HateXplain (Mathew et al., 2020). Mentioned HateXplain-based models are fine-tuned on HateXplain dataset (Mathew et al., 2020).

**Implementation Details.** Hugging Face transformers library (Wolf et al., 2020) is used to get embeddings from pre-trained BERT (bert-base-uncased) (Devlin et al., 2019) and BiRNN (Schuster and Paliwal, 1997). The model is trained for 200 epochs with a learning rate of 0.001, using Adam optimizer. The experimental results in Table 1 show that our model achieves remarkable performance comparing to benchmarks with explaining occurring phenomena. We utilized a single layer for each type of GNN, with a maximum tokenization length of 512 in the tokenizer and length of BERT embeddings ( $F$ ) set to 128.

### 4.2 Experimental Results

Table 1 shows the performance of various models in detecting hate speech, highlighting accuracy and Macro F1 metrics. Traditional models




Figure 3: Explanation Graph

like CNN-GRU and BiRNN show lower performance, with BiRNN-HateXplain offering slight improvements. BERT-based models perform better, particularly BERT-HateXplain. However, our proposed models, XG-HSI-BiRNN and XG-HSI-BERT, significantly outperform all others, with XG-HSI-BERT achieving the highest accuracy (0.741) and Macro F1 (0.747). These results demonstrate the superior effectiveness of our dual GNN approach in hate speech detection.

## 5 Graph Explanation Case Study

For a given post, "*How is all that awesome Muslim diversity going for you native germans? You have allowed this yourselves. If you do not stand and fight against this. You get what you asked for what you deserve!*", the predicted classification was offensive towards Islam. As per the explainer (Figure 4), the neighbouring and self-tokens<sup>5</sup> helped to classify this as offensive to Muslims are *fight, Muslim diversity, brooks, ##rish, donald, syrian, schultz, typed*. The text's association of "*Muslim diversity*" with potential blame and its confrontational tone in phrases like "*stand and fight against this,*" combined with neighbouring tokens like syrians, brooks, syrians denoted negative sentiment. More detailed analysis is added in Appendix A.

## 6 Discussion

We believe that our study not only addresses the immediate challenge of identifying and explaining hate speech directed at Islam but also recognizes the broader impact of hate speech propagation on

<sup>5</sup>Each sentence was tokenized, and then we collected embeddings from BERT as features. Those numbers in the figure denotes to particular token, used in tokenization.

online platforms. The proliferation of Islamophobic language fosters intolerance, division, and hostility within communities, perpetuating harmful stereotypes and prejudices. By leveraging GNNs in our XG-HSI framework, we not only detect hate speech but also provide explanations for its occurrence, shedding light on the underlying factors driving such behaviour. GNNs excel in capturing complex relationships and patterns within data, enabling them to effectively identify instances of hate speech and elucidate the contextual nuances surrounding them. By leveraging the inherent structure of social networks and textual data, our approach offers a comprehensive understanding of how hate speech propagates in online discourse.

In future research, exploring the integration of multimodal data sources, such as images and videos, could enhance the robustness of hate speech detection models, particularly in detecting nuanced forms of Islamophobic content. Additionally, investigating the dynamic nature of online communities and incorporating temporal aspects into GNN architectures could provide deeper insights into the evolution of hate speech propagation and enable more proactive interventions to counter its spread.

## 7 Conclusion

Identifying and addressing Islamophobic hatred on social media is crucial for achieving harmony and peace. This research presents a novel method using GNNs to detect hate speech towards Islam. Empirical findings demonstrate that our model achieves exceptional performance, significantly outperforming all others, with XG-HSI-BERT achieving the highest accuracy (0.741) and Macro F1 (0.747). Explainability aspect of this approach is also very promising, as it provides insights into both correlations and causation. This further highlights the potential of GNNs in combating online hate speech and fostering a safer, more inclusive online environment.

## Limitations

The limitations of our study include the use of only one dataset, which, while sufficient for this initial exploration, should be expanded upon in future research to validate and extend our findings. Additionally, while Graph Neural Networks (GNNs) are known to be computationally intensive, especially with large-scale datasets, the relatively limited number of hate speech keywords suggests that GNNs

may still be highly effective. Furthermore, more efficient GNN training methods are now available such as G3 (Wan et al., 2023) and Graphite (Gong et al., 2022), which address some of the computational challenges in future applications.

## Ethical Implications

Our work on using GNNs to detect hate speech targeting Islam carries significant ethical responsibilities. We focus on minimizing biases in the model to ensure fair treatment of all groups, emphasizing the need for transparency in how the model arrives at its decisions. By using interpretable GNN methods, we strive to provide clear explanations for the model's classifications, allowing for greater accountability. We also acknowledge the potential risks of misuse and take steps to prevent these, adhering to ethical guidelines that respect privacy and avoid unjust censorship.

## Societal Implications

The societal impact of our work lies in its potential to create a safer online environment by effectively identifying and mitigating Islamophobic content. By enhancing the detection accuracy and providing clear explanations for the identified hate speech, our model contributes to fostering more inclusive and respectful online communities. Additionally, our work highlights the importance of combating digital hate speech, which can lead to real-world harm. We aim to empower platforms and policy-makers with tools that uphold freedom of expression while curbing harmful rhetoric, thus promoting social harmony and understanding.

## Potential Risks

The application of our model presents several risks. One major concern is the potential for model misclassification, which could lead to false positives or negatives, impacting users unfairly. Additionally, there is a risk of over-reliance on automated systems, which might not capture nuanced contexts and could inadvertently suppress legitimate speech. Annotation errors can also induce bias (Sap et al., 2019), but as we used a previously peer-reviewed benchmark dataset, we hope those type of concerns are already addressed.

## Acknowledgements

I express my sincere gratitude to the Computational Intelligence and Operations Laboratory (CIOL) for

all their support. I also extend my appreciation to all the reviewers for their valuable suggestions to improve the work.

## References

- Ashwitha A, Shruthi G, Shruthi H R, Makarand Upadhyaya, Abhra Pratip Ray, and Manjunath T C. 2021. [Sarcasm detection in natural language processing](#). *Materials Today: Proceedings*, 37:3324–3331.
- Diah Ikawati Ayuningtias, Oikurema Purwati, and Pratiwi Retnaningdyah. 2022. Making muslim the enemy: A transitivity analysis on anti-islam hate speech. *Research on Humanities and Social Sciences*.
- Necva Bölücü and Pelin Canbay. 2021. [Hate speech and offensive content identification with graph convolutional networks](#). In *Fire*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Rui Cao, Roy Ka-Wei Lee, and Tuan-Anh Hoang. 2020. [Deepbate: Hate speech detection via multi-faceted text representations](#). *Proceedings of the 12th ACM Conference on Web Science*.
- Thomas Davidson, Dana Warmusley, Michael W. Macy, and Ingmar Weber. 2017. [Automated hate speech detection and the problem of offensive language](#). In *International Conference on Web and Social Media*.
- Gretel Liz De la Peña Sarracén and Paolo Rosso. 2022. Convolutional graph neural networks for hate speech detection in data-poor settings. In *Natural Language Processing and Information Systems*, pages 16–24, Cham. Springer International Publishing.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Javier Ferrando, Gerard I. Gállego, Ioannis Tsiamas, and Marta R. Costa-jussà. 2023. [Explaining how transformers use context to build predictions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5486–5513, Toronto, Canada. Association for Computational Linguistics.
- Zhangxiaowen Gong, Houxiang Ji, Yao Yao, Christopher W. Fletcher, Christopher J. Hughes, and Josep Torrellas. 2022. [Graphite: optimizing graph neural networks on cpus through cooperative software-hardware techniques](#). In *Proceedings of the 49th Annual International Symposium on Computer Architecture, ISCA '22*, page 916–931, New York, NY, USA. Association for Computing Machinery.
- William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 1025–1035, Red Hook, NY, USA. Curran Associates Inc.
- L. Hebert, L. Golab, and R. Cohen. 2022. [Predicting hateful discussions on reddit using graph transformer networks and communal context](#). In *2022 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, pages 9–17, Los Alamitos, CA, USA. IEEE Computer Society.
- Thomas N. Kipf and Max Welling. 2017. [Semi-supervised classification with graph convolutional networks](#). In *International Conference on Learning Representations*.
- György Kovács, Pedro Alonso, and Rajkumar Saini. 2021. [Challenges of hate speech detection in social media: Data scarcity, and leveraging external resources](#). *SN Computer Science*, 2(2).
- Jessica Lin. 2022. [Leveraging world knowledge in implicit hate speech detection](#). In *Proceedings of the Second Workshop on NLP for Positive Impact (NLP4PI)*, pages 31–39, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Zhibin Lu, Pan Du, and Jianyun Nie. 2020. [Vgcnbert: Augmenting bert with graph embedding for text classification](#). *Advances in Information Retrieval*, 12035:369 – 382.
- Anqi Mao, Mehryar Mohri, and Yutao Zhong. 2023. [Cross-entropy loss functions: Theoretical analysis and applications](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 23803–23828. PMLR.
- Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2020. [Hatexplain: A benchmark dataset for explainable hate speech detection](#). In *AAAI Conference on Artificial Intelligence*.

- Marzieh Mozafari, Reza Farahbakhsh, and Noël Crespi. 2020. [Hate speech detection and racial bias mitigation in social media based on bert model](#). *PLoS ONE*, 15.
- Ankit Pal, Muru Selvakumar, and Malaikannan Sankarabubbu. 2020. [Multi-label text classification using attention-based graph neural network](#). *ArXiv*, abs/2003.11644.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Anchal Rawat, Santosh Kumar, and Surender Singh Samant. 2024. [Hate speech detection in social media: Techniques, recent trends, and future challenges](#). *WIREs Computational Statistics*, 16(2).
- Koustuv Saha, Eshwar Chandrasekharan, and Munmun De Choudhury. 2019. [Prevalence and psychological effects of hateful speech in online college communities](#). In *Proceedings of the 10th ACM Conference on Web Science, WebSci '19*. ACM.
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. [The risk of racial bias in hate speech detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678, Florence, Italy. Association for Computational Linguistics.
- Simone Scardapane. 2024. [Alice’s adventures in a differentiable wonderland – volume i, a tour of the land](#).
- Anna Schmidt and Michael Wiegand. 2017. [A survey on hate speech detection using natural language processing](#). In *SocialNLP@EACL*.
- Mike Schuster and Kuldip Paliwal. 1997. [Bidirectional recurrent neural networks](#). *Signal Processing, IEEE Transactions on*, 45:2673 – 2681.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph attention networks](#).
- Lilapati Waikhom and Ripon Patgiri. 2021. [Graph neural networks: Methods, applications, and opportunities](#).
- Xinchen Wan, Kaiqiang Xu, Xudong Liao, Yilun Jin, Kai Chen, and Xin Jin. 2023. [Scalable and efficient full-graph gnn training for large graphs](#). *Proc. ACM Manag. Data*, 1(2).
- Azmine Toushik Wasi. 2024. [HRGraph: Leveraging LLMs for HR data knowledge graphs with information propagation-based job recommendation](#). In *Proceedings of the 1st Workshop on Knowledge Graphs and Large Language Models (KaLLM 2024)*, pages 56–62, Bangkok, Thailand. Association for Computational Linguistics.
- Azmine Toushik Wasi, Raima Islam, Mst Rafia Islam, Taki Hasan Rafi, and Dong-Kyu Chae. 2024a. [Exploring bengali religious dialect biases in large language models with evaluation perspectives](#).
- Azmine Toushik Wasi, Taki Hasan Rafi, and Dong-Kyu Chae. 2024b. [Diaframe: A framework for understanding bengali dialects in human-ai collaborative creative writing spaces](#). In *Companion of the 2024 Computer-Supported Cooperative Work and Social Computing (CSCW Companion '24)*, page 7, New York, NY, USA. ACM.
- Azmine Toushik Wasi, Taki Hasan Rafi, Raima Islam, and Dong-Kyu Chae. 2024c. [BanglaAutoKG: Automatic Bangla knowledge graph construction with semantic neural graph filtering](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2100–2106, Torino, Italia. ELRA and ICCL.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. [How powerful are graph neural networks?](#) In *International Conference on Learning Representations*.
- Rex Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. 2019. [GNNExplainer: generating explanations for graph neural networks](#). Curran Associates Inc., Red Hook, NY, USA.
- Yufeng Zhang, Xueli Yu, Zeyu Cui, Shu Wu, Zhongzheng Wen, and Liang Wang. 2020. [Every document owns its structure: Inductive text classification via graph neural networks](#). In *Annual Meeting of the Association for Computational Linguistics*.

## A Extended Explanation of Case Study

The task of detecting and classifying offensive content, especially hate speech, is inherently complex due to the nuanced and often implicit nature of such language. In the example provided—"How is all that awesome Muslim diversity going for you native Germans? You have allowed this yourselves. If you do not stand and fight against this. You get what you asked for, what you deserve!"—the model identified the post as offensive toward Islam. This classification was aided by analysing specific tokens and their relationships within the text using a Graph Neural Network (GNN) framework, particularly with the GNNExplainer (Ying et al., 2019) method.

As discussed above in Section 1, GNNs excel in tasks where the relationships between data points are as critical as the data points themselves. In the context of hate speech detection, GNNs can capture the intricate web of semantic and syntactic connections between words, phrases, and even larger text segments. This capability allows the model to consider not just isolated words but also the context in which they appear, making it particularly powerful for understanding language that may be implicitly biased or offensive.

In this example presented in Figure 4, the GNNExplainer was employed to determine which tokens—both in isolation and in combination with their neighbouring tokens—contributed to the model's decision to classify the post as offensive. The key tokens identified, such as "fight," "Muslim diversity," and "Syrian," are not inherently negative but, when analysed in context, reveal an underlying sentiment of hostility and blame. The phrase "stand and fight against this" suggests a confrontational stance, while the juxtaposition of "Muslim diversity" with a directive to "stand and fight" subtly frames the diversity as a threat. The mention of "Syrian" further adds to the narrative by invoking a specific group, which, in the context of the surrounding words, contributes to a negative sentiment.

GNN-based explainers are particularly effective because they allow us to visualize and interpret the model's decision-making process by highlighting the most influential tokens and their connections. This interpretability is crucial in sensitive applications like hate speech detection, where understanding why a model made a certain decision can help in refining the model, addressing potential biases,




Figure 4: Explanation Graph

and ensuring that it aligns with ethical guidelines. Moreover, by using a GNN-based approach, the model can weigh the significance of different parts of the text more effectively than traditional linear models. The graph structure allows the model to account for the interactions between words and their broader context, providing a more holistic understanding of the text. This is particularly important in hate speech detection, where context often determines whether a statement is offensive.