

# Cyberbullying Detection via Aggression-Enhanced Prompting

Aisha Saeid , Anu Sabu , Girish A. Koushik 

Ferrante Neri  & Diptesh Kanojia 

 NICE Research Group &  Institute for People-Centred AI,  
School of Computer Science & Electronic Engineering,  
University of Surrey, UK

{a.saeid,as05318,g.koushik,f.neri,d.kanojia}@surrey.ac.uk

## Abstract

Detecting cyberbullying on social media remains a critical challenge due to its subtle and varied expressions. This study investigates whether integrating aggression detection as an auxiliary task within a unified training framework can enhance the generalisation and performance of large language models (LLMs) in cyberbullying detection. Experiments are conducted on five aggression datasets and one cyberbullying dataset using instruction-tuned LLMs. We evaluated multiple strategies: zero-shot, few-shot, independent LoRA fine-tuning, and multi-task learning (MTL). Given the inconsistent results of MTL, we propose an enriched prompt pipeline approach in which aggression predictions are embedded into cyberbullying detection prompts to provide contextual augmentation. Preliminary results show that the enriched prompt pipeline consistently outperforms standard LoRA fine-tuning, indicating that aggression-informed context significantly boosts cyberbullying detection. This study highlights the potential of auxiliary tasks, such as aggression detection, to improve the generalisation of LLMs for safety-critical applications on social networks.

## 1 Introduction

Digital communication platforms and social media networks have fundamentally transformed human interaction, enabling expression across geographic and cultural divides. However, these platforms have also become hotspots for harmful behaviours such as cyberbullying and online harassment (Hinduja and Patchin, 2010; Kowalski et al., 2014). Cyberbullying, in particular, is a pressing issue due to its profound psychological and emotional impact on victims (Wang et al., 2019). Unlike traditional bullying, which could be confined to specific times and locations, cyberbullying victims can be targeted unchecked, as observed across multiple digital platforms, significantly impacting the young

adult populace (Slonje et al., 2013).

Detecting cyberbullying is challenging due to its linguistic complexity, ranging from overt aggression to subtle, passive-aggressive expressions that often elude conventional moderation systems with implicit targeted meaning (Xu et al., 2012). Traditional keyword-based and rule-driven detection methods struggle to adapt to evolving slang and hidden messaging patterns in social media (Waseem et al., 2017). Recent advances in Large Language Models (LLMs), such as the Gemma series (Gemma Team, 2024), have contributed to improvements in natural language processing (NLP), particularly in understanding context, sentiment, and linguistic subtleties. However, when applied to cyberbullying detection, these models may face challenges due to domain-specific vocabulary, implicit language use, and the scarcity of annotated data tailored to harmful online content (Yi et al., 2025).

This study investigates whether integrating aggression-labelled datasets with cyberbullying datasets within a unified training framework can enhance the generalisation and detection capabilities of large language models (LLMs). Moreover, we investigate architectural modifications, including targeted Transformer-heads adaptations (Houlsby et al., 2019), designed to focus model attention on aggression signals, thus improving the detection of both direct and indirect cyberbullying. A core challenge addressed here is the detection of nuanced cyberbullying, especially in low-resource settings with limited annotated data. Although LLMs are versatile, their performance decreases without task-specific supervision, particularly for subtle, context-dependent aggression (Chang et al., 2024). Cyberbullying-only datasets often lack coverage of the entire aggression spectrum, resulting in models that miss harmful content or generate excessive false positives, undermining moderation trust (Salminen et al., 2020b). Furthermore, con-

ventional attention mechanisms struggle to prioritise aggression signals amid noisy social media text (Khan et al., 2022).

Guided by the question of **how joint training in aggression and cyberbullying data can enhance LLM effectiveness**, our work *examines how Transformer-head modifications can optimise for cyberbullying prediction using aggression-related features*. We evaluated model performance in zero-shot, few-shot, and fine-tuning scenarios, and introduce a novel prompt enrichment pipeline that leverages aggression predictions to augment cyberbullying detection.

The remainder of this article is structured as follows. Section 2 reviews related work on cyberbullying detection, aggression modelling, and the adaptation of LLMs. Section 3 outlines the methodology, including datasets, experimental design, model architectures, and evaluation metrics. Section 4 presents comparative evaluations across zero-shot, few-shot, independent LoRA fine-tuning, and multi-task learning baselines. The discussion addresses the effectiveness of the model, ethical considerations, limitations, and directions for future research. Finally, Section 5 concludes the paper by summarising the key contributions and their implications for advancing research in cyberbullying detection, and outlines potential directions for future work to further enhance the effectiveness and generalisability of LLMs in this domain.

## 2 Related Work

Early cyberbullying detection systems relied on rule-based and lexicon-based methods, which struggled with ambiguous language and evolving slang, often producing false positives and missing actual abuse (Maity et al., 2023; Muminovic, 2025). Machine learning (ML) models such as Support Vector Machines (SVMs) improved detection by incorporating sentiment and linguistic features, but still required extensive feature engineering and failed to generalize well across platforms (Yi and Zubiaga, 2022). The advent of deep learning and LLMs like BERT, RoBERTa, and Gemma advanced the field by enabling better contextual understanding. Fine-tuned transformer models now demonstrate notable gains in precision and recall for cyberbullying detection using task-specific data (Gutiérrez-Batista et al., 2024; Philipo et al., 2025; Gemma Team, 2024).

**Role of Aggression in Cyberbullying** Recent studies show that aggression signals improve both the accuracy and interpretability of cyberbullying detection, helping models better identify hostile intent compared to general toxicity models. (Han et al., 2024). Researchers have emphasised that aggression manifests across a spectrum from overt insults to subtle forms such as sarcasm and passive aggression, which often remain undetected in models trained solely on cyberbullying datasets (Salmi nen et al., 2020a). Recent empirical work by (Qian et al., 2019) demonstrated that aggression-labelled datasets significantly improve models’ ability to disambiguate between harmful and non-harmful content in noisy social media environments.

**Zero-shot and Few-shot Learning** Parallel to the development of richer datasets, zero-shot and few-shot learning paradigms have emerged as compelling solutions to data scarcity challenges, enabling models to generalise to unseen tasks. Minimal supervision-based techniques have shown notable success in sentiment analysis and hate speech detection, where labelled data are often sparse or unevenly distributed (Brown et al., 2020; Schick and Schütze, 2021). Pre-trained LLMs such as Gemma are inherently suited to zero-shot transfer, yet their performance in cyberbullying detection without extensive fine-tuning remains inconsistent (Yi et al., 2025). The effectiveness of these models improves when few-shot examples incorporate domain-specific cues like aggression indicators, as they provide contextual anchors that guide model predictions toward the correct semantic space (Vidgen et al., 2019). However, most existing research does not systematically explore how multi-level aggression cues, when embedded in zero- and few-shot learning settings, can enhance cyberbullying detection.

Studies in (Zampieri et al., 2020; Ranasinghe and Zampieri, 2023) demonstrate that aggression detection models can identify overt and covert hostility with reasonable accuracy. Furthermore, prompt engineering techniques have recently emerged as powerful strategies for injecting auxiliary knowledge into LLM inputs, enhancing task-specific performance without extensive re-training (Zhou et al., 2022). However, limited research has explored using aggression detection outputs to enrich prompts specifically for improving cyberbullying detection, a gap that this study aims to address. Building upon these observations,

this study proposes an integrated framework that leverages aggression-informed prompt enrichment and parameter-efficient adaptations to enhance cyberbullying detection performance, which is detailed in the following sections.

### 3 Methodology

This section outlines the methodology used to evaluate the effectiveness of integrating aggression detection into the classification pipeline for enhancing cyberbullying detection using instruction-tuned LLMs. Our approach leverages standard classification output from the Low-rank adaptation (LoRA) adapter in a hierarchical manner for enriching cyberbullying detection.

Our analysis focuses on three fundamental questions: (i) how well LLMs generalise to cyberbullying detection under varying supervision levels, (ii) what role parameter-efficient fine-tuning (LoRA) plays in performance improvement, and (iii) whether auxiliary contextual signals, specifically *predicted aggression labels*, can enhance model reasoning and prediction accuracy. To answer these questions, we conduct a systematic evaluation using five publicly available aggression datasets and one cyberbullying dataset. The methodology comprises two main components: (1) A comparative baseline evaluation using zero-shot, few-shot, LoRA-based supervised fine-tuning, and joint multi-task learning (MTL), which leverages hard parameter sharing using joint loss weight updates to both adapters. (2) An enriched prompt pipeline (EPP) that incorporates predicted aggression labels into cyberbullying classification.

#### 3.1 Datasets

Table 1 presents the datasets used for the aggression (D1–D5) and cyberbullying (D6) detection tasks. These identifiers (D1–D6) are used throughout the paper to refer to each dataset. The aggression datasets were collected from social media platforms, including Twitter and online forums, and annotated for varying levels of aggressive behaviour (Rawat et al., 2023; Nafis et al., 2023; Kumar et al., 2018; Samghabadi et al., 2020; Kumar et al., 2020). The cyberbullying dataset (D6) was obtained from a publicly available Kaggle repository and consists of tweets annotated for various forms of online abuse (Fati et al., 2025).

#### 3.2 Comparative Evaluation: Baselines

The first phase of experimentation focuses on establishing baselines using different learning strategies that reflect varying degrees of task supervision and model adaptation complexity. These strategies were applied in a selection of instruction-tuned LLMs, each independently assessed to measure its response to different configurations.

**Zero-shot** In this setting, pre-trained instruction-tuned LLMs are evaluated directly on the cyberbullying and aggression detection tasks without task-specific fine-tuning.

**Few-shot** This setup introduces a limited set of manually labelled examples to train the model. The objective is to determine the extent to which minimal supervision improves classification, simulating real-world low-resource scenarios and testing the model’s rapid learning capacity.

**Supervised Fine-Tuning (SFT w/ LoRA)** LoRA is employed to adapt LLMs by injecting trainable low-rank matrices into existing model weights using SFT. Each LoRA is fine-tuned independently on the aggression and cyberbullying tasks. This approach offers a parameter-efficient method for task-specific tuning, ideal for scenarios where computational resources or data volume are constrained.

**Joint Multi-Task Learning (MTL)** This strategy fine-tunes the model concurrently on both tasks, aggression and cyberbullying detection, sharing representational layers to potentially capture generalisable features. It tests whether task co-learning facilitates transfer learning benefits or introduces representational conflicts. We leverage multi-task learning by updating model weights using the joint loss (sum) from different heads to both LoRA adapters, and training them with individual classification heads for each task<sup>1</sup>, using embeddings from the final LLM layer.

These baseline strategies are evaluated using four standard metrics: accuracy, precision, recall, and F1-score. Given the imbalanced, multi-class nature of the cyberbullying detection problem, F1-score is designated as the primary

<sup>1</sup>Transformer Heads Library

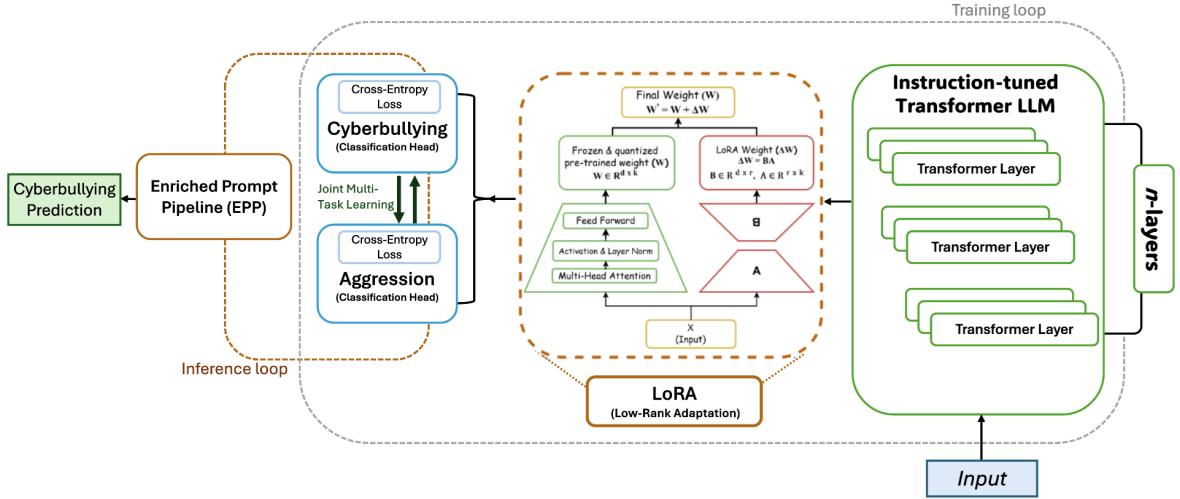


Figure 1: An overview of the proposed system architecture for cyberbullying detection. The diagram illustrates the training loop, which incorporates a LoRA adapter and a joint multi-task learning framework for aggression and cyberbullying, and the inference loop that utilises Enriched Prompt Pipeline (EPP).

performance measure. This baseline evaluation provides the foundational evidence to motivate more advanced integration strategies.

### 3.3 Enhanced Prompt-Pipeline (EPP)

To further enhance model understanding and classification accuracy, we propose a new methodology that augments model inputs for cyberbullying detection with contextual information derived from aggression classification. This pipeline is based on the hypothesis that aggression cues, *whether overt, covert, or absent, can provide auxiliary semantic signals that help LLMs better detect subtle or implied forms of cyberbullying*. The proposed pipeline is structured into three sequential steps:

**Step 1) Aggression Prediction** Each input text is first passed through an aggression detection model fine-tuned using LoRA. The model assigns the post to one of three categories: Overtly Aggressive (OAG), Covertly Aggressive (CAG), or Not-Aggressive (NAG). These predictions serve as summarised emotional cues, offering a linguistic profile of the text’s tone or intent.

**Step 2) Prompt Enrichment** The predicted aggression label is embedded directly into the classification prompt for cyberbullying detection. This enrichment mechanism is designed to introduce semantic context that primes the model’s attention to potentially harmful expressions. A typical enriched prompt follows the format:

*“This post was predicted as [Overtly Aggressive, Covertly Aggressive, Not-Aggressive]. Based on this, classify the following content for cyberbullying.”* This strategy transforms the task from isolated content classification into a contextual reasoning problem, where aggression becomes a guiding signal.

**Step 3) Cyberbullying Detection** The enriched prompt and the original user-generated post are then fed into a second LoRA-fine-tuned LLM trained for cyberbullying classification (as shown in Figure 1). By conditioning the model on the aggression label, we aim to improve its ability to detect indirect or context-dependent forms of cyberbullying.

### 3.4 Experiment Setup

To support the methodology described above, all experiments were carried out using three large language models tuned to instructions: Gemma-2-2B<sup>2</sup>, Gemma-2-9B<sup>3</sup>, and Gemma-3-4B<sup>4</sup>. These models were selected primarily to represent a range of parameter sizes and fine-tuning capabilities. For standard LoRA fine-tuning, each model was trained for 1 epoch, reflecting a lightweight and computationally efficient adaptation process. In the joint multi-task learning configuration, training duration was extended to between 3 and 6 epochs, allowing the models to jointly optimise across tasks. Training

<sup>2</sup>[google/gemma-2-2b-it](https://google/gemma-2-2b-it)

<sup>3</sup>[google/gemma-2-9b-it](https://google/gemma-2-9b-it)

<sup>4</sup>[unsloth/gemma-3-4b-it-unsloth-bnb-4bit](https://unsloth/gemma-3-4b-it-unsloth-bnb-4bit)

**Table 1:** Description of aggression and cyberbullying datasets

	<b>Aggression Datasets (D1-D5)</b>	<b>Cyberbullying Dataset (D6)</b>
<b>Dataset Name</b>	D1: Political Aggression Dataset D2: Code-Mixed Aggression and Offensive Language Dataset (TRAC - 2) D3: Aggression Identification D4: Trac - 2020 Submission D5: Aggression Identification	Cyberbullying Classification
<b>Number of instances</b>	D1: 1,999 samples D2: 10,000 samples D3: 12,246 samples D4: 6,000 samples D5: 10,000 samples	47,000 samples
<b>Categories</b>	- Not Aggressive (NAG) – Label: 0 - Covertly Aggressive (CAG) – Label: 1 - Overtly Aggressive (OAG) – Label: 2	- Ethnicity/Race - Religion - Gender/Sexual - Not Cyberbullying
<b>Language</b>	D1: English D2: Hindi-English (Code-Mixed) D3: English with Hindi transliteration D4: English, Hindi, Bengali D5: English	English
<b>Splits</b>	80% Training 10% Test 10% Validation	75% Training 25% Test and validation (with 2,000 instances in validation)

was performed with a batch size of 8 and a learning rate set to  $1e-4$ . The LoRA rank parameter was fixed at  $r = 8$ , enabling low-rank updates without full-scale retraining. Few-shot evaluations were conducted using  $k = 3$  examples per class, simulating limited supervision scenarios. Execution times varied depending on model size and architecture. Smaller models completed training in approximately 45 minutes, while larger models required up to 9 hours. All experiments were executed using two GPUs equipped with 24GB VRAM. Model evaluation employed four performance metrics: accuracy, precision, recall, and F1-score. Given the inherent class imbalance in cyberbullying datasets, the macro-averaged F1-score was treated as the primary performance indicator throughout all evaluation phases (Luo et al., 2025).

## 4 Results And Discussion

The selected approaches encompass a range of generalisation, supervision, and efficiency scenarios to provide a comprehensive evaluation: Zero-shot learning evaluates a model’s baseline capacity to perform a task with no task-specific supervision. Few-shot learning tests the model’s adaptability to small, annotated datasets. LoRA enables efficient adaptation by injecting trainable low-rank updates. MTL leverages shared representation learning across related tasks. The Enriched Prompt Pipeline introduces auxiliary contextual knowledge

to improve downstream task comprehension.

### 4.1 Zero-shot vs. Few-shot Performance

In the zero-shot setting, models are tested without task-specific examples to assess generalisation from instruction tuning alone. Few-shot learning introduces limited supervision to evaluate the minimum data required for meaningful performance gains. As shown in Table 2, few-shot learning consistently yields higher F1 scores than zero-shot learning across all three models. For aggression detection, the gemma-3-4b-it-unsloth-bnb-4bit model increases from an F1-score of 0.53 (zero-shot) to 0.63 (few-shot), while the Gemma-2-2B and Gemma-2-9B models show smaller increases from 0.54 to 0.56 and 0.57 to 0.60, respectively. These results suggest that models are better able to identify explicit hostility or implicit aggression patterns when supervised in a limited but focused manner. Cyberbullying detection also shows a substantial increase in F1 scores. As detailed in the table, the Gemma-2-2B model’s F1-score increases from 0.63 to 0.83, and Gemma-2-9B shows a similar increase from 0.79 to 0.83. This indicates that even a few examples help models capture context-dependent and often nuanced instances of cyberbullying, such as sarcasm, coded language, or indirect threats. The results presented in Table 2 thus emphasise the critical role of task-specific samples in enabling LLMs to understand and differentiate harmful content ef-

**Table 2:** *macro-F1* score comparison across models and methods for aggression and cyberbullying detection. The table groups results from zero-shot, few-shot, LoRA, Multi-Task Learning (MTL), and Enriched Prompt Pipeline (EPP) evaluations. **Bold** indicates the best performing method within a specific comparison for a given task, while the *italicized* value is noted for comparison against LoRA.

Model	Aggression Detection					Cyberbullying Detection				
	Zero-shot	Few-shot	LoRA	MTL	EPP	Zero-shot	Few-shot	LoRA	MTL	EPP
Gemma-2-2B	0.54	0.56	<b>0.67</b>	0.51	<b>0.67</b>	0.63	<b>0.83</b>	0.84	0.90	<b>0.99</b>
Gemma-2-9B	0.57	0.60	0.65	0.53	0.65	0.79	<b>0.83</b>	<b>0.93</b>	0.89	<b>0.99</b>
Gemma-3-4B	0.53	<b>0.63</b>	0.50	0.49	0.50	0.34	0.57	0.84	0.76	0.86

fectively.

#### 4.2 Few-shot vs. LoRA SFT

LoRA enables efficient task-specific fine-tuning by updating fewer parameters, making it a viable approach for resource-constrained settings. As detailed in Table 2, this approach results in a higher F1 score for aggression detection compared to few-shot learning for two of the models: for example, Gemma-2-2B increases from 0.56 to 0.67 and Gemma-2-9B from 0.60 to 0.65. However, gemma-3-4b-it-unsloth-bnb-4bit drops from 0.63 (few-shot) to 0.50 (LoRA), likely due to overfitting or misalignment in instruction-heavy or compact models.

In the cyberbullying task, LoRA fine-tuning also results in higher F1 scores for all models. Gemma-2-9B, for instance, shows a score increase from 0.83 (few-shot) to 0.93 (LoRA), while gemma-3-4b-it-unsloth-bnb-4bit score increases from 0.57 to 0.84. These results indicate that LoRA is beneficial for tasks involving rich semantic complexity, as it enables the models to refine their internal representations without extensive re-training. The results in Table 2 thus highlight the effectiveness of LoRA for adapting models to challenging, multi-dimensional language classification problems, especially in cases where few-shot learning reaches its performance ceiling.

#### 4.3 LoRA SFT vs. MTL

Multi-task learning (MTL) investigates whether jointly training on aggression and cyberbullying detection tasks enables models to learn more generalizable linguistic features, based on shared traits such as hostile intent and emotional intensity. As shown in Table 2, a comparison with LoRA reveals that MTL consistently yields a lower F1 score for aggression detection across all models. For ex-

ample, Gemma-2-2B shows a decrease from 0.67 (LoRA) to 0.51 (MTL), and Gemma-2-9B from 0.65 to 0.53. This drop in performance suggests that joint optimization may cause task interference, especially when there is a disparity in task difficulty or dataset size, potentially leading to a trade-off where one task improves at the expense of the other.

The results of cyberbullying detection under MTL are mixed. Gemma-2-2B shows an increase in its F1 score from 0.84 to 0.90, suggesting some benefit from joint learning. However, other models like gemma-3-4b-it-unsloth-bnb-4bit and Gemma-2-9B show a decrease in performance even with more epochs. For instance, Gemma-2-9B score decreases from 0.93 (LoRA) to 0.89 (MTL), as highlighted by the italicized value. These outcomes illustrate that MTL’s effectiveness is model- and task-dependent, requiring careful tuning to avoid negative transfer.

#### 4.4 LoRA SFT vs. Proposed EPP

We introduce the Enriched Prompt Pipeline (EPP) to address limitations of LoRA and MTL by enhancing cyberbullying detection through context engineering (aggression-informed prompts) by sequentially prompting adapters. This approach augments model input with predicted aggression labels, improving classification accuracy without increasing model complexity or training time. As shown in Table 2, EPP shows a substantial increase in F1 scores for cyberbullying detection: for instance, Gemma-2-2B increases from 0.84 (LoRA) to 0.99, and Gemma-2-9B also shows an increase from 0.93 to 0.99. Even the previously less stable gemma-3-4b-it-unsloth-bnb-4bit model shows an increase from 0.84 to 0.86.

As expected, the aggression results in Table 2 are identical for both the LoRA and EPP methods. This is because no retraining was performed on aggression tasks in the EPP setup. Instead, aggression

labels were solely used to enrich the cyberbullying prompts, aligning with the study’s main objective to enhance cyberbullying detection using auxiliary aggression signals.

## 5 Conclusion and Future Work

This study focused on improving cyberbullying detection using LLMs, with aggression detection leveraged as an auxiliary signal to enhance performance. We evaluated multiple learning paradigms, including zero-shot, few-shot, LoRA fine-tuning, joint MTL, and enriched prompting with three instruction-tuned LLMs. Our proposed enriched prompt pipeline, which incorporates aggression-related cues into model inputs, showed substantial gains, demonstrating its utility in improving model generalization. This work indicates that lightweight, context-sensitive prompt augmentation is a promising approach for tasks that require socially sensitive NLP.

Our future work will focus on assessing a wider range of instruction-tuned LLMs, including models such as Llama and Mistral. Although we encountered technical challenges related to accessing and deploying these models during this study, we plan to address these in future experiments.

## References

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language models are few-shot learners](#). *arXiv preprint arXiv:2005.14165*.

Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kajie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, and Yidong Wang. 2024. [A survey on evaluation of large language models](#). *ACM Transactions on Intelligent Systems and Technology*, 15(3):1–45.

Suliman Mohamed Fati, Mohammed A Mahdi, Mohamed AG Hazber, Shahanawaj Ahamad, Sawsan A Saad, Mohammed Gamal Ragab, and Mohammed Al-Shalabi. 2025. Enhancing multi-class cyberbullying classification with hybrid feature extraction and transformer-based models. *Computer Modeling in Engineering & Sciences (CMES)*, 143(2).

Gemma Team. 2024. Gemma: Open models for responsible ai. <https://ai.google.dev/gemma>. Google DeepMind.

Karel Gutiérrez-Batista, Jesica Gómez-Sánchez, and Carlos Fernandez-Basso. 2024. Improving automatic cyberbullying detection in social network environments by fine-tuning a pre-trained sentence transformer language model. *Social Network Analysis and Mining*, 14(1):136.

Huimin Han, Muhammad Asif, Emad Mahrous Awwad, Nadia Sarhan, Yazeed Yasid Ghadi, and Bo Xu. 2024. Innovative deep learning techniques for monitoring aggressive behavior in social media posts. *Journal of Cloud Computing*, 13(1):19.

Sameer Hinduja and Justin W. Patchin. 2010. [Bullying, cyberbullying, and suicide](#). *Archives of Suicide Research*, 14(3):206–221.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for nlp](#). *Preprint*, arXiv:1902.00751.

Umair Khan, Salabat Khan, Atif Rizwan, Ghada Atteia, Mona M. Jamjoom, and Nagwan Abdel Samee. 2022. [Aggression detection in social media from textual data using deep learning models](#). *Applied Sciences*, 12(10):5083.

Robin M. Kowalski, Susan P. Limber, and Patricia W. Agatston. 2014. Bullying in the digital age: A critical review and meta-analysis of cyberbullying research among youth. *Psychological Bulletin*, 140(4):1073–1137. Available at: <https://pubmed.ncbi.nlm.nih.gov/24512111/>.

Ritesh Kumar, Atul Kr Ojha, Shervin Malmasi, and Marcos Zampieri. 2018. Benchmarking aggression identification in social media. In *Proceedings of the first workshop on trolling, aggression and cyberbullying (TRAC-2018)*, pages 1–11.

Ritesh Kumar, Atul Kr Ojha, Shervin Malmasi, and Marcos Zampieri. 2020. Evaluating aggression identification in social media. In *Proceedings of the second workshop on trolling, aggression and cyberbullying*, pages 1–5.

Kai Luo, Ce Zheng, and Zhenyu Guan. 2025. Reinforced multi-modal cyberbullying detection with subgraph neural networks. *International Journal of Machine Learning and Cybernetics*, 16(3):2161–2180.

Krishnan Maity, Prince Jha, Raghav Jain, Sriparna Saha, and Pushpak Bhattacharyya. 2023. “explain thyself bully”: Sentiment aided cyberbullying detection with explanation. In *International Conference on Document Analysis and Recognition*, pages 132–148. Springer.

Amel Muminovic. 2025. [Moderating harm: Benchmarking large language models for cyberbullying detection in youtube comments](#). *International Journal of Computer Applications*, 187(25):1–9.

Nazia Nafis, Diptesh Kanodia, Naveen Saini, and Rudra Murthy. 2023. Towards safer communities: Detecting aggression and offensive language in code-mixed tweets to combat cyberbullying. In *The 7th Workshop on Online Abuse and Harms (WOAH)*, pages 29–41.

Adamu Gaston Philipo, Doreen Sebastian Sarwatt, Jianguo Ding, Mahmoud Daneshmand, and Huansheng Ning. 2025. Assessing text classification methods for cyberbullying detection on social media platforms. *IEEE Transactions on Information Forensics and Security*.

Jing Qian, Anna Bethke, Yinyin Liu, Elizabeth Belding, and William Yang Wang. 2019. [A benchmark dataset for learning to intervene in online hate speech](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4755–4764, Hong Kong, China. Association for Computational Linguistics.

Tharindu Ranasinghe and Marcos Zampieri. 2023. [A text-to-text model for multilingual offensive language identification](#). In *Findings of the Association for Computational Linguistics: IJCNLP-AACL 2023*. Accepted to Findings of IJCNLP-AACL 2023.

Akash Rawat, Nazia Nafis, Dnyaneshwar Bhadane, Diptesh Kanodia, and Rudra Murthy. 2023. Modelling political aggression on social media platforms. In *Proceedings of the 13th workshop on computational approaches to subjectivity, sentiment, & social media analysis*, pages 497–510.

Joni Salminen, Matthias Hopf, Sami A. Chowdhury, Soon-Gyo Jung, Bernard J. Jansen, and Haewoon Kwak. 2020a. [Developing an online hate classifier for multiple social media platforms](#). *Human-centric Computing and Information Sciences*, 10(1):1.

Joni Salminen, Matthias Hopf, Sami A. Chowdhury, Milica Milenovic, Soon gyo Jung, Haewoon Kwak, and Bernard J. Jansen. 2020b. [Developing an online hate classifier for multiple social media platforms](#). *Human-centric Computing and Information Sciences*, 10(1):1.

Nilofar Safi Samghabadi, Parth Patwa, Srinivas Pykl, Prerana Mukherjee, Amitava Das, and Thamar Solorio. 2020. Aggression and misogyny detection using bert: A multi-task approach. In *Proceedings of the second workshop on trolling, aggression and cyberbullying*, pages 126–131.

Timo Schick and Hinrich Schütze. 2021. [Exploiting cloze questions for few shot text classification and natural language inference](#). *arXiv preprint arXiv:2001.07676*.

Robert Slonje, Peter K. Smith, and Ann Frisén. 2013. [The nature of cyberbullying, and strategies for prevention](#). *Computers in Human Behavior*, 29(1):26–32.

Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019. Challenges and frontiers in abusive content detection. In *Proceedings of the third workshop on abusive language online*. Association for Computational Linguistics.

Peng Wang, Xinqiang Wang, Xiaoyan Yang, and Li Lei. 2019. [Affective and stress consequences of cyberbullying: The mediating role of social anxiety](#). *Symmetry*, 12(9):1536.

Zeerak Waseem, Thomas Davidson, Dana Warmsley, and Ingmar Weber. 2017. [Understanding abuse: A typology of abusive language detection subtasks](#). *Preprint*, arXiv:1705.09899. To appear in the proceedings of the 1st Workshop on Abusive Language Online.

Jun-Ming Xu, Kwang-Sung Jun, Xiaojin Zhu, and Amy Bellmore. 2012. Learning from bullying traces in social media. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 656–666.

Peiling Yi and Arkaitz Zubiaga. 2022. Cyberbullying detection across social media platforms via platform-aware adversarial encoding. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, pages 1430–1434.

Peiling Yi, Arkaitz Zubiaga, and Yunfei Long. 2025. [Detecting harassment and defamation in cyberbullying with emotion-adaptive training](#). *arXiv preprint arXiv:2501.16925*.

Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. [Semeval-2020 task 12: Multilingual offensive language identification in social media \(offenseval 2020\)](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1425–1447, Barcelona (online). International Committee for Computational Linguistics.

Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2022. [Large language models are human-level prompt engineers](#). *arXiv preprint arXiv:2211.01910*.