Sentiment and sarcasm: Analyzing gender bias in sports through social media with deep learning

Sethulakshmi Praveen
CSE Department
IIIT Sri City
AP, India
sethu7praveen@gmail.com

Balaji TK
CSE Department
IIIT Sri City
AP, India
balaji.tk@iiits.in

Sreeja S R
CSE Department
IIIT Sri City
AP, India
sreeja.sr@iiits.in

Annushree Bablani
CSE Department
IIIT Sri City
AP, India
annushree.bablani@iiits.in

Abstract

Gender bias continues to be a pervasive issue, especially in public discourse surrounding high-profile events like the Olympics. Social media platforms, particularly Twitter, have become a central space for discussing such biases, making it crucial to analyze these conversations to better understand public attitudes. Sentiment analysis plays a key role in this effort by determining how people feel about gender bias. However, sarcasm often complicates sentiment analysis by distorting the true sentiment of a tweet, as sarcastic expressions can mask negative or positive sentiments. To address this, the study introduces a novel framework called SENSA (SENtiment and Sarcasm Analysis), designed to detect both sentiment and sarcasm in tweets related to gender bias. The framework leverages the R2B-CNN model for robust sarcasm and sentiment classification. Using approximately 5,000 tweets related to gender bias from 2010 to August 30, 2024, SENSA applies advanced sarcasm detection to account for shifts in sentiment caused by sarcastic remarks. The R2B-CNN model demonstrates high accuracy of 92.32% along with achieving 92.75% precision and 92.53% F1-score for sarcasm detection, and a 93.67% accuracy, 92.33% precision and 92.33% F1-score for sentiment classification. SENSA provides a comprehensive understanding of gender bias discussions on social media by capturing both sentiment and sarcasm to reveal deeper insights into public perceptions.

1 Introduction

Gender bias refers to the unequal treatment or perception of individuals based on their biological gender, often resulting in unfair advantages or disadvantages (Heise et al., 2019). People encounter this bias in various aspects of daily life, from being stereotyped to facing real inequality and discrimination. For example, nursing is often viewed as a female-dominated field, and men who pursue this profession may be labeled as “not man enough". Conversely, mechanical engineering is typically male-dominated, and women in the field may face negative stereotypes (Ohse, 2014). Such biases can manifest in trivial comments like calling someone a “girl" for not performing well in sports or in more serious issues such as wage inequality.

The effects of gender bias can be deeply personal, impacting individuals’ self-worth and opportunities. It can result in people being unable to advance in their careers or have their achievements ignored simply because of their gender (Herrmann et al., 2019). This bias extends beyond just male and female, transgender and non-binary individuals also face these challenges. The consequences of gender bias are damaging to society as it fosters an environment of inequality and inferiority, leading to lost opportunities and creating unnecessary divisions.

Social media platforms, especially Twitter, amplify these issues, providing a space for people to express their thoughts, often anonymously (Oz et al., 2024). Gender bias discussions often come to light on social media because people feel more comfortable addressing such sensitive issues behind a screen. This makes social media an essential platform for gathering data on gender bias, as users frequently express opinions and frustrations about these topics.

Sarcasm also plays a significant role in online discourse about gender bias. Sarcasm, where the literal meaning of words contrasts with the intended message, often complicates the sentiment of a statement (Strozzo, 2023). For example, a tweet like “Oh great, another woman in engineering!" might seem neutral or positive, but the underlying sarcasm conveys negativity. Twitter is rife with sarcastic remarks, especially around social issues like gender bias, making it an excellent platform for analyzing public sentiment.

The 2024 Paris Olympics highlighted gender biases, particularly around the misgendering of contestants, following the incident involving a boxer Imane Khalef (IndiaToday, 2024), which sparked debates on social media. Many users voiced differing opinions, with some supporting the athlete and others questioning her gender identity. This situation highlights the need to analyze how social media conversations reveal attitudes toward gender bias, motivating the study to explore these dynamics in depth. The objective of this study is to explore public sentiment and the role of sarcasm in these discussions, particularly in the context of high-profile events like the Olympics. To address this, a novel framework, SENSA (SENtiment and Sarcasm Analysis), is proposed, incorporating the R2B-CNN model for sarcasm and sentiment analysis. This framework effectively analyzes tweets related to gender bias, focusing on sarcasm and sentiment to uncover underlying public attitudes. The analysis provides insights into how gender biases are discussed and perceived, contributing to a deeper understanding of public opinion and supporting efforts toward more inclusive and equal treatment across all genders.

The contributions of this study are:

- A novel framework, SENSA, is proposed for analysing gender bias in sports.
- Collected and curated tweets related to gender bias for analysis.
- Introduced the R2B-CNN model for combined sentiment and sarcasm analysis within the framework.
- Presented the in-depth analysis of sarcasm and sentiment over gender bias.

The paper is structured as follows: Section 2, reviews the literature, Section 3 explains the methodology of this study, Section 4 discusses the results, and Section 5 concludes the study.

2 Literature Survey

In recent years, several studies have explored the analysis of gender bias and sentiment across various domains on social media platforms. The study (Jain and Mondal, 2024) examines interactions between Indian politicians and journalists on Twitter, with a focus on gender-based differences. Using a gender dataset (Panda et al., 2020), different politician accounts are identified and cross-matched with the MyNeta platform. The interaction data is categorized based on gender combinations between politicians and journalists. Findings indicated a significant gender bias, with male politicians receiving more mentions and their tweets gaining more popularity compared to female politicians. The study further noted that male politicians consistently garnered more attention in terms of follower count, although no significant emotional differences are found across gender categories through emotion analysis.

In the study (Mertens et al., 2019), they talk largely about gender bias in political interactions on digital platforms by considering how politicians present themselves on Twitter and how they are approached by others. The study (Aalberg and Jenssen, 2007) also revealed that gender stereotypes lead to changes in the assessment of politicians, their communication and party support. The authors analyzed gender differences in Twitter communication with politicians, using sentiment analysis and topic modelling on the Twitter data (Stier et al., 2018). A sentiment measure based on the Lexicoder Sentiment Dictionary (LSD) shows how positive or negative words vary across tweets, while a second measure quantifies personal versus professional language using LIWC dictionaries. Results reveal that female politicians receive more personal and emotional tweets compared to male counterparts, especially in conservative parties. However, sentiment in politicians' tweets largely aligns with party identity rather than gender. Structural Topic Modeling (STM) further confirms significant differences in the communication directed toward male and female politicians.

In (Betti et al., 2023), authors analysed gender bias in English songs. The study exploits WEAT (Caliskan et al., 2017), SC-WEAT (Charlesworth et al., 2021) and SWEAT (Bianchi et al., 2021). Authors in (Betti et al., 2023) fine-tuned BERT classifier on WASABI dataset (Meseguer-Brocal et al., 2017) for sexist content detection and also analyzed language bias using word embeddings. The study's results revealed significant gender-based biases in song lyrics. Male solo artists are found to produce a higher proportion of sexist content compared to female solo artists or groups, with the BERT classifier detecting sexist language patterns. The Word Embedding Association Tests (WEAT, SC-WEAT, and SWEAT) further confirmed that male artists'

Table 1: Literature Survey

Study	Description	Methodology	Results	Dataset
(Jain and Mondal, 2024)	Examines interactions between Indian politicians and journalists on Twitter, focusing on gender-based differences.	Genderize.io and MyNeta platform.	Male politicians received more mentions and popularity in terms of followers compared to female politicians.	gender dataset (Panda et al., 2020)
(Mertens et al., 2019)	Analyzes gender bias in political interactions on Twitter.	Lexicoder Sentiment Dictionary (LSD) and LIWC dictionaries. Structural Topic Modeling (STM)	Female politicians received more personal and emotional tweets, especially from conservative parties.	Twitter data (Stier et al., 2018)
(Betti et al., 2023)	Focuses on sexism in English song lyrics using machine learning and word embeddings.	BERT.	Male solo artists showed higher proportions of sexist content compared to female artists.	WASABI dataset (Meseguer-Brocal et al., 2017)
(Kalra and Zubiaga, 2021)	Identifies sexism in tweets and Gabs using deep neural networks.	BERT with MultiCNN.	accuracy - 0.749.	Twitter

lyrics are more strongly associated with negative gender stereotypes. Table 1 gives a brief overlook of the discussed papers.

While these studies contribute to understanding gender bias in various contexts, a notable research gap exists in the analysis of gender bias in sports, specifically during major events like the Olympics. Sentiment analysis of tweets related to gender bias, particularly with sarcasm, has been largely overlooked. Addressing this gap is crucial to avoid reinforcing harmful stereotypes and fostering a more gender-equal society. To meet this challenge, the novel SENSA framework is introduced, integrating advanced sarcasm and sentiment analysis techniques to explore these discussions in depth, particularly in the context of high-profile sporting events.

3 Methodology

The study proposes a novel SENSA framework designed to analyze sentiment shifts concerning sarcasm in user tweets related to gender bias. This framework processes tweets by first normalizing them through preprocessing. After preprocessing, the tweets are fed into a sentiment model to determine the expressed sentiment. Subsequently, the tweets are passed to a sarcasm detection model to assess the presence of sarcasm. If sarcasm is detected, the sentiment undergoes a shift; for instance, a positive sentiment is adjusted to negative, and a negative sentiment is transformed to positive. Finally, the overall sentiment is analyzed after accounting for any shifts caused by sarcasm. The graphical abstract of this proposed framework is presented in Figure 1.

3.1 Dataset

The study uses the data collected from Twitter using the Twitter Selenium Scraper tool [1]. The search queries focus on keywords such as “misgendering Olympics", “equality Olympics", “gender roles Olympics", “transgender Paris Olympics", and “gender bias Olympics 2024". The resulting gender bias tweet dataset includes features like the Name, Handle, Verified, Content, Timestamp, Comments, Retweets, Likes, Analytics, Hashtags, Mentions, Emojis, Profile Image, Tweet URL, and

[1] https://github.com/godkingjay/selenium-twitter-scraper

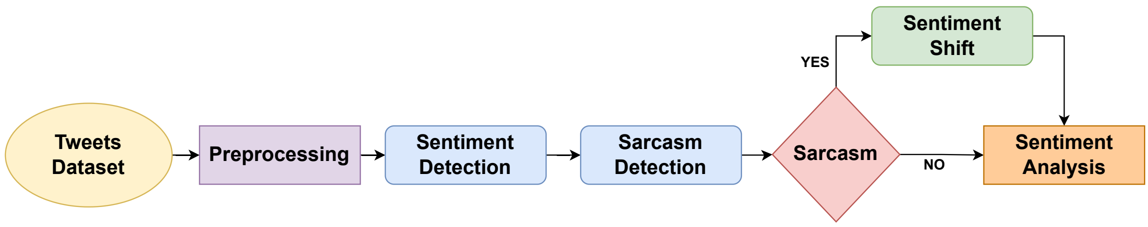

Figure 1: The graphical abstract of the SENSA framework.

Tweet ID.

To train the model for sarcasm and sentiment analysis, two additional datasets are utilized due to the lack of labeled datasets specifically addressing gender bias in this context. The sarcasm dataset (Handoyo et al., 2021) is employed to train the sarcasm detection model. This dataset includes 19,548 entries labeled as sarcastic (1) and 19,576 labeled as non-sarcastic (0). Additionally, a Kaggle sentiment dataset [2] is used for fine-tuning sentiment classification. The Kaggle sentiment dataset comprises 17,409 entries for each class, including Positive, Negative, and Neutral. Both datasets are split into training and testing sets in an 80:20 ratio, with 80% allocated for training and 20% for testing.

3.2 Preprocessing

Data preprocessing is essential, as the datasets used for training and testing may contain unwanted content that could negatively impact the model's performance. The preprocessing steps for this study involve removing unicode characters, duplicates, user mentions (@), and hashtags. Duplicates are eliminated from the test data based on the Tweet_ID, user handle, and content. Since emojis significantly contribute to understanding the sarcasm in a sentence, they are replaced with textual representations using Python's emoji library [3] to facilitate better processing by the model.

3.3 Implementation of R2B-CNN model

The proposed framework utilizes the R2B-CNN model for identifying sentiment and sarcasm in the tweet. The R2B-CNN model takes RoBERTa embeddings as input, which are passed through a RoBERTa layer followed by two Bi-GRU layers and a convolutional layer. The final layer is a Dense layer that outputs whether the input is sarcastic or not for the sarcasm analysis model. The R2B-CNN sentiment analysis model follows the same architecture, with the only difference being in the final Dense layer, which classifies the input as Positive, Negative, or Neutral. The architectural representation of the R2B-CNN model is presented in Figure 2.

The RoBERTa model in the R2B-CNN model excels at capturing contextual information within the text, providing a solid foundation for understanding general language. When this is layered with a Bi-GRU (Bidirectional Gated Recurrent Unit), which processes both past and future tokens, the model enhances its ability to capture sequential dependencies. This, when again stacked with another BiGRU layer, adds the sequential understanding of the model, allowing it to learn more deep patterns. Further incorporating a convolutional layer helps in identifying local patterns and n-grams through convolutional filters. This hybrid approach thus combines RoBERTa's comprehensive language understanding with BiGRU's refined sequential context and CNN's extraction of localized features. As a result, the model achieves improved feature representation and classification accuracy. By integrating these different processing techniques, the model becomes more flexible and robust, capable of adapting to various types of textual data and capturing diverse patterns that a single model type might miss.

The layered architecture of fine-tuned R2B-CNN starts by tokenizing data using the RoBERTa tokenizer. The tokenizer outputs two values, namely input IDs and an attention mask. The tokenizer also sets all fields to a particular maximum length by padding if the length is shorter or truncating if the length is longer. Finally, the tokenizer outputs the data in the form of tensors, which are passed through the RoBERTa main layer with default dimensions, producing a pooled output of 768 dimensions. The output is then processed through two BiGRU layers, each with 768 dimensions. This

[2] https://www.kaggle.com/datasets/jp797498e/twitter-entity-sentiment-analysis

[3] https://pypi.org/project/emoji/

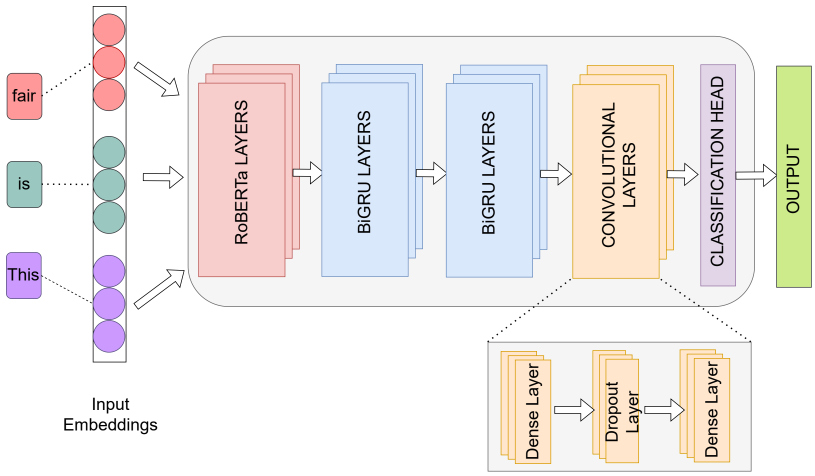

Figure 2: The R2B-CNN Architecture.

is followed by a dense layer with a ReLU activation function and a dropout layer. Another dense layer with a ReLU activation function processes the data before it reaches the classification head. The classification head uses a softmax activation function to output probabilities ranging from 0 to 1 for sarcasm model. The argmax function is applied to determine the binary classification as either 0 or 1, representing non-sarcastic and sarcastic, respectively. The model operates in the same way for sentiment classification. However, the final classification head outputs three probabilities. These probabilities can be argmaxed to determine the classification as 0, 1, or 2, where 0 represents Negative, 1 represents Neutral, and 2 represents Positive. Both models are set to run for 100 epochs with an early stopping patience of 5. The input is fed in batches with batch size 64. Adam optimizer is used with 5e-5 learning rate and Categorical Crossentropy as the loss function. Thus this novel method is used to analyze sarcasm and sentiment in the collected tweets on gender bias. The hyperparameters used to train the proposed R2B-CNN model are presented in Table 2.

Table 2: Parameters used in training the DL models.

Parameter	Final Model
Batch size	64
Activation	ReLU*/softmax**
Optimizer	Adam
Learning rate	5e-5
Epochs	100
Early stopping	5
Loss function	Categorical Crossentropy

*ReLU activation function is used in the dense layers.
**Softmax activation function is used in the classification layer.

Table 3: Performance analysis of sarcasm detection models

Model	Precision-micro	Recall-micro	F1-score	Accuracy
R-GR	92.70%	92.30%	92.02%	92.16%
R-2GR	92.30%	92.10%	92.01%	92.04%
R-2BGR	92.30%	92.10%	92.01%	92.04%
R-BGR-3D	92.30%	91.60%	92.00%	91.93%
R-BGR	92.60%	92.00%	91.90%	92.02%
R-3D	92.18%	**92.34%**	92.22%	92.02%
R2B-CNN	**92.75%**	92.31%	**92.53%**	**92.32%**

4 Result and Analysis

The performance of the R2B-CNN model for sarcasm and sentiment detection is evaluated against several baseline models. For sarcasm detection, the comparison includes R-GR (RoBERTa-GRU) (Tan et al., 2023), R-2GR (RoBERTa-2xGRU), R-2BGR (RoBERTa-2xBiGRU) (Li et al., 2021), R-BGR-3D (RoBERTa-BiGRU-CNN) (Kamal et al., 2023), R-BGR (RoBERTa-BiGRU) (Xie et al., 2024; Fan et al., 2022), and R-3D (RoBERTa-CNN) (Kumar and Reddy, 2023). In the sentiment detection category, the models considered are R-BGR-3D (Kamal et al., 2023) and R-3D (Kumar and Reddy, 2023). This evaluation offers valuable insights into the strengths and weaknesses of each model in detecting sarcasm and sentiment. The metrics considered for this comparison analysis are F1-score, precision, recall, and accuracy.

The proposed R2B-CNN model outperformed its baseline models by achieving the highest precision at 92.75% and an F1-score of 92.53% on doing sarcasm classification. This indicates its effectiveness in accurately identifying sarcastic content. R-GR achieved a precision of 92.70%, closely following R2B-CNN, but its F1-score is lower at

Table 4: Performance analysis of sentiment detection models

Model	Precision-micro	Recall-micro	F1-score	Accuracy
R-GR	91.08%	91.00%	91.09%	91.18%
R-2GR	91.18%	90.99%	90.79%	90.39%
R-BGR	91.01%	91.00%	91.03%	91.09%
R-2BGR	90.31%	90.39%	90.30%	90.43%
R-3D	84.98%	85.00%	84.99%	84.36%
R-BGR-3D	92.32%	92.00%	92.33%	91.94%
R2B-CNN	**92.33%**	**92.32%**	**92.33%**	**93.67%**

92.02%. Other models like R-2GR and R-2BGR recorded similar performances with precision values of 92.30% and F1-scores around 92.01%. R-BGR-3D and R-3D also showed competitive results, but their overall metrics are slightly lower than those of R2B-CNN, indicating that the latter provides a superior capability for sarcasm detection in tweets related to gender bias. Table 3 showcases the performance metrics of different sarcasm detection models.

The performance analysis of models trained for sentiment analysis is presented in Table 4. The R2B-CNN model excelled with an accuracy of 93.67%, making it the top-performing model in this category. Its precision and F1-score both recorded at 92.33%, indicating a strong balance between precision and recall. The R-BGR-3D model followed closely, achieving 92.32% precision and 91.94% accuracy. The R-GR model demonstrated solid performance with a precision of 91.08%, a recall of 91.00%, and an accuracy of 91.18%. In contrast, the R-2GR model showed slightly lower results with a precision of 91.18%, a recall of 90.99%, and an accuracy of 90.39%. The R-BGR model maintained a precision of 91.01%, a recall of 91.00%, and an accuracy of 91.09%. The R-2BGR model recorded a precision of 90.31%, a recall of 90.39%, and an accuracy of 90.43%. Lastly, the R-3D model lagged behind with an accuracy of only 84.36%, accompanied by a precision of 84.98% and a recall of 85%.

Overall, the results highlight that the R2B-CNN model stands out in sarcasm detection, while the R2B-CNN excels in sentiment classification. This comparative analysis underscores the effectiveness of specialized models tailored for specific tasks, enhancing the understanding of public sentiment and sarcasm in discussions of gender bias in the context of high-profile events like the Olympics. The ROC vs Precision-recall curve is presented in Figure 3a and 3b.

The collected tweets are tested with the proposed model R2B-CNN to see the sarcasm trends in the data. When analyzing the sarcasm trend Figure 4 unveils a fascinating trend in sarcasm during the year 2024, particularly highlighting the surge in sarcastic comments as the 2024 Paris Olympics unfolded. Notably, sarcasm reached its peak after July 2024, with August recording a staggering 43 sarcastic tweets. This trend reveals an intriguing phenomenon, as the Olympic Games progressed, users expressed more sarcastic sentiments, suggesting that the excitement and intensity of the event fueled a wave of irony and critique among spectators, especially towards the end of the games.

To be exact, there are 630 sarcastic comments and 4352 non-sarcastic comments. Some of the sample texts are classified as sarcastic and non-sarcastic as given in Table 5. Here, 1 indicates sarcastic and 0 indicates non-sarcastic, respectively. From this a tweet such as “Ah, the great Olympic debate of 2024! Here’s a thought: if we’re this riled up about who punches whom for a shiny medal, imagine the uproar if we started questioning the gender of the pigeons in pigeon racing. “Is that a he-pigeon or a she-pigeon¿‘ is initially labeled as positive during sentiment analysis. However, after sarcasm analysis, it is determined that the sentence is sarcastic. Thus, instead of labeling it as positive, the correct sentiment label would be negative. Hence sentiment shift is required when the tweet is labeled as sarcastic.

The proposed R2B-CNN sentiment model categorized a total of 651 tweets as Positive, while around 1,210 as Negative. The analysis also revealed that out of the collected dataset labelled Positive or Negative after sentiment analysis, 1,657 tweets are classified as non-sarcastic, whereas 204 are identified as sarcastic. This sharp contrast in sentiment is compelling, after accounting for sarcasm, the final counts showed 743 tweets classified as positive and 1,118 as negative. This shift underscores the profound impact of sarcasm on public sentiment on gender bias, particularly in events like the Olympics. The graphical representation of this sentiment shift is presented in Figure 5.

Hence, the analyses of both sarcasm and sentiment shifts reveal a prevailing negative sentiment among users concerning gender bias in the Olympics. Even when accounting for instances where positive sentiments are initially misclassified as negative due to the sarcastic nature of certain statements, it is evident that the overall sentiment

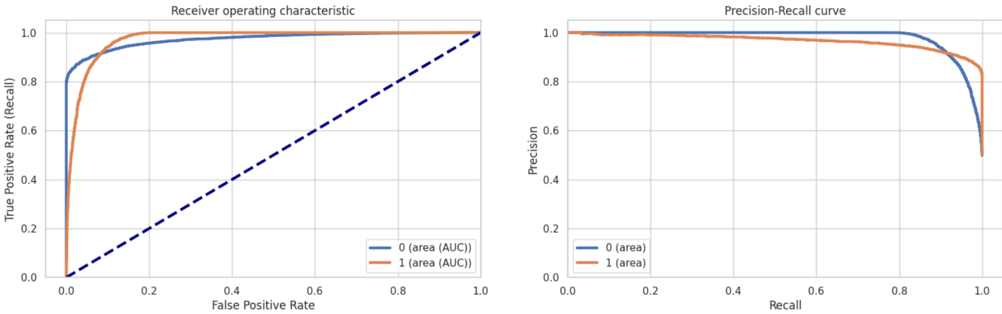

(a)

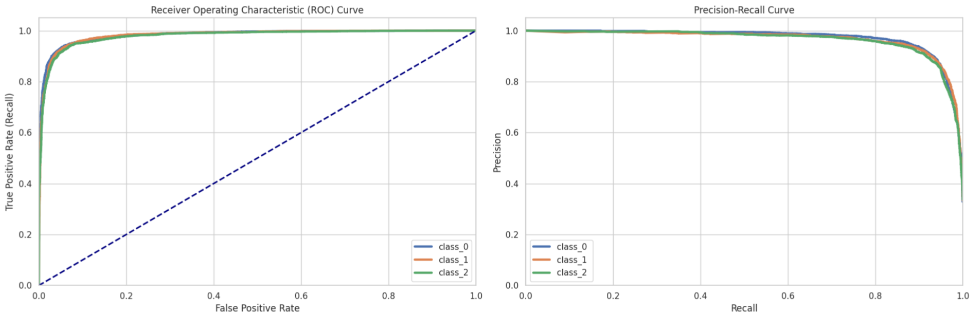

(b)

Figure 3: Comparison of ROC and Precision-Recall curves for R2B-CNN model on a)Sarcasm and b)Sentiment tasks.

Table 5: Sample tweets classified by the proposed R2B-CNN sarcasm model.

Tweets	**Class**
At the 2024 Paris Olympic Games, Team China featured 405 athletes, with 269 women making up 66 percent of the team. As we celebrate the first-ever gender parity in Olympic history, we #ChangeTheGame through more #WomenInSport and act together for #GenderEquality.	0
Paris 2024 achieves Gender Equality at the Olympics	0
Ah, the great Olympic debate of 2024! Here's a thought: if we're this riled up about who punches whom for a shiny medal, imagine the uproar if we started questioning the gender of the pigeons in pigeon racing. 'Is that a he-pigeon or a she-pigeon?'	1
Paris Olympics 2024 proved us all that as a female boxer in Olympics, u need to artificially increase your testosterone levels in order to have fair competition with opponents like imane khalif who are naturally blessed with high level of testosterone(or imane increased it art..)	1

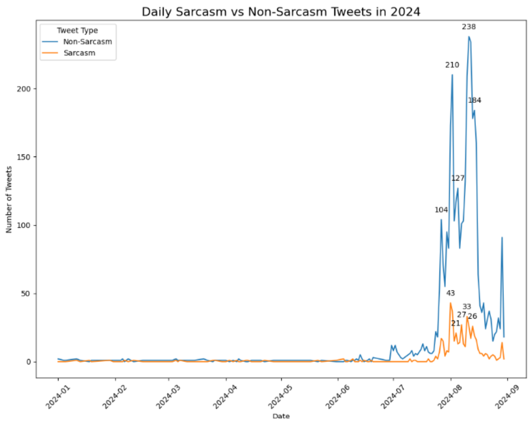

Figure 4: Line graph of sarcastic tweets from January to September 2024.

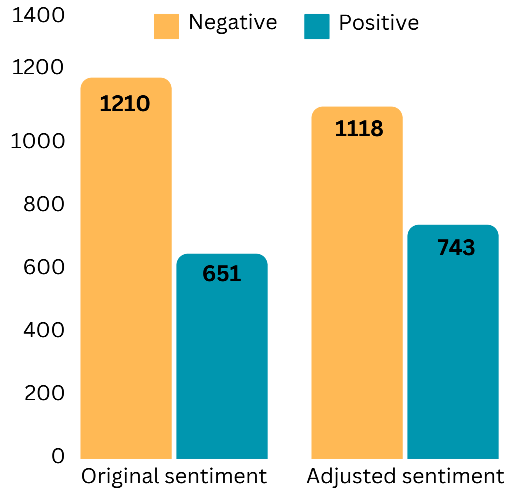

Figure 5: Graphical representations of sentiment shifts concerning sarcasm in tweets.

remains predominantly negative. This reinforces the notion that, despite variations in expression, the majority sentiment surrounding gender bias issues in the Olympic context consistently trends toward negativity.

The SENSA framework offers several advantages, however, the framework has some shortcomings. The first is its large size. RoBERTa itself is already a sizable model, and when combined with two BiGRU layers and a CNN with multiple layers, it requires significant computational power and resources. Moreover, the multiple layers in the model necessitate the use of dropout layers to prevent overfitting, further adding to the complexity of the model. A second and very major shortcoming of the framework is nothing but the unavailability of a proper training dataset. Even though the analysis has been done on tweets related to gender bias the training dataset is on miscellaneous tweets. This, even though not a lot, might affect the correctness of the classification.

5 Conclusion

Gender bias continues to be a pressing issue in contemporary society, particularly, it disturbs the ecosystem in sports events. The primary objective of this study is to analyze public sentiment and sarcasm in tweets related to gender bias in a sports context. To accomplish this, the novel SENSA framework integrates sentiment analysis with sarcasm detection, offering a comprehensive understanding of public discourse. The analysis is performed on 5,000 tweets dating from 2010 to August 30, 2024. The R2B-CNN model demonstrated strong performance in sarcasm detection, with a precision of 92.75%, recall of 92.31%, F1-score of 92.53%, and accuracy of 92.32%. For sentiment analysis, it achieved a precision of 92.33%, recall of 92.32%, F1-score of 92.33%, and accuracy of 93.67%. This showcases the model's effectiveness in accurately identifying sarcasm and adjusting sentiment classifications.

The study revealed a prevailing negative sentiment surrounding gender bias discussions, indicating that the majority of users express concern regarding this issue, especially in relation to high-profile sporting events. This analysis on gender bias highlights the significance of understanding sarcasm in shaping public sentiment. Future work can involve expanding the dataset to include a wider range of events and utilizing the SENSA framework to explore other dimensions of public discourse, such as intersectionality and media representation, further enhancing the understanding of societal attitudes toward gender biases. Additionally, incorporating real-time analysis could provide more immediate insights into evolving public sentiments.

References

Toril Aalberg and Anders Todal Jenssen. 2007. Gender stereotyping of political candidates. *Nordicom Review*, 28(1):17–32.

Lorenzo Betti, Carlo Abrate, and Andreas Kaltenbrunner. 2023. Large scale analysis of gender bias and sexism in song lyrics. *EPJ data science*, 12(1):10.

Federico Bianchi, Marco Marelli, Paolo Nicoli, and Matteo Palmonari. 2021. Sweat: Scoring polarization of topics across different corpora. *arXiv preprint arXiv:2109.07231*.

Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.

Tessa ES Charlesworth, Victor Yang, Thomas C Mann, Benedek Kurdi, and Mahzarin R Banaji. 2021. Gender stereotypes in natural language: Word embeddings show robust consistency across child and adult language corpora of more than 65 million words. *Psychological Science*, 32(2):218–240.

Guowei Fan, Tao Wu, and Yamei Lei. 2022. Research on rumor detection based on roberta-bigru model. In *International Conference on Artificial Intelligence and Security*, pages 194–204. Springer.

Alif Tri Handoyo, Derwin Suhartono, et al. 2021. Sarcasm detection in twitter–performance impact while using data augmentation: Word embeddings. *arXiv preprint arXiv:2108.09924*.

Lori Heise, Margaret E Greene, Neisha Opper, Maria Stavropoulou, Caroline Harper, Marcos Nascimento, Debrework Zewdie, Gary L Darmstadt, Margaret Eleanor Greene, Sarah Hawkes, et al. 2019. Gender inequality and restrictive gender norms: framing the challenges to health. *The Lancet*, 393(10189):2440–2454.

Julia Herrmann, Karoline Koeppen, and Ursula Kessels. 2019. Do girls take school too seriously? investigating gender differences in school burnout from a self-worth perspective. *Learning and Individual Differences*, 69:150–161.

IndiaToday. 2024. Olympic gender-row boxer Imane Khelif receives a hero's welcome in Algiers | Paris Olympics 2024 News - Times of India — timesofindia.indiatimes.com. https://tinyurl.com/bu5duay8. [Accessed 25-09-2024].

Brisha Jain and Mainack Mondal. 2024. Uncovering gender bias within journalist-politician interaction in indian twitter. In *Proceedings of the 16th ACM Web Science Conference*, pages 296–301.

Amikul Kalra and Arkaitz Zubiaga. 2021. Sexism identification in tweets and gabs using deep neural networks. *arXiv preprint arXiv:2111.03612*.

Ashraf Kamal, Padmapriya Mohankumar, and Vishal Kumar Singh. 2023. Financial misinformation detection via roberta and multi-channel networks. In *International Conference on Pattern Recognition and Machine Intelligence*, pages 646–653. Springer.

P. Saidesh Kumar and P. Vijayapal Reddy. 2023. Document clustering using roberta and convolution neural network model. *International Journal of Intelligent Systems and Applications in Engineering*, 12(8s):221–230.

Lizhou Li, Shuang Zheng, and Quanmin Wang. 2021. Roberta and stacked bidirectional gru for fine-grained chinese named entity recognition. In *Proceedings of the 2021 6th International Conference on Mathematics and Artificial Intelligence*, pages 95–100.

Armin Mertens, Franziska Pradel, Ayjeren Rozyjumayeva, and Jens Wäckerle. 2019. As the tweet, so the reply? gender bias in digital communication with politicians. In *Proceedings of the 10th ACM Conference on Web Science*, pages 193–201.

Gabriel Meseguer-Brocal, Geoffroy Peeters, Guillaume Pellerin, Michel Buffa, Elena Cabrio, Catherine Faron Zucker, Alain Giboin, Isabelle Mirbel, Romain Hennequin, Manuel Moussallam, et al. 2017. Wasabi: a two million song database project with audio and cultural metadata plus webaudio enhanced client applications.

Dawn M Ohse. 2014. *Applying the SCM and BIAS Map models to theories of sexual harassment*. Southern Illinois University at Carbondale.

Mustafa Oz, Saif Shahin, and Scott B Greeves. 2024. Platform affordances and spiral of silence: How perceived differences between facebook and twitter influence opinion expression online. *Technology in Society*, 76:102431.

Anmol Panda, A'ndre Gonawela, Sreangsu Acharyya, Dibyendu Mishra, Mugdha Mohapatra, Ramgopal Chandrasekaran, and Joyojeet Pal. 2020. Nivaduck-a scalable pipeline to build a database of political twitter handles for india and the united states. In *International Conference on Social Media and Society*, pages 200–209.

Sebastian Stier, Arnim Bleier, Malte Bonart, Fabian Mörsheim, Mahdi Bohlouli, Margarita Nizhegorodov, Lisa Posch, Jürgen Maier, Tobias Rothmund, and Steffen Staab. 2018. Systematically monitoring social media: The case of the german federal election 2017. *arXiv preprint arXiv:1804.02888*.

Samuel L Strozzo. 2023. *Comedic Sarcasm: Meaning Generation by Second Audience Teens*. Ph.D. thesis, University of Leicester.

Kian Long Tan, Chin Poo Lee, and Kian Ming Lim. 2023. Roberta-gru: A hybrid deep learning model for enhanced sentiment analysis. *Applied Sciences*, 13(6):3915.

Yi Xie, Kaiyu Li, Zhaomin Liu, Shicheng Li, and Di Liu. 2024. Sentiment analysis of microblog public opinion based on deep learning. *Mathematical Modeling and Algorithm Application*, 2(3):21–25.