

# Nullpointer at CheckThat! 2024: Identifying Subjectivity from Multilingual Text Sequence

Md. Rafiul Biswas<sup>1,†</sup>, Abrar Tasneem Abir<sup>2,†</sup> and Wajdi Zaghouani<sup>3,\*,†</sup>

<sup>1</sup>Hamad Bin Khalifa University, Doha, Qatar

<sup>2</sup>Carnegie Mellon University in Qatar, Education City, Doha, Qatar

<sup>3</sup>Northwestern University in Qatar, Education City, Doha, Qatar

## Abstract

This study addresses a binary classification task to determine whether a text sequence, either a sentence or paragraph, is subjective or objective. The task spans five languages—Arabic, Bulgarian, English, German, and Italian—along with a multilingual category. Our approach involved several key techniques. Initially, we preprocessed the data through parts of speech (POS) tagging, identification of question marks, and application of attention masks. We fine-tuned the sentiment-based Transformer model 'MarieAngeA13/Sentiment-Analysis-BERT' on our dataset. Given the imbalance with more objective data, we implemented a custom classifier that assigned greater weight to objective data. Additionally, we translated non-English data into English to maintain consistency across the dataset. Our model achieved notable results, scoring top marks for the multilingual dataset (Macro F1-0.7121) and German (Macro F1-0.7908). It ranked second for Arabic (Macro F1-0.4908) and Bulgarian (Macro F1-0.7169), third for Italian (Macro F1-0.7430), and ninth for English (Macro F1-0.6893).

## Keywords

subjectivity, natural language processing, sentiment, fact checking, news articles, text sequence

## 1. Introduction

The concepts of objectivity and subjectivity are crucial in shaping methodologies, interpretations, and the perceived validity of findings in many natural language processing (NLP) applications, such as sentiment analysis and information extraction [1, 2]. Objectivity analysis relies on data that can be measured, observed, and verified by others and is achieved through careful experimental designs, standard procedures, and statistical analysis. In an ideal sense, objective analysis is supposed to be free from individual biases, emotions, and personal judgments, thereby ensuring that the results are universally valid and replicable [3].

Subjectivity, on the other hand, refers to perspectives, interpretations, or analyses that are influenced by personal experiences, feelings, beliefs, or biases [4]. Subjective analysis is inherently shaped by the individual's background, cultural context, and personal viewpoints. While often perceived as less reliable or credible in scientific contexts, subjectivity is an unavoidable aspect of human cognition and can provide valuable insights, particularly in fields such as humanities, social sciences, and qualitative research where personal interpretation and contextual understanding are essential [5].

Identifying whether a text sequence expresses personal opinions, emotions, or factual information is essential for enhancing the accuracy and relevance of automated systems in diverse fields such as social media monitoring, customer feedback analysis, and news content categorization. In data analysis, the tension that arises from the interaction of objectivity and subjectivity frequently affects decision-making procedures and the dissemination of findings. The challenge lies in creating systems that can accurately classify text sequences—whether sentences or paragraphs—as either subjective, reflecting personal opinions or sentiments, or objective, presenting factual information devoid of personal bias [6]. In an effort to improve the acceptability and credibility of work, researchers may strive for objectivity, occasionally avoiding or hiding choices that might be viewed as subjective. Subjective opinions can,

CLEF 2024: Conference and Labs of the Evaluation Forum, September 09–12, 2024, Grenoble, France

\*Corresponding author.

✉ rafiulbiswas@gmail.com (Md. R. Biswas); abir@cmu.edu (A. T. Abir); wajdi.zaghouani@northwestern.edu (W. Zaghouani)

🆔 0000-0002-5145-1990 (Md. R. Biswas); 0000-0002-2375-4134 (A. T. Abir); 0000-0003-1521-5568 (W. Zaghouani)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

for example, slightly skew the analysis's ostensibly objective results in the data selection, analytical method selection, and result interpretation processes. Thus, there is a high chance that the dataset contains a relatively higher number of objective values compared to subjective values.

Task 2 in CheckThat Lab at CLEF 2024 [7] classifies text as either subjective or objective. This binary classification task requires systems to accurately identify the nature of a text sequence. The task is offered in multiple languages: Arabic, Bulgarian, English, German, and Italian, providing a comprehensive multilingual evaluation of the systems' capabilities. The challenge of multilingual and cross-linguistic text classification is compounded by the inherent linguistic and cultural differences that influence the expression of subjectivity and objectivity.

This study presents an approach to a binary classification task aimed at discerning subjective from objective text across multiple languages. By leveraging advanced NLP techniques and Transformer models, we aim to enhance the accuracy and robustness of subjective-objective text classification. The implications of this research extend to improving automated news analysis, enhancing content recommendation systems, and promoting a comprehension understanding in various languages.

## 2. Related Works

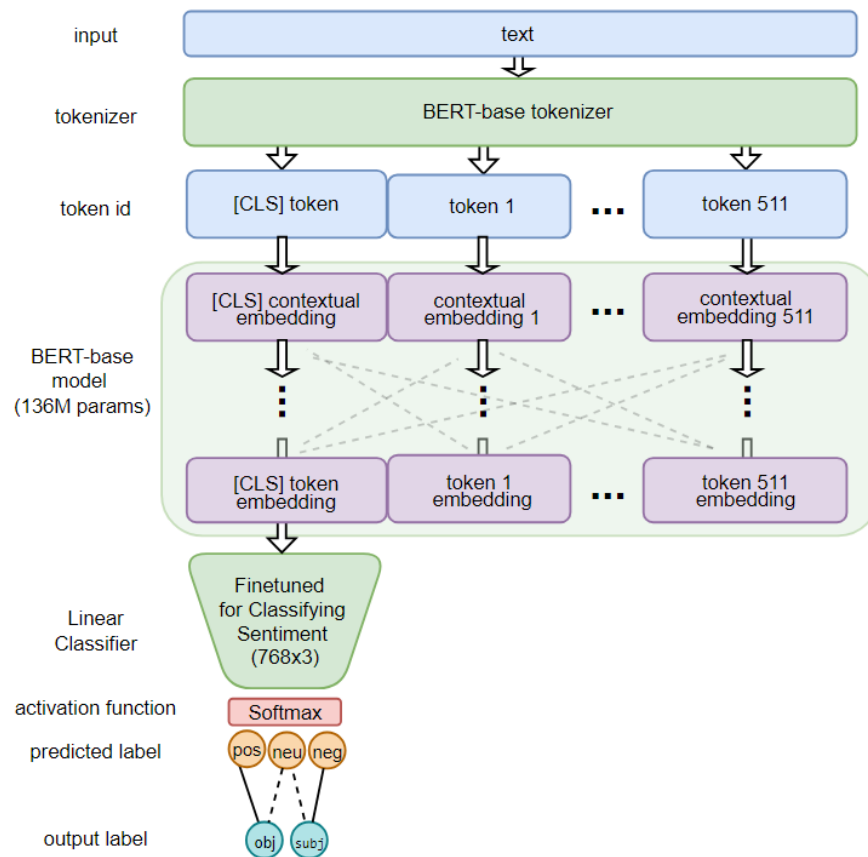
The task of classifying text as subjective or objective has been studied extensively in natural language processing. Early work by Wiebe et al. [8] laid the foundations for subjectivity analysis, proposing a scheme for annotating subjective elements in text. They developed a system called OpinionFinder [1] which performed subjectivity analysis using various lexical and syntactic features. More recently, deep learning approaches have been applied to this task with great success. Nakov et al. [9] provide a thorough overview of modern approaches to sentiment analysis, including detecting subjectivity. They highlight the effectiveness of leveraging pre-trained language models like BERT [10] and fine-tuning them for the target task. Several studies have specifically examined subjectivity classification in a multilingual setting. Balahur et al. [11] constructed a multilingual dataset for subjectivity classification in English, Spanish, French and German. They experimented with various machine translation approaches to make the problem cross-lingual. Similarly, Mihalcea et al. [12] generated subjectivity datasets for English and Romanian, using English tools and manually translating the subjective sentences into Romanian. The CLEF [13](Conference and Labs of the Evaluation Forum) has run workshops on automatic identification and verification of claims in political debates, speeches, and news articles since 2018 [14]. The CheckThat! shared task at CLEF focuses on detecting checkworthy claims across various languages including Arabic [15], which is one of the languages in the current study. In terms of methodology, fine-tuning pre-trained Transformer models has proven very effective for subjectivity and sentiment tasks. Xu et al. [16] fine-tuned BERT for sentiment classification and demonstrated its strong performance on multiple benchmarks. Exploring multi-task learning, Yu and Jiang [17] showed that jointly learning sentiment and subjectivity through a shared BERT encoder led to improvements on both tasks.

## 3. System Overview

This works system for subjectivity classification comprises several key components, including data preprocessing, model selection, and training strategies. This section provides an overview of each component and the techniques employed (see Figure 1).

### 3.1. Data Preprocessing

The first step in the pipeline is data preprocessing, which involves cleaning and transforming the raw text data into a suitable format for model. The preprocessing steps include:



**Figure 1:** Diagram for classification of subjectivity in text sequence

- **Demojization:** We convert emoji characters into their text descriptions using a demojizer<sup>1</sup> to ensure consistent input to the model.
- **Removing users and links:** We remove user mentions and URLs from the text, as they are not relevant for subjectivity classification.
- **Handling poorly formatted TSV files:** Some of the provided TSV files were poorly formatted, so we use a custom dataset class to handle the processing instead of relying on the pandas library.

We also experiment with additional preprocessing techniques such as part-of-speech (POS) tagging and attention masking, but find that they do not significantly improve the performance of the model.

### 3.2. Model Selection

For the subjectivity classification task, we choose to fine-tune pre-trained Transformer models that have been previously trained on sentiment analysis tasks. Specifically, we use the 'MarieAngeA13/Sentiment-Analysis-BERT' model, which is a BERT-based model fine-tuned for sentiment analysis. We find that this approach of using a model already fine-tuned for a related task (i.e., multi-task learning) yields better results compared to fine-tuning a pre-trained model from scratch. The code and data can be found in the GitHub repository <https://github.com/AbRAR-Abir/CLEF2024task02>.

### 3.3. Training Strategies

The training was conducted on a remote Dell server running the latest Ubuntu 22 OS with 512 GB RAM and 24-core CPU. The server was equipped with NVIDIA A100 GPU with 80 GB GPU memory. We employ several training strategies to improve the performance of the model listed below.

<sup>1</sup><https://carpedm20.github.io/emoji/docs/api.html#emoji.demojize>

- **Label mapping:** The pre-trained sentiment analysis model is designed for three-class prediction (positive, neutral, negative), while our subjectivity classification task requires only two classes (subjective and objective). We experimented with different label mappings and found that mapping subjective to negative sentiment and objective to positive sentiment yielded the best results.
- **Confidence weighting:** For the English dataset, we incorporate the confidence level information provided in the dataset (in the 'solved\_conflict' column where 1[true] means conflict was resolved i.e., higher annotation confidence and vice versa). We assign 20% higher weight for the training losses- coming from the annotations with higher confidence (i.e, 1.2 weight)- before passing the losses to the loss function so that backpropagation prioritizes minimizing loss for higher confidence annotation compared to their counterparts.
- **Hyperparameter tuning:** We experiment with different hyperparameter settings and find that a batch size of 16, learning rate of  $2e-5$ , and training for 20 epochs yields the best performance.

### 3.4. Language Adaptation

To handle the multilingual nature of the task, we employ machine translation to convert non-English data into English. We use the Google Translator API through the deep translator library for this purpose. While we also experiment with fine-tuning language-specific pre-trained models for non-English languages, we find that translating the training and test datasets to English and using the English model yields better performance. These preprocessing, model selection, and training strategies form the core of the subjectivity classification system. In the following sections, we detail our experimental setup and present the results of the approach.

## 4. Results

This section presents the results of subjectivity classification system across various languages and datasets. We first describe the dataset characteristics and then provide a detailed analysis of the model's performance using different evaluation metrics. Finally, we compare our results with those of other participating teams in the CheckThat! Lab at CLEF 2024.

### 4.1. Dataset Description

The dataset for the Subjectivity Subtask consists of sentences from news articles in five languages: Arabic, Bulgarian, English, German, and Italian. Additionally, there is a multilingual dataset that combines all five languages. Table 1 shows the distribution of objective and subjective sentences in the training and test sets for each language. Across all languages, the percentage of objective sentences is higher than that of subjective sentences, with the imbalance being more pronounced in the training sets. This imbalance poses a challenge for subjectivity classification systems, as they need to learn from skewed data distributions. For the Arabic language, the training set comprises 1185 sentences, with 905 being objective (76.37%), and the test set includes 748 sentences, of which 425 are classified as objective (56.81%). In Bulgarian, the training set contains 729 sentences, where 406 are objective (55.69%), and the test set consists of 250 sentences, with 143 objective sentences (57.2%). The English dataset includes 830 sentences for training, with 532 labeled as objective (64.09%), and the test set has 484 sentences, with 362 objectives (74.79%). For the German language, the training set comprises 800 sentences, of which 492 are objective (61.5%), and the test set contains 337 sentences, with 226 being objective (67.07%). For the Italian language, the training set includes 1613 sentences, with 1231 objectives (76.31%), and the test set comprises 513 sentences, with 377 objectives (73.4%). The multilingual dataset combines all five languages and comprises 5159 sentences in the training set, of which 3568 are objective (69.16%). The test set contains 500 sentences, evenly split with 250 objective sentences (50%) and 250 subjective sentences (50%). This comprehensive dataset provides a robust foundation for developing and evaluating systems that distinguish between subjective and objective statements in news articles across multiple languages.

**Table 1**  
Training and Test Data Distribution

| Language     | Dataset      | OBJ (N) (%)  | SUBJ (N) (%) |
|--------------|--------------|--------------|--------------|
| Arabic       | Train (1185) | 905 (76.37)  | 280 (23.63)  |
|              | Test (748)   | 425 (56.81)  | 323 (43.18)  |
| Bulgarian    | Train (729)  | 406 (55.69)  | 323 (44.23)  |
|              | Test (250)   | 143 (57.2)   | 107 (42.8)   |
| English      | Train (830)  | 532 (64.09)  | 298 (35.9)   |
|              | Test (484)   | 362 (74.79)  | 122 (25.2)   |
| German       | Train (800)  | 492 (61.5)   | 308 (38.5)   |
|              | Test (337)   | 226 (67.07)  | 111 (32.93)  |
| Italian      | Train (1613) | 1231 (76.31) | 382 (23.68)  |
|              | Test (513)   | 377 (73.4)   | 136 (26.5)   |
| Multilingual | Train (5159) | 3568 (69.16) | 1591 (30.83) |
|              | Test (500)   | 250 (50)     | 250 (50)     |

## 4.2. Performance Metrics

We evaluate our subjectivity classification model using various performance metrics, including macro-averaged F1-score, precision, recall, and accuracy. Table 2 presents the results for each language and the multilingual dataset. Our model achieves the best performance on the German dataset, with an F1 Macro score of 0.79 and an accuracy of 0.81, indicating high prediction correctness. The multilingual dataset obtains good performance, with an F1 Macro score of 0.71, an F1 SUBJ of 0.69, and an accuracy of 0.71. The model also shows good performance for the Italian language with an F1 Macro score of 0.74 and strong subjective class metrics, with an F1 SUBJ of 0.64. On the other hand, the model struggles the most with the Arabic dataset, obtaining an F1 Macro score of 0.49 and an accuracy of 0.52. The performance is relatively lower than in other languages, which shows the difficulty in identifying subjective data. The model performs well in Bulgarian, achieving an F1 Macro score of 0.72 and high subjective class performance with an F1 SUBJ of 0.69. For English, the performance is moderate to good, with an F1 Macro score of 0.68. The model handles subjective data in English relatively better, with an F1 SUBJ of 0.54, precision (P SUBJ) of 0.52, and recall (R SUBJ) of 0.64. The overall accuracy for English is 0.64.

In summary, the model shows the highest performance in German, followed by Italian and Bulgarian, with Arabic being the most challenging language for the model. The performance in English is moderate, and the overall multilingual performance is strong, suggesting the model’s effectiveness across multiple languages but with some variability in specific language performance.

**Table 2**  
Performance metrics across different languages

| Language     | F1 Macro | P Macro | R Macro | F1 SUBJ | P SUBJ | R SUBJ | Accuracy |
|--------------|----------|---------|---------|---------|--------|--------|----------|
| Arabic       | 0.49     | 0.49    | 0.50    | 0.37    | 0.43   | 0.33   | 0.52     |
| Bulgarian    | 0.72     | 0.72    | 0.72    | 0.69    | 0.66   | 0.72   | 0.72     |
| English      | 0.68     | 0.43    | 0.50    | 0.54    | 0.52   | 0.64   | 0.64     |
| German       | 0.79     | 0.78    | 0.81    | 0.73    | 0.67   | 0.80   | 0.81     |
| Italian      | 0.74     | 0.73    | 0.77    | 0.64    | 0.57   | 0.73   | 0.78     |
| Multilingual | 0.71     | 0.72    | 0.71    | 0.69    | 0.76   | 0.63   | 0.71     |

- **F1 Macro:** The macro-averaged F1 score, which is the harmonic mean of precision and recall across all classes.
- **P Macro:** The macro-averaged precision.
- **R Macro:** The macro-averaged recall.
- **F1 SUBJ:** The F1 score for subjective classification.

- **P SUBJ**: The precision for subjective classification.
- **R SUBJ**: The recall for subjective classification.
- **Accuracy**: The overall accuracy of the model.

### 4.3. Comparison with Other Teams

We compare the performance of our subjectivity classification model with that of other participating teams in the CheckThat! Lab at CLEF 2024. Table 3 shows the official results for each language and the multilingual dataset. Our team achieves the highest rank in the German and multilingual categories, with Macro F1 scores of 0.7908 and 0.7121, respectively. We also secure the second position in Arabic and Bulgarian. For Arabic, our model achieved second place with a Macro F1 score of 0.4908 and a SUBJ F1 score of 0.37. In Bulgarian, our model also secured second place with a Macro F1 score of 0.7169 and a SUBJ F1 score of 0.69. In Italian, our model ranks third with a Macro F1 score of 0.7430 and a SUBJ F1 score of 0.64. In the English category, our model ranks ninth with a Macro F1 score of 0.6893 and a SUBJ F1 score of 0.54.

These results showcase the competitiveness of our approach in the shared task, especially in the German and multilingual categories. They also indicate areas for improvement, particularly in English, where our model’s performance is lower than other teams. Overall, our team’s participation in the ArAIEval shared task demonstrated strong performance across multiple languages, securing top ranks in several categories and showcasing our model’s capabilities in multilingual data and subjective data evaluation.

**Table 3**

Official results for six test languages in Subtask2 CheckThat! Lab at CLEF 2024

| Language            | Team               | Rank | Macro F1 | SUBJ F1 |
|---------------------|--------------------|------|----------|---------|
| <b>Arabic</b>       | IAI Group          | 1    | 0.4947   | 0.46    |
|                     | <b>Nullpointer</b> | 2    | 0.4908   | 0.37    |
|                     | Baseline           | 3    | 0.4852   | 0.40    |
|                     | JUNLP (last)       | 7    | 0.3623   | 0.00    |
| <b>Bulgarian</b>    | Baseline           | 1    | 0.7531   | 0.73    |
|                     | <b>Nullpointer</b> | 2    | 0.7169   | 0.69    |
|                     | Hybrinfox          | 3    | 0.7147   | 0.65    |
|                     | JUNLP (last)       | 5    | 0.3639   | 0.00    |
| <b>English</b>      | Hybrinfox          | 1    | 0.7442   | 0.60    |
|                     | <b>Nullpointer</b> | 9    | 0.6893   | 0.54    |
|                     | Baseline           | 11   | 0.6346   | 0.45    |
|                     | IAI Group (last)   | 15   | 0.4491   | 0.39    |
| <b>German</b>       | <b>Nullpointer</b> | 1    | 0.7908   | 0.73    |
|                     | IAI Group          | 2    | 0.7302   | 0.66    |
|                     | Baseline           | 3    | 0.6994   | 0.63    |
|                     | Hybrinfox (last)   | 4    | 0.6968   | 0.57    |
| <b>Italian</b>      | JK_PCIC_UNAM       | 1    | 0.7917   | 0.69    |
|                     | <b>Nullpointer</b> | 3    | 0.7430   | 0.64    |
|                     | Baseline           | 4    | 0.6503   | 0.52    |
|                     | IAI Group (last)   | 5    | 0.5862   | 0.49    |
| <b>Multilingual</b> | <b>Nullpointer</b> | 1    | 0.7121   | 0.69    |
|                     | Hybrinfox          | 2    | 0.6849   | 0.63    |
|                     | Baseline           | 3    | 0.6697   | 0.66    |
|                     | IAI Group (last)   | 4    | 0.6292   | 0.67    |

## 5. Discussion

Our system leveraged state-of-the-art pre-trained language models, specifically BERT, which we fine-tuned for subjectivity classification task. Through extensive experiments, we demonstrated the ef-

fectiveness of our approach, achieving competitive performance in various languages. Our system ranked first in the German and multilingual categories, second in Arabic and Bulgarian, and third in Italian. These results highlight the robustness of our model and its ability to generalize across different languages. We also investigated the impact of various preprocessing techniques, such as part-of-speech tagging and attention masking, on the performance of our system.

Furthermore, our analysis of the dataset characteristics revealed the challenges posed by the imbalance between objective and subjective sentences across all languages. This imbalance underscores the need for developing strategies to handle skewed data distributions effectively.

Our work contributes to the growing body of research on subjectivity classification and multilingual natural language processing. The insights gained from our experiments can inform future research directions and help develop more robust and accurate systems for subjectivity analysis across diverse languages.

However, our study also has some limitations. The performance of our system in English was relatively lower compared to other languages, indicating room for improvement. Future work could explore more advanced techniques, such as domain adaptation and transfer learning, to enhance the model's performance in English and other languages. Moreover, the scope of our study was limited to the dataset provided by the CheckThat! Lab. Further research could investigate the generalizability of our approach to other datasets and domains, such as social media and customer reviews.

## 6. Conclusion

In conclusion, our subjectivity classification system, Nullpointer, demonstrates the potential of leveraging pre-trained language models and multilingual approaches for identifying subjective and objective statements in news articles. As the volume of online content continues to grow, the ability to automatically distinguish between subjective and objective information becomes increasingly crucial. Our work contributes to this important research area and paves the way for more advanced and reliable subjectivity analysis systems in the future.

## 7. Acknowledgments

We acknowledge Qatar National Research Fund grant NPRP14C0916-210015 from the Qatar Research Development and Innovation Council (QRDI) for funding this research.

## References

- [1] T. Wilson, P. Hoffmann, S. Somasundaran, J. Kessler, J. Wiebe, Y. Choi, C. Cardie, E. Riloff, S. Patwardhan, Opinionfinder: A system for subjectivity analysis, in: Proceedings of HLT/EMNLP 2005 Interactive Demonstrations, 2005, pp. 34–35.
- [2] A. Gelman, C. Hennig, Beyond subjective and objective in statistics, *Journal of the Royal Statistical Society Series A: Statistics in Society* 180 (2017) 967–1033.
- [3] R. A. Hackett, Decline of a paradigm? bias and objectivity in news media studies, *Critical Studies in Media Communication* 1 (1984) 229–259.
- [4] J. Kocoń, M. Gruza, J. Bielaniewicz, D. Grimling, K. Kanclerz, P. Miłkowski, P. Kazienko, Learning personal human biases and representations for subjective tasks in natural language processing, in: 2021 IEEE International Conference on Data Mining (ICDM), IEEE, 2021, pp. 1168–1173.
- [5] U. Müller, J. Carpendale, M. Bibok, T. Racine, Subjectivity, identification and differentiation: Key issues in early social development, *Monographs of the Society for Research in Child Development* (2006) 167–179.
- [6] M. Othman, H. Hassan, R. Moawad, A. M. Idrees, Using nlp approach for opinion types classifier (2015).

- [7] J. M. Struß, F. Ruggeri, A. Barrón-Cedeño, F. Alam, D. Dimitrov, A. Galassi, M. Siegel, M. Wiegand, Overview of the CLEF-2024 CheckThat! lab task 2 on subjectivity in news articles, 2024.
- [8] J. Wiebe, R. Bruce, T. O'Hara, Development and use of a gold-standard data set for subjectivity classifications, in: Proceedings of the 37th annual meeting of the Association for Computational Linguistics, 1999, pp. 246–253.
- [9] P. Nakov, A. Ritter, S. Rosenthal, F. Sebastiani, V. Stoyanov, Semeval-2016 task 4: Sentiment analysis in twitter, in: Proceedings of the 10th international workshop on semantic evaluation (SemEval-2016), 2016, pp. 1–18.
- [10] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, 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), 2019, pp. 4171–4186.
- [11] A. Balahur, R. Steinberger, E. van der Goot, B. Pouliquen, M. Kabadjov, Opinion mining on newspaper quotations, in: Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology-Volume 03, IEEE Computer Society, 2009, pp. 523–526.
- [12] R. Mihalcea, C. Banea, J. Wiebe, Learning multilingual subjective language via cross-lingual projections, in: Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics, 2007, pp. 976–983.
- [13] A. Barrón-Cedeño, F. Alam, T. Chakraborty, T. Elsayed, P. Nakov, P. Przybyła, J. M. Struß, F. Haouari, M. Hasanain, F. Ruggeri, X. Song, R. Suwaileh, The clef-2024 checkthat! lab: Check-worthiness, subjectivity, persuasion, roles, authorities, and adversarial robustness, in: N. Goharian, N. Tonelotto, Y. He, A. Lipani, G. McDonald, C. Macdonald, I. Ounis (Eds.), Advances in Information Retrieval, Springer Nature Switzerland, Cham, 2024, pp. 449–458.
- [14] P. Nakov, A. Barr'on-Cede no, G. Da San Martino, F. Alam, R. M'iguez, T. Caselli, M. Kutlu, W. Zaghouni, C. Li, S. Shaar, et al., The clef-2022 checkthat! lab on fighting the covid-19 infodemic and fake news detection, in: European Conference on Information Retrieval, Springer, 2022, pp. 416–428.
- [15] F. Alam, F. Dalvi, S. Shaar, N. Durrani, H. Mubarak, A. Nikolov, G. Da San Martino, A. Ali, F. Sajjad, T. Caselli, et al., Fighting the covid-19 infodemic in social media: a holistic perspective and a call to arms, in: Proceedings of the International AAAI Conference on Web and Social Media, volume 15, 2021, pp. 913–922.
- [16] H. Xu, B. Liu, L. Shu, P. S. Yu, Bert post-training for review reading comprehension and aspect-based sentiment analysis, 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), 2019, pp. 2324–2335.
- [17] J. Yu, J. Jiang, Adapting bert for target-oriented multimodal sentiment classification, in: Proceedings of the 28th International Joint Conference on Artificial Intelligence, 2019, pp. 5408–5414.