

OneLove beyond the field - A few-shot pipeline for topic and sentiment analysis during the FIFA World Cup in Qatar

Christoph Rauchegger¹ and Sonja Mei Wang² and Pieter Delobelle³

¹ Technische Universität Berlin, Germany

² University of Wuppertal, Germany

³ Department of Computer Science, KU Leuven; Leuven.AI, Belgium

Abstract

The FIFA World Cup in Qatar was discussed extensively in the news and on social media. Due to news reports with allegations of human rights violations, there were calls to boycott it. Wearing a OneLove armband was part of a planned protest activity. Controversy around the armband arose when FIFA threatened to sanction captains who wear it. To understand what topics Twitter users Tweeted about and what the opinion of German Twitter users was towards the OneLove armband, we performed an analysis of German Tweets published during the World Cup using in-context learning with LLMs. We validated the labels on human annotations. We found that Twitter users initially discussed the armband’s impact, LGBT rights, and politics; after the ban, the conversation shifted towards politics in sports in general, accompanied by a subtle shift in sentiment towards neutrality. Our evaluation serves as a framework for future research to explore the impact of sports activism and evolving public sentiment. This is especially useful in settings where labeling datasets for specific opinions is unfeasible, such as when events are unfolding.

1 Introduction

In December 2010, it was announced that Qatar will host the FIFA World Cup in 2022 (BBC News, 2010). This announcement was met with concerns and criticism, including the effect of high temperatures during June and July (Matzarakis and Fröhlich, 2015), and accusations of bribery (Panja and Draper, 2020). Important political topics were allegations of human rights violations, such as withholding wages from migrant workers and their unexplained deaths (Heerdt and Roorda, 2023; Human Rights Watch, 2023), as well as abuse endured by LGBT people (Crane, 2022; Human Rights Watch, 2022).

As one of the biggest sport events of 2022, the FIFA World Cup in Qatar was followed by football

fans and social media users who expressed their enthusiasm, but also critique towards this sporting event on the microblogging platform X, formerly known as Twitter (Pak and Paroubek, 2010; Fan et al., 2020). FIFA claims that there were over five billion “engagements” with the World Cup on social media (FIFA, 2023). The FIFA World Cup is one of the most followed sporting events worldwide – a mega-event (Müller et al., 2023).

While calls for boycotts became public before the World Cup and initiatives like “Boycott Qatar 2022” were established (Sportschau, 2022), there were planned protest activities during the World Cup, such as the OneLove captain’s armband. Ten European teams announced their intention to wear the colored armband (KNVB Media, 2022). Announcements of sanctions made national teams cancel these plans, with threats of receiving a yellow card for wearing it becoming known on November 21st, 2022 (DW, 2022). Teams released statements expressing that it had been their intention to wear the armband to stand for inclusion, diversity and mutual respect, but chose not to in order to avoid negatively impacting the players (Ramsay and Nabbi, 2022). Instead of wearing the armband, the German national team held their hands in front of their faces to make a statement (ZDF heute, 2022). The sanctions against the German team’s protest activity and the backing down of teams were discussed on social media, such as Twitter.

We conducted an analysis of 132k Tweets about the FIFA World Cup in Qatar to analyze topics dominating discussions on this event and opinions on one particular topic, namely the OneLove armband. By conducting this analysis, we also evaluate if our proposed method can replace supervised pipelines to analyze and measure public opinions. Supervised machine learning models play a crucial role in measuring sentiment online (Dang et al., 2020) by classifying user-generated texts (Lheureux, 2023; Rustam et al., 2021; Scott et al., 2021). However,




Figure 1: Topic modeling and sentiment analysis pipeline.

training these models requires labeled data, which involves significant time and resources. Thus, some projects are limited to using existing models, for instance when there is no annotation budget, or when monitoring events or crises, where there is no time to create a labeled dataset (Bruyne et al., 2024). Another limitation is that most studies of sentiment of specific events on social media happen after the event has passed, instead of while an event is unfolding, which is possible with our proposed method.

We specifically aim to answer the following research questions, both related to the World Cup in Qatar and for monitoring opinions on social media:

- RQ1** Which topics related to the OneLove armband were discussed on Twitter? (§ 3.1)
- RQ2** What is the sentiment on Twitter towards FIFA’s ban of the OneLove armband? (§ 3.2)
- RQ3** How suitable is a zero-shot or few-shot language model-based pipeline for measuring public sentiment from Twitter/X? (§ 4)

2 Background and Related Work

Numerous studies have conducted analyses on Twitter data for FIFA World Cups, including Meier et al. (2021), Patel and Passi (2020), Nuñez Franco (2023), Hassan and Wang (2023), and Fan et al. (2020). Controversies surrounding countries hosting these events were a focus of previous research. Meier et al. (2021) who analyzed Twitter data on the 2018 FIFA World Cup in Russia describe that it was controversial due to the annexation of Crimea in March 2014 and subsequent military conflicts in Ukraine. The authors argue that mega-events are used by host countries to portray a positive national image for domestic and global audiences. They also mention that in the past, countries have come under scrutiny when hosting such events, as civil society has increasingly focused on human rights and civil liberties of host countries, and demanded accountability and political reforms.

Brannagan and Giulianotti (2018) predicted that Qatar intended to use the global attention received from hosting the FIFA World Cup to show Qatar’s pursuits of peace, security and integrity, and bolster their attractiveness for international tourists. However, the criticism due to the news coverage of alleged human rights violations lead to calls for boycotts and planned protest activities. Hassan and Wang (2023) found that political discussions surrounding the FIFA World Cup in Qatar initially dominated but then gradually declined, shifting the focus to sporting achievements and cultural exchange. Findings from the sentiment analysis conducted by Nuñez Franco (2023) include that more than 66% of Tweets showed a positive sentiment and that the majority of data contained neutral hashtags, with only around 6% of Tweets containing a negative hashtags.

Similarly to (Meier et al., 2021), we argue that a sentiment analysis can be an indicator of whether the public is aware of contested issues around mega sports events such as FIFA World Cups and whether protest activities to raise awareness of human rights issues are supported. Further analysis, such as a qualitative content analysis of non-supportive Tweets, can inform future protest activities.






Differently from previous Twitter analyses on the FIFA World Cup in Qatar, we focus not only on finding out what topics were discussed, but also how users perceived the protest activity of wearing the OneLove armband, a much discussed protest activity. In addition, we analyze German Tweets, while previous research has focused on English Tweets.

3 Method

All analyzed Tweets were collected by the authors through the formerly freely available, official Twitter API. This API allowed access to all public Tweets of the last 30 days. To find Tweets with certain characteristics, we used the following filter:

```
-is:retweet -is:reply -is:quote lang:de
-liveticker -newsticker (#WM2022 OR
```

Table 1: Performance of different sentiment analysis models. All models were evaluated on our manually labeled test set, as discussed in § 3.2. We link to the publicly available models and indicate if they are multilingual (🌐) or only for German (🇩🇪).

Model	Predicted labels	Acc [%]	F ₁ [%]
 nlpstown/bert-base-multilingual-uncased-sentiment	1 to 5 stars	46.4	19.6
 oliverguhr/german-sentiment-bert	Positive, negative, neutral	62.6	43.9
 mistralai/Mistral-7B-Instruct-v0.2	In favor, against or neutral	64.9	47.3
+ 3-shot prompt (examples)		67.1	50.7
+ translation to English		69.8	54.7
+ Chain of Thought (CoT) reasoning (Kojima et al., 2022)		74.3	61.5
+ all three techniques		76.1	64.2
 gpt-3.5-turbo ¹ (3-shot prompt, translation, CoT)	In favor, against or neutral	59.9	39.9
 gpt-4-turbo ¹ (3-shot prompt, translation, CoT)	In favor, against or neutral	80.2	70.3

#FIFAWorldCup OR FIFA OR WM OR (WM (Katar OR Qatar)) OR ((Fußball OR Fussball)(Qatar OR Katar OR Weltmeisterschaft))

This string contains the German and English spelling variant of Qatar and Football. Weltmeisterschaft (abbr. WM) is the translation for "World Cup" in German. These hashtags and words were selected to ensure that Tweet topics are connected to the World Cup. Tweets that only contain content about football or only Qatar should be excluded through this search. Retweets, replies and quotes were excluded from the collection of data, as context needs to be taken into account when interpreting such Tweets (see Limitations). The language was set to German and liveticker and newsticker were excluded to avoid Tweets only about news regarding this event, as they are not relevant for sentiment analysis. The data set consists of 132,150 Tweets that were Tweeted in a time period from November 20 to December 18, 2022.

3.1 Topic modeling

To get an overview of important topics in our data set, we performed unsupervised topic modeling. We used BERTopic (Grootendorst, 2022) with a multilingual model and a CountTokenizer with n-grams of length 1 to 3, to also include bi-grams like ‘one love’ and tri-grams like ‘One Love Binde’. The BERTopic library only includes stopwords in English, so we iteratively wrote our own short list to improve the clusters (see Appendix A).

We validated the discovered topics using a separate test set, where we manually annotated 600 Tweets belonging to seven topics: irrelevant, game Tweet, news, boycott, human rights, OneLove Binde and general politics. The highest amount of Tweets were news and reporting of the game progress, which are not interesting for our analysis.

The OneLove armband was relevant to Twitter

users and our topic model found this topic and the issues related to this. As mentioned before, due to the media attention and controversy surrounding this protest activity, we decided to focus on Tweets about the OneLove armband. We manually annotated 200 Tweets that contain the word OneLove or a spelling variation (one-love and one love) as being for, against or neutral towards wearing the armband. Another of the authors annotated 100 Tweets to calculate Inter-Annotator Agreement (IAA). Of the 200 annotated Tweets on OneLove, 88 were annotated as for OneLove and 20 against OneLove. We did this by selecting the topics from our BERTopic pipeline that were associated with the Tweets that we manually labeled as ‘OneLove Binde’ (see Figure 2).

The topics found by our topic modelling pipeline align well with different aspects related to the OneLove armband, such as the penalty for wearing the armband despite the ban (a yellow card), discussions about politics in sports and various OneLove-related topics.

3.2 Few-shot sentiment analysis

We used sentiment analysis to analyze the Tweets on the OneLove topic, which is a commonly used method to extract subjective information, such as opinions and attitudes (Medhat et al., 2014; Mäntylä et al., 2018). However, our task is to gauge support for wearing the OneLove armband, which is distinct—and often the opposite—of support for the ban that FIFA implemented. This task, often referred to as stance detection (Alturayef et al., 2023), is different than what publicly available sentiment analysis models are trained on, so we perform an evaluation of different language models.

¹GPT-3.5 and GPT-4 are not fully supported with select-queries in Guidance, so unconstrained generation is used and this affects the accuracy.




Figure 2: Timeline of discussed topics related to OneLove. The topics are found by topic modeling (§ 3.1) and linked to OneLove with manual annotation of a subset of Tweets.

In addition, we focus on evaluating zero-shot and few-shot methods to answer RQ3. Because of this focus, our approach could be directly applicable for new studies following other events, which in turn also allows for a faster reaction to develop a pipeline while an event is still unfolding.

We first create a set of gold labels by having two native German speakers annotate 148 randomly sampled Tweets using Doccano¹. Our annotation labels are ‘in favor of the OneLove Binde’ ($n = 63$), ‘against the OneLove Binde’ ($n = 15$) and ‘neutral’ ($n = 70$). Some of the Tweets did not directly contain an opinion in the text, but used tags like #WMderSchande (World Cup of Shame) and #BoycottQatar. Other Tweets were commenting on how the OneLove armband is only a replacement for a real rainbow-colored armband and a compromise, these Tweets were also annotated as being for the armband. The inter-annotator agreement is $\kappa = 0.68$, which indicates significant agreement (Viera et al., 2005).

We evaluate two already finetuned BERT-based models and validate them by calculating macro-averaged accuracy and micro-averaged F_1 scores between the predicted labels and the gold labels. However, both models are finetuned for sentiment analysis on different domains, mostly reviews with a star-based rating system. As a consequence, the observed error rates are too high when applied to our dataset, as shown by the F_1 scores of 20% and 44% in Table 1.

To address this, we test a similar setting using generative LMs, where we tested Mistral, GPT-3.5 and GPT-4. For the publicly available model Mistral 7B (Jiang et al., 2023), we evaluated several variations of a prompt: (i) only the base prompt with the coding instructions, (ii) with 3 examples for 3-shot classification, (iii) with additional instructions to translate the Tweet to English, (iv)

¹<https://doccano.github.io/doccano/>

Chain of Thought (CoT, “let’s think step by step”) reasoning (Kojima et al., 2022) and (v) a combination of all prompts. We also test the conversation-tuned GPT-3.5 and GPT-4 models by OpenAI using all aforementioned prompting techniques, as they proved to increase the performance on the Mistral model. We also note that the models struggle most with the ‘against OneLove binde’ class, for instance GPT-4 has an AUC score of 0.697 for this class, while the others are slightly higher with 0.769 (in favor) and 0.747 (neutral).

Based on this evaluation (see Table 1), we find that an optimized prompt with examples, Chain of Thought reasoning and a translation all improve the Mistral-based labeling. Interestingly enough, a BERT-based classifier performs almost as well (62.6% accuracy) as Mistral without prompting techniques (64.9%), at a lower inference cost. Nevertheless, by optimizing the prompt Mistral achieves 76.1% accuracy, which outperforms GPT-3.5 and is only slightly worse than the relatively expensive GPT-4 model, so we use Mistral with all prompting techniques to classify all Tweets on OneLove.

4 Discussion

Opinions on the OneLove armband. We first analyzed the topics found in the Tweets, where we found multiple topics related to OneLove (see Figure 2). There are more Tweets about the game itself, but ‘one love’ and ‘boycott’ are the first political topics we found with our topic modeling. The range of topics is broad, and include for instance discussions of the consequences of wearing the OneLove armband, women’s rights and politics in sports. Initially, the discussion was about OneLove and related topics, such as women’s rights, homosexuality and homophobia, as well as the mention of ‘legal action’. Interestingly enough, most of the Tweets after the ban, on November 21, 2022, shifted to focus on politics in sports in general, as opposed to the armband itself or banning it.

The sentiment of the Tweets slightly shifted towards a more neutral stance, however there are more supportive Tweets for wearing the OneLove armband (see Figure 3). We also observe an initial spike of Tweets as a reaction on the ban of the OneLove armband, but this quickly ebbed away. The overall sentiment correlates with some surveys, e.g. Dörner (2022), although we observe more Tweets taking a neutral stance.




Figure 3: Sentiment of Tweets related to OneLove.

Few-shot analysis pipeline. In this paper, we also presented a few-shot pipeline to cluster Tweets using topic modeling and analyze them using zero-shot and few-shot language models. We found that both aspects of the pipeline work sufficiently well to use for events and crises: First, the topics discovered by BERTopic aligned with our manual labels (see § 3.1). Second, large language models, such as Mistral-7B, are surpassing finetuned BERT models (see § 3.2) and allow for in-context learning with few examples (Kojima et al., 2022), meaning classification pipelines are quicker to set up.

Comparison to Surveys Some social media users do not post their opinion publicly or post nothing, even if they have an opinion. The imagined audience, lurking behavior and deletion of Tweets impact the availability of social media data (Litt and Hargittai, 2016; Gong et al., 2016; Almuhimedi et al., 2013; Zhou et al., 2016). Our analysis can be used as an additional data point alongside other sources of information. Many surveys conducted by newspapers and other media outlets observed similar trends in the fraction of people supporting a boycott as our analysis does. A 2022 survey conducted by a local newspaper shows that 46% of German citizens who responded intended to boycott the 2022 FIFA World Cup, while 28% were against a boycott (Dörner, 2022). Another survey’s results showed that 72% of respondents were in favor of a boycott (NDR, 2022). There was also a non-representative survey which German government officials responded to and approximately half favored a boycott (Merkur, 2022).

While surveys with fixed responses can give an overview of agreement and disagreement with an issue, an advantage of social media data compared to such surveys is that it is possible to analyze how individuals form or explain their stance on a topic. Topic modeling and sentiment analysis can also be combined with a qualitative analysis of Tweets

if particular topics are of interest. Further analyses could focus on whether there are any group differences among those supporting or boycotting protest activities, as previous research has focused on inferring demographic data from information provided on Twitter (Sloan et al., 2013; Culotta et al., 2015; Sloan et al., 2015).

5 Conclusion

The FIFA World Cup 2022 in Qatar was accompanied by controversies, among them allegations of human rights violations of migrant workers and LGBT people. Football teams planned to wear the OneLove armband as a protest activity, however the FIFA threatened to sanction wearing the armband. We conducted topic modeling and found that Twitter users talked about the ban of the OneLove armband for a few days after the ban. We then focused on Tweets about OneLove and the subsequent ban of the OneLove armband, to gauge whether Tweets were for or against wearing the armband. Our analysis shows that there was more support for wearing the armband than not, although the support did fade over time, which could indicate that Twitter users perceived the protest as unsuccessful. We identified a shift from the armband-specific discourse to a broader discussion on politics in sports. A purpose of this protest activity was to raise awareness of human rights violations. If a sentiment analysis reveals negative public perception towards a particular protest activity, this suggests a potential need for reassessment and modification of the protest strategy which is possible while an event takes place through our proposed few-shot pipeline.

Limitations

The data set we analyzed consisted of German Tweets, which means that results only allow us a glimpse into German-speaking Twitter users’ opinions on the FIFA World Cup in Qatar. X, formerly known as Twitter, has users from certain demographics and is not used by every German-speaking person, thus the results are not generalizable to all German-speakers. We did not try to infer demographics from our data. Nonetheless, our work contributes to research on social media data centered on languages other than English and goes beyond Germany’s borders, as German is also spoken in other countries, such as Austria and Switzerland.

We elected to not include replies, which means there might be posts in threads that are for or

against the OneLove armband that we did not cover with our analysis. Since the added context (the original Tweet) might make interpreting these Tweets more difficult², we decided to leave out these replies. Future work could focus to include these replies with the original Tweet.

Finally, we evaluated multiple models to classify Tweets. De Bruyne et al. (2021) showed that different BERT-based models performed differently with regard to emotion detection or sentiment analysis, which might be the case for the models we tested as well. Biases in the model might affect the classification as well (Talat et al., 2021). However, an analysis of these issues for the models we used is out of scope for this work.

Ethical considerations

We only used publicly available data and did not interact with human subjects, which means our work did not classify as human subjects research by our IRBs. Since this research did not seek informed consent, data were taken from a public online venue. On X, formerly known as Twitter, users have different privacy options, such as posting public and private Tweets. We collected only public Tweets which were easily accessible for all users through the search function. Following the conditions of the Twitter API, we cannot publish the full dataset that we used, instead we publish a “dehydrated” version which requires access to the API to receive full information on each Tweet. For our analysis and collection of data, we only used the Tweet, no further information on the person who posted it was collected.

Acknowledgments

Pieter Delobelle received funding from the Flemish Government under the “Onderzoeksprogramma Artificial Intelligence (AI) Vlaanderen” programme and was supported by the Research Foundation - Flanders (FWO) under EOS No. 30992574 (VeriLearn). He also received a grant from “Interne Fondsen KU Leuven/Internal Funds KU Leuven”.

References

Hazim Almuhammedi, Shomir Wilson, Bin Liu, Norman Sadeh, and Alessandro Acquisti. 2013. [Tweets are forever: a large-scale quantitative analysis of deleted tweets](#). In *Proceedings of the 2013 Conference on*

²Is positive or negative sentiment towards the ban or the OneLove Binde or towards the original Tweet?

Computer Supported Cooperative Work, CSCW '13, page 897–908, New York, NY, USA. Association for Computing Machinery.

Nora Alturayef, Hamzah Luqman, and Moataz Ahmed. 2023. A systematic review of machine learning techniques for stance detection and its applications. *Neural Computing and Applications*, 35(7):5113–5144.

BBC News. 2010. [England miss out to Russia in 2018 World Cup vote](#).

Paul Michael Brannagan and Richard Giulianotti. 2018. Soft power and soft disempowerment: Qatar, global sport and football’s 2022 world cup finals. In *Leveraging Mega-Event Legacies*, pages 89–105. Routledge.

Luna De Bruyne, Toni G.L.A. van der Meer, Orphée De Clercq, and Véronique Hoste. 2024. [Using state-of-the-art emotion detection models in a crisis communication context](#). *Computational Communication Research*, 6(1):1.

Jonathan Crane. 2022. [HRW report alleges abuse of LGBTQ qataris](#).

Aron Culotta, Nirmal Kumar, and Jennifer Cutler. 2015. [Predicting the demographics of Twitter Users from website traffic data](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 29(1).

Nhan Cach Dang, María N Moreno-García, and Fernando De la Prieta. 2020. Sentiment analysis based on deep learning: A comparative study. *Electronics*, 9(3):483.

Luna De Bruyne, Orphee De Clercq, and Veronique Hoste. 2021. Emotional robbert and insensitive bertje: Combining transformers and affect lexica for dutch emotion detection. In *Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 257–263.

Jan Dörner. 2022. [Katar: So viele Deutsche wollen die WM boykottieren](#).

DW. 2022. [OneLove campaign hit by threat of FIFA sanctions](#).

Minghui Fan, Andrew Billings, Xiangyu Zhu, and Panfeng Yu. 2020. Twitter-based BIRGing: Big data analysis of english national team fans during the 2018 FIFA World Cup. *Communication & Sport*, 8(3):317–345.

FIFA. 2023. [One month on: 5 billion engaged with the FIFA World Cup Qatar 2022™](#).

Wei Gong, Ee-Peng Lim, Feida Zhu, and Pei Hua Cher. 2016. On unravelling opinions of issue specific-silent users in social media. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 10, pages 141–150.

- Maarten Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794*.
- Ahmed A. M. Hassan and Jia Wang. 2023. The qatar world cup and twitter sentiment: Unraveling the interplay of soft power, public opinion, and media scrutiny. *International Review for the Sociology of Sport*, 0(0).
- Daniela Heerdt and Lucas Roorda. 2023. Lessons learned in Qatar: The role of the Netherlands and its businesses in addressing human rights abuses in mega-sporting events. *Netherlands International Law Review*, pages 1–46.
- Human Rights Watch. 2022. Qatar: Security forces arrest, LGBT people.
- Human Rights Watch. 2023. Qatar: Six months post-World Cup, migrant workers suffer.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- KNVB Media. 2022. Ten countries embrace the Dutch OneLove campaign.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. <https://arxiv.org/abs/2205.11916>.
- Yasemin Lheureux. 2023. Predictive insights: leveraging twitter sentiments and machine learning for environmental, social and governance controversy prediction. *Journal of Computational Social Science*, pages 1–22.
- Eden Litt and Eszter Hargittai. 2016. The imagined audience on social network sites. *Social Media+ Society*, 2(1):2056305116633482.
- Mika V Mäntylä, Daniel Graziotin, and Miikka Kuutila. 2018. The evolution of sentiment analysis—a review of research topics, venues, and top cited papers. *Computer Science Review*, 27:16–32.
- Andreas Matzarakis and Dominik Fröhlich. 2015. Sport events and climate for visitors—the case of FIFA World Cup in Qatar, 2022. *International Journal of Biometeorology*, 59:481–486.
- Walaa Medhat, Ahmed Hassan, and Hoda Korashy. 2014. Sentiment analysis algorithms and applications: A survey. *Ain Shams engineering journal*, 5(4):1093–1113.
- Henk Erik Meier, Michael Mutz, Julia Glathe, Malte Jetzke, and Martin Hölzen. 2021. Politicization of a contested mega event: The 2018 FIFA World Cup on twitter. *Communication & Sport*, 9(5):785–810.
- Merkur. 2022. Sollte die deutsche Politik die Katar-WM boykottieren? Exklusive Bundestags-Umfrage zeigt gespaltenes Bild.
- Martin Müller, David Gogishvili, Sven Daniel Wolfe, Christopher Gaffney, Miriam Hug, and Annick Leick. 2023. Peak event: the rise, crisis and potential decline of the Olympic Games and the World Cup. *Tourism Management*, 95:104657.
- NDR. 2022. Umfrage zur WM in Katar: Mehrheit im Norden für Boykott.
- Luis Fernando Nuñez Franco. 2023. On sentiment analysis of Twitter content related to the FIFA World CUP 2022 in Qatar.
- Alexander Pak and Patrick Paroubek. 2010. Twitter as a corpus for sentiment analysis and opinion mining. In *LREc*, volume 10, pages 1320–1326.
- Tariq Panja and Kevin Draper. 2020. U.S. says FIFA officials were bribed to award World Cups to Russia and Qatar.
- Ravikumar Patel and Kalpdrum Passi. 2020. Sentiment analysis on twitter data of world cup soccer tournament using machine learning. *IoT*, 1:218–239.
- George Ramsay and Zayn Nabbi. 2022. England’s Harry Kane and several other European captains told not to wear ‘OneLove’ armband at World Cup.
- Furqan Rustam, Madiha Khalid, Waqar Aslam, Vaibhav Rupapara, Arif Mehmood, and Gyu Sang Choi. 2021. A performance comparison of supervised machine learning models for Covid-19 tweets sentiment analysis. *PloS one*, 16(2):e0245909.
- Kristen Scott, Pieter Delobelle, and Bettina Berendt. 2021. Measuring shifts in attitudes towards covid-19 measures in belgium. *Computational Linguistics in the Netherlands Journal*, 11:161–171.
- Luke Sloan, Jeffrey Morgan, Pete Burnap, and Matthew Williams. 2015. Who tweets? deriving the demographic characteristics of age, occupation and social class from twitter user meta-data. *PloS one*, 10(3):e0115545.
- Luke Sloan, Jeffrey Morgan, William Housley, Matthew Williams, Adam Edwards, Pete Burnap, and Omer Rana. 2013. Knowing the tweeters: Deriving socio-logically relevant demographics from twitter. *Sociological Research Online*, 18(3):74–84.
- Sportschau. 2022. Erneut Fan-Protteste gegen WM in Katar,.
- Zeerak Talat, Smarika Lulz, Joachim Bingel, and Isabelle Augenstein. 2021. Disembodied machine learning: On the illusion of objectivity in nlp. *arXiv preprint arXiv:2101.11974*.
- Anthony J Viera, Joanne M Garrett, et al. 2005. Understanding interobserver agreement: the kappa statistic. *Fam med*, 37(5):360–363.

ZDF heute. 2022. [DFB-Team hält sich den Mund zu.](#)

Lu Zhou, Wenbo Wang, and Keke Chen. 2016. Tweet properly: Analyzing deleted tweets to understand and identify regrettable ones. In *Proceedings of the 25th International Conference on World Wide Web*, pages 603–612.

A Stopwords

- der
- dem
- das
- und
- er
- ein
- den
- ist
- es
- die
- ich
- man
- zu
- nicht
- so
- wie
- was
- auch
- aber
- wenn
- als
- noch
- mal
- sich
- dass
- nur
- oder
- dann

B Prompts

```
1 [INST] Rate German Tweets as 'supportive', 'against' or 'neutral' of wearing the
   One Love bracelet or one love binde in German. These Tweets were posted
   during the world cup in Qatar.
2
3 Coding instructions:
4 - Critique for the FIFA policy of banning the bracelet or wishing players could
   wear it should be interpreted as support for the one love binde.
5 - Similarly, if support for the one love binde is slightly implied, that is
   sufficient for it to be supported.
6 - If the tweet is factual without any sentiment, consider it neutral.
7 - If the tweet uses negative connotations with the one love binde and what it
   stands for (the LGBT community), the tweet is against.
8
9 Tweet: "lul und ich werde angefeindet weil ich den tv anmache hugo lloris
   weigert sich bei der anstehenden fußball weltmeisterschaft in katar die one
   love binde zu tragen er wolle die kultur des gastgebers respektieren so die
   begründung des frankreich kapitäns boykottqatar" [/INST]
10 Support for wearing the one love binde: supportive
11 [INST] Tweet: "dfb torwart neuer hält an one love binde bei wm in katar fest" [/
   INST]
12 Support for wearing the one love binde: neutral
13 [INST] Tweet: "soso die scheiß will wohl die onelove binde verbieten ihr habt
   echt nicht mehr alle latten am zaun infantino und seine geldsäcken sollte man
   in die wüste schicken fifaworldcup qatarworldcup qatar dopa" [/INST]
14 Support for wearing the one love binde: against
15 [INST] Tweet: "{tweet}" [/INST]
16 English translation: {translation}
17 [INST] Now look at the previous tweets and analyze the similarities with the
   latest tweet that we want to label and if that tweet expresses support for
   wearing the one love binde. Let's think step by step:[/INST] {reason}
18 [INST] Therefore the tweet expressed the following sentiment towards wearing the
   one love binde:[\INST] {decision}
```

Listing 1: Prompt for Mistral-7B