

Taking sides: Public Opinion over the Israel-Palestine Conflict in 2021

Arsal Imtiaz, Danish Khan, Hanjia Lyu, Jiebo Luo

University of Rochester, Rochester, New York

{aimtiaz2, dkhan5, hlyu5}@ur.rochester.edu, jluo@cs.rochester.edu

Abstract

The Israel-Palestine Conflict, one of the most enduring conflicts in history, dates back to the start of 20th century, with the establishment of the British Mandate in Palestine and has deeply rooted complex issues in politics, demography, religion, and other aspects, making it harder to attain resolve. To understand the conflict in 2021, we devise an observational study to aggregate stance held by English-speaking countries. We collect Twitter data using popular hashtags around and specific to the conflict portraying opinions neutral or partial to the two parties. We use different tools and methods to classify tweets into pro-Palestinian, pro-Israel, or neutral. This paper further describes the implementation of data mining methodologies to obtain insights and reason the stance held by citizens around the conflict.

Introduction

With the increasing usage of social media in the political context, Twitter maintains its stature as one of the top platforms to express political opinions (Jungherr 2014). It is oriented towards facilitating digital political advertising in addition to playing an active role in how candidates use them to communicate with voters (Kreiss and McGregor 2018). It has also allowed users to stay updated with daily news and politicians to effectively mobilize protesters (Morselli, Passini, and McGarty 2021).

The Israel-Palestine conflict is one of the longest-running and most controversial conflicts in the world.¹ The spark between the Arabs and Jews in Palestine can be traced back to the establishment of British Mandate in 1917 during the British Occupation of Palestine (Tsimhoni 1984). Since then, multiple violent confrontations took place between Jews and Arabs in Palestine, costing thousands of lives from both sides and making it an important issue for the world to address (Ginat 2018).

Housing policy and design have always favored the shaping of the Jewish state and colonial domination in Israel further (Haas 2021). This framework used by the Zionist Israeli regime to counter insurgent practises of Palestinians in the

country sparked the conflict in 2021.² It resulted in a number of events which escalated the conflict as a Human Rights issue by the United Nations.

This study attempts to understand the response to the current Israel-Palestine conflict by English-speaking countries. The importance lies in the dissection of controversies around the conflict. The Council on Foreign Relations warned that a third intifada could break out and that renewed tensions will escalate into large-scale violence.³ Being updated about daily happenings and opinions around the conflict might untangle controversies and prevent the said intifada. The proposed methodology in the study provides an aggregated stance of English-speaking countries and gathers insights into the topics discussed by people pertaining to these countries on the grounds of their support towards Israel/Palestine.

The study makes use of modern day techniques in data mining and linguistics to study opinions based on individuality and nationality. The unresolved status of the subject instigates the need to perform a study with modern day techniques in Data Science and related fields.

Related Work

Twitter as a Political Medium

There is an ever growing use of Twitter globally and politics plays a key role in this growth. With the rapid increase of the use of social media in the 2010s, analysis of data within these platforms has become one of the adapted methodologies to gather insights. Political participation in Twitter is performed both by politicians and citizens (Polat 2005). The methodology mentioned in Sabatovych (2019) provides insights into “social media revolution” and “democratic revolution” by studying the Maidan Revolution using Twitter data of Ukrainians. This study also aims to study public opinion in 2021 toward the Israel-Palestine conflict to gather insights about the issue.

Stance Detection

Stance detection is the task of automatically determining from text whether the author of the text is in favor of, against,

¹https://www.business-standard.com/article/international/israel-palestine-conflict-history-guide-to-world\%E2%80%99s-longest-running-jews-muslim-arab-war-121051300886_.html

²<https://www.vox.com/22440330/israel-palestine-gaza-airstrikes-hamas-updates-2021>

³<https://www.cfr.org/global-conflict-tracker/conflict/israeli-palestinian-conflict>

or neutral toward a proposition or target (Mohammad, Sobhani, and Kiritchenko 2017). Stance on social media data can be modeled using various online signals. These signals can be categorized into two main types, namely, (1) content signals, such as the text of the tweets, and (2) network signals, such as the users' connections and interactions on their social networks (ALDayel and Magdy 2021). In this paper, we study the former and attempt to determine stance on tweets pertaining to a bounded geographical location. In the socio-political domain, ALDayel and Magdy (2021) show that there is an orthogonal relation between sentiment and stance but studies such as Li and Caragea (2019) show that sentiment can be used as a proxy to determine stance. In this study, we implement and compare transformer-based stance detection models and the models relying on sentiment to detect stance.

Israel-Palestine Conflict and Social Media

Although there is limited research on the Israel-Palestine conflict using data mining, the studies conducted on this subject have proposed efficient methods for obtaining insights. Zeitzoff (2011) collected data from social media and implemented vector autoregression to measure conflict dynamics between the Israeli Air Force (IAF) and Hamas during the Gaza conflict of 2008. Alagha and Dahrooj (2019) also conducted a study on the Israel-Palestine conflict by using modern NLP tools, such as ArkTweetNLP and LingPipe, to analyze individual and national opinions on the Israel-Palestine conflict in the past. The study presented in this paper uses a similar methodology but focuses on the context of the Israel-Palestine conflict in 2021 and considers a larger dataset. We also make use of visualization tools and methods to obtain inferences regarding support.

Problem Statement

The objective of this paper is to study stance in English-speaking countries around the Israel-Palestine conflict in 2021. It further attempts to reason support towards Israel/Palestine in English-speaking countries with the help of techniques and methods in Natural Language Processing.

Methodology

This section of the paper discusses the pipeline for the study. As shown in Figure 1, the data acquired is passed through multiple stages presented in the pipeline to obtain semantic understanding and insights around the conflict. The stages of the pipeline are explained in detail in the subsections below.

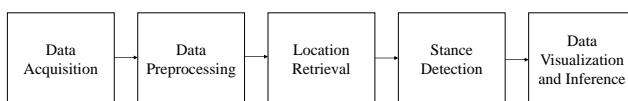


Figure 1: Pipeline of the analysis.

Data Acquisition

Social media platforms such as Facebook, Twitter, Reddit, *etc.*, can be used to convey different forms of messages, and hence, their role in propagating political information is not surprising (Zhuravskaya, Petrova, and Enikolopov 2020). The findings of Bode (2016) suggested that “Twitter users gain in political knowledge by the virtue of their use of that medium”, which supports the argument regarding the political nature of Twitter.

For this study, we have considered Twitter data because of its speed, immediacy, ease of use via hashtags, and its global reach (Schreiner 2018). Twitter API⁴ is used to scrape english tweets around the timeline in the format shown in Table 1. The API, under academic access allows us to scrape tweets at a faster rate with a higher number of API calls, and provides more information about each tweet and its associated user.

Field	Description
text	Tweet text content
language	Tweet language
tweet_id	Unique tweet ID
source	Device used for posting the tweet
conversation_id	Unique ID for every conversation
retweet_count	Number of retweets
reply_count	Number of replies to the tweet
like_count	Number of likes of the tweet
quote_count	Number of times the tweet has been quoted
author_id	Unique user ID
created_at	Datetime of the posted tweet

Table 1: Summary of the tweet fields.

Data acquisition was carried out in primarily two phases with respect to the timeline shown in Figure 2. First, a list of popular hashtags regarding the conflict is curated. A small sample of tweets of roughly 1,500-2,000 tweets per day are extracted over the entire timeline of the conflict. The hashtags specific to the study such as #freepalestine, #antisemitism, #istandwithpalestine, #istandwithisrael, #jerusalemfightsback, *etc.* are used. The hashtags acquired in the first phase are then used as parameters of acquiring the larger dataset.

The API configuration is set to scrape 200 tweets every ten minutes of the entire timeline. A total of 493,877 tweets ranging from April 13th, 2021 to May 23rd, 2021 are collected.

Data Preprocessing

Data preprocessing is an essential step to formulate the data in a form suitable for stance detection. In this stage, we will discuss data cleaning, demojizing emoticons, part-of-speech subject estimation, *etc.*.

Data cleaning for our data includes removal of stop-words, hyperlinks, HTML escape sequences, white-spaces, *etc.* with regular expressions and NLTK libraries in Python. Hashtags on the other hand are retained to preserve meaning

⁴<https://developer.twitter.com/en/docs/twitter-api>

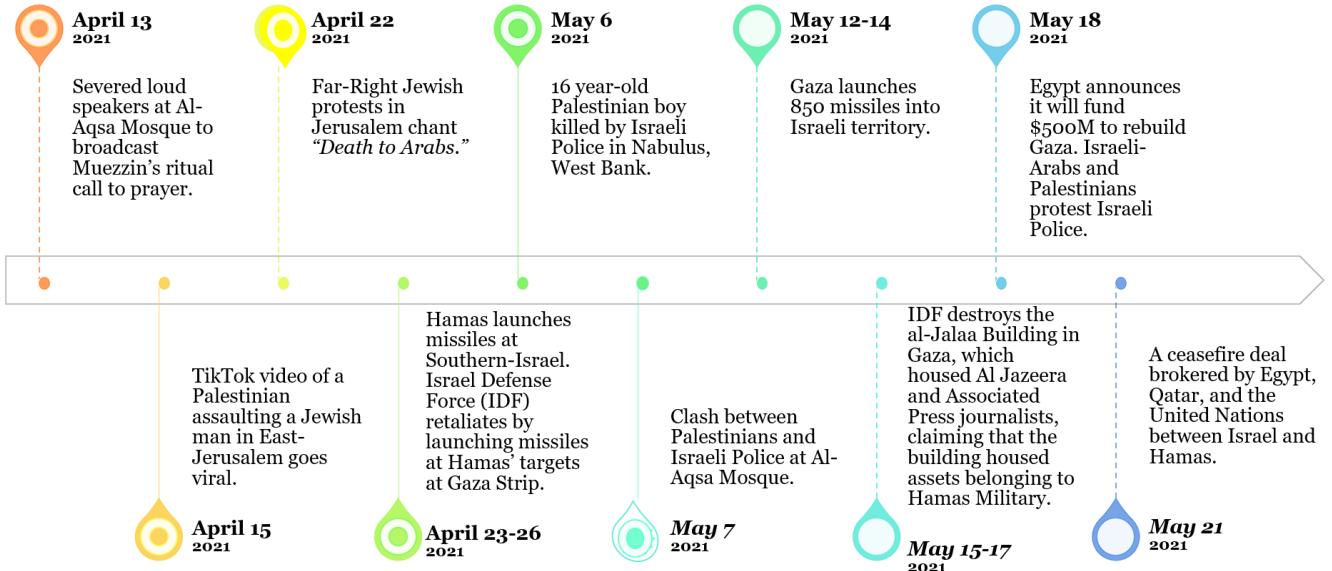


Figure 2: The timeline of the Israel-Palestine Conflict in 2021.

in the tweets (Almatrafi, Parack, and Chavan 2015). For instance, “Save children in #Gaza” would have lost its context if "#Gaza" is removed entirely from the tweet.

Soranaka and Matsushita (2012) stressed on the possibility that users may emphasize emotions by the use of emoticons along with co-occurring emotional words or complement the emotional words presented in the tweet. Hence, we demojize emoticons in the tweet to retain meaning. Emojo,⁵ a python library is used to convert emoticons to text such as happy, sad, angry, excited, etc.

Gimpel et al. (2011) recognised part-of-speech tagging as one of the most fundamental parts in a linguistic pipeline that often leads to a richer linguistic analysis. Therefore, we use SpaCy part-of-speech tagging across our data to estimate the subject of the tweet.

Location Retrieval

Only an estimated 2% of the tweets are geo-tagged with latitude and longitude coordinates, and there is a need for methods to assign geo-locations to each tweet and its associated user (Schlosser, Toninelli, and Cameletti 2021). Primarily we intend to extract user location by referring to the “geo” field of each tweet. If the tweet is geo-tagged, it is stored as a location in the location field. For the situations where the tweet is not geo-tagged, the user’s location, as mentioned in their bio is fetched using user fields from the API’s .json file. If the user has not mentioned their location in their bio, their location is stored as NULL. Python’s GeoPy library⁶ is used to process the locations to a universal country format for the entirety of data.

⁵<https://github.com/carpedm20/emoji/>

⁶<https://geopy.readthedocs.io/en/stable/>

Stance Detection

Twitter provides a great corpus of text for stance detection and opinion mining based on its sheer usage (Pak and Paroubek 2010). In this study, we have implemented stance detection to classify tweets into three classes - pro-Palestine, pro-Israel, and neutral via two different approaches. The first approach uses pre-trained sentiment analysis tools such as VADER (Hutto and Gilbert 2014), LIWC,⁷ PySentimiento (Pérez, Giudici, and Luque 2021) as a part of the model. These tools provide an array of features on the data, and are clubbed with logistic regression/multi-layer feedforward neural network for text classification.

The structure of the multi-layer feedforward neural network is described in Table 3. The input of the network is configured to the outputs of VADER, LIWC and PySentimiento. There exists three hidden layers with 512, 256 and 64 neurons respectively. The output layer of the model has 3 neurons computing the probability for each class using a softmax function. The model is configured to have a learning rate of 0.001, batch size of 20, categorical cross entropy as the loss function, Adam as the optimizer and is trained for 500 epochs.

The second approach attempts to solve the problem of stance detection using pre-trained transformer-based models. We have fine-tuned an XLNet (Yang et al. 2019) model on our corpus of tweets to determine stance. The pre-trained model used for this study is *xlnet-base-cased*. The model is configured to have three outputs for stance detection and is trained for 10 epochs with a batch size of 32.

For training and testing models, we have manually annotated 979 tweets into three classes - pro-Palestine (n=510), pro-Israel (n=290) and neutral (n=179). The labeled data is then split to training and testing sets in the ratio of 3:1.

⁷<https://www.liwc.app/>

Model	Accuracy	Precision			Recall			F1-Score		
		Pro-Palestine	Neutral	Pro-Israel	Pro-Palestine	Neutral	Pro-Israel	Pro-Palestine	Neutral	Pro-Israel
Vader with Logistic Regression	0.63	0.70	0.44	0.58	0.69	0.19	0.81	0.70	0.27	0.68
Vader with Neural Network	0.63	0.71	0.32	0.57	0.77	0.17	0.64	0.74	0.22	0.60
LIWC with Logistic Regression	0.66	0.76	0.28	0.72	0.76	0.28	0.73	0.76	0.28	0.72
LIWC with Neural Network	0.60	0.61	0.39	0.63	0.78	0.19	0.54	0.68	0.25	0.58
PySentimiento with Logistic Regression	0.65	0.74	0.28	0.69	0.74	0.26	0.71	0.74	0.27	0.70
PySentimiento with Neural Network	0.67	0.65	0.62	0.71	0.89	0.19	0.60	0.76	0.20	0.65
XLNet	0.82	0.84	0.59	0.91	0.82	0.56	0.97	0.83	0.57	0.94

Table 2: Model comparison: accuracy, precision, recall and F1-score.

Layer	Activation Function	Output Shape
Dense	relu	(None, 512)
Dense	relu	(None, 256)
Dense	relu	(None, 64)
Dense	softmax	(None, 3)

Table 3: Multi-Layer Feedforward Neural Network Structure.

Tweet Metric	Value
Average number of words per tweet	21.63
Median number of words per tweet	21.63
Minimum number of words	1
Maximum number of words	133
Standard deviation of the number of words	12.702
Total number of tweets	493,877
Number of pro-Israel tweets	14,982
Number of neutral tweets	250,109
Number of pro-Palestine tweets	228,813

Table 4: Descriptive statistics of the tweets.

The models are passed through the testing set and their accuracy, precision, recall, F1-score are reported in Table 2. Our results confirm the findings of ALDayel and Magdy (2021) that stated the relation between sentiment and stance as orthogonal.

Data Visualization and Inferences

People are generally more engaged when information is represented in a space and grid form in comparison to representation in the form of lists and statistics. Visualizing opinions in a political context may enhance intentional and deliberate cognition, and interpretations, whilst mitigating framing effects and subconscious processes (Faridani et al. 2010). In this study, we have made use of *Tableau* and *Wordart*⁸ to generate bar plots and wordclouds.

Results

This section provides statistical information about the data in Table 5 and attempts to reason stances held by English-speaking countries. Given the performances reported in Table 2, all the stance-related statistics in Table 4 are based on the stance determined by the XLNet model.

We have set a threshold of 10,000 tweets per country to consider their stance in our study. Table 5 provides the to-

tal number of tweets by the four most actively opinionated English-speaking countries in this conflict. Figure 3 shows a clear support for Palestine in the discussed conflict for these countries.

Country	No. of Tweets
United States	38,418
Pakistan	23,346
United Kingdom	22,611
India	19,543

Table 5: Ranking of countries in terms of expressing opinions around the recent conflict.



Figure 3: Number of pro-Israel and pro-Palestine tweets of the top four countries.

Wordclouds are generated to visualize topics that Palestine/Israel supporters are talking about. We attempt to reason their opinions and juxtapose the events that took place. The highly frequent words attained from the wordclouds can be queried in the data to understand and reason support among pro-Israeli/pro-Palestinian supporters. The inferences attained from pro-Palestinian tweets can be assumed to sound the opinion of the masses but the inference attainable from pro-Israeli tweets poses a higher possibility of opinions to be specific and deep rooted in culture. We consider the case of United States as an example in this study. The same methodology can be applied to other countries to attain inferences as well.

With the United States being the major contributor of the data, the finding presents some interesting insights. According to Figure 4, *freepalestine*, *Gaza*, *Savesheikhjarrah*, etc. are some of the frequent words that have occurred amongst Palestinian supporters in the United States. On querying tweets with the words obtained above, there is a clear expression of contempt toward the events within the timeline.

⁸<https://wordart.com/>



Figure 4: Pro-Palestinian tweets in the United States.

For example, there have been tweets that depict the condemnation of children's death in Gaza.

Upon further querying and studying tweets, a number of Palestinian supporters expressed discontentment toward the President and the Vice President of the United States. The supporters were not able to accept with the fact that US President approved sale of weapons worth \$735M to Israel.⁹ The sale of weapons to Israel was a counterproductive move in the efforts of waning and stopping the conflict.

Similarly, opinions around Israeli supporters can also be drawn with this methodology. Israeli supporters expressed opinions about topics in the timeline that are generally associated with negative connotations. Figure 5 shows a few topics - *Hamas, terrorist, rocket, attack, etc.* that are frequent among pro-Israeli tweets. On querying, there is a strong possibility that the condemnation rises from the hate of terrorism - *Hamas* in this case.



Figure 5: Pro-Israeli tweets in the United States.

Conclusion

In this study, we propose a methodology to present insights into the public opinion toward the conflict at hand. It possesses the ability to obtain inferences about stance toward Israel and Palestine in accordance to real time events and evidence with respect to individual nations. With the acquisition of data in such a distinct domain, we present data in visuals encompassing the trends of support in the conflict with events and occurrences around the conflict in the United States. Our study supports the claim by Hosch-Dayican et al. (2016) of using Twitter more as an outlet for expressing discontent than as a medium for negative campaigning. However, the basis of our study on English-speaking countries excludes the valued opinions of non-English speaking countries, especially near Israel-Palestine which would have an impact on the study. Inclusion of the said tweets would have provided greater clarity on the situation based on multi-lingual stance detection.

From a technical standpoint, we find that the transformer-based model - XLNet outperforms off-the-shelf hybrid implementations of sentiment analysis tools in political domains specific to conflicts between parties, which is in line with the findings of Bestvater and Monroe (2022) that political opinions are highly complex and sentiment analysis methods may not capture stance. There exists a class imbalance issue in the data, with pro-Israel tweets being very few in comparison to pro-Palestine tweets, which can be addressed in future work. Multi-lingual models can be considered in the future to take other popular languages such as Arabic, Hebrew, etc. into account. This would help us process and understand data around the conflict from a global perspective.

Broader Impact and Ethics Statement

The study possesses the potential of gathering numerous insights on a broader scale. Considering specific time frames around heated events provides detailed voicing of people irrespective of their involvement in the conflict. It provides a mechanism into understanding why people support Israel/Palestine in the conflict. However, patterns and insights attained from the conflict can also be used in a detrimental way to filter out tweets by organizations or governments.

According to the Twitter Developer Agreement and Policy, the organization prohibits the use of data in any way that would breach user privacy. For non-commercial purposes such as academic research, Twitter provides some liberty in distribution of content but limits the scope to distributing User IDs, Tweet IDs and Direct-Message IDs only. Making the data with a broader array of information such as tweets, geo-data and other personal information puts the licensee and the organization in jeopardy for potential breach of privacy. Sensitive Twitter data is collected in the study for gathering insights at an aggregate level without disclosing any personal or sensitive information about users.

References

Alagha, I.; and Dahrooj, O. 2019. Mult-Level Analysis of Political Sentiments Using Twitter Data: A Case Study of

⁹<https://www.washingtonpost.com/politics/2021/05/17/power-up-biden-administration-approves-735-million-weapons-sale-israel-raising-red-flags-some-house-democrats/>

- the Palestine-Israeli Conflict. *Jordanian Journal of Computers and Information Technology* 05: 1.
- ALDayel, A.; and Magdy, W. 2021. Stance detection on social media: State of the art and trends. *Information Processing & Management* 58(4): 102597. ISSN 0306-4573.
- Almatrafi, O.; Parack, S.; and Chavan, B. 2015. Application of Location-Based Sentiment Analysis Using Twitter for Identifying Trends towards Indian General Elections 2014. In *Proceedings of the 9th International Conference on Ubiquitous Information Management and Communication, IMCOM '15*. New York, NY, USA: Association for Computing Machinery. ISBN 9781450333771.
- Bestvater, S. E.; and Monroe, B. L. 2022. Sentiment is Not Stance: Target-Aware Opinion Classification for Political Text Analysis. *Political Analysis* 1–22.
- Bode, L. 2016. Political News in the News Feed: Learning Politics from Social Media. *Mass Communication and Society* 19(1): 24–48.
- Faridani, S.; Bitton, E.; Ryokai, K.; and Goldberg, K. 2010. Opinion Space: A Scalable Tool for Browsing Online Comments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10*, 1175–1184. New York, NY, USA: Association for Computing Machinery. ISBN 9781605589299.
- Gimpel, K.; Schneider, N.; O'Connor, B.; Das, D.; Mills, D.; Eisenstein, J.; Heilman, M.; Yogatama, D.; Flanigan, J.; and Smith, N. 2011. Part-of-Speech Tagging for Twitter: Annotation, Features, and Experiments. volume 2, 42–47.
- Ginat, A. 2018. British Mandate for Palestine .
- Haas, O. 2021. *Housing Policy, Design, And Struggle: The Colonial Production of Space of Israel/Palestine, One New City at a Time*. Ph.D. diss., Dept. of Environmental Studies, York University, Toronto, Ontario.
- Hosch-Dayican, B.; Amrit, C.; Aarts, K.; and Dassen, A. 2016. How Do Online Citizens Persuade Fellow Voters? Using Twitter During the 2012 Dutch Parliamentary Election Campaign. *Social Science Computer Review* 34(2): 135–152.
- Hutto, C.; and Gilbert, E. 2014. VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. *Proceedings of the International AAAI Conference on Web and Social Media* 8(1): 216–225.
- Jungherr, A. 2014. Twitter in Politics: A Comprehensive Literature Review. *SSRN Electronic Journal* .
- Kreiss, D.; and McGregor, S. C. 2018. Technology Firms Shape Political Communication: The Work of Microsoft, Facebook, Twitter, and Google With Campaigns During the 2016 U.S. Presidential Cycle. *Political Communication* 35(2): 155–177.
- Li, Y.; and Caragea, C. 2019. Multi-Task Stance Detection with Sentiment and Stance Lexicons. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 6299–6305. Hong Kong, China: Association for Computational Linguistics.
- Mohammad, S. M.; Sobhani, P.; and Kiritchenko, S. 2017. Stance and Sentiment in Tweets. *ACM Trans. Internet Technol.* 17(3). ISSN 1533-5399.
- Morselli, D.; Passini, S.; and McGarty, C. 2021. Sos Venezuela: an analysis of the anti-Maduro protest movements using Twitter. *Social Movement Studies* 20(5): 509–530.
- Pak, A.; and Paroubek, P. 2010. Twitter as a Corpus for Sentiment Analysis and Opinion Mining. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*. Valletta, Malta: European Language Resources Association (ELRA).
- Polat, R. K. 2005. The Internet and Political Participation: Exploring the Explanatory Links. *European Journal of Communication* 20(4): 435–459.
- Pérez, J. M.; Giudici, J. C.; and Luque, F. 2021. pysentimiento: A Python Toolkit for Sentiment Analysis and SocialNLP tasks.
- Sabatovich, I. 2019. Do social media create revolutions? Using Twitter sentiment analysis for predicting the Maidan Revolution in Ukraine. *Global Media and Communication* 15(3): 275–283.
- Schlosser, S.; Toninelli, D.; and Cameletti, M. 2021. Comparing Methods to Collect and Geolocate Tweets in Great Britain. *Journal of Open Innovation: Technology, Market, and Complexity* 7: 44.
- Schreiner, T. 2018. Information, Opinion, or Rumor? The Role of Twitter During the Post-Electoral Crisis in Côte d'Ivoire. *Social Media + Society* 4(1): 2056305118765736.
- Soranaka, K.; and Matsushita, M. 2012. Relationship between Emotional Words and Emoticons in Tweets. In *2012 Conference on Technologies and Applications of Artificial Intelligence*, 262–265.
- Tsimhoni, D. 1984. The Status of the Arab Christians Under the British Mandate in Palestine. In *Middle Eastern Studies*, 166–192.
- Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J. G.; Salakhutdinov, R.; and Le, Q. V. 2019. XLNet: Generalized Autoregressive Pretraining for Language Understanding. *CoRR* abs/1906.08237.
- Zeitzoff, T. 2011. Using Social Media to Measure Conflict Dynamics: An Application to the 2008–2009 Gaza Conflict. *Journal of Conflict Resolution* 55(6): 938–969.
- Zhuravskaya, E.; Petrova, M.; and Enikolopov, R. 2020. Political Effects of the Internet and Social Media. *Annual Review of Economics* 12(1): 415–438.