

Factoring Hate Speech: A New Annotation Framework to Study Hate Speech in Social Media

Gal Ron^{1†}, Effi Levi^{2†}, Odelia Oshri¹, Shaul R. Shenhav¹

¹Department of Political Science, The Hebrew University of Jerusalem

²Institute of Computer Science, The Hebrew University of Jerusalem

{gal.ron2|odelia.oshri|shaul.shenhav}@mail.huji.ac.il

efle@cs.huji.ac.il

Abstract

In this work we propose a novel annotation scheme which factors hate speech into five separate discursive categories. To evaluate our scheme, we construct a corpus of over 2.9M Twitter posts containing hateful expressions directed at Jews, and annotate a sample dataset of 1,050 tweets. We present a statistical analysis of the annotated dataset as well as discuss annotation examples, and conclude by discussing promising directions for future work.

1 Introduction

Social media has come to constitute a space for the propagation of hostility (see ElSherief et al., 2018, p. 1) and provides fertile grounds for the radicalization of individuals in support of violent extremist groups (Reynolds and Tuck, 2016; Mitts, 2019). Much research has been devoted to automatically identifying hate speech in online forums (Mathew et al., 2021; Kim et al., 2022; Wiegand et al., 2022) along with factors that facilitate its propagation (Scharwächter and Müller, 2020; Newman et al., 2021), and classifying different forms of hateful and abusive content (Davidson et al., 2017; Founta et al., 2018). However, the very concept of “hate” and its presence in written text is somewhat illusive and amorphous (Fortuna et al., 2020); therefore, some efforts have been made to define different typologies of hate speech (Waseem et al., 2017; Founta et al., 2018; Davidson et al., 2017; Mathew et al., 2021).

In this work, we address this important subject through the specific case of hate speech directed towards Jews in Twitter posts. Rather than attempting to classify hate speech into different types, we present a novel annotation scheme which factors hate speech into a comprehensive set of separate

discursive aspects that often appear in hate speech, capturing its intricate and diverse nature.

This factorization aims to achieve several goals: first, make the annotation process more focused and accurate, by decomposing the amorphous and ambiguous concept of “hate” into more specific and narrowly defined discourse aspects, rendering the annotation process more objective. Second, it allows exploring and analyzing hate speech across these different aspects, hopefully leading to a deeper understanding of its complexities, variety, and nuance. Furthermore, this set of aspects defines various possible distinct combinations, each of which encodes a different and unique configuration of hate speech. Although this annotation scheme was designed to capture and characterize hate speech directed towards Jews, with the exception of one group-specific aspect, it is general enough to be applied to any other group-directed hate speech.

We constructed a corpus of Twitter conversations in English containing over 2.9M tweets, collected through Twitter API v2. In order to evaluate our annotation scheme on real Twitter posts, we used it to annotate a sample of 1,050 tweets taken from the corpus. We present a quantitative analysis of the annotated dataset, as well as a qualitative one (through the use of some examples). We conclude by discussing several directions to extend and develop our work.

Content Warning: This document contains some examples of hateful content. This is strictly for the purpose of enabling this research. Please be aware that this content could be offensive and cause you distress.

2 Tweets Corpus

The tweets were extracted through the Twitter API v2 using the tweepy python module ¹. We applied

[†]Both authors contributed equally to this work.

¹<https://github.com/tweepy/tweepy>

	Single Tweets	Conversations		Total	Conversation Length	
	# Tweets	# Conversations	# Tweets	# Tweets	Mean	STD
Neutral	601,917	109,172	1,005,095	1,607,012	9.21	71.73
Racial	527,541	97,730	788,576	1,316,117	8.07	127.43
Total	1,129,458	206,902	1,793,671	2,923,129		

Table 1: Complete tweet corpus statistics. “Single Tweets”: tweets that were posted as new tweets rather than as a reply, and were not replied to.

for, and were granted, an Academic Research Access to the API², which offers a full-archive access to public data posted on the platform, going back to April 2006.

To increase the likelihood of retrieving tweets that contain expressions of hate towards Jews, we used two types of keyword-based filters in our queries: *neutral* keywords and *racial* keywords. The neutral stop list – containing 14 words – was compiled from keywords referenced in previous studies (Gunther et al., 2021; Chandra et al., 2021). The racial stop list – containing 28 words and expressions – was compiled using the Hatebase database, a multilingual lexicon for racial terms³, by extracting from the database all the English terms pertaining to Jews and Judaism that had at least one sighting.

Following preliminary experimentation with the API directed at increasing the chances of retrieving a conversation (thread) containing Jew-related hate expressions, we decided to focus on collecting conversations which stemmed from a “source” tweet (a new post rather than a reply to another tweet) adhering to our keyword filters. We devised the following 2-step process. Given a specific date (24-hour interval) and a specific keyword filter:

1. Query the API for English “source” tweets containing any of the keywords in the filter, posted within the specified date.
2. For every tweet extracted in step 1: if the tweet was replied to, query the API for all the available tweets in the resulting conversation (some tweets, such as deleted or private tweets, were not available for extraction).

For each date between July 1st 2018 and June 30th 2022 (defining a period of exactly 4 years), we applied the procedure with the *racial* keywords filter and collected as many tweets as possible. Then,

²No longer available, as of May 2023

³<https://hatebase.org>

we applied the same procedure with the *neutral* keyword filter to collect a similar number of tweets from the same date⁴. This was done in order to keep the corpus as balanced as possible between the two types of keyword filters.

The result is a large corpus of Twitter conversations started between July 1st 2018 and June 30th 2022, segmented by the type of the filter applied to the conversation’s source tweet (*neutral* or *racial* Jews-related keywords). Aside from the text itself, the corpus includes additional meta-data for each tweet: tweet ID, conversation ID, posting date, reply-to ID (if the tweet was written as a reply to another tweet), tweet statistics (retweets, replies, likes, quotes and views), place & country (if available), author ID (a unique identifier for the author of the tweet) and author statistics (followers, verified status). Note that the “conversation ID” and “reply-to ID” fields allow a complete hierarchical reconstruction of a conversation given any tweet from that conversation.

Statistics for the corpus are given in Table 1.

3 Hate Speech Annotation

3.1 Annotation Scheme

As discussed in Section 1, we have devised a novel annotation scheme with the goal of factoring hate speech into several separate aspects. The scheme encodes five different discursive categories, which are designed to capture the main recurring aspects of hate speech as employed and defined in previous studies (Kaid and Holtz-Bacha, 2007; Davidson et al., 2017; Arango et al., 2022; Khurana et al., 2022), as well as the discursive elements of hate speech that are described in the United Nations’ “Strategy and Plan of Action on Hate Speech”⁵. An aggregative definition, as suggested here, enables us to identify hate speech towards the target

⁴preliminary experiments showed that tweets containing the *neutral* keywords are significantly more abundant compared to the *racial* keywords

⁵<https://www.un.org/en/hate-speech/understanding-hate-speech/what-is-hate-speech>

	Single Tweets	Conversations		Total
	# Tweets	# Conversations	# Tweets	# Tweets
Neutral	263	82	263	526
Racial	262	82	262	524
Total	525	164	525	1,050

Table 2: Statistics for the sample dataset

group in a broad yet nuanced sense, while also differentiating between various forms of expression. Importantly, our annotation scheme aims to capture clear expressions of hate rather than mere hateful terms. As such, tweets containing the racial keywords which were used in our queries to the Twitter API were annotated under the relevant category only when it was plausible to suspect that these words were indeed employed to express hate.

Category	# Tweets
Contempt	5
Abuse	181
Call for Anti-Group Action	31
Prejudice	12
Holocaust Denial	4

Table 3: Annotation statistics for the sample dataset

The five categories are:

1. **Contempt** – speech that conveys a strong disliking of, or negative attitudes towards the targeted group, and does so in a neutral tone or form of expression.
2. **Abuse** – speech that demeans, degrades, vilifies, mocks, humiliates, or conveys general hostility that is expressed using emotionally-charged language.
3. **Call for Anti-Group Action** – an incitement of violence and/or discrimination against the target group.
4. **Prejudice** – the expression of negative thoughts/beliefs regarding the targeted group on the basis of the group’s characteristics, and/or (negative) monolithic references to the targeted group.
5. **Holocaust Denial** – the only category specific to our target group (Jews), this includes derecognition of the holocaust, or statements that recognize the fact that the holocaust happened

but degrade from its scope, mock it (and/or the people it hurt), and belittle its significance.

These discursive categories not only encompass substantive elements of hate, such as contempt and prejudice, but also address the manner in which negative discourse is conveyed, including abusive language and incitement of violence.

All the categories, except the fifth (Holocaust Denial), are general and may naturally be applied to other groups besides our target group (Jews). In addition, while the categories encode separate aspects of hateful discourse, they may conjointly characterize the same expression; for example, a post can be abusive while also expressing prejudice. Consequently, the annotation scheme defines a *multi-label* classification task.

3.2 Annotated Sample Dataset

For the purpose of conducting a preliminary analysis of our annotation scheme over real tweet data, we annotated a sample dataset of 1,050 tweets from the corpus described in Section 2. These tweets were sampled by iterating the dates backwards, starting from June 30th 2022. For each date, one conversation with a length of $2 \leq k \leq 10$ tweets was randomly selected, then k additional single tweets (1-tweet conversations) were randomly selected from the same date. This procedure was performed separately for each of the two keyword filter types (*neutral* and *racial*). The result is a collection of 1,050 tweets from conversations started between June 4th 2022 and June 30th 2022 (statistics are given in Table 2). Each tweet was encoded with a subset of the five possible categories (described in Section 3.1), including the empty set (none of the categories).

Table 3 shows the annotation statistics for the sample dataset. Note that despite the use of keyword filters to retrieve the tweets from the Twitter API, all five categories are generally sparse in the dataset. The most common category is Abuse, with 181 instances (out of the 1,050 tweets in the dataset). Given that half of these tweets were col-

	Contempt	Abuse	Call for Anti-Group Action	Prejudice
Abuse	-0.0316			
Call for Anti-Group Action	-0.0121	0.2332		
Prejudice	0.1227	0.1644	0.0342	
Holocaust Denial	-0.0043	0.0536	-0.0108	0.1388

Table 4: Inter-category Pearson’s correlations in the sample dataset

lected using the *racial* keywords filter, this is consistent with previous findings that hate speech is highly likely to contain racial slurs (Davidson et al., 2017); intuitively, it is the most “direct” way to express hate (among the five categories). It is important to note, however, that the mere presence of a racial slur does not automatically merit an Abuse annotation, as evident in the following example:

- (1) In the span of 5 min I have been both called a "toxic fan" for not liking the Kenobi show and a "Zionazi" by a fan of the show. Maybe you should reconsider who are the "toxic" ones (None)

The second most common category is Call for Anti-Group Action, followed by Prejudice and Contempt, with the most uncommon category being Holocaust Denial.

Table 4 displays the inter-category correlations (measured in Pearson’s r) in the sample dataset. In general, no considerable correlations were found between any two categories. Abuse was found to be somewhat correlated with Call for Anti-Group Action ($r = 0.2332$), and to a lesser degree with Prejudice ($r = 0.1644$). Naturally, calls for violence and prejudiced expressions are often accompanied by abusive language, for example:

- (2) Ahhhhhh the good old days , yids were bums then still bums now. (Abuse, Prejudice)

A minor correlation between Prejudice and Contempt ($r = 0.1227$) is possibly an indication that prejudiced perspectives serve as a kind of "rationale" for hate that does not always require the more emotional use of abusive language. This is demonstrated in the following example:

- (3) Perhaps America is just too fat, spoiled and lazy not to be noosed by the Jews into another low road oblivion that profits the jews. It

seems incapable of tweaking itself or nurturing and governing through its’ higher itself. Better perhaps to help it rot? (Prejudice, Contempt)

Another minor correlation between Prejudice and Holocaust Denial ($r = 0.1388$) may be attributed to the fact that both categories are closely associated with conspiracy theories.

4 Conclusion

In this paper, we address the task of capturing and characterizing hate speech directed towards Jews in Twitter posts. For that purpose, we devised a novel annotation scheme that encodes five different aspects of hate speech, four of which are not specific to our target group (Jews), allowing us to factorize the generally amorphous concept of “hate” into more concretely defined aspects. We utilized the Twitter API v2 to collect and assemble a corpus of Twitter conversations in English containing over 2.9M tweets, using two types of keyword filters (*neutral* and *racial*) to maximize the likelihood of retrieving tweets that contain expressions of hate towards Jews. We then used our annotation scheme to annotate a sample of 1,050 tweets, and demonstrated its potential contribution through select examples. We intend to make all of these resources (tweet corpus, annotation guidelines and sample dataset) available to the research community.

We are currently engaged in an ongoing effort to train additional annotators and use our assembled tweet corpus to produce a large and comprehensive annotated dataset. We are also working – in parallel – on assembling and annotating a similar corpus for Muslim-related hateful expressions.

Another direction we are currently pursuing is taking advantage of the fact that the corpus is comprised of complete Twitter conversations, to annotate expressions of hate in the context of the conversation which the tweet is a part of (rather than just based on the content of the tweet itself).

For example, replying to a hateful post with strong agreement may be considered as hate speech only if the context (preceding posts) is taken into account. Using the hierarchical structure of the conversation will allow not only encoding such cases, but also modelling the dynamics of hate speech as it progresses through the conversation and over time.

In addition, we plan to utilize the tweet corpus to explore *counter-hate* speech (Benesch et al., 2016; Wright et al., 2017; Garland et al., 2020). As these types of expressions are reactive by nature, complete Twitter conversations are instrumental in addressing and analyzing them. Augmenting hate-speech annotated Twitter conversations with counter-hate annotation will allow us to explore the inter-changing dynamics of hate and counter-hate speech, as well as which kinds of counter messages are tailored to the different hate categories. We might find for example that the effective counter messages for abusive speech are those that attack the user, while most effective counter messages for prejudice deliver data and facts to contradict the prejudiced beliefs.

Limitations

There are three main limitations to our work. One limitation results directly from the single target group included in our analysis (Jews). While our annotation scheme was designed to be as general as possible (with the exception of the Holocaust Denial category), applying it to a single target group does not allow us to evaluate the extent of its generalizability to other target groups.

A second limitation has to do with messages that support and fuel hate, without containing actual hateful content (expressing agreement with another hateful message). While such messages may spread hate, they would not be encoded in our annotation setup, since the message context (e.g., the surrounding conversation) is not taken into account during the annotation process.

Thirdly, our annotation scheme does not currently account for how the annotated hate is perceived by the message’s readers. This information may lie in the reaction incurred by the hateful message – the tweet’s replies, as well as the its meta-data (number of likes, quotes, etc.).

By applying the annotation to other target groups, and by annotating complete conversations (thus capturing the context of the tweets), future research could address the three limitations.

5 Acknowledgments

We wish to thank Prof. Peter J. Loewen, Prof. Ron Levi, Prof. Christopher Cochrane, and Mr. Thomas Bergeron for their valuable contribution.

This research was partially funded by the University of Toronto - Hebrew University of Jerusalem Research and Training Alliance, and by the Knapp Family Foundation Doctoral Fellowship at the Vidal Sassoon International Center for the Study of Antisemitism.

References

- Aymé Arango, Jorge Pérez, Bárbara Poblete, Valentina Proust, and Magdalena Saldaña. 2022. Multilingual resources for offensive language detection. *WOAH 2022*, page 122.
- Susan Benesch, Derek Ruths, Kelly P Dillon, Haji Mohammad Saleem, and Lucas Wright. 2016. Considerations for successful counterspeech. *Dangerous speech project*.
- Mohit Chandra, Dheeraj Pailla, Himanshu Bhatia, Aadilmehdi Sanchawala, Manish Gupta, Manish Shrivastava, and Ponnurangam Kumaraguru. 2021. “subverting the jewtocracy”: Online antisemitism detection using multimodal deep learning. In *13th ACM Web Science Conference 2021*, pages 148–157.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Proceedings of the international AAAI conference on web and social media*, volume 11, pages 512–515.
- Mai ElSherief, Shirin Nilizadeh, Dana Nguyen, Giovanni Vigna, and Elizabeth Belding. 2018. Peer to peer hate: Hate speech instigators and their targets. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 12.
- Paula Fortuna, Juan Soler, and Leo Wanner. 2020. Toxic, hateful, offensive or abusive? what are we really classifying? an empirical analysis of hate speech datasets. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 6786–6794, Marseille, France. European Language Resources Association.
- Antigoni Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. Large scale crowdsourcing and characterization of twitter abusive behavior. In *Proceedings of the international AAAI conference on web and social media*, volume 12.
- Joshua Garland, Keyan Ghazi-Zahedi, Jean-Gabriel Young, Laurent Hébert-Dufresne, and Mirta Galesic. 2020. [Countering hate on social media: Large scale](#)

- classification of hate and counter speech. In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 102–112, Online. Association for Computational Linguistics.
- Jikeli Gunther, Deepika Awasthi, David Axelrod, Daniel Miebling, Pauravi Wagh, and Weejoeng Joeng. 2021. Detecting anti-jewish messages on social media. building an annotated corpus that can serve as a preliminary gold standard. In *Workshop Proceedings of the 15th International AAAI Conference on Web and Social Media*.
- Lynda Lee Kaid and Christina Holtz-Bacha. 2007. *Encyclopedia of political communication*. SAGE publications.
- Urja Khurana, Ivar Vermeulen, Eric Nalisnick, Marloes Van Noorloos, and Antske Fokkens. 2022. [Hate speech criteria: A modular approach to task-specific hate speech definitions](#). In *Proceedings of the Sixth Workshop on Online Abuse and Harms (WOAH)*, pages 176–191, Seattle, Washington (Hybrid). Association for Computational Linguistics.
- Youngwook Kim, Shinwoo Park, and Yo-Sub Han. 2022. [Generalizable implicit hate speech detection using contrastive learning](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6667–6679, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2021. Hatexplain: A benchmark dataset for explainable hate speech detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14867–14875.
- Tamar Mitts. 2019. From isolation to radicalization: Anti-muslim hostility and support for isis in the west. *American Political Science Review*, 113(1):173–194.
- Benjamin Newman, Jennifer L Merolla, Sono Shah, Danielle Casarez Lemi, Loren Collingwood, and S Karthick Ramakrishnan. 2021. The trump effect: An experimental investigation of the emboldening effect of racially inflammatory elite communication. *British Journal of Political Science*, 51(3):1138–1159.
- Louis Reynolds and Henry Tuck. 2016. The counter-narrative monitoring & evaluation handbook. *Institute for Strategic Dialogue*.
- Erik Scharwächter and Emmanuel Müller. 2020. [Does terrorism trigger online hate speech? on the association of events and time series](#). *The Annals of Applied Statistics*, 14(3).
- Zeerak Waseem, Thomas Davidson, Dana Warmusley, and Ingmar Weber. 2017. Understanding abuse: A typology of abusive language detection subtasks. *arXiv preprint arXiv:1705.09899*.
- Michael Wiegand, Elisabeth Eder, and Josef Ruppenhofer. 2022. [Identifying implicitly abusive remarks about identity groups using a linguistically informed approach](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5600–5612, Seattle, United States. Association for Computational Linguistics.
- Lucas Wright, Derek Ruths, Kelly P Dillon, Haji Mohammad Saleem, and Susan Benesch. 2017. Vectors for counterspeech on twitter. In *Proceedings of the first workshop on abusive language online*, pages 57–62.