

Decoding implicit hate speech: The example of antisemitism

Becker, Matthias J.; Troschke, Hagen

Erstveröffentlichung / Primary Publication

Sammelwerksbeitrag / collection article

Empfohlene Zitierung / Suggested Citation:

Becker, M. J., & Troschke, H. (2023). Decoding implicit hate speech: The example of antisemitism. In C. Strippel, S. Paasch-Colberg, M. Emmer, & J. Trebbe (Eds.), *Challenges and perspectives of hate speech research* (pp. 335-352). Berlin <https://doi.org/10.48541/dcr.v12.20>

Nutzungsbedingungen:

Dieser Text wird unter einer CC BY Lizenz (Namensnennung) zur Verfügung gestellt. Nähere Auskünfte zu den CC-Lizenzen finden Sie hier: <https://creativecommons.org/licenses/by/4.0/deed.de>

Terms of use:

This document is made available under a CC BY Licence (Attribution). For more information see: <https://creativecommons.org/licenses/by/4.0>

Recommended citation: Becker, M. J., & Troschke, H. (2023). Decoding implicit hate speech: The example of antisemitism. In C. Strippel, S. Paasch-Colberg, M. Emmer, & J. Trebbe (Eds.), *Challenges and perspectives of hate speech research* (pp. 335–352). Digital Communication Research. <https://doi.org/10.48541/dcr.v12.20>

Abstract: This article deals with the problem and the different levels of implicit in the context of antisemitic hate speech. By means of authentic examples stemming from social media debates, we show how alluded antisemitic concepts can be inferred on the basis of conservative interpretation. We point out the role of different sources of knowledge in reconstructing the ideas implied in a statement as well as potential sources of error in the interpretation process. The article thus focuses on the methods and findings of the interdisciplinary research project “Decoding Antisemitism” conducted at the Center for Research on Antisemitism (ZfA) at TU Berlin. By doing so, it explains the procedure of qualitative content analysis as a fruitful approach to understand the actual repertoire of antisemitic hate speech online.

License: Creative Commons Attribution 4.0 (CC-BY 4.0)

Matthias J. Becker & Hagen Troschke

Decoding Implicit Hate Speech

The example of antisemitism

1 Why focus on implicit?

Anyone examining forms of hate speech, either on- or offline, will soon come into contact with utterances that communicate hateful ideas, but without clearly assignable word material. In this article, we will demonstrate how implicit hate speech functions using the example of antisemitism.¹ Writers do not necessarily have to refer to Jews or Israel, nor do they have to reproduce antisemitic stereotypes, such as greed or infanticide, or express invective, threats, or death wishes—all of which should be classed as explicit antisemitism. Writers communicate antisemitic ideas through an enormous spectrum of language use patterns both at the word and sentence levels.² This spectrum of linguistic variations of meaning—among them the

1 The prerequisite here is a viable definition of antisemitism that takes into account both historical and contemporary manifestations of hostility toward Jews. As a basis for our work, we use the internationally recognised IHRA definition (IHRA, 2016). It nevertheless had to be scientifically specified by assigning the various antisemitic tropes to the enumeration of aspects of antisemitism given with the definition.

2 On reasons for communicating antisemitism in implicit forms, see Troschke and Becker (2019, pp. 152–154) and Becker (2020).

dominant field of implicit patterns (i.e., semantically ambiguous or underspecified)—has to be taken into consideration in order to be able to make reliable statements about the actual presence of antisemitism online.

Two recent studies underscore the importance of taking implicit statements into account when measuring antisemitism online. A corpus analysis of the comments sections of the news websites *Zeit Online* (German) and *The Guardian* (British), focusing on the use of Nazi comparisons in Middle East discourse (with such comparisons understood as a form of current antisemitism), showed that only one out of 304 Nazi comparisons in reader comments was *explicit* (i.e., equating Israel and Nazi Germany based on the classic pattern *X is like Y*; Becker, 2021). All other Nazi comparisons were characterized by varying degrees of implicitness, either through incomplete comparisons, innovative metaphors, or onomastic or open allusions, which require world knowledge to extrapolate them.³

A similar picture emerged in the context of the Berlin-based research project “Decoding Antisemitism.”⁴ The project’s first “Discourse Report,” which primarily refers to British news websites, shows that the commenting readership tended to communicate antisemitic stereotypes implicitly (Becker, Troschke et al., 2021).⁵ The examination of approximately 1,200 comments from web debates on the Jewish billionaire and philanthropist George Soros found that—provided that there is a qualitative approach—roughly 15 per cent of the examined comments contained antisemitic statements. Prior to the qualitative analysis, we conducted searches with relevant words that represent antisemitic concepts. On this basis

3 Other corpus analyses of the German-speaking internet confirm the finding that the linguistic reproductions of antisemitic attributions are largely implicit (Schwarz-Friesel, 2020). With regard to implicit language usage patterns outside of the internet, see, for example, Schwarz-Friesel and Reinharz (2017).

4 From 2020, the three-year pilot project “Decoding Antisemitism: An AI-Driven Study on Hate Speech and Imagery Online” is being carried out at the Center for Research on Antisemitism at TU Berlin in cooperation with King’s College London. In this project, which is funded by the Alfred Landecker Foundation, antisemitism in comments sections on British, French, and German mainstream news websites and social media platforms are qualitatively and quantitatively analyzed (see <https://www.alfredlandecker.org/en/article/decoding-antisemitism>).

5 The report focuses on web debates triggered by the media, *The Guardian*, *The Independent*, and the *Daily Mail*, about Jewish billionaire and philanthropist George Soros, the EHRC report on antisemitism in the British Labour Party, and the exclusion of Jeremy Corbyn.

alone, only a small fraction of the antisemitic statements were identified. Thus, only by cataloguing current language usage patterns based on qualitative analysis can we refine quantitative studies and increase the degree to which they represent a research object.⁶

This article explores how the problem of implicitness can be grasped by research practice, using findings from the Decoding Antisemitism project's analysis of implicitly produced antisemitic meanings in social media comments. It presents various forms of implicitness and discusses how these can be included in the analysis. The article then lays out the role played by different sources of knowledge for understanding or even inferring the subject matter of a comment and refers to sources of error for the interpretation process. Finally, we illustrate our interpretive approach using examples from social media. In this way, we also show where the limits of the interpretation of implicit statements may lie.

2 Knowledge areas for extrapolating the implicit

How do we deal with implicit, ciphered statements? To fully extrapolate the meanings conveyed by implicit statements or statements containing implicitness, we need to distinguish *three areas of knowledge*: First, the interpreters require the necessary *language knowledge* to recognize and understand even the most delicate nuances in a statement. Second, the *context* of a statement (e.g., within a thread) and its potential impact must be taken into account. Third, relevant *world knowledge* of the broader context is required, including general knowledge about the cultural space (including society, politics, history), discourses, and conventions, as well as specific knowledge of the subject whose implicit mediation is to be investigated.⁷ In our case, this is the indispensable knowledge of historical and contemporary antisemitism in all its manifestations.

6 See also the project's second "Discourse Report," which compares results from qualitative and quantitative analyses of commentary sections discussing the escalation phase of the Middle East conflict in May 2021 (Becker, Allington et al., 2021).

7 For *world knowledge*, see, for example, Plümacher (2006) and Schwarz-Friesel (2013, pp. 37–41).

The necessary *language knowledge* is available to all members of a language community with sufficient linguistic competence. However, there are differences between language producers and recipients in terms of the precision with which the structure of meanings can be understood.

Context knowledge results from a reception of the context that is equally accessible to all. In our case, it includes the trigger for the comments, whether it is media articles, posts, videos, initial comments, or associated comments. Here, too, despite the same contextual information, all recipients differ in the extent to which they are able to fully integrate this knowledge into the process of forming conclusions.

In terms of *world knowledge*, major differences can exist between recipients regarding a particular subject. These differences are greater the further the object is away from the center of general cultural knowledge. The more extensive the object-specific world knowledge, the quicker meanings based on implicit acceptances or assumptions can be fully identified. This is therefore an important prerequisite for the analysis. In this particular case, this knowledge is constituted by our research on historical and contemporary forms of antisemitism and antisemitic discourses and our insights stemming from past analyses. This knowledge is required to assign meaning to the most linguistically explicit antisemitic statements and even more so in the case of implicit comments. The latter sometimes only refer to fragments of an antisemitic concept. However, the whole of these fragments must be known to the interpreters in order for them to extrapolate the respective concept from the reference to a part of it.

3 Securing interpretations

The application of the knowledge from these three areas determines the extent to which all meanings and nuances are identified when categorizing texts. Insufficient knowledge as well as incomplete or faulty reasoning processes at the linguistic and conceptual levels can lead to antisemitic meanings not being recognized or being interpreted in a text without sufficient evidence. The effect of over-interpretation can, however, also be the result of over-sensitivity caused by priming: The continuous examination of the subject matter in the coding process can lead to false presumptions of the (not reliably verifiable)

presence of the subject in the texts. Distortions are possible in two directions: overlooking antisemitic meanings or false positives.

This means it is crucial that statements should be categorized conservatively. Conservative attribution means deciding in favor of the more likely meaning in situations where there are at least two possible interpretations of an ambiguous utterance in order to prevent false positives. Conflicting valid interpretations are set out, the probabilities of the correctness of different interpretations are weighed against each other, and the use of world knowledge to enrich each interpretation is undertaken with caution. On the one hand, the clearly defined interpretation scheme (compiled in a guidebook⁸) allows one to arrive at meaningful interpretations; on the other hand, it avoids over-interpreting a statement (false positives).

When the same text corpus is categorized by several coders looking for certain content, their respective levels of topic-related world knowledge are the element most likely to produce differences in their conclusions. The extent of world knowledge thus has a decisive impact on the comparability of reasoning and coding processes between coders and the resulting categorizations of texts. Therefore, the level of this world knowledge has to be raised collectively in advance.

In order to minimize deviations in interpretation—and thereby categorization—as much as possible, it is important to define the procedure with a comprehensive (and continuously refined) guidebook containing coding instructions. This allows for the orderly presentation of all conceptual and linguistic-semiotic phenomena (along with a listing of numerous representative and distinguishing examples) for the benefit of facilitating the general understanding of the phenomenon in question. This guidebook is constructed using both definitions drawn from existing research literature and those developed inductively in the course of engagement with the empirical material itself. It serves as a common knowledge base, both with regard to the object of investigation and all the areas where implicitness is found.

In our research project, we created distinctions between antisemitic and non-antisemitic attributions as follows: All anti-Jewish stereotypes used explicitly against Soros, for example, were coded as antisemitic, since it can be assumed

8 The guidebook contains the key elements of the resources used by human coders to analyze comment threads: stereotypes and linguistic and image-analytical categories defined and substantiated with explicit and implicit examples.

that Soros' Jewish identity is widely known. Even if a particular writer unintentionally expressed himself or herself in this way, there is high potential that a negative attribution will be associated with Soros' Jewishness by a large portion of the web comment's readership. If, on the other hand, Soros' actions as an investment banker are demonized, this, as well as any form of criticism of Soros and his practices, however harsh, is not regarded as antisemitic, even if the underlying worldview, which we cannot infer directly from the comment, may be antisemitic. Our understanding of antisemitism comprises antisemitic concepts, insults (whether specifically antisemitic or aimed at Jewish identity), and various speech acts that express the wish to harm Jews or Israelis based on their Jewish or Israeli identity. For example, "Soros is the evil of the world/the evilest person" is an antisemitic utterance. "Soros is an evil banker," however, is an example of a strongly negative, but non-antisemitic, evaluation, explicitly linked to his professional background as a banker.

In the first phase of the coding process, consensual validation across coders—reaching intersubjective agreement via discussion about categorizations—should take place in a succession of small steps. Once a common understanding has been established, checks on intercoder agreement can take place at longer intervals. Furthermore, regularly collating intercoder reliability—and resolving disagreement and, if necessary, henceforth adjusting guidebook instructions—assures the quality of the coding.

On the basis of the guidebook, we have developed an extensive code system using the analytical tool MAXQDA that includes antisemitic concepts (at the content level, e.g., stereotypes) as well as phenomena at the linguistic and semiotic levels through which the concepts are communicated. In this way, an utterance can be coded with regard to its conceptual as well as structural particularities. We can then record the linguistic—and possibly semiotic or semiotically accompanied—expressions of content and examine how they are combined.⁹

9 For methodological literature on qualitative content analysis, see, for example, Mayring (2015) and Kuckartz (2018).

4 Interpretative approach

In this section, we show how an antisemitic concept can be communicated via different types of utterance. We will then demonstrate how implicitness can be realized at the different levels and how these levels can interact. Subsequently, we will move on to authentic corpus examples from our research and, based on the interpretation of these, we will reconstruct our interpretative journey through the three aforementioned areas of knowledge, providing an insight into our conservative approach and the use of our guidebook.

In order to extrapolate its meaning, we break down a comment into units of meaning, determine which propositions are present, track how it is embedded in the context, and identify linguistic forms of implicitness, as well as any informational gaps that have to be filled by conclusions. We then summarize the conclusions regarding the individual components of the comment and how they relate to each other.

Language and world knowledge are relevant for every interpretation. An interpretation without world knowledge is impossible for the identification of antisemitism. Certainly, an explicit attribution can be understood from a linguistic point of view (and without being augmented by world knowledge); however, the utterance still needs to be situated within the ideology of antisemitism. Context knowledge is very often a relevant source for interpretation, but there are statements that can be comprehensively interpreted without contextual information because their meaning is independent from the context and would be the same in other contexts.

The possibilities for disguising a concept within subtle words, that is to communicate it implicitly, are numerous and can be found on several levels that are becoming increasingly complex. They begin at the word level when a writer uses *acronyms* or *puns*, for example. Changes to the surface of the word add another unit of meaning without having to do so explicitly. Another possibility involves using *allusions*, where the surface of the lexeme used remains intact, but—due to the apparent conflict between the meaning of the allusion and the utterance into which it has been transplanted—an indirectly communicated meaning is constituted. At the word group or sentence level, implicit antisemitism can be communicated by means of so-called *indirect speech acts*, in which what is meant results from the combined evaluation of all the communicated

units. These phenomena represent only a small part of how implicitness is constituted. In what follows, we will give examples of the various levels.

4.1 Word level: Using synonyms

Let us begin with a simple concept: The word *money* is the linguistic expression of the concept of a payment tool and means of storing value. However, when analyzing authentic language data, it quickly becomes clear that concepts are reproduced linguistically with a high degree of variation and that the standardized expression is only used sometimes. This makes the linguistic variance for antisemitic concepts all the greater since, for the most part, they do not have standardized linguistic expressions in the language community. Speakers who refer to the concept MONEY¹⁰ can alternatively use words from other concept areas, such as *dough*, *bread*, *loot*, *wedge*, *moolah*, and *lolly*, as synonyms. Therefore, various signifiers exist for the concept being conveyed. When extrapolating the intended meaning, the reader or listener can use the mental lexicon entry for MONEY: Which synonyms exist within the language community, and is there a match here? Alternatively, the context helps to make the extrapolation of the intended object more precise. This already shows how interactions between different knowledge bases have to be reconciled by the recipient faced with a situational speech act.

The actions of “asking for money” or “demanding money” can also be paraphrased in English and French, as in German, with the metaphorical phrase, “holding out one’s hand.” The connection is underpinned by a reasoning process based on language knowledge through a conventional metaphor (Skirl, 2009). Metaphors can exist at the word or sentence level and serve to concretize abstract facts. This means that they have a function that promotes knowledge. At the same time, however, they can also manipulate by conveying a controversial thought in an indirect and partly elaborated way, thus giving it the status of being sayable.

The conceptual connection between Jews and AVARICE has a prominent position in antisemitic thinking. Writers with this mindset might speak directly of

10 Since stereotypes are phenomena that exist on the conceptual (i.e., mental) level and can be reproduced using language, stereotypes are given in small caps on the following pages in accordance with the conventions of cognitive linguistics.

“greedy and money-grabbing Jews”—but they could also conceptualize Jews using the metaphorical phrase mentioned above and thus involve themselves in the discourse in a more subtle way. Due to the conventional status of the metaphor in combination with the social condemnation of antisemitism in post-war Germany, the statement “Jews are always holding out their hands” would certainly also face sanction in numerous contexts of expression.

4.2 *Sentence level: Indirect speech acts, irony, and rhetorical questions*

A further increase in implicitness is achieved through changes at the sentence level. Speech act theory deals with speech acts of a direct and indirect nature (Searle, 1969, 1975). Indirect speech acts are statements that, word-for-word, do not express what is meant. In such cases, there are two levels: one part that is expressed literally (the literal or secondary act), and another part that is intended (the primary act). Irony is an example of an indirect speech act. The association of Jews with money and greed can be expressed in the context of an ironic statement, such as: “Yes, yes, I know, Jews have never made much of money ...” To decipher the irony, knowledge shared with the writer must be used. In this case, it is initially language knowledge: The emphasized (exaggerated) affirmation that opens the statement, its reinforcement by the generic statement, and the omission points that leave the issue open can serve as indicators that in this statement the assertion of the direct speech act is negated by an indirect one.

However, irony can also be deciphered through world knowledge, and it can be concluded in this case too that by using such information, which includes knowledge of the corresponding antisemitic stereotype, the stereotype is indirectly affirmed. The corresponding knowledge base and its contrast with the information in the statement indicate that there is an implicature to be drawn here (Levinson, 1983). The implicature is a conclusion relating to the actual meaning of the sentence: what is intended. It is up to the recipient to determine here that in the statement the assertion of the direct speech act is negated by an indirect one and then to infer what is actually being meant. The use of irony enables the writer to present the insinuation of *AVARICE* in a persuasive way. In addition, he or she can avoid being pinned down to the meaning when threatened by sanctions.

Parts of the meaning of the statement can thereby be deleted in terms of information—they can be withdrawn or denied.

A rhetorical question is another indirect speech act. If someone writes, “Do Jews always hold out their hands?”, he or she can withdraw from the threat of sanctions by claiming to have innocently asked a genuine question. In addition to this slightly encrypting function, rhetorical questions serve to determine and thus emphasize in an elaborated form what is being asked (Lee-Goldman, 2006).

The presence of irony as well as rhetorical questions and the need for corresponding conclusions can be extrapolated from the language, context, and world knowledge. If someone wants to increase the degree of implicitness in statements linking Jews to money by, for example, placing a question pronoun in the subject position, as in “Who is (once again) holding out their hand?”, context knowledge can be a prerequisite, in this case, of an anaphoric reference, for understanding who is being referred to here (in a roundabout way through an indirect speech act).

We may encounter rhetorical questions like these in German contexts of expression when talking about the Nazi past and the culture of remembrance, for example, in comments sections on the internet that refer to these subjects. Simply drawing on language and context knowledge in the inference process may potentially leave the recipient unsatisfied about the meaning of the statement because even in light of the comment’s trigger—the article itself—there may be several people or groups or no named person who could fit the role of demanding money. In these cases, it is only possible to identify who is meant by augmenting the context information and interpreting it using world knowledge.

4.3 *Changing stereotypes*

After the end of National Socialism, antisemitic thinking did not simply disappear in Germany. Instead, previous stereotypes were updated. The stereotype avarice, at least in its explicit form of presentation, was no longer widespread in the public communication space (Bergmann & Erb, 1986). It was not only the form of expression that changed but also the concept of the stereotype. The insinuation of general avarice was joined by that of the instrumentalization of the holocaust (and of antisemitism in general), the allegation that Jews would capitalize on particular issues and subsequently use the Holocaust to make themselves

rich. It is this expression of specifically German secondary antisemitism that the abovementioned rhetorical question refers to and that allows the implicature to be drawn that Jews should be used as the subject of the sentence.¹¹ The presupposition in brackets “once again” also makes the assumption that this is not a one-off incident but a routine relationship between both sides, the victims or their descendants and the (later) society of the perpetrators, in which the former are said to harass and take advantage of the latter.

5 Corpus examples

Based on this chain of examples, it is clear how the three areas of knowledge (partly connected with each other) are used to (gradually or even fully) extrapolate the levels of meaning behind language usage patterns in terms of their complexity and ambiguity. We will now present a range of web comments from our recent corpora to illustrate how, in our research, we draw on the aforementioned areas of knowledge to show how chains of inference are constructed.

Below the line of a BBC documentary uploaded to YouTube that critiques the usage of antisemitic conspiracy theories in discussions about George Soros, a commenter reproduces an antisemitic concept by using, among other things, a semiotic marker:

“Evil \$oros hands.”

The finding, resulting from language knowledge, that Soros’ name was not spelled correctly and that a dollar sign was inserted instead, leads to the question as to why this was done. Here, by drawing an implicature, an informational gap has to be filled: The signifiers are merged into a compound. It is also a pun since changes are simultaneously made to the surface of the name Soros. The compound brings together the concept areas (the individual Soros with that of

11 Experience abroad shows that the manifestation of a secondary guilt-deflecting or exonerating antisemitism is largely unknown. There, the statements discussed here could not be fully interpreted, which again illustrates how access to cultural and milieu-specific world knowledge is fundamental for the extrapolation of language and thinking patterns.

money). This leads to the conclusion that Soros is associated with money and that an affinity for or the pursuit of money is alleged to be a fixed trait of his character. In this way, he is assigned the stereotype GREED—in addition to the explicit attribution of EVIL.

In a *Daily Mail* article, Soros is described as a philanthropist. In a comment that refers to this, the user's view is that this attribution should be corrected to say:

“Philanthrocapitalist not philanthropist. Huge difference.”

The play on words, “philanthrocapitalist,” represents the (linguistically unsuccessful) attempt to portray Soros as a special “friend of capital” or “capitalism” and evokes the stereotype GREED by insinuating that Soros prioritizes the opportunities of profiteering that capitalism offers above caring about the well-being of people. Language knowledge is sufficient for decoding the compound word. However, to be able to understand that the attribution is directed at Soros, context knowledge in the form of the anaphorical connection between “philanthropist” and Soros has to be included.

Context knowledge can be linked to a comment using, among other things, anaphors. Due to their obviousness, in many cases, anaphoric connections apparently do not need to be mentioned. However, given the fact that in many current research projects on hate speech the testing of automatic recognition takes place on the basis of machine learning, the proportion of anaphoric connections in the construction of meaning of a hate comment becomes more important: The more the antisemitic meaning is based on this contextual information, the more difficult it is for algorithms to correctly categorize the text if they are unable to take that information into account.

In the context of an announcement by Soros that he will fund universities, one comment reads:

“He will finance the Far Left Globalism Marxist Indoctrination and students brainwash ...!”

The pronoun “he” refers to Soros. Without this knowledge, the attribution made in the comment could not be assigned the intended Jewish object, and the specifically antisemitic nature of this post would not be recognized. At the same

time, the topos of the comment would also independently reflect an antisemitic perspective, namely the insinuation that there is a system of global reach and absolute power said to be based on left-wing political ideologies. The connection to universities is established through “students.” The accompanying attributes convey the image of a university teaching in the service of an ideology supported by Soros. This draws on the antisemitic topos of POWER and CONSPIRACY, which claims that Jews are behind ideologies or social developments. This topos can also be expressed in insinuations of influence on media, politics, and socialization factors, resulting in developments that are beneficial to them.

Context knowledge may also be required for a multiple-step interpretation. Referring to warnings made by Soros in an interview about the situation in the European Union (EU), a user remarks, “Its days [that of the EU] were numbered regardless of COVID-19” and thus predicts its imminent end. Another user replies with “Hopefully so are his!” and thus takes up the prediction, supplies it with a wish, and turns it against Soros with a change of object—namely wishing him an early death. In order to infer the antisemitic death wish, the anaphoric connections from “are” to “were” in the reference comment and from “his” to “Soros” in the article must be identified and linked.

The antisemitic content of a text can also be extrapolated without any context knowledge, as in the following comment also posted in response to the above-mentioned documentary by the BBC on antisemitic conspiracy theories:

“WWG1WGA IGWT !”

Clearly world knowledge is required to decode this. The first acronym stands for the emblematic slogan, “Where we go one we go all,” of the adherents of the meta-conspiracy theory QAnon, which is linked to a number of antisemitic ideas and acts as an allusion to this. The slogan is intended to create a sense of community and strengthen the solidarity of its supporters. This function is supported by the subsequent acronym, which stands for “In God we trust” and provides the aspect of confidence in and the support of a higher power for one’s own cause—and thus its legitimacy. The exclamation mark emphasizes that the slogan is to be understood here as an appeal or a commitment to QAnon. The affirmation of antisemitic conspiracy theories is an act of antisemitic communication. However, this affirmation can be extrapolated in another way by adding context knowledge. Since the

comment is a reaction to the repudiation of conspiracy theories, the implicature can be drawn that the allusion must be understood as an opposing position that serves to demarcate a common bogeyman—in this case, Soros.¹²

In reference to an article that discusses criticism of Soros on Facebook, a user states:

“Soros, Zuckerberg, all good hearted Christian names. They want you to believe that the meek shall inherit the earth”

Since the user patently falsely identifies prominent Jewish people as Christians, the reader can (assuming his or her world knowledge allows him or her to recognize this fact) thus draw the implicature that, in the user’s opinion, the named people are also not “good hearted” (a trait here assigned to Christians) but rather are to be identified as wicked. It is an indirect speech act in the form of irony. Subsequently, it is claimed that it is in their interests to lull people into false belief: From our world knowledge about Christian theology and the relationship of this assumption and hope for the “meek” to actual history, we know that this is not true nor will be true. Therefore, here too, the opposite must be concluded in order to understand the attribution. The conclusion here is thus that both individuals are eager to make people believe in an error, thereby putting them (or keeping them) in a state of defenselessness, which makes them incapable of acting in the face of the world’s challenges; this will then allow the former to achieve their true goals more easily. We know from the implicature in the first sentence that they are said to pursue these goals with bad intentions. The reverse implies that Soros and Zuckerberg, rather than striving for the (Christian) ideal, actually want a world order of hardness and ruthlessness in which people are subjugated to their power. Correspondingly, the stereotypes of DECEIT, HYPOCRISY, and GREED FOR POWER are found here.

In response to the BBC documentary, an accusation that regularly crops up that its content or even the BBC itself was influenced by Soros (or Jews in general).

12 Since the web comment refers to the context of a BBC documentary on antisemitism, it can be assumed that precisely those aspects of the QAnon conspiracy theories are activated that are clearly antisemitic – and not the aspects that are not inherently antisemitic.

In the following example, the user rejects the connection between antisemitism and conspiracy theories in relation to Soros and indirectly claims that the BBC is spreading falsehoods for money:

“This is nothing to do with anti-Semitism. I suggest you (The BBC film makers) look into this more carefully; those who still know what the truth is and haven’t taken your 30 pieces of silver.”

The user formulates a construction of opposites between “those who still know what the truth is and haven’t taken [any money]” and those who are allegedly bribed. From the proposition, “taken your 30 pieces of silver,” a form of financial influence can be generally inferred. Adding world knowledge, on the other hand, enables it to be classified into an antisemitic world view since, according to the Gospels, Judas betrayed Jesus and obtained this sum of money. In this respect, it is an allusion to a core concept of anti-Judaism: the betrayal of God’s son by a Jew, which all Jews have been accused of since then. The rejection of antisemitism goes hand in hand with the invocation of one of the oldest antisemitic attributions. At the same time, the *INFLUENCE ON THE MEDIA* stereotype is activated.

It is not always possible to assign a valid interpretation to a text. It is more often the case that, though conclusive interpretations are possible, the probability of their correctness appears to be too low to be able to categorize them on this basis. This includes the following reaction to the abovementioned BBC documentary:

“TRUMP-PENCE 20-24 America.Freedom.Constitution.”

Due to the world knowledge that Trump supports conspiracy theories, the implicature can be drawn that this sudden reference to a second term in office propagated by the user is based on the idea or wish that Trump and Pence will fight the imagined conspiracies. Alternative interpretations here would be that the user, inspired by the trigger, changes the topic and does not want to refer to the context or that it is simply a multiple post that does not establish a connection to the context.

One of the posts about Soros’s support for universities reads:

“He needs to go”

It remains unclear whether this refers to wishing that Soros should withdraw from the public eye or that he should die. For this reason and in contrast to utterances like “The end is near ...” or “Soros should be neutralized,” we do not assign this text to the death wish category.

6 Conclusion

These examples demonstrate the ways in which antisemitic stereotypes can appear in a variety of forms and show how categorizing them using only a few indicators or merely quantitatively on the basis of, for example, key words, collocations, or n-grams, is only able to partially capture (antisemitic) hate speech. The differentiated code system we work with opens up the possibility of, but also demonstrates the need for, determining the existence of antisemitic concepts more precisely than would be the case, for example, with a code system that only differentiates antisemitic and non-antisemitic texts or the different forms of antisemitism. At the same time, it makes coding easier as individual concepts are constantly in view instead of having to be repeatedly recalled in relation to a general category. The advantages such an approach provides for qualitative analysis are also applicable to the subsequent step of machine learning. The data provided might theoretically allow algorithms not only to recognize antisemitic concepts in the course of the learning process in accordance with the categorization of the concepts in the training data, but even to learn to differentiate them (to a certain extent).

The linguistic level in the code system supports the visualization process when making interpretations, as this has to be substantiated specifically on the basis of the language usage. Both the conceptual and the linguistic levels play a double role: They support the interpretation process using the existing categories in a close examination of the indicators, and they enable more detailed analysis results.

Matthias J. Becker is project lead of the “Decoding Antisemitism” project and a postdoctoral researcher at the Center for Research on Antisemitism (ZfA) at TU Berlin, Germany. <https://orcid.org/0000-0003-2847-4542>

Hagen Troschke is a researcher in the “Decoding Antisemitism” project. <https://orcid.org/0000-0003-4098-4115>

References

- Becker, M. J. (2020). Antisemitism on the Internet: An underestimated challenge requiring research-based action. *Justice*, (64), 32–40. <https://www.ijl.org/justicem/no64/index.html#32>
- Becker, M. J. (2021). *Antisemitism in reader comments: Analogies for reckoning with the past*. Palgrave Macmillan.
- Becker, M. J., Allington, D., Ascone, L., Bolton, M., Chapelan, A., Krasni, J., Placzynta, K., Scheiber, M., Troschke, H., & Vincent, C. (2021). *Decoding antisemitism: An AI-driven study on hate speech and imagery online*. Discourse Report 2. Technische Universität Berlin: Zentrum für Antisemitismusforschung. <https://decoding-antisemitism.eu/publications/second-discourse-report/>
- Becker, M. J., Troschke, H., & Allington, D. (2021). *Decoding antisemitism: An AI-driven study on hate speech and imagery online*. First Discourse Report. Technische Universität Berlin: Zentrum für Antisemitismusforschung. <https://decoding-antisemitism.eu/publications/first-discourse-report/>
- Bergmann, W., & Erb, R. (1986). Kommunikationslatenz, Moral und öffentliche Meinung. Theoretische Überlegungen zum Antisemitismus in der Bundesrepublik Deutschland [Communication latency, morality and public opinion. Theoretical reflections on antisemitism in the Federal Republic of Germany]. *Kölner Zeitschrift für Soziologie und Sozialpsychologie*, 38, 223–246.
- IHRA – International Holocaust Remembrance Alliance. (2016). *Working definition of antisemitism*. <https://www.holocaustremembrance.com/resources/working-definitions-charters/working-definition-antisemitism>
- Kuckartz, U. (2018). *Qualitative Inhaltsanalyse. Methoden, Praxis, Computerunterstützung (Grundlagentexte Methoden)* [Qualitative content analysis. Methods, practice, computer support (Basic texts on methods)]. Beltz.
- Lee-Goldman, R. (2006). *A typology of rhetorical questions. Syntax and semantics circle*. UCB.
- Levinson, S. C. (1983). *Pragmatics*. Cambridge University Press.
- Mayring, P. (2015). *Qualitative Inhaltsanalyse. Grundlagen und Techniken* [Qualitative content analysis. Basics and techniques] (12th ed.). Beltz.

- Plümacher, M. (2006). "Weltwissen." Ein sprachwissenschaftlicher Terminus phänomenologisch betrachtet ["World knowledge." A linguistic term phenomenologically considered]. In D. Lohmar & D. Fonfara (Eds.), *Interdisziplinäre Perspektiven der Phänomenologie. Neue Felder der Kooperation: Cognitive Science, Neurowissenschaften, Psychologie, Soziologie, Politikwissenschaft und Religionswissenschaft* [Interdisciplinary perspectives on phenomenology. New fields of cooperation: Cognitive science, neuroscience, psychology, sociology, political science and religious studies] (pp. 247–261). Springer.
- Schwarz-Friesel, M. (2013). *Sprache und Emotion* [Language and emotion] (2nd ed.). UTB.
- Schwarz-Friesel, M. (2020). "Antisemitism 2.0" – The spreading of Jew-hatred on the World Wide Web. In A. Lange, K. Mayerhofer, D. Porat, & L. H. Schiffman (Eds.), *Comprehending and confronting antisemitism. A multi-faceted approach* (pp. 311–338). De Gruyter.
- Schwarz-Friesel, M., & Reinharz, J. (2017). *Inside the antisemitic mind: The language of Jew-hatred in contemporary Germany*. Brandeis University Press.
- Searle, J. R. (1969). *Speech acts. An essay in the philosophy of language*. Cambridge University Press.
- Searle, J. R. (1975). Indirect speech acts. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics. Vol. 3, Speech acts* (pp. 59–82). Academic Press.
- Skirl, H. (2009). *Emergenz als Phänomen der Semantik am Beispiel des Metaphernverstehens. Emergente konzeptuelle Merkmale an der Schnittstelle von Semantik und Pragmatik* [Emergence as a phenomenon of semantics exemplified by metaphor comprehension. Emergent conceptual features at the interface of semantics and pragmatics]. Narr.
- Troschke, H., & Becker, M. J. (2019). Antisemitismus im Internet. Erscheinungsformen, Spezifika, Bekämpfung [Antisemitism on the Internet. Manifestations, specifics, combating]. In G. Jikeli & O. Glöckner (Eds.), *Das neue Unbehagen. Antisemitismus in Deutschland und Europa heute* [The new unease. Antisemitism in contemporary Germany and Europe] (pp. 151–172). Olms.