Title: Hate Speech Operationalization: A Preliminary Examination of Hate Speech Indicators and Their Structure

Jana Papcunová¹, Marcel Martončík², Denisa Fedáková³, Michal Kentoš⁴, Miroslava Bozogánová⁵, Ivan Srba⁶, Robert Moro⁷, Matúš Pikuliak⁸, Marián Šimko⁹, Matúš Adamkovič¹⁰

¹Corresponding author, Institute of Social Sciences, CSPA, Slovak Academy of Sciences, Slovakia; Department of Psychology, Faculty of Arts, Pavol Jozef Šafárik University in Košice, Slovakia, jpapcunova@saske.sk, <https://orcid.org/0000-0002-9358-1419>

²Institute of Psychology, Faculty of Arts, University of Prešov, Slovakia
martoncik@protonmail.ch, <https://orcid.org/0000-0003-4869-6900>

³Institute of Social Sciences, CSPA, Slovak Academy of Sciences, Slovakia, dfedak@saske.sk, <https://orcid.org/0000-0001-6816-1802>

⁴Institute of Social Sciences, CSPA, Slovak Academy of Sciences, Slovakia, kentos@saske.sk, <https://orcid.org/0000-0003-2341-2457>

⁵Institute of Social Sciences, CSPA, Slovak Academy of Sciences, Slovakia, bozoganova@saske.sk, <https://orcid.org/0000-0003-4295-1464>

⁶Kempelen Institute of Intelligent Technologies, Slovakia, ivan.srba@kinit.sk, <https://orcid.org/0000-0003-3511-5337>

⁷Kempelen Institute of Intelligent Technologies, Slovakia, robert.moro@kinit.sk, <https://orcid.org/0000-0002-3052-8290>

⁸Kempelen Institute of Intelligent Technologies, Slovakia, matus.pikuliak@kinit.sk, <https://orcid.org/0000-0003-1353-9462>

⁹Kempelen Institute of Intelligent Technologies, Slovakia, marian.simko@kinit.sk, <https://orcid.org/0000-0002-2306-4408>

¹⁰Institute of Social Sciences, CSPA, Slovak Academy of Sciences, Slovakia, matho.adamkovic@gmail.com, <https://orcid.org/0000-0002-9648-9108>

“This is an original manuscript of an article published by Springer Nature in *Complex & Intelligent Systems* (2021), on 11 October 2021, available online: <https://doi.org/10.1007/s40747-021-00561-0>

Funding: This study was supported by the SAS PhD Student Program Grant under Contract no. APP0175, the Slovak Research and Development Agency, under Contract nos. APVV-18-0140 & APVV-20-0319, and VEGA 2/0068/19 and by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No. 952215

Conflicts of interest/Competing interests: None declared

Availability of data and material: The materials, data, and analytic code are freely available at <https://osf.io/9fykd/>

Acknowledgement: We would like to thank Maria Bielikova for the contribution to the paper and all additional experts and colleagues who provided feedback and/or helped with the preparation of the materials. Thank you for your valuable input. We also would like to express our sincere gratitude to anonymous reviewers for their valuable feedback on the manuscript.

Author contribution: (CRediT taxonomy): JP: conceptualization, writing – original draft, investigation, writing – review & editing; MM: writing – original draft, writing – review & editing; DF: conceptualization, supervision, writing – original draft. MK: conceptualization, investigation, data curation; MB: conceptualization, writing – review & editing; MA: conceptualization, formal analysis, methodology, writing – original draft, writing – review & editing; IS: conceptualization, writing – original draft, writing – review & editing; RM: conceptualization, writing – original draft, writing – review & editing; MP: writing – original draft, writing – review & editing; MS: conceptualization, writing – review & editing. All the authors have revised the manuscript, made a substantial contribution to the work, and approved it for publication.

Abstract

Hate speech should be tackled and prosecuted based on how it is operationalized. However, the existing theoretical definitions of hate speech are not sufficiently fleshed out or easily operable. In order to overcome this inadequacy, and with the help of interdisciplinary experts, we propose an empirical definition of hate speech by providing a list of 10 hate speech indicators and the rationale behind them (the indicators refer to specific, observable, and measurable characteristics that offer a practical definition of hate speech). A preliminary exploratory examination of the structure of hate speech, with the focus on comments related to migrants (one of the most reported grounds of hate speech), revealed that two indicators in particular, *denial of human rights* and *promoting violent behavior*, occupy a central role in the network of indicators. Furthermore, we discuss the practical implications of the proposed hate speech indicators – especially (semi-)automatic detection using the latest natural language processing (NLP) and machine learning (ML) methods. Having a set of quantifiable indicators could benefit researchers, human right activists, educators, analysts, and regulators by providing them with a pragmatic approach to hate speech assessment and detection.

Keywords: hate speech, hate speech indicators, hate speech structure, migrants, comments

1. Introduction

According to a Eurobarometer survey (2016), 75% of people who follow or participate in online debates have witnessed or experienced abuse, threat, or hate speech. A more recent study by Eurostat (2020) points to concerns about the online behaviors of children and young people and the harmful consequences – greater dependency, anxiety, and aggression. Both governmental and civil society organizations have put substantial effort into tackling hate speech (e.g., ECRI general policy recommendation no. 15 on combating hate speech, 2015; The European Commission Code of Conduct on countering illegal hate speech online, 2016; the “No hate speech movement” of the Council of Europe, 2013; etc.). Nonetheless, numerous scientific studies have indicated that the incidence of hate content is continually increasing (Mathew, 2019; Sambaraju, 2020; Salminen et al., 2020).

Successfully combating hate speech requires a proper multidisciplinary understanding of the phenomenon. Intuitively, people generally have some inner understanding of what hate speech is and when they cross the line. But the extent to which they assess the content and form can vary greatly (MacAvaney et al., 2019). Hate speech should be tackled and prosecuted, based on its *operationalization*. As stated by Brown (2017b, p. 45), “How can we ban it if nobody agrees about what it is?” A number of ad-hoc definitions have been introduced in recent years by legal entities, scholars, and social networks. Producing an accurate definition of online hate speech is, however, a challenging task given its multidimensional and complex nature. For example, legal definitions, which excessively concentrate on explicit hate, could make hate speech detection ineffective and targeted groups may feel that there is no one to protect them.

Providing a *precise, universal, operable definition* of online hate speech that is generally accepted by multiple stakeholders – scholars and practitioners, who may come from different, unrelated fields – is even more challenging. It is nonetheless extremely important – such a definition would make it easier for stakeholders to share know-how, annotate data consistently, compare automatic detection methods, study, and compare the prevalence of hate speech across different platforms/countries, and would generally make communication easier, in the knowledge that others are likely to have some understanding of what we mean when we talk about online hate speech. Nevertheless, there is as yet no such theoretical definition (Fortuna & Nunes, 2018; Srba et al., 2021) and none is likely to appear and become widely accepted in the near future.

In order to address this problem and overcome the existing ambiguity, and lack of a unified and operational theoretical definition, we adopt a slightly different approach – we define hate speech empirically by providing a list of *hate speech indicators* (i.e., specific, observable, measurable characteristics that offer a practical means of defining the concept) and the rationale behind them. To the best of our knowledge, thus far such an approach has not been used to operationalize hate speech. Nevertheless, it has been used for other types of undesired user behavior and content (e.g., cyberbullying or fake news). In addition, some of the literature (e.g., Waseem et al., 2017; Zampieri et al. 2019) already relies on some indicator precursors to overcome the ambiguity in definitions of hate speech and other relevant concepts.

The goal of this paper is therefore to operationalize hate speech using indicators and explore their use and the implications thereof, using a social science and computer science approach.

The main contributions of our approach are:

1. Instead of grounding hate speech in a psychological narrative, we take a pragmatic approach (inspired by diagnostic manuals like the Diagnostic and Statistical Manual of Mental Disorders DSM-5 (APA, 2013) or the International Classification of Diseases ICD-11 (WHO, 2018)) and 1) provide a list of hate speech indicators and the rationale behind them using a multidisciplinary approach, and 2) depict the structure of the indicators to reveal the core/peripheral indicators and the strength of the relationship between them.
2. Online hate speech is a research field that is open to diverse knowledge-based fields, and cross-disciplinary approaches are both key and often recommended (Laaksonen et al., 2020; Paz et al, 2020; Srba et al., 2021; Waqas, 2019). While social science researchers lack knowledge on automatic hate speech detection, most data analysts have

no social science background. This paper takes into account the multiplicity of both computer science and social science approaches.

3. Although there are automatic methods of hate speech detection, using advanced NLP and ML methods, much could be done to aid social scientists interested in under-taking large-scale studies and in-depth research on hate speech. Having a set of quantifiable indicators could prove to be a beneficial pragmatic approach: a) for measuring changes/assessments; b) enabling international comparisons; c) improving communication between scholars; d) recognizing the extent of hate speech on the local/national level; e) as a basis for legal proceedings; and from the computer science perspective, for f) the annotation of extensive new hate speech datasets and for training novel explainable detection methods that provide explanations, or justifications, on what aspects of the text (e.g., a comment) lead to the decision of the automated method.

2. Background and related work

2.1 Definitions of hate speech

Definitions of hate speech can largely be found in four different contexts: 1) legal, 2) lexical, 3) scientific, and 4) practical; and these differ in scope and content. Below is a brief overview of each.

1) The purpose of legal definitions is straightforward: to identify messages that violate existing legal norms and require government regulation, i.e., messages that are publicly shared, incite, promote, or justify hatred, discrimination, or hostility toward a specific group and/or individual, based on certain attributes, such as race or ethnic origin, religion, disability, gender, age, sexual orientation/gender identity (Article 20, International Covenant on Civil and Political Rights, 1966; Council of Europe, 1997; ECRI, 2016; Code of Conduct on Countering Illegal Hate Speech, European Commission, 2016). There is no universally accepted legal definition of hate speech and different states, the primary duty bearers under international human rights, fall under different jurisdictions (Brown, 2017a, Walker, 1994 in Calvert, 1997). What is considered “hateful” can be controversial and disputed. State actors, including Governments, legislatures, State authorities, and courts, therefore need a common understanding of the phenomenon before they can take relevant action to address it (Guterres, 2019).

2) Five of the most well-known online dictionaries define hate speech as follows. Hate speech is: 1) “public speech that expresses hate or encourages violence towards a person or group based on something such as race, religion, sex, or sexual orientation” (The Cambridge Dictionary); 2) “speech expressing hatred of a particular group of people” (The Merriam Webster); 3) “speech that attacks a person or group on the basis of race, religion, gender, or sexual orientation” (Collins Dictionary); 4) “a statement expressing hatred for a particular group of people” (The MacMillan Dictionary); 5) “speech or writing that attacks or threatens a particular group of people, especially on the basis of race, religion or sexual orientation” (Oxford Learner's Dictionary). All these definitions share the same characteristics. They define hate speech as a message (written or oral) that expresses hatred and encourages violence toward a specific group of people. These simple definitions concentrate on the meaning of the words that make up the collocation – hate and speech.

3) Scientific definitions go much further, beyond the apparent meaning of the constituent words. The term hate speech is used in different disciplines such as economics, philosophy, sociology,

psychology, or computer science, and, despite the lack of a definitional consensus, we can trace some of the common constituent characteristics. As Parekh (2012) states, when defining hate speech, it is necessary to distinguish hate speech from related terms that do not fit the definition, such as expressing dislike, lack of respect, a demeaning view of others, disapproval, the use of abusive or insulting speech, and speech that does “not call for action” (p. 40). These subtle distinctions could be resolved by having a proper definition of hate speech.

According to Parekh (2012, p. 40-41) hate speech “expresses, encourages, stirs up, or incites hatred against a group of individuals distinguished by a particular feature or set of features such as race, ethnicity, gender, religion, nationality, and sexual orientation. [...] and is often (but not necessary) expressed in offensive, angry, abusive, and insulting language”. Moreover Cohen-Almagor (2011) states that it is intended to dehumanize, harass, intimidate, debase, degrade, victimize, or incite brutality against the targeted groups. For an extensive philosophical analysis of the term hate speech, see Brown (2017a, 2017b).

Parekh (2012) further recognizes three main aspects of the construct of hate speech: 1) it is directed at specific individuals or groups of individuals; 2) undesirable attributes are ascribed to the target group, thereby creating stigmatization; 3) it leads to discrimination because the ascription of undesirable attributes encourages the target to be seen as undesirable.

The third aspect of Parekh's definition is highlighted by Gelber (2019) as well - the inequality is a core element of hate speech. In the subordinate–superior relationship, hate speech puts the creator in a position of authority and the target in the position of the subordinate, which encourages the target to be seen as inferior and legitimizes discriminatory behavior (Gelber, 2019).

According to Gelber (2019), it is important also to note that the emotion of hatred does not have to be evoked for hate speech to cause harm. Hate is not and does not have to be a key element of the definition (see also Brown, 2017a, “myth of hate”). On the other hand, as Brown (2017) points out, if manifestations of hate speech did not include hate, then this concept would be grossly misleading, but it is not. In summary, a lot of speech evokes the strong emotion of hate, but not necessarily all of it.

Last but not least, it should be noted that hate speech need not relate only to a minority or disadvantaged social group as was previously thought by classic scholars such as Jacobs and Potter (1998) or Walker (1994) as cited in (Bilewicz & Soral, 2020). In past centuries hate speech was specifically directed at minorities (e.g., Afro-Americans, the Roma community).

In the computer science research literature (see our previous work (Srba et al., 2021) for an overview of the computer-science perspective on hate speech), the most common focus related to hate speech is its (semi-)automatic detection. With the help of natural language processing (NLP) and machine learning (ML), the models can classify a particular piece of content (usually textual) according to the binary presence of hate speech or a particular aspect thereof (e.g., whether it attacks an individual or a group), see the survey by Fortuna & Nunes (2018) for more information. In computer science (and especially data science), there are two specific factors influencing the employed definitions of hate speech (Srba et al., 2021): 1) adoption of how the data are already annotated (either by other researchers or by social media platforms); 2) the observable factors that can be identified from the content available in the dataset (i.e., without accessing more intangible clues such as the authors' thinking, intent, or attitude that can be studied in other disciplines, e.g. psychology).

Like many other disciplines, computer science lacks a generally accepted definition of hate speech. In addition, it is not clear how hate speech relates to other similar/related concepts (e.g., toxic language) or superior concepts (e.g., abusive/offensive language) (Srba et al., 2021). The boundaries between these concepts are often blurred and there are few works that specifically tackle the distinctions. One example is Davidson et al. (2017), who attempt to distinguish hate speech from other instances of offensive language in their work on the automated detection of hate speech.

4) Practical definitions can be found in the way online platforms (e.g., social media) and existing detection tools define hate speech. Facebook, Twitter, and YouTube are social media platforms that mediate online communication and have developed their own definitions of hate speech to ensure users stick to the rules and as part of their internal regulatory policies (e.g., terms of service or community standards). They have signed a Code of Conduct on the regulation of illegal hate speech with the European Commission (2016). Facebook's community standards (2021) define hate speech as "a direct attack against people on the basis of what they call protected characteristics: race, ethnicity, national origin, disability, religious affiliation, caste, sexual orientation, sex, gender identity and serious disease".

The practical tools for identifying and moderating hate speech, such as profanity filters, content moderation filters, or human driven approaches, have been widely studied (Fortuna & Nunes, 2018; Kalsnes, 2020; Mathew, 2019, Srba et al., 2021, Švec et al., 2018). In the context of social media platforms, the findings on utilization of practical tools are somewhat puzzling as they either show that social platforms perform too much or too little content moderation and lack transparent decision-making processes (Gillespie, 2018).

Capturing the complexity of hate speech in a unified, precise, and operable definition (either theoretical or more empirically oriented) could potentially lead to a better understanding of hate speech and help reduce its occurrence online. Fortuna & Nunes (2018) reviewed all four types of (legal, lexical, scientific, and practical) hate speech definitions. Following a content analysis, they proposed their own theoretical definition (probably one of the most fleshed out ones) (Fortuna & Nunes, 2018, p. 5):

Hate speech is language that attacks or diminishes, that incites violence or hate against groups, based on specific characteristics such as physical appearance, religion, descent, national or ethnic origin, sexual orientation, gender identity or other, and it can occur with different linguistic styles, even in subtle forms or when humour is used.

There are other useful and inspiring attempts at a universal definition reached through a process of characterizing the toxicity of comments (Guimarães et al, 2020; Chua & Banerjee, 2020), or utilizing context information (Wilson & Land, 2020; Gao & Huang, 2017).

Nevertheless, when it comes to their practical application (e.g., during data annotation for machine learning), these theoretical definitions are still too vague and difficult to capture. In addition, many local factors (e.g., cultural, platform, task at hand) make it difficult or even impossible to produce a universal theoretical definition. Indeed, we argue that the potential for finding one seems unlikely at the moment.

Consequently, our aim in this paper is to define hate speech empirically by identifying specific, measurable, observable hate speech indicators.

2.2 Indicators of hate speech

An indicator is “a sign that shows you what something is like or how a situation is changing” (Oxford Learner's Dictionary). Previous studies have looked at indicators of cyberbullying (Tudkuea & Laeheem, 2014), operational indicators of fake news (Molina et al., 2019), and behavioral indicators related to sexting (Hudson, 2014). Indicators are also frequently used in psychopathology as a means of conceptualizing and diagnosing mental health problems. We propose that operationalizing hate speech in the form of indicators could fulfil the same purpose – provide objective, measurable cues to help “diagnose” hate speech and find effective interventions.

Some existing research papers on hate speech deal with the content characteristics that describe hate speech and can be considered precursors to hate speech indicators. However, these are not primarily aimed at identifying hate speech indicators using an adequate methodology and so do not assess them satisfactory.

First, Fortuna & Nunes (2018) identified four dimensions (that can be considered analogous to indicators) that allowed them to compare definitions of hate speech: 1) hate speech has specific targets; 2) hate speech is to incite violence or hate; 3) hate speech is to attack or diminish; and 4) hate speech can be expressed through (and hidden in) humor.

Secondly, some computer science studies have gone further than producing a binary classification and have distinguished different aspects of abusive/offensive language (as a superior concept to hate speech). Waseem et al. (2017) proposed a typology of abusive language based on defining two dimensions: 1) directed versus generalized, and 2) explicit versus implicit abusive language. Zampieri et al. (2019) similarly annotated and automatically classified offensive language according to directness (targeted insult, untargeted) and target identification (individual, group, other). By emphasizing these critical aspects of offensive/abusive language, these studies reduce the ambiguity in definitions of hate speech and related concepts. Ousidhoum et al. (2019) focused specifically on hate speech and manually annotated a multilingual dataset of Tweets in three languages, distinguishing five aspects: directness, hostility type (e.g., abusive, hateful, offensive), target attribute (i.e., what is being targeted, e.g., origin, gender, religion), target group (i.e., who is being targeted, e.g., individual, women), and annotator sentiment (i.e., the emotion felt by the annotator after reading the Tweet, e.g., disgust, fear, anger).

Other things resembling indicators can be identified in data annotation methodologies as well. The manual annotation of hate speech is a subjective process and annotators need extensive cultural and societal knowledge (Waseem, 2016). Human annotators must therefore be provided with detailed instructions on how to recognize and label hate speech consistently. For this purpose, Waseem & Hovy (2016) proposed a list of 11 criteria to guide annotators, examples items are hate speech uses sexist or racial slurs; criticizes a minority and uses straw man arguments; and defends xenophobia or sexism.

Finally, in NLP and ML models, training of detection models commonly involves a feature engineering step. Features utilized to train models (especially those with a strong prediction capability) can be considered a kind of indicator as well. Fortuna & Nunes (2018) analyzed a number of detection methods and identified two categories of features: general features commonly used in text classification approaches and specific hate speech detection features. The latter category, which is more closely related to the concept of indicators, was further divided into several subcategories of features that are intrinsically related to the characteristics

of hate speech: 1) othering language, 2) perpetrator characteristics, 3) objectivity-subjectivity of the language, 4) declarations of superiority of the ingroup, 5) a focus on particular stereotypes, and 6) intersectionism of oppression.

To conclude, despite the increasing understanding in the social sciences of the factors that may underlie online hateful behavior (Azucar, 2018; Bliuc, 2018), we found no studies attempting to systematically develop a set of indicators covering the main aspects of hate speech. In order to fill this gap and conduct a systematic study of hate speech indicators, we aim to provide an empirical definition of hate speech, while taking into account both social science and computer science approaches (with the help of existing research and interdisciplinary experts).

3. Choosing the hate speech indicators

The primary goal of our research was to propose an empirical definition of hate speech. As stated by Chaffee (1991) “an empirical definition of a concept enables us to conclude whether an event we encounter is an instance of that concept or not.” This social science insight has extremely important relevance for computer science and supervised machine learning – in practice, it could improve the accuracy and effectiveness of hate speech detection and removal.

The concept of hate speech is a social construct. It represents agreement between people about what speech is hateful and what is not. Therefore, we started our search for an empirical definition by holding focus group discussions with two groups: group 1 comprising 12 participants selected and assembled by the authors of this manuscript; group 2 consists of a team of researchers – psychologists with different levels of seniority – primarily the authors of this manuscript. The participants (group 1), who varied in gender, age and educational background, discussed and commented on the concept of hate speech, based on their personal experiences. Research ethics were concerned with respecting the research participants throughout the study. The focus groups of researchers took the form of a structured informal discussion described below. Eight focus group sessions were conducted over two months, April–May 2020, with the team of researchers (group 2) in one-hour regular discussion sessions. All eight focus groups were conducted utilizing a dual moderator model (Krueger & Casey, 2009; Morgan, 1998) until the mutual agreement on the selected indicators was reached. First the discussion focused on hate speech indicators in general, but as the sessions continued the focus shifted more towards the hate speech indicators relating to migrants. The focus on migrants was selected on the basis of current reports by: 1) the European Commission (2020) stating that xenophobia, anti-migrant hatred (15%), and anti-gypsyism (9.9%) are among the most reported grounds for hate speech; 2) UNHCR (2020) about powerful narratives denigrating refugees and turning them into objects of fear; 3) the ECRI report on Slovakia (2020) pointing to escalation in hate speech against Jews, Muslims, migrants, Roma, and black people.

The most significant outcome of the focus group discussions and subsequent non-systematic review of the hate speech literature was the creation of a list of 14 hate speech indicators relating to migration and migrants (see Figure 1). The indicators cover the various features of hate speech in a form that is easily accessible to policy makers, educators, researchers, and human rights activists. Crucially, the occurrence of each indicator is easily quantifiable.

Figure 1. Initial 14 hate speech indicators identified via the literature review and focus group discussions.

The list of 14 indicators (including general information about the purpose of the study and a description of each indicator including a representative example) was then sent to experts who were asked to evaluate the content validity of the indicators and for any additional comments. The experts were all based in Slovakia and specialized in asylum and migration law (Human Rights League), sociology (Centre for the Research of Ethnicity and Culture), leading media (*SME*), a non-profit organization (DigiQ), IT, social work and human rights protection, and psychology. Representatives of minorities in Slovakia (Islamic Foundation Slovakia) were also sent a copy. An online consultation was held with one of the experts in relation to her comments and recommendations. Altogether, 18 experts responded and evaluated the validity of the indicators. Their task was to indicate (on a 5-point scale where 0 = does not indicate hate speech at all; 4 = fully indicates hate speech) the intensity of each indicator. The intraclass correlation coefficient (ICC; indicates the level of agreement between the raters) shows very good agreement between the experts' evaluations, $ICC = .80, p < .001$. The experts concluded that denial of other rights is the most intensive hate speech indicator whereas the weakest hate speech indicators were manipulation, slant, and use of tendentious analogies (for more details see <https://osf.io/mctkv/>). After checking the descriptives of the expert ratings, we further discussed and re-evaluated the content of the indicators and merged those we deemed very similar, reducing the list to the final 10 indicators. For an overview of the process see Figure 2.

Figure 2. Flowchart of the process of selecting the indicators

4. List of hate speech indicators and the rationale

The list of 10 indicators and the rationale behind them was arrived at through a process of explication, identifying the theoretical and empirical definitions of hate speech and then evaluating and modifying them.

Below we provide a description of each indicator. The purpose of the indicators is to provide a simple but comprehensive picture of the hate speech phenomenon, which can then be used to help to identify, evaluate, or summarize the results achieved along the way in tackling hate speech, and may help in monitoring hate speech trends in different regions and cultures. The indicators can also be used for setting goals and benchmark criteria. The indicators are relevant, unambiguous, comprehensible, interpretable, and comparable over time.

1. *Sexist language*. Sexist hate speech relates to utterances which spread, incite, promote, or justify hatred based on sex, gender, or sexual preference, with the aim of humiliating and objectifying the target, destroying their reputation, and making them vulnerable and fearful. It also includes second sexism – discrimination against men and boys. The feminist aspect is even more prominent in relation to migration. Almost half of migrants are women and girls (Gender Statistics, 2019). Negative depictions often appear in the media, where female migrants often face double discrimination – as women and as migrants.

2. *Attacking minorities as a “traditionally disadvantaged group”*. This indicator includes any sort of racial offense, xenophobia, antisemitism, discrimination on the grounds of the person’s age, disability, ethnicity, national or religious group aimed at inciting intolerance or racial and ethnic hatred.

3. *Denial of fundamental human rights*. This includes any calls for the exclusion or segregation of a person or group of people, being denied the right to promote one's language and traditions, peaceful assembly etc.

4. *Promoting violent behavior*. Comments aimed at inciting others to commit acts of violence or that can reasonably be expected to have that effect. This also includes excusing violence, promoting negationism, and condoning terrorism.

5. *Problematic hashtags, nicknames, symbols*. This can also mean referring to an organization that has committed hate crimes.

6. *Ad hominem attacks*. These are personal attacks on a person's character or motive, based on feelings of prejudice, rather than facts, logic. They include accusations of lying, ignorance, stupidity, and condescension.

7. *Negative stereotypes of a minority*. These are negative expectations about the out-group, and appear alongside negative emotions (e.g., fear, anger) towards the out-group. They promote the view that the targeted group is inferior or subordinate and are closely linked and related to Indicator no. 2.

8. *Texts containing ambiguous statements, irony, sarcasm*. Making generalized negative statements about minority groups, ridicule, intimidation, hostility. Mocking the concept, events or victims of hate crimes, even if no real person is discussed.

9. *Manipulative texts/misinterpretation of the truth*. Any misinterpretations made purposefully, with the intention of fooling or harming someone. For instance, denial of historical events.

10. *Slurs and vulgarisms*. Direct, explicit verbal or nonverbal attacks, using insulting labels to refer to an individual or group based on their race, ethnicity, national origin, religious affiliation, or other attribute.

Our proposed list of indicators is similar to work by Waseem & Hovy (2016) based on critical race theory. They provided a list of criteria for human annotators identifying and annotating hate speech based on Twitter feeds. It differs from our indicators in that we relied on the views of the multidisciplinary experts involved (i.e., we did not base our list on theory alone) and, while some of the indicators are the same (e.g., hate speech is the use of a sexist or racial slur, attacks a minority, uses problematic hashtags, negatively stereotypes a minority), others are slightly different. In addition, following the preliminary exploratory examination of the structure of hate speech (cf. Section 5), it seems that the denial of human rights and promotion of violent behavior (not present in the work of Waseem & Hovy (2016)) plays a central role in the network of indicators.

5. Using indicators to examine the structure of hate speech

To obtain better insights into hate speech, we examined the internal structure and dynamics of hate speech indicators via a network analysis. In recent years, the network approach to the conceptualization of social and behavioral phenomena has appeared in the behavioral sciences and has become increasingly popular, especially in psychopathology (e.g., Robinaugh et al., 2019). The network approach shifts the focus away from latent constructs (i.e., unobservable entities that cause a set of observable behaviors) and toward complex systems for grouping

variables. In this perspective, a phenomenon occurs owing to a causal system of mutually interacting variables. The four network theory principles for conceptualizing mental disorders have been summarized by Borsboom (2017) and can be applied to hate speech contexts and the indicators we propose. That is, hate speech can be regarded as a specific combination of indicators of varying intensities that mutually interact and reinforce/weaken each other (the indicators theoretically correspond to the nature of the construct, i.e., the indicators proposed above).

The full procedure of hate speech structure examination was as follows:

1. A sample of around 240 comments about migration/migrants was obtained via a qualitative analysis of public social media and local news (Facebook, *SME*, *HNonline*, *Magazin1*) or created by some of the members of the research team. The comments differed in the level of hate expressed: some slightly misinterpreted historical facts, whereas others contained vulgarisms or explicitly called for violent acts.

To prevent the range restriction phenomenon, neutral comments were also selected to capture the whole spectrum of speech (ranging from neutral comments to very hateful comments).

2. Members of the research team were then asked to rate each comment in terms of the indicators described above and to assess the hate speech in the comment overall. Each comment was independently rated by four members of the research team, consisting primarily of the authors of this manuscript.

3. The process of selecting comments was stopped once to check how well the comments represented the whole spectrum of hate speech severity (based on the ratings). Sexism, irony, and manipulative texts from the spectrum of positive, neutral, and very hateful comments seemed to be underrepresented, so we added another 60 comments.

4. Before proceeding to the analysis, we screened the comments one last time and excluded those ($N = 13$) not explicitly related to migration. The resulting dataset consisted of 283 representative comments. The final set of comments (in Slovak) is available at <https://osf.io/452pb/>.

5. Once the selection/extraction of the comments had been finalized, the level of agreement between the raters for each indicator was calculated. Despite all the ICCs being significant, the average level of agreement was $ICC = .64$, indicating moderate reliability, according to the rule of thumb. However, given that some indicators, like ad hominem attacks, manipulative statements, or ambiguous or ironic statements, are naturally dependent on the person interpreting them, rater agreement can in fact be better than appears at first glance.

6. The rating for each indicator was then averaged (individual ratings can be found in the dataset at <https://osf.io/g3b87/>).

7. Furthermore, we checked how well the average perception of hate speech correlated with the general impression of hate speech, the average severity of the indicators and their product (the sum score of the indicators), as well as how strongly the indicators were correlated. Excluding sexism, there was a moderate to strong relationship between all the indicators and the general impression of hate speech. All the indicators correlated positively, with the magnitude of the correlations ranging from .02 to .72; the average correlation between the indicators was $r = .32$.

8. A network analysis approach was used to examine the structure of the hate speech indicators. It tells us how the indicators (nodes) of a phenomenon interconnect (edges), which play a more central role (strength and expected influence: how strongly an indicator is related to other indicators; closeness and betweenness: how well an indicator connects with other indicators), and which do not. For the purposes of this demonstration, a network consisting of the average evaluation of each of the hate speech indicators was computed. Here we depict the structure of hate speech (only for the 283 comments which featured at least one hate speech indicator; the network of all the comments is available in the supplementary materials), show the estimates of the centrality measures, and very briefly discuss the stability of the network. For more details on the technical side, please see the supplementary materials at <https://osf.io/g3b87/>.

For descriptive purposes, the correlation matrix and the corresponding network of binary relationships are presented in Table 1 and Figure 3a respectively. A regularized network of the hate speech indicators is depicted in Figure 3b. Please note that Figures 3a and 3b are visualizations of the networks – although these are intuitive to comprehend, visualizations can be misleading and so interpretations should be based on the statistical estimates. Regarding the estimated centrality measures (Figure 4), the core indicator of hate speech appears to be: (a) denial of fundamental human rights accompanied by (b) promoting violent and aggressive behavior, and (c) the use of vulgarisms. In contrast, a rather peripheral role in the network is played by the use of: (a) ambiguous statements, irony, or sarcasm and (b) the use of problematic hashtags or nicknames. Paradoxically, after controlling for the other variables in the network, some of the correlations between denial of human rights and the other indicators ended up negative (that is why the expected influence drops so substantially compared to the strength estimate).

Table 1. Correlation matrix of the hate speech indicators, general hate speech impression, and a sum score of the indicators

	1	2	3	4	5	6	7	8	9	10	11	12
1	-											
2	.25	-										
3	.13	.45	-									
4	.17	.66	.72	-								
5	.14	.37	.32	.55	-							
6	.14	.43	.48	.55	.48	-						
7	.20	.73	.31	.54	.28	.32	-					
8	.34	.25	.16	.27	.19	.17	.18	-				
9	.04	.35	.13	.26	.09	.18	.46	.13	-			
10	.14	.47	.40	.61	.51	.55	.36	.13	.07	-		
11	.29	.83	.60	.84	.52	.54	.72	.38	.38	.60	-	
12	.38	.81	.65	.87	.62	.65	.73	.42	.42	.70	.92	-

Note: 1: Sexism, 2: Attacks on minorities, 3: Denial of human rights, 4: Promoting violent behavior, 5: Problematic hashtags and nicknames, 6: Ad hominem attacks, 7: Negative stereotypes, 8: Ambiguous statements and irony, 9: Manipulative text, 10: Slurs and vulgarisms, 11. General hate speech impression, 12. Hate speech score (sum of nos 1-10)

The results are, to our knowledge, the first attempt to represent the structure of hate speech and as such, should be interpreted with caution. The attribute of promoting violent and aggressive behavior features in both legal and scientific definitions. This indicator seems to distinguish

hate speech from distinct but related behaviors, such as dislike, lack of respect, or disapproval since, as Parekh (2012) states, it “calls for action”. Our analysis shows that insulting and abusive language, previously considered an important but unnecessary component by scholars, seems to be another core indicator of hate speech. Since both indicators are easy to operationalize, they provide clearer guidance for the effective identification of hate speech. Some scholars (Gelber, 2019; Parekh, 2012) consider inequality or the subordinate-superior relationship to be a defined component of hate speech. We suggest that this vaguely defined component manifests in specific types of speech as the third core indicator, denial of fundamental human rights. Denial of fundamental human rights places the person using hate speech in the position of a superior who determines who should and should not be considered a full-fledged human being. Again, focusing on the list of fundamental human rights and manifestations of their denial could facilitate the identification (in a public context) and measurement (in an academic context) of hate.

Figure 3. Visualization of hate speech structure

Figure 3a. A network of binary correlations between the hate speech indicators

Figure 3b. A regularized network of partialized correlations between the hate speech indicators

1: Sexism, 2: Attacks on minorities, 3: Denial of human rights, 4: Promoting violent behavior, 5: Problematic hashtags and nicks, 6: Ad hominem attacks, 7: Negative stereotypes, 8: Ambiguous statements and irony, 9: Manipulative text, 10: Slurs and vulgarisms

Figure 4. Centrality measures of a regularized network of the hate speech indicators (ordered by strength)

Note: The x-axis represents the standardized score for each variable (y-axis) in the index. A basic (simplified) description of the indices is presented in the body of the paper above (The structure of hate speech, point no. 8). For more detailed information on the strength, closeness, and betweenness please see the tutorial paper by Constantini et al., 2015, for more information on expected influence please see Robinaugh et al., 2016.

6. Practical implications

This paper presents a new approach to understanding hate speech online. Although there are automatic ways of detecting hate speech using advanced ML and NLP methods, still much could be done to aid social scientists interested in under-taking large-scale studies and in-depth research on hate speech. Our research into the investigation of hate speech indicators and their use in combating online hate speech incorporates a social science and computer science perspective, while offering some practical implications.

Having a set of quantifiable indicators could benefit researchers, human right activists, educators, analysts, and regulators by providing them with a pragmatic approach to hate speech assessment and detection. This pragmatic approach could benefit:

1. Scholars and researchers seeking to improve the way they communicate about the phenomenon of online hate speech. Without a clear concept and understanding of the concept, how can scholars and researchers know what speech should be targeted?
2. The non-profit sector, when conducting comparative international studies, measuring changes/assessments in the local community, or fostering debates, discussions, or the exchanging of ideas on how to protect disadvantaged client groups.
3. Regulators who need to identify the extent of hate speech on the local/national level, provide recommendations and mechanisms for fighting against and preventing hate speech, and for facilitating further efforts and initiatives.
4. Social media moderators and online portals who could use it as a basis for hate speech detection and moderation and cultivating discussion.
5. Legal proceedings. As Pejchal (2018) stresses in a study analyzing legal instruments in post-communist Slovakia and the Czech Republic, there is a need to support and create

a favorable environment for discussing democracy and the values it stands for – free speech versus hate speech.

While the proposed indicators may prove beneficial in the manual detection of hate speech in the above-mentioned scenarios, their greatest potential lies in automatic hate speech detection. First, many existing datasets (e.g., Davidson et al., 2017; Golbeck et al., 2017; Zampieri et al., 2021) contain a label that a post is an instance of hate speech or offensive language, but often no finer-grained description or explanation is given. Although the annotators had definitions or guidelines to follow (see, e.g., Waseem and Hovy, 2016), these are not captured in the data. Thus, the same label in two different datasets may have a different meaning, which can hamper the generalizability of the trained models and prevent the use of datasets for transfer learning or even a simple comparison of the different models' performances. Having a dataset annotated with the proposed indicators would alleviate these issues.

Moreover, the indicators can be used as input features or weak labels to train hate speech detection models. Such models would have the added benefit of explainability (e.g., the comment is labeled as hate speech because it promotes violent behavior and negatively stereotypes a minority), which is important for moderating social media content.

However, for the indicators to be used as weak labels or as explanations, they need to be automatically quantifiable from an input text (e.g., a social media post). In Table 2 we show how the proposed indicators occur in existing datasets or lexicons, which could be used to train models to automatically detect or quantify them. Given the available datasets and the current level of maturity of existing NLP approaches, the proposed indicators vary in readiness for automatic detection. *Problematic hashtags, nicknames and symbols* as well as *slurs and vulgarisms* can be detected with a high level of precision using simple word lexicons. However, the cultural sensitivity of certain word usages (words that are problematic in some cultures, but not in others) remains an open problem. Curated lexicons of problematic hashtags are still missing, although existing lists of hate speech words could be used for this purpose. Language aimed at the *denial of fundamental human rights* or *promoting violent behavior* can be detected with machine learning methods based on sufficiently large datasets, since it is often characterized by a relatively homogenous set of features. While existing datasets often contain a label pertaining to violence and aggressiveness (Kumar et al., 2018; Basile et al., 2019; Poletto et al., 2017; Sanguinetti et al., 2018), the denial of fundamental human rights is rarely explicitly labeled, despite our analysis showing it is a strong indicator of hate speech.

Table 2. Map of proposed indicators to existing NLP datasets on hate speech and related phenomena, such as offensive and abusive language and the assessment of the current level of readiness for the automatic detection of each indicator given the availability of data and the capability of existing approaches.

Indicator	Occurs in existing datasets	Level of readiness for automatic detection
Sexist language	Waseem, 2016; Waseem and Hovy, 2016; Basile et al., 2019; Ousidhoum et al., 2019; Chung et al., 2019; Fersini et al., 2018a;	Sexist language can to some extent be detected using machine learning methods and large enough datasets.

	Fersini et al., 2018b (Italian, Spanish)	However, certain types of sexist language are still hard to detect, e.g. latent sexist language, ambiguous statements or sexist language that draws on a broader context and not individual phrases or sentences.
Attacking minorities as a “traditionally disadvantaged group”	Waseem, 2016 (racism); Waseem and Hovy, 2016 (racism); Basile et al., 2019; Ousidhoum et al., 2019 (labeled as origin)	Language attacking minorities can be detected using machine learning methods and large enough datasets, since it is often characterized by a relatively homogenous set of features. However, certain types of attack are still hard to detect, e.g. ambiguous statements or attacks that draw on a broader context and not individual phrases or sentences.
Denial of fundamental human rights	Partially in Sanguinetti et al., 2018 (Italian language dataset) as intensity label.	Language containing the denial of fundamental human rights can be detected using machine learning methods and large enough datasets, since it is often characterized by a relatively homogenous set of features.
Promoting violent behavior	Coded as aggressiveness in Kumar et al., 2018 Basile et al., 2019; Poletto et al., 2017; Sanguinetti et al., 2018 (last two in Italian) Coded also as intensity in Sanguinetti et al., 2018 (in Italian)	Language promoting violent behavior can be detected using machine learning methods and large enough datasets, since it is often characterized by a relatively homogenous set of features.

Problematic hashtags, nicknames, symbols	<p>As far as we know, there is no curated list of problematic hashtags.</p> <p>However, there is a list of hate speech vocabulary, which could be used for this purpose: https://hatebase.org/</p>	Problematic hashtags, nicknames and symbols can be detected with a high level of precision using simple word lexicons. Issues with the culturally sensitive usage of certain words (words that are problematic in some cultures, but not in others) remain an open problem.
Ad hominem attacks	<p>Habernal et al., 2018</p> <p>Name calling label in Da San Martino et al., 2020</p> <p>Name calling and Smears labels in Dimitrov et al., 2021</p>	Certain types of <i>ad hominem</i> attacks can be easy to detect, especially attacks where other indicators are present, such as slurs, stereotypical language. But there are context-sensitive <i>ad hominem</i> attacks whose detection requires a natural language understanding level not yet achieved by the current systems. E.g. "He spends all his time in a library" might be an <i>ad hominem</i> attack aiming at a person's lack of practical skills, but in other contexts it could in fact be praise.
Negative stereotypes of a minority	Explicitly present only in Italian language datasets: Poletto et al., 2017; Sanguinetti et al., 2018; Bosco et al., 2018	Stereotypical language can be detected using machine learning methods and large enough datasets, since it is often characterized by a relatively homogenous set of features (e.g. specific words and phrases). Issues remain with the sheer amount of different stereotypes that exist in different cultures.

Texts with ambiguous statements, irony, sarcasm	Irony: Van Hee et al., 2018 Humor and offense: Meany et al., 2021	Despite the existing research, automatic irony or sarcasm detection is still a challenge. Current systems do not possess the level of natural language understanding to reliably detect such phenomena.
Manipulative texts/misinterpretation of the truth	Da San Martino et al., 2020; Dimitrov et al., 2021	Despite the existing research, automatic detecting a manipulative style or misinterpretations is still a challenge. Current systems do not possess the level of natural language understanding to reliably detect such phenomena.
Slurs and vulgarisms	Profanity filter, e.g. https://github.com/rominf/profanity-filter/blob/master/profanity_filter/data/en_profane_words.txt List of ethnic slurs: https://en.wikipedia.org/wiki/List_of_ethnic_slurs There is also a list of hate speech vocabulary, which might be used for this purpose: https://hatebase.org/ Offensive language datasets, e.g. Davidson et al., 2017; Zampieri et al., 2020; Rosenthal et al., 2021.	Slurs and vulgarisms can be detected with a high level of precision using simple word lexicons. Issues with the culturally sensitive usage of certain words (words that are slurs in some cultures, but not in others) remain an open problem.

Sexist language or language attacking minorities as a “traditionally disadvantaged group” can also be detected with machine learning methods and there are datasets available containing these labels (e.g., Waseem and Hovy, 2016; Basile et al., 2019; Ousidhoum et al., 2019; Chung et al., 2019). However, certain types of these are still hard to detect, e.g. latent sexist language, ambiguous statements or language drawing on a broader context and not just individual phrases or sentences. *Stereotypical language* can be detected with machine learning methods as well, but it is explicitly labeled in only a few datasets (Poletto et al., 2017, Sanguinetti et al., 2018; Bosco et al., 2018). Another issue is the sheer amount of different stereotypes in different cultures.

Certain types of *ad hominem* attacks are relatively easy to detect, especially attacks in which other indicators feature, such as slurs and stereotypical language. The existing datasets that contain name calling or smears as a type of *ad hominem* reflect this (Da San Martino et al., 2020; Dimitrov et al., 2021). But there are context-sensitive *ad hominem* attacks whose detection requires a natural language understanding level not yet achieved by current systems.

For instance, “He spends all his time in a library” might be an ad hominem attack aiming at a person’s practical skills, but in other contexts it could in fact be praise.

Despite the existing research, automatic *irony or sarcasm* detection or automatic detection of a *manipulative style or misinterpretations* is still far from a solved problem. Although there are labeled datasets on irony (Van Hee et al., 2018), humor (Meany et al., 2021), and manipulative (propaganda) techniques (Da San Martino et al., 2020; Dimitrov et al., 2021), the current systems do not possess a level of natural language understanding that would enable them to reliably detect such phenomena. Yet our analysis shows that the use of ambiguous statements, irony or sarcasm is a peripheral indicator of hate speech.

7. Limitations of the present study and suggestions for future research

The present study has several limitations that can be divided into two main areas – robustness of procedure and technicalities related to the analysis. The list of indicators is probably not complete, and one can expect it to change as more ideas or empirical evidence come to light. This is not surprising given that even relatively well-established concepts in diagnostic manuals (e.g., PTSD, compare DSM-IV and the DSM-5) are continually evolving. It could be argued that the procedure employed here relied heavily on experts’ judgement (both in the selection of the indicators as well as their evaluation). Although expert judgement is essential at all stages of the scientific process (Brownstein et al., 2019), there is some evidence (e.g., Cassidy & Buede, 2009) that it may be less accurate than intuitively expected. To prevent systematic mistakes in the conceptualization of a phenomenon, it would perhaps be better to use a pragmatic data-driven approach (e.g., Yarkoni & Westfall, 2017) combined with an expert consensus. This could be a future area of research. Furthermore, the validity and hence generalizability of the set of indicators is likely to be dependent on the way the topic (e.g., LGBTI, gender, political issues) interacts with the cultural setting. To tackle this issue, replication studies on different topics in different cultures would be a very welcome addition. Applied to the present research, the results are probably limited by the specifics of Slovak culture/language and the polarizing topic of migration which may have distorted the ratings.

From the technical perspective, there are two major limitations to this study. First, the input number for calculating the network is taken from the averaged ratings of the hate speech indicators obtained from a fixed set of independent raters. Although this is a legitimate practice (Shrout & Fleiss, 1979), a larger pool of expert raters would make the results more robust. Second, given the small sample (i.e., the number of comments), the stability of the observed estimates and the robustness of the whole network is moderate at best (for the results of simulations see <https://osf.io/g3b87/>). As noted above, the present network primarily serves as a probe into the topic and much more extensive research and replication is needed in order to obtain credible information about the structure of hate speech.

From the computer science perspective, the labeled dataset of hate speech comments is too small to be used in training automated detection models. However, this was beyond the scope of this work and a collection of a larger dataset is planned in the future. One limitation for future research and, more importantly, for practical deployment is that, as discussed above, not all of the proposed indicators are easily quantifiable (detected) by automatic means using ML and NLP methods (given the current state-of-the-art and available resources). However, the indicators, identified in our analysis as most significant for the operationalization of hate speech, can be detected with machine learning methods given large enough datasets, since they are characterized using a relatively homogeneous set of features. Lastly, one could object that our proposed indicators replace a simple binary classification of whether a comment is or is not

hate speech with a set of classification/regression problems. But as argued above, the main advantage is that using indicators leads to unambiguous, interpretable, and comparable labels, which can then be used to detect and explain hate speech. Also, as shown above, most indicators can already be found in some form or other in the existing datasets, which points to their relevance and feasibility.

8. Conclusion

Hate speech is prevalent in all societies, across various social media platforms, and continues to present a challenge to practitioners and regulators. The latest report from the Organization for Security and Co-operation in Europe OSCE (2020) includes hate crimes and bias-motivated crimes but fails to provide information about discrimination or hate speech owing to the lack of consensus on whether these acts should be criminalized. The more vague a theory is, the less it helps us describe, explain, predict, and control a phenomenon (Fried, 2020). This applies to weak, narrative theories, which can be contrasted with strong theories that enable phenomena to be precisely formalized (Fried, 2020). A good hate speech theory would be useful in providing us with a recipe for accurately identifying forms of hate speech. Our goal was to operationalize hate speech by proposing indicators and examining their structure. We defined hate speech in the context of migrants as any text that promotes *violent behavior, denies human rights, contains slurs, vulgarisms, or ad hominem attacks, uses negative stereotypes, or purposefully manipulates the truth or historical facts*. The indicators are interchangeable, and may not all be present, but we consider the *promotion of violent behavior, use of vulgarisms, and denial of human rights* to be the core aspects of hate speech. The proposed indicators can be directly used in practice (i.e., a more precise assessment of hate speech severity by rating the various indicators) and further developed by other researchers. The grey zone deserves special attention and, from the pragmatic perspective, would benefit from a more nuanced conceptualization of the hate speech phenomenon.

For future research, if we are to fully grasp the phenomenon, resources must be shared between researchers, and studies should be conducted/replicated across different cultures and topics, as there is a pressing need for more evidence-based information. In this regard, we plan to use the proposed methodology and indicators to collect a larger sample of labeled comments, of such a volume that it can also be used to train automatic hate speech detection models. To the best of our knowledge, our work provides the first (albeit small) public dataset of labeled hate speech comments in the Slovak language, and the extended dataset will be a significant contribution to the limited NLP resources available in Slovak (and other Slavic languages for that matter). Lastly, we plan to examine the methods for the automatic detection (or quantification) of the proposed indicators using the available resources and state-of-the-art NLP approaches as discussed in this paper.

References

- APA (2013). Diagnostic and statistical manual of mental disorders, Fifth edition. Arlington, VA: American Psychiatric Association.
- Azucar, D., Marengo, D., & Settanni, M. (2018). Predicting the Big 5 personality traits from digital footprints on social media: A meta-analysis. *Personality and individual differences*, 124, 150-159.
- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, Manuela Sanguinetti (2019). SemEval-2019 Task 5: Multilingual Detection of Hate Speech Against Immigrants and Women in Twitter. In:

Proceedings of the 13th International Workshop on Semantic Evaluation. Association for Computational Linguistics, pp. 54-63. doi: 10.18653/v1/S19-2007

Bilewicz, M., & Soral, W. (2020). Hate Speech Epidemic. The Dynamic Effects of Derogatory Language on Intergroup Relations and Political Radicalization. *Political Psychology*, 41(S1), 3–33. doi:10.1111/pops.12670

Bliuc, A. M., Faulkner, N., Jakubowicz, A., & McGarty, C. (2018). Online networks of racial hate: A systematic review of 10 years of research on cyber-racism. *Computers in Human Behavior*, 87, 75-86.

Bosco, C., Felice, D. O., Poletto, F., Sanguinetti, M., & Maurizio, T. (2018). Overview of the evalita 2018 hate speech detection task. In *Evalita 2018-Sixth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian*. Vol. 2263, pp. 1-9. CEUR. URL: <https://iris.unito.it/retrieve/handle/2318/1686264/465071/paper010.pdf>

Borsboom, D. (2017). A network theory of mental disorders. *World Psychiatry*, 16(1), 5–13. doi:10.1002/wps.20375

Brown, A. (2017). What is Hate Speech? Part 2: Family Resemblances. *Law and Philosophy*, 36(5), 561–613. doi:10.1007/s10982-017-9300-x

Brown, A. (2017a). What is hate speech? Part 1: The Myth of Hate. *Law and Philosophy*, 36(4), 419–468. doi:10.1007/s10982-017-9297-1

Brownstein, N. C., Louis, T. A., O'Hagan, A., & Pendergast, J. (2019). The role of expert judgment in statistical inference and evidence-based decision-making. *The American Statistician*, 73(sup1), 56-68. doi: 10.1080/00031305.2018.1529623

Calvert, C. (1997). Hate Speech and Its Harms: A Communication Theory Perspective. *Journal of Communication*, 47(1), 4–19. doi:10.1111/j.1460-2466.1997.tb02690.x

Cassidy, M. F., & Buede, D. (2009). Does the accuracy of expert judgment comply with common sense:caveat emptor. *Management Decision*, 47(3), 454–469. doi:10.1108/00251740910946714

Chua, A.Y.K., & Banerjee, S. (2020). "The topic of terrorism on Yahoo! Answers: questions, answers and users' anonymity", *Aslib Journal of Information Management*, Vol. 72 No. 1, 1-16. doi:10.1108/AJIM-08-2019-0204

Yi-Ling Chung, Elizaveta Kuzmenko, Serra Sinem Tekiroglu, Marco Guerini (2019). CONAN - COunter NArratives through Nichesourcing: a Multilingual Dataset of Responses to Fight Online Hate Speech. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 2819–2829. Association for Computational Linguistics. Doi: 10.18653/v1/P19-1271

CoE-ECRI – Council of Europe - European Commission against Racism and Intolerance: ECRI Report on the Slovak Republic (sixth monitoring cycle); Adopted on 1 October 2020; Published on 8 December 2020, 8 December 2020

Cohen-Almagor, R. (2011). Fighting Hate and Bigotry on the Internet. *Policy & Internet*, 3(3), 89–114. doi:10.2202/1944-2866.1059

Community standards, Facebook (2021) Available at https://www.facebook.com/communitystandards/hate_speech/

Costantini, G., Epskamp, S., Borsboom, D., Perugini, M., Möttus, R., Waldorp, L. J., & Cramer, A. O. (2015). State of the aRt personality research: A tutorial on network analysis of personality data in R. *Journal of Research in Personality*, 54, 13–29. doi:10.1016/j.jrp.2014.07.003

- Council of Europe. (2020). Committee of Experts on Combating Hate Speech. Retrieved from <https://www.coe.int/en/web/committee-on-combatting-hate-speech>
- Council of Europe. (1997). Recommendation No. R (97) 20 of the Committee of Ministers to Member States on “Hate Speech” Retrieved from <https://rm.coe.int/1680505d5b>
- Council of Europe. (2013). No hate speech movement. Retrieved from <https://www.coe.int/en/web/no-hate-campaign>
- Davidson, T., Warmley, D., Macy, M., & Weber, I. (2017). Automated hate speech detection and the problem of offensive language. arXiv preprint arXiv:1703.04009
- Giovanni Da San Martino, Alberto Barrón-Cedeño, Henning Wachsmuth, Rostislav Petrov, Preslav Nakov (2020). SemEval-2020 Task 11: Detection of Propaganda Techniques in News Articles. In: Proceedings of the Fourteenth Workshop on Semantic Evaluation, pp. 1377–1414. International Committee for Computational Linguistics. URL: <https://aclanthology.org/2020.semeval-1.186>
- Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, Giovanni Da San Martino (2021). SemEval-2021 Task 6: Detection of Persuasion Techniques in Texts and Images. In: Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021), pp. 70–98. Doi: 10.18653/v1/2021.semeval-1.7
- ECRI. (2016). General Policy Recommendation No. 15 on Combating Hate Speech. Retrieved from <https://coe.int>
- EU Commission. (2016). Eurobarometer, S. 452, Media pluralism and democracy. Retrieved from <https://ec.europa.eu>
- Eurostat. (2020). Being young in Europe today - digital world. Retrieved from <https://ec.europa.eu/eurostat>
- EU Commission. (2016). Code of conduct on countering illegal hate speech online. Retrieved from <https://ec.europa.eu>
- EU Commission (2020). Countering illegal hate speech online. 5th evaluation of the Code of Conduct. Retrieved from https://ec.europa.eu/info/sites/default/files/codeofconduct_2020_factsheet_12.pdf
- Fersini, E., Nozza, D., & Rosso, P. (2018a). Overview of the evalita 2018 task on automatic misogyny identification (AMI). EVALITA Evaluation of NLP and Speech Tools for Italian, 12, 59. URL: <https://boa.unimib.it/retrieve/handle/10281/219426/452274/10281-219426.pdf#page=68>
- Fersini, E., Rosso, P., & Anzovino, M. (2018b). Overview of the Task on Automatic Misogyny Identification at IberEval 2018. In: Proceedings of the Third Workshop on Evaluation of Human Language Technologies for Iberian Languages (IberEval 2018) co-located with 34th Conference of the Spanish Society for Natural Language Processing (SEPLN 2018). CEUR-WS. URL: <http://ceur-ws.org/Vol-2150/overview-AMI.pdf>
- Fried, E. I. (2020). Lack of Theory Building and Testing Impedes Progress in The Factor and Network Literature. Psychological Inquiry, 31(4), 271–288. doi:10.1080/1047840x.2020.1853461
- Fortuna, P., & Nunes, S. (2018). A survey on automatic detection of hate speech in text. ACM Computing Surveys (CSUR), 51(4), 1-30. doi:10.1145/3232676
- Gao, L. & Huang, R. (2017). Detecting Online Hate Speech Using Context Aware Models. arXiv:1710.07395
- Gelber, K. (2019). Differentiating hate speech: A systemic discrimination approach. Critical Review of International Social and Political Philosophy, 1–22. doi:10.1080/13698230.2019.1576006

Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. New Haven, CT: Yale University Press

Jennifer Golbeck, Zahra Ashktorab, Rashad O. Banjo, Alexandra Berlinger, Siddharth Bhagwan, Cody Buntain, Paul Cheakalos, Alicia A. Geller, Quint Gergory, Rajesh Kumar Gnanasekaran, Raja Rajan Gunasekaran, Kelly M. Hoffman, Jenny Hottle, Vichita Jienjittlert, Shivika Khare, Ryan Lau, Marianna J. Martindale, Shalmali Naik, Heather L. Nixon, Piyush Ramachandran, Kristine M. Rogers, Lisa Rogers, Meghna Sardana Sarin, Gaurav Shahane, Jayanee Thanki, Priyanka Vengataraman, Zijian Wan, and Derek Michael Wu (2017). A Large Labeled Corpus for Online Harassment Research. In *Proceedings of the 2017 ACM on Web Science Conference (WebSci '17)*. Association for Computing Machinery, New York, NY, USA, 229–233. DOI:10.1145/3091478.3091509

Guimarães, S.S., Reis, J.C.S., Ribeiro, F.N., & Benevenuto, F. (2020). Characterizing Toxicity on Facebook Comments in Brazil. In *Proceedings of the Brazilian Symposium on Multimedia and the Web (WebMedia '20)*. Association for Computing Machinery, New York, NY, USA, 253–260. doi: 10.1145/3428658.3430974

Guterres, A. (2019). United Nations Strategy and Plan of Action on Hate Speech. Taken from: [https://www.un.org/en/genocideprevention/documents/U_\(20Strategy\)](https://www.un.org/en/genocideprevention/documents/U_(20Strategy).). doi:10.1163/2210-7975_hrd-9915-20190006

Ivan Habernal, Henning Wachsmuth, Iryna Gurevych, Benno Stein (2018). Before Name-Calling: Dynamics and Triggers of Ad Hominem Fallacies in Web Argumentation. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 386–396. Association for Computational Linguistics. Doi: 10.18653/v1/N18-1036

Hudson, H. K., Fetro, J. V., & Ogletree, R. (2014). Behavioral Indicators and Behaviors Related to Sexting Among Undergraduate Students. *American Journal of Health Education*, 45(3), 183–195. doi:10.1080/19325037.2014.901113

Chaffee, S. H. (1991). *Explication (communication concepts 1)*. Newbury Park.

Kalsnes, B., & Ihlebæk, K. A. (2020). Hiding hate speech: political moderation on Facebook. *Media, Culture & Society*, 43(2), 326–342. doi:10.1177/0163443720957562

Krueger, R. A., & Casey, M. A. (2009). *Developing a questioning route. Focus groups: a practical guide for applied research*, 4.

Ritesh Kumar, Atul Kr. Ojha, Shervin Malmasi, Marcos Zampieri (2018). Benchmarking Aggression Identification in Social Media. In: *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*. Association for Computational Linguistics. 11 p. URL: <https://aclanthology.org/W18-4401/>

Laaksonen, S.-M., Haapoja, J., Kinnunen, T., Nelimarkka, M., & Pöyhtäri, R. (2020). The Datafication of Hate: Expectations and Challenges in Automated Hate Speech Monitoring. *Frontiers in Big Data*, 3. doi:10.3389/fdata.2020.00003

MacAvaney, S., Yao, H. R., Yang, E., Russell, K., Goharian, N., & Frieder, O. (2019). Hate speech detection: Challenges and solutions. *PloS one*, 14(8), e0221152. doi:10.1371/journal.pone.0221152

Mathew, B., Saha, P., Tharad, H., Rajgaria, S., Singhanian, P., Maity, S. K., ... & Mukherjee, A. (2019, July). Thou shalt not hate: Countering online hate speech. In *Proceedings of the International AAAI Conference on Web and Social Media (Vol. 13, pp. 369-380)*.

J.A. Meaney, Steven R. Wilson, Luis Chiruzzo, Adam Lopez, Walid Magdy (2021). SemEval-2021 Task 7: HaHackathon, Detecting and Rating Humor and Offense. In: *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pp. 105–119, Association for Computational Linguistics. URL: <https://aclanthology.org/2021.semeval-1.9.pdf>

Molina, M. D., Sundar, S. S., Le, T., & Lee, D. (2019). “Fake News” Is Not Simply False Information: A Concept Explication and Taxonomy of Online Content. *American Behavioral Scientist*, 65(2), 180–212. doi:10.1177/0002764219878224

Morgan, D. L., Krueger, R. A., & Scannell, A. U. (1998). Planning focus groups. Sage.
Council of Europe (2013) The No Hate Speech Movement. Retrieved from: <https://coe.int>

OSCE. (2020). Annual Report 2019, Retrieved from: <http://osr.org>

Ousidhoum, N., Lin, Z., Zhang, H., Song, Y., & Yeung, D. (2019). Multilingual and Multi-Aspect Hate Speech Analysis. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 4667–4676.
<https://doi.org/10.18653/v1/d19-1474>

Parekh, B. (2012). Is There a Case for Banning Hate Speech? In M. Herz & P. Molnar (Eds.), *The Content and Context of Hate Speech: Rethinking Regulation and Responses* (pp. 37-56). Cambridge: Cambridge University Press. doi:10.1017/CBO9781139042871.006

Paz, M. A., Montero-Díaz, J., & Moreno-Delgado, A. (2020). Hate Speech: A Systematized Review. *SAGE Open*, 10(4), 215824402097302. doi:10.1177/2158244020973022

Pejchal, V. (2018). Hate speech regulation in post-communist countries: migrant crises in the Czech and Slovak republics. *International Journal for Crime, Justice and Social Democracy*, 7(2), 58. doi:10.5204/ijcjsd.v7i2.500

Poletto, F., Stranisci, M., Sanguinetti, M., Patti, V., Bosco, C. (2017). Hate Speech Annotation: Analysis of an Italian Twitter Corpus. In: Proceedings of the Fourth Italian Conference on Computational Linguistics (CLiC-it 2017). CEUR-WS, Vol. 2006. URL: <http://ceur-ws.org/Vol-2006/paper024.pdf>

Robinaugh, D. J., Millner, A. J., & McNally, R. J. (2016). Identifying highly influential nodes in the complicated grief network. *Journal of abnormal psychology*, 125(6), 747. doi:10.1037/abn0000181

Robinaugh, D. J., Hoekstra, R. H. A., Toner, E. R., & Borsboom, D. (2019). The network approach to psychopathology: a review of the literature 2008–2018 and an agenda for future research. *Psychological Medicine*, 50(3), 353–366. doi:10.1017/s0033291719003404

Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Marcos Zampieri, Preslav Nakov (2021). A Large-Scale Semi-Supervised Dataset for Offensive Language Identification. arXiv preprint arXiv:2004.14454

Salminen, J., Hopf, M., Chowdhury, S.A. et al. (2020). Developing an online hate classifier for multiple social media platforms. *Hum. Cent. Comput. Inf. Sci.* 10, 1. doi:10.1186/s13673-019-0205-6

Sambaraju, R., & McVittie, C. (2020). Examining abuse in online media. *Social and Personality Psychology Compass*, 14(3), e12521. doi:10.1111/spc3.12521

Shrout, P. E., & Fleiss, J. L. (1979). Intraclass in assessing rater reliability.

Manuela Sanguinetti, Fabio Poletto, Cristina Bosco, Viviana Patti, Marco Stranisci (2018). An Italian Twitter Corpus of Hate Speech against Immigrants. In: Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018). European Language Resources Association (ELRA). URL: <https://aclanthology.org/L18-1443>

Shrout, P. E., & Fleiss, J. L. (1979). Intraclass correlations: Uses in assessing rater reliability. *Psychological Bulletin*, 86(2), 420–428. doi:10.1037/0033-2909.86.2.420

Srba, I., Lenzini, G., Pikuliak, M., & Pecar, S. (2021). Addressing Hate Speech with Data Science: An Overview from Computer Science Perspective. *Hate Speech - Multidisziplinäre Analysen Und Handlungsoptionen*, 317–336. doi:10.1007/978-3-658-31793-5_14

Švec, A., Pikuliak, M., Šimko, M., & Bielíková, M. (2018). Improving Moderation of Online Discussions via Interpretable Neural Models. Proceedings of the Second Workshop on Abusive Language Online - ALW '18 at EMNLP, 2015, 60–65.

The UN Refugee Agency, (2020), Countering toxic narratives about refugees and migrants. Retrieved from: <https://www.unhcr.org/>

Tudkuea, T., & Laeheem, K. (2014). Development of Indicators of Cyberbullying among Youths in Songkhla Province. *Asian Social Science*, 10(14). doi:10.5539/ass.v10n14p74

UN General Assembly, International Covenant on Civil and Political Rights, 16 December 1966, United Nations, Treaty Series, vol. 999, p. 171.

Cynthia Van Hee, Els Lefever, Véronique Hoste (2018). SemEval-2018 Task 3: Irony Detection in English Tweets. In: *Proceedings of The 12th International Workshop on Semantic Evaluation*, pp. 39–50, Association for Computational Linguistics. Doi: 10.18653/v1/S18-1005

Waqas, A., Salminen, J., Jung, S., Almerexhi, H., & Jansen, B. J. (2019). Mapping online hate: A scientometric analysis on research trends and hotspots in research on online hate. *PLOS ONE*, 14(9), e0222194. doi:10.1371/journal.pone.0222194

Waseem, Z. (2016). Are You a Racist or Am I Seeing Things? Annotator Influence on Hate Speech Detection on Twitter. *Proceedings of the First Workshop on NLP and Computational Social Science*, 138–142. <https://doi.org/10.18653/v1/w16-5618>

Waseem, Z., & Hovy, D. (2016). Hateful Symbols or Hateful People? Predictive Features for Hate Speech Detection on Twitter. *Proceedings of the NAACL Student Research Workshop*. doi:10.18653/v1/n16-2013

Waseem, Z., Davidson, T., Warmesley, D., & Weber, I. (2017). Understanding abuse: A typology of abusive language detection subtasks. *arXiv preprint arXiv:1705.09899*

Wilson, R.A., & Land, M.K. (2021). Hate Speech on Social Media: Content Moderation in Context (September 10, 2020). 52 *Connecticut Law Review* 1029. Available at SSRN: <https://ssrn.com/abstract=3690616>

World Health Organization. (2018). International classification of diseases for mortality and morbidity statistics (11th Revision). Retrieved from <https://icd.who.int/browse11/l-m/en>

World Bank. 2019. The Little Data Book on Gender 2019. World Bank, Washington, DC. World Bank. <https://openknowledge.worldbank.org/handle/10986/31689>

Yarkoni, T., & Westfall, J. (2017). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, 12(6), 1100–1122. <https://doi.org/10.1177/1745691617693393>

Zampieri, M., Malmasi, S., Nakov, P., Rosenthal, S., Farra, N., & Kumar, R. (2019). Predicting the Type and Target of Offensive Posts in Social Media. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1*, 1415–1420. <https://doi.org/10.18653/v1/n19-1144>