

CHAPTER 4

The words of hate speech

A lexical study of homotransphobia in an Italian Twitter corpus

Elena Sofia Safina

University of Naples Federico II

Abstract

The everyday discursive production of dehumanising representations and stereotypical beliefs regarding the LGBTQIA+ community undermines the self-respect of both individuals and the target group by damaging their social agency and entitlement dimensions. This chapter proposes a quantitative and qualitative analysis of the TWEER corpus, which consists of 5660 Italian tweets on queer topics. The aims of the work are quantifying the presence of hate speech online and describing the main linguistic features that characterise such language in Italian. Quantitative analysis consists in a manual annotation of the corpus based on a fine-grained scheme comprising six labels (Sanguinetti et al. 2018): hate speech, intensity of hate, aggressiveness, offensiveness, irony,

How to cite this book chapter:

Safina, Elena Sofia. 2024. 'The words of hate speech: A lexical study of homotransphobia in an Italian Twitter corpus'. In *An Investigation of Hate Speech in Italian: Use, Identification and Perception*, edited by Silvio Cruschina and Chiara Gianollo, 111–157. Helsinki: Helsinki University Press. DOI: <https://doi.org/10.33134/HUP-27-4>.

and stereotype. We found that hate speech covers 13 per cent of the entire corpus, but only 6 per cent of those tweets contained an explicit inciting of hatred, while most hateful tweets were superficially polite, containing dangerous prejudices against LGBTQIA+ individuals. We then propose a lexical-semantic study on a sub-corpus which focuses on isolating the most representative meaning clusters in explicitly hateful texts. By analysing each lexical word, we found three main clusters, namely references to politics, health, and ethics, while mentions of sexual identity issues were far rarer, confirming that even explicit hate relies on a heteronormative matrix rather than an impulsive intolerance of certain kinds of sexual orientations or gender identities.

Keywords: corpus analysis, sentiment analysis, hate speech, homotransphobia

4.1 Introduction

From a linguistic point of view, hate speech has been a focus of research particularly in computational linguistics and natural language processing (NLP; Basile et al. 2020). The main focus of the computational studies was first on hate speech detection (Warner and Hirschberg 2012), based on sentiment analysis annotations (Patti, Bosco, and Damiano 2017). More recent studies have positively evaluated detection systems implemented with modern deep learning tools (Chakravarthi et al. 2022), such as fine-tuned large language models like HateBERT (Caselli et al. 2021) and RoBERTa (Nozza 2022), which have demonstrated high performance in detecting homotransphobia in YouTube comments. Many other subtasks have been developed over the years, such as target detection (Silva et al. 2016), author profiling (Mishra and Del Tredici 2018), and automatic user detection (Musto et al. 2019). These aims first required a search for the linguistic indicators of hatred, primarily identified as foul language and explicit incitement to physical violence evidenced by words such as ‘hitting’, ‘eliminating’, or ‘fighting’.

Although the indicators of offensiveness and aggressiveness are inherently objective measures for categorising hate speech, it is equally true that the co-occurrence of aggressiveness and offensiveness is quite rare in corpora, and offensive language is also a factor of ambiguity for machine performance (Davidson et al. 2017; Malmasi and Zampieri 2017; Pamungkas, Cignarella, and Basile 2018). For instance, offensive language can also occur in harmless texts with an ironic or expressive function, which inevitably leads to confusion in the detection task. Moreover, this type of co-occurrence is common to any content that could be described as hate speech and cannot therefore constitute a distinctive feature of a particular discriminatory discourse, such as homotransphobic speech.

One of the limitations that NLP studies have encountered over the years is the ineffectiveness of generic hate speech detection systems (Chakravarthi et al. 2022; Nozza 2022), since each hate type has its own linguistic peculiarities, especially at the lexical and semantic levels, which may elude recognition by a general classifier. Some studies have therefore looked at the detection of specific types of hate speech in Italian, such as the automatic recognition of misogyny (Attanasio and Pastor 2020; Fersini, Nozza, and Rosso 2020; Muti and Barrón-Cedeño 2020). However, according to more recent studies (Chakravarthi et al. 2022; Nozza 2022; Locatelli, Damo, and Nozza 2023), homotransphobic speech still receives little attention from NLP researchers compared with other types of hate.

Another notable example is the manual annotation of a xenophobic and racist hate speech corpus (Poletto et al. 2017; Sanguinetti et al. 2018) based on a rich and fine-grained annotation model. The authors of these papers proposed a scheme that aimed both to detect hate speech and to describe additional strategies of linguistic expression of hatred, such as aggressiveness, offensiveness, irony, or stereotype.

Hate speech has been studied from a more qualitative perspective in social sciences (Leonard et al. 2022), philosophy of language (Bianchi 2017), and Critical Discourse Analysis (CDA),

with a focus on the structural and semantic features of hate messages (Assimakopoulos, Baider, and Millar 2017). Work within the CDA framework has specifically helped to identify some of the referential strategies and frames evoked in institutional homotransphobic discourse (Reddy 2002), in the press (Mongie 2016), and in computer-mediated communication (Socciarelli 2019). These works have also contributed to highlighting the importance of sociocultural context analysis in hate speech studies.

Qualitative analysis has proved to be essential, even combined with some NLP tasks: an important work by Locatelli, Damo, and Nozza (2023) applied a multimodal process of annotation to a homotransphobic Twitter corpus by integrating hate and topic detection tasks. Starting from a qualitative review of the main themes involved in homotransphobia, the researchers managed not only to quantify the presence of hate speech at the cross-linguistic level but also to create a taxonomy of the most frequent topics for each language.

Lavender and queer linguistics (Liddicoat 2009; Norocel 2011; Peterson 2013) can also be considered benchmarks in the study of homotransphobia, since these approaches analyse the discursive construction of the heteronormative model and consider it the conceptual matrix of homotransphobic discourse (Bucholtz and Hall 2004; Coates 2013), even if the explicit incitement of hate is not the main focus of those works.

Relying on the strands of computational and queer linguistics research, this work addresses two different methodologies and goals. Indeed, [Section 4.2](#) specifically aims at detecting the presence of hate speech in a Twitter corpus; thus, we provide a sentiment analysis based on a six-label annotation scheme to quantify the percentage distribution of different hateful features, such as aggressiveness, offensiveness, irony, and stereotype throughout the corpus.

Since quantitative methodology does not allow for an in-depth study of tweets, [Section 4.3](#) instead looks more closely and qualitatively at the words adopted by haters, interfacing with a very narrow subcorpus. The main goal of [Section 4.3](#) is to isolate the

meanings that differentiate homotransphobic speech from other types of hate, creating a spectrum of types and degrees of hate speech. Therefore, based on the extraction of all lexical words of the subcorpus, this section reviews the main semantic clusters involved in Italian homotransphobia, going beyond a description of slurs and searching for apparently neutral meanings that could foster dangerous prejudices against the LGBTQIA+ community.

4.2 Quantitative computational analysis

4.2.1 Corpus construction and description

This section describes the sentiment analysis of online texts automatically collected from Twitter (now X). As the main goal was to detect hate speech against LGBTQIA+ people, we decided to name the corpus ‘TWEER’, a portmanteau of ‘tweet’ and ‘queer’.

The corpus was built between June and July 2019, and contains 5660 tweets in Italian about queer topics. The corpus is the outcome of three data-filtering operations on a bigger dataset called TWITA (Basile, Lai, and Sanguinetti 2018), kindly made available for this research by the Computer Science Department at the University of Turin. This larger corpus consisted of 500 million tweets, and had already been used for hate speech detection purposes, particularly for xenophobia and racism-related hate speech (Poletto et al. 2017; Sanguinetti et al. 2018).

For the data filtering, we first selected specific production periods by including only those tweets produced by users in June–July 2018 and March–April 2019. The first period covered the Pride months, and while the second related to the World Congress of Families XIII (WCF), held in Verona on 29–31 March 2019.¹ The choices were therefore based on the hypothesis that a greater concentration of political events could have aroused frequent

1 According to its official website, WCF (2019) is a large, international public event that aims to unite leaders, organisations, and families to affirm, celebrate, and defend the natural family as the only stable foundation of society.

discussion by users on Twitter about homotransphobia, whether in support or against such demonstrations.

We opted for a keyword-based approach, selecting a set of keywords associated with the queer target, such as *gay*, *omosessual**, *lgbt*, *lesbica* (lesbian), *bisex*, *bisessual** (bisexual), *trans*, *transessual** (transsexual), *transgender*, *queer*, *gender*, *genderfluid*.² Along with these keywords, we selected another set of typical Italian homotransphobic slurs, such as all the terms that correspond to ‘faggot’ in Italian: *froci*, *finocchi*, *culatton**, *ricchion**, *checc**, *succhiacazz**, *ciucciacazz**, *rottinculo*, *rotto in culo*, *piglianculo*, *piglia in culo*, and in regional variants such as *caghin** (Sardinia), *buliccio* (Liguria), *busone* (Emilia-Romagna) *bucaiol** (Tuscany), *garrusu* (Sicily) (see [Chapter 5](#) in this volume for the variable intrinsic offensiveness of these terms). Finally, we added a set of neutral keywords represented by single words or phrases, with and without hashtags, which described the main queer topics on social media, such as #loveislove, #famigliarcobaleno (rainbow family), #unionicivili (civil unions), #wcf, #wcfverona, #congressomondialedellefamiglie (World Congress of Families), #contro-natura (unnatural), and #pride.

Because of the huge number of tweets obtained, we used a third random filter to reduce the corpus to 6000 tweets, and after off-topic tweets were removed 5660 tweets were left. The final version of the corpus was manually annotated by the author according to the scheme and guidelines described in the next section.

4.2.2 Annotation scheme: tagset design and issues

The annotation task was completed manually by using a tagset consisting of six labels that had been already used by Sanguinetti et al. (2018) (see also [Chapter 5](#) in this volume for an adaptation of

2 Here and throughout the chapter, an asterisk (*) indicates that we also included inflected and derived forms of the word. For example, *omosessual** includes the singular form *omosessuale* (homosexual) but also the plural *omosessuali* and the derived form *omosessualità* (homosexuality).

this tagset). The annotation model relies on a set of variables that the European Court of Human Rights considers in the analysis of hate speech cases; the model attempts to encompass all those variables in a single coherent framework. The tagset includes, besides a hate speech label, labels for aggressiveness, offensiveness, irony, stereotype, and intensity of hate (Sanguinetti et al. 2018: 2800).

The hate speech tag presents a binary choice of values (yes/no). Confirming the presence of hate speech in the tweet depends on the co-occurrence of two factors: the *target*, thus a reference to the LGBTQIA+ community, and the *action*, meaning the illocutionary force of the utterance, thus the intention of spreading, inciting, promoting, or justifying hatred or violence towards the queer target, or a message that aims to dehumanise, delegitimise, hurt, or intimidate the target (Sanguinetti et al. 2018: 2800).

- (1) It's only fair that the government is against lesbians.³

If the 'yes' label applies, one of five degrees of intensity must be selected: the degree of intensity can be between 1 and 4, or 0 to indicate the absence of hate speech. Indeed, intensity is the only hate speech-dependent tag, while the other four categories are more descriptive and are independent of each other. The two lower degrees of intensity (1 and 2) describe implicit discrimination, while the higher degrees (3 and 4) describe explicit hatred. Definitions and examples are given in the following.

Degree 1: There is no explicit incitement to violence, but the text negatively depicts the queer target. It could be a derogative judgement against a single person or the whole social minority, designed to promote prejudices or to discredit the target:

- (2) Trans people didn't even have the decency, they always parade their obscenity!!

Degree 2: There is still no explicit incitement, but the tweet aims to dehumanise or delegitimise the targeted group by questioning

3 All the examples provided in this section have been created by the author to clarify the meaning of each label of the tagset.

their fundamental rights, which are described as a threat to the in-group's rights:

- (3) The government only cares about faggots and their civil unions, who is caring about Italian workers?

Degree 3: There is explicit incitement towards discriminatory or violent acts, but users do not refer to themselves as the direct promoter of the violent actions:

- (4) The church should refuse to let these homosexual perverts in on Sundays!!!

Degree 4: There is explicit incitement towards discriminatory or violent acts, and they are promoted by the author in person:

- (5) As soon as I find that slutty lesbian, I swear I'll smash her head!!!

Turning to the independent categories, the aggressiveness label has three possible values (weak, strong, or absent) and refers to the user's willingness to be aggressive or violent through the justification of discriminatory acts against the target (weak label, as in (6)) or by promoting violent actions against the target (strong label, as in (7)):

- (6) It's normal that a gay couple has been hurt on the street! They were kissing each other!!!!

- (7) I want all those faggots out of my neighbourhood!

The offensiveness label could be considered complementary to the previous label as it takes into account the target's, rather than the hater's, point of view. Relying on the same three values, offensiveness focuses on the lexicon employed in the tweet. For example, if the tweet contains a negative representation of the target by means of the expression of negative qualities, it receives the weak label (8), while if the message features highly disparaging lexical items, it receives the strong label (9):

- (8) Bisexual people don't exist at all. They're just moody.
- (9) Here we go again! Another bastard faggot paedophile in the Church!

While the previous tags are useful in detecting homotransphobic speech, the following tags were added to the tagset by Sanguinetti et al. (2018) to investigate the implicit strategies that may express hateful content. The irony tag, which has only two values ('yes' and 'no'), indicates the presence of any kind of ironic, satirical, or sarcastic expression in the tweet. This linguistic feature is quite important because it could mitigate the spread of hate speech. It is also particularly challenging for automatic hate speech detection systems, because 'sometimes, the presence of figurative language even baffles human annotators. Moreover, external world knowledge is often required in order to infer whether an utterance is ironic' (Pamungkas, Cignarella, and Basile 2018: 204):

- (10) Yes dude, gay families exist, and I am a flying unicorn.

Finally, the stereotype label (which can have the value 'yes' or 'no') is associated with all tweets that contain false prejudicial beliefs about the LGBTQIA+ community, which are disseminated in order to justify existing discrimination or to lay the foundations for new discriminatory phenomena:

- (11) Children would grow up very bad with two mums or dads.

Based on this annotation scheme, the following section will present a quantitative analysis of the results of the manual annotation task.

4.2.3 Results and discussion

We conducted a quantitative analysis of the distribution and frequency of the labels in the annotated TWEER corpus by using the R statistical tool. What emerges from the distribution of hate

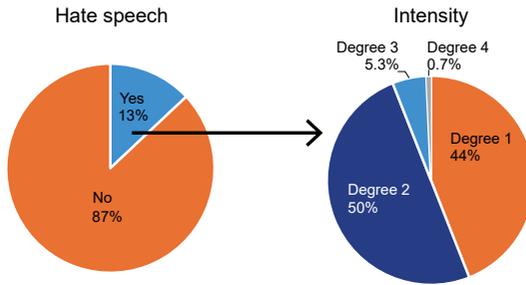


Figure 4.1: Distribution of hate speech and intensity labels in TWEER.

speech labels ([Figure 4.1](#)) is that hate speech covers 13 per cent of the entire corpus, or 742 tweets. To better understand this data, it is important to compare the frequency of hate speech with the distribution of the intensity label. Indeed, less than 6 per cent of the tweets labelled as hate speech (44 tweets) explicitly incite violent or discriminatory actions (degrees 3 and 4), while a larger number of tweets convey implicit and mitigated hate speech (degrees 1 and 2).

Turning to aggressiveness and offensiveness, we found two opposing trends: although each category has been tagged in less than 10 per cent of total tweets ([Figure 4.2](#)), there are more tweets labelled with weak aggressiveness than tweets labelled strongly aggressive, while tweets with strong offensiveness are more frequent than tweets labelled weakly offensive.

However, given that the offensiveness category is independent of the hate speech category, this unexpected trend of strong

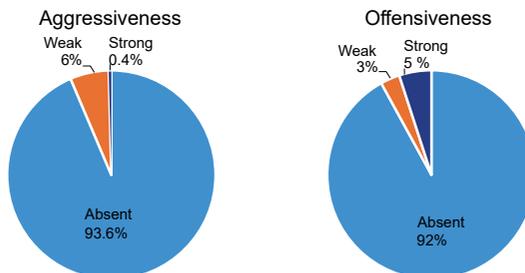


Figure 4.2: Distribution of aggressiveness and offensiveness labels in TWEER.

lexical offensiveness is understandable as a typical linguistic feature of social networks rather than a homotransphobic trait. In fact, 20 per cent of tweets labelled as offensive refer to totally harmless messages (hate speech=No). Rösner and Krämer (2016) describe the absence of a traceable and verified identity on social networks as the cause of the *online disinhibition effect*, the feeling of anonymity and deindividuation that may lead to the extended use of uncivil language; however, offensive words are sometimes used for benign purposes, such as conveying irreverent and ironic meanings, or highlighting the emphasis of the utterance by using insults as filler words.

With regard to the distribution of the irony tag (Figure 4.3), the corpus does not show frequent use of mitigation strategies (only 5 per cent of total tweets). Conversely, stereotype is the most frequent label in TWEER, accounting for 12 per cent of total tweets.

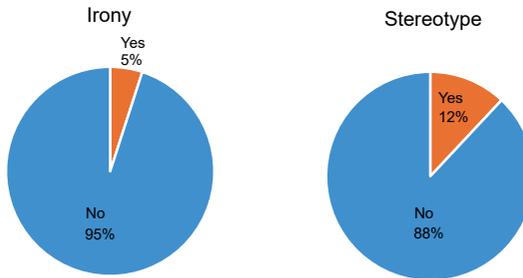


Figure 4.3: Distribution of irony and stereotype labels in TWEER.

In summary, the four labels are distributed in descending order as follows: stereotype (12 per cent of tweets), offensiveness (8 per cent of tweets), aggression (6 per cent of tweets), and irony (4 per cent of tweets). Most hate speech is conveyed in a moderate, mild, and polite form, relying on the effectiveness and immediacy of stereotypes. This is confirmed by the data, given the low component of aggression and offensiveness and the low levels of explicit hate messages (degrees 3 and 4) in the corpus. Similar results are reported by Locatelli, Damo, and Nozza (2023), who found fewer

negative sentiments but more prejudicial homotransphobic views in their dataset made up of 25,000 tweets in Italian.

The quantitative analysis illustrates how the hate speech detection task is difficult and treacherous, partly due to the use of implicit strategies by users, but also due to the significant degree of variability in the linguistic structures that spread hate online, which are difficult to trace by rigid measurement metrics such as those used in this study. The limits of the annotation scheme and of this first quantitative analysis will be further explored in [Section 4.4](#).

In conclusion, the significant distribution of stereotype labels, especially in explicit harmful tweets (hate speech=Yes; intensity degrees 3 and 4) leads us to our next research aim: a more in-depth analysis of stereotypes and the lexical and semantic description of online homotransphobia.

4.3 Lexical and semantic analysis

4.3.1 *Corpus construction and description*

This section describes a lexical study of a restricted corpus of tweets with the aim of investigating which meanings contribute to the construction of homotransphobic discourse. Starting from the TWEER corpus (see [Section 4.2](#)), we collected each lexical word contained in tweets with degrees 3 and 4 under the intensity label in order to avoid any kind of ambiguity in the interpretation of those texts.⁴ For this more qualitative analysis, we explored a subcorpus made up of 43 tweets, 38 with intensity degree 3, and 5 with intensity degree 4, ultimately comprising 665 types and 1246 tokens.⁵ In terms of lexical words, the subcorpus contains items from the lexical classes of nouns, adjectives, and verbs.

4 By lexical words we mean all words with descriptive-referential, as opposed to purely grammatical, content.

5 By 'type' we mean each word of the corpus with descriptive-referential meaning, while 'tokens' refers to all the occurrences of each type-word in the corpus.

The research goal of this study is to identify which meanings are ‘activated’ and ‘salient’ (Arduini and Fabbri 2013) in homo-transphobic discourse, to describe which semantic spheres are involved in stereotypical representations of the queer community, and to detect possible characteristic clusters among lexical words that realise these semantic spheres. As we mentioned in [Section 4.1](#), there are several studies in queer and lavender linguistics that investigate reference strategies and LGBTQIA+ framing (Motschenbacher and Stegu 2013; Motschenbacher 2019). These problem-oriented works focus on how the words used to refer to the target such as ‘gay’, ‘queer’, ‘transgender’, and ‘lesbian’ are qualified by frequent co-occurrences or collocations (Socciarelli 2019). By analysing frequency, keyness, and collocations, researchers were able to establish not only which words qualified the target but also which semantic frames (Fillmore 1985) are usually employed in referring to that target, intended as sets of words associated with stable cognitive structures.

In his work on the Corpus of Contemporary American English (COCA), Motschenbacher (2018) identified six ‘basic sexual usage categories’ intended to represent the main clusters of meanings associated with the queer target, namely Identity, Gender, Partner, Relationship, Desire, and Practice. However, what we see from his analysis is that homosexual or transgender targets are often addressed within a political frame rather than within an affective or erotic frame. In fact, the author found more co-occurrences of politics and social identity words than expressions of desire, sexuality, and erotic practices. Only two out of eight target words, ‘bisexual’ and ‘homosexual’, concerned desire and sexual practices.

We adopt a similar methodology, but considering every lexical word contained in the tweets, even if it does not qualify queer people in co-text, in order to describe the whole semantic environment in which homotransphobia is observed in terms of argumentations and specific topics.

This study does not therefore use tokens as measuring units, but semantic families that consist of one word and all its inflected and derived forms. For example, all the occurrences of *pedofilo*

(paedophile), *pedofili* (paedophiles), and *pedofilia* (paedophilia) have been grouped into the unit *pedofilo** (paedophile). This measuring unit has the advantage of treating related meanings together, thereby directly individuating core meanings in homo-transphobic speech.

Firstly, we use the AntConc software to compile a frequency wordlist. The most frequent lexical words are listed in [Table 4.1](#).⁶ The tables presented in this chapter will have a column on the right reporting the semantic family of reference and a column on the left indicating the number of occurrences of all the inflected and derived forms relating to that semantic family. The asterisk (*) indicates the presence of inflected and derived forms related to the semantic family, while a double asterisk (**) signals that we found different spellings of the same unit of meaning due to informal writing or to the hashtag function on Twitter. For example, in the nine occurrences of the unit of meaning ‘Matteo Salvini’, the presentation of the name varied with respect to the lack of blank spaces between the first and last name (‘MatteoSalvini’), lower-case letters (‘matteosalvini’), and reference by last name (‘Salvini’), but we considered all those forms as the same unit of meaning. The subcorpus contains 222 semantic families.

It is important to highlight that the words *gay*, *lgbt*, *famiglie*, *froci*, *checche*, *gender*, *omosessual**, and *ricchione* in the list had already functioned as keywords for the construction of the TWEER corpus (see [Section 4.2.1](#)), which probably explains why their frequency is higher, as evidenced by ‘gay’ being the most frequent word. We looked at the whole set of words from a problem-oriented, corpus-based approach. Corpus-based studies typically analyse corpus data to validate, refute, or refine a hypothesis, while the corpus-driven approach claims that the corpus itself should be considered the source of the hypothesis about language

6 Because of the very narrow dimensions of the subcorpus, we provide a list of every word that occurred more than once in the dataset. Thus, in this qualitative analysis, the frequency criterion only serves as an auxiliary descriptive data point.

(McEnery and Hardie 2012: 5–6). More precisely, the problem-oriented element in the corpus-based approach indicates that the analysis of language features is informed by critical knowledge about a specific social problem—in our case, homotransphobia (Motschenbacher 2019).

With this in mind, we derive four main clusters of meanings typically associated with homotransphobic hate (tables 4.2, 4.3, 4.4, and 4.5), and two more clusters containing general offensive and violent language (tables 4.6 and 4.7). These clusters can be interpreted as abstract models that are useful in schematising homotransphobic discourse.

Table 4.1: Most frequent words in the TWEER subcorpus.

Frequency	Semantic family
25	gay*
9	matteo salvini** (Italian Politician)
8	lgbt
7	bambini* (children); inesistenti* (non-existent); legge (law)
6	froci* (faggots); italiani* (Italians); roma (Rome); virginia raggii** (Italian politician)
5	famiglie* (families); natura (nature); zingari (gypsies); sessuali* (sexual)
4	gender*; pedofilo* (paedophiles); uomo (man)
3	africani* (African people); delinquere (to commit a crime); diritti (rights); donna (woman); fermate [imperative mood] (stop); ius (Latin = right); lobby; negri (niggers); stop; vivere (to live)
2	basta [hortative] (enough!); checche (faggots); civili (civil); coglione (asshole); depravata (depraved); eros; etici* (ethical); farmaco (drug, medication); fontana (Italian politician Lorenzo Fontana); governolega (Lega Government); inutile (useless); immorali* (immoral); liberando (dalla cacca) (to break free from the shit); madre* (mother); merda (shit); movimento ([political] movement); nazioni (nations); palle [informal] (balls, testicles); popoli* (people/folk); razza (race); repubblica (republic); ricchione (faggot); rottoinculo* ([literal] broken ass, wreck); spazziamo (via) (let's sweep away)

First cluster: Politics

Confirming the findings of Motschenbacher's (2018) study, mentioned above, Politics is the most significant cluster, as references to the political discourse were the most frequent in the subcorpus (36 per cent of total semantic families). As we see from [Table 4.2](#), the politics cluster comprises all those words that describe Italian and international political topics and major figures: first, we find references to leading Italian politicians, especially from the right, such as Francesco Storace, Lorenzo Fontana, and Matteo Salvini, the most cited referent in the subcorpus, since he was in government during the data-collection phase, representing his political party Lega Nord, addressed as 'governolega' in the corpus.

Predictably enough, the main figures from the left are the recipients of hate content, such as in (14), while those from the right are considered both as role models and as reliable political partners, as in (12) and (13):⁷

- (12) SEMPRE PIU' PUTIN! COSI' SI FA. Mondiali, i cosacchi controlleranno i gay: 'Effusioni in pubblico segnalate alla polizia.'

'MORE AND MORE PUTIN! THIS IS WHAT YOU DO. World Cup, the Cossacks will control the gays: "Public displays of affection reported to the police"'

- (13) Quando salvini caccerà gli extracomunitari e i gay, in italia non ci sarà più delinquenza e ci sarà lavoro per tutti gli italiani.

'As soon as Salvini expels non-EU citizens and gays, there will be no more crime in Italy and there will be jobs for all Italians.'

7 All the following examples are from the TWEER corpus.

- (14) @virginiaraggi @Roma Un privato può affittare la casa a chi vuole. Chi discrimina in realtà è lei. E la smetta di sculettare a favore di gay e immigrati.

‘@virginiaraggi @Roma A private individual can rent the house to whoever he wants. The one who’s actually discriminating is you. And stop strutting around in favour of gays and immigrants.’

We also find references to female left-wing personalities, including Monica Cirinnà, a senator of the Italian Republic and representative of the political campaign for civil rights for the LGBTQIA+ community, and Virginia Raggi, the mayor of Rome in 2019. The tweet referring to Raggi in (14) displays not only homotransphobic discrimination but also sexist vilification of her professional role. Along with Italian politicians, we also found references to powerful groups, such as Bilderberg (*bidelberg*^{**} in TWEER), and to famous journalists, such as Bianca Berlinguer and Ezio Mauro. There were three further subclusters relating to politics:

a) Law and criminality: This category is particularly interesting because, alongside neutral institutional terms such as *governo* (government), *ministro* (minister), and *repubblica* (republic), homotransphobia occurs in a criminal frame, such as in (15), with terms with negative meanings such as *delinquere* (to commit a crime), *criminale* (criminal), *mafioso* (mafia man), *sentenza* (conviction), *polizia* (police).

- (15) @RadioSpada Tu sei un coglione non un ministro, lasciamo vivere i bambini come madre natura comanda e sterminiamo la #LGBT sporca e maledetta assicurazioni a delinquere di stampo mafioso!!!⁸

‘@RadioSpada You’re an asshole, not a minister, let the children live as Mother Nature intends and exterminate the dirty and cursed #LGBT, criminal mafia conspiracy.’

8 In the example in (15), *assicurazioni* (insurances, assurances) is a malapropism for *associazioni* (associations).

b) Freedom and fundamental rights: As is usual in the argumentative strategies of many in-groups (Van Dijk 2004), especially in highly polarised societies, homotransphobic speech portrays the perpetrators of hate as the victims by depicting the in-group as the target of discrimination through the use of words such as *discrimina** (discrimination), *odiano* (they hate), and *razzista* (racist). From the perspective of the in-group, the use of expressions such as *ideologia* (ideology), *cambiare* (to change), *controlleranno* (they will control), *libertà* (freedom), *ribellarci* (to rise up/rebel against something), and *dittatura* (dictatorship) is a specific argumentative strategy intended to frame straight and cisgender people as being subordinate to a superior out-group power, namely the ‘LGBT lobby’ (16).

(16) Eliminare e cancellare le leggi gayste e il gaysmo di stato subito.

‘You must remove pro-gay laws and the National Gay cult, now.’

c) Geography and immigration: The final subcluster is the most relevant, because it demonstrates typical hate speech behaviour, namely the assimilation strategy (Van Dijk 2004). Assimilation has been described as a rhetorical strategy that aims to dehumanise and objectify specific groups of people, such as social minorities. Along with the construction of a social dichotomy through the use of deictic ‘us’, sometimes replaced by generalisations such as the ‘people’ or the ‘Nation’, assimilation tends to depict the out-group, ‘them’, using figures or demonstratives, thus erasing the minority’s human traits or cultural peculiarities (Orrù 2017: 35). In our case, alongside the individual dehumanisation of the queer minority, we find *assimilation of minorities*, since the hateful message addresses a large and heterogeneous group of people, which is perceived and represented as a compact group with no internal differences. We found this assimilation of minorities in 13 out of 43 tweets. In one case (17), the author represented the enemy of the in-group through the juxtaposition of ethnic,

religious, political, sexual, and even professional groups, talking about a ‘long list’ of hated subjects with ‘no distinctions’ needed.

- (17) Mi viene un rutto ogniqualvolta leggo Salvini si/no/ma. Salvini sta cambiando il vento che aveva già affondato l’Italia. Non è ora di ‘distinguo’. Spazziamo via zecche rosse, froci/e, giornalai superpagati, toghe indegne. Lista lunga. Poi vediamo. Chi dubita è dall’altra parte.

‘I’m going to burp every single time I read Salvini yes/no/maybe. Salvini is changing the wind that had already sunk Italy. No more time to ‘distinguish’. Let’s sweep away red ticks ([figurative] communists), fags, overpaid news-agents ([ironic-derogative] journalists), unworthy robes ([figurative] magistracy). Long list. Then we’ll see. Anyone with doubts is on the other side.’

Of the assimilated minorities, the group with immigrant status is referred to most often. As we can see from [Table 4.2](#), there are several references to the migration frame, such as *immigrati* (immigrants), *extracomunitari* (non-EU citizens), *profughi* (asylum seekers), and *rifugiati* (refugees). Moreover, we find specific minorities addressed by their geographic or ethnic origins, *africani* (Africans) and *zingari* (gypsies), by religion, *musulmani* (Muslims), *ebrei* (Jewish), and also by the racist slur *negri* (niggers).

We can easily contextualise these frequent references to immigrants in Italian political debate since the early 2010s (Orrù 2017), where the representation of migrants arriving by sea as a *wave of invasion* has gained significant ground in the media agenda. In addition to this, by being deprived of a specific identity, immigrants and sexual minorities are merely characterised by *otherness*. Thus, just as immigration is believed to lead to an inevitable drift towards a loss of security and national identity, homosexual orientation is held to result in social disorder and the dissolution of traditional values.

Table 4.2: List of words in the Politics cluster.

Fre- quency	Semantic family
Politicians and journalists	
9	matteo salvini** (Italian politician)
5	virginia raggi** (Italian politician)
2	fontana (Italian politician Lorenzo Fontana); governolega (Lega Nord government)
1	cirinnà (Italian senator Monica Cirinnà); eziomauro (Italian journalist Ezio Mauro); luigidimaio (Italian politician Luigi Di Maio); berlinguer (Italian journalist Bianca Berlinguer); piddini (Democratic Party voters); putin (Russian president Vladimir Putin); storace (Italian politician Francesco Storace); bidelberg** (Bilderberg group)
Law and criminality	
7	legge* (law)
3	delinquere* (to commit a crime); diritti (rights); ius (+ iusgenus) (Latin = right); lobby*
2	civili (collocates of 'unioni') (civil unions); movimento (movement); repubblica (republic)
1	condanna (criminal sentence); contratto (contract); criminale (criminal); firma (signature); governo (government); illegale (illegal); leader; lista (list); mafioso (mafia man); ministro (minister); rubacchiavano (they were sneaking); segretario (secretary); toghe (robes – figurative use for magistracy); vietati (forbidden)
Freedom and fundamental rights	
1	agevolazioni (benefits); cambia* ([someone/something] changes); comandare (to command); contrapporre (to counterpose); controlleranno (they will control); discrimina* ([someone/something] discriminates); diffondere (to spread); dittatura (dictatorship); giornalisti (newsagents); ideologia (ideology); imporre (to impose); lavoro (job); libertà (freedom); lotta ([political] struggle); mantenere (to maintain); odiano (they hate); parere (opinion); perbenismo (self-righteousness); prevaricato (overlooked); privato (private); razzista (racist); ribellarsi (to rise up/rebel against); scandalizzato (shocked); segnalate [imperative mood] (report them!); sinistra (left-wing); sostenuto (endorsed)

Fre- quency	Semantic family
Geography and immigration	
6	italiani* (Italian people); roma (Rome)
3	africani* (African people); negri (niggers); zingari (gypsies)
2	nazioni (nations); popoli* (people/folk)
1	ebrei (Jewish); extracomunitari (non-EU citizens); immigrati (immigrants); invasione (invasion); islamici (Islamic); malta (Malta); musulmani (Muslims); olandesi (Dutch people); paese (country); provincia (Italian regional district); rifugiati (refugees); straniera (foreign); taranto (Taranto, city in Apulia); toscani (Tuscan people); venezia (Venice)

Second cluster: Nature

The second cluster accounts for 16 per cent of the lexical words in the subcorpus, namely those words related to what we consider as ‘human experience’ (Table 4.3). This cluster is generically named ‘Nature’ since it includes not only human-referring terms but also references to the animal world and abstract concepts about life and experience. We found two subclusters within the ‘Nature’ group, used to represent two typical attitudes in general and homotransphobic hate speech. First, we find many occurrences of words that refer to animals being used to dehumanise the hated target, a strategy also noticeable in other types of hate speech (De Mauro 2016; De Smedt et al. 2018). Second, we see the old-fashioned belief that homosexuality is a medical condition that modifies physical and mental traits in human beings or produces deviant social behaviours.

a) Human beings and animals: This subcluster concerns the *anthropological homophobia* paradigm (Rossi Barilli 1999) which has spread through the Western world since the twentieth century, and is based particularly on the dichotomy between the civilised, ordered, sober, new bourgeois class with their moral concerns, and brutal, savage, exotic societies. According to this paradigm, homosexuality afflicted savage, poor, and uneducated individuals, those unable to escape their tribal impulses. In our case,

the animal frame served as a dehumanising device to associate homosexual identities with the sexual practices of wild animals (Locatelli, Damo, and Nozza 2023), by metonymically reducing their sexual orientation to ‘savage’ anal penetration, as in (18), for example. The anthropological category of wildness also deprives homosexuals of their entitlement to fundamental cultural rituals, such as having a proper burial, as in (19).

(18) @bravimabasta @USERNAME Il tuo culo lo hai da tempo regalato al primo mandrillo che hai trovato, almeno taci, ricchione di merda! [...]

‘@bravimabasta @USERNAME You gave your ass long ago to the first lecher (lit. mandrill) you found, at least shut up, you bloody faggot! [...]

(19) [...] I froci nei cimiteri sono vietati. Li diano in pasto agli squali, almeno servono a qualcosa alla fine della loro inutile esistenza.

‘[...] Fags in cemeteries are prohibited. Feed them to the sharks, at least they’ll serve some purpose at the end of their useless existence.’

b) Health and disease: In this subcluster, we find body-related terms, such as *sangue* (blood) or *scalpo* (scalp); words referring to different conditions, ranging from the hypernyms *disturbi* (disorders), and *patologie* (pathologies), to the hyponyms *anoressia* (anorexia), *psicosi* (psychosis), *ansia* (anxiety), and *disforia* (dysphoria), up to the extreme and negatively connotated *impazzire* (going crazy) or *pazzi* (fools). We also found some references to neurodivergence, with terms such as *autismo* (autism) or *ritardi mentali* (mental retardation). It is important to underline that the most frequent unit of meaning in the Nature cluster is *pedofilo** (including paedophile[s] and paedophilia), confirming the violent and ongoing stigma of purported paedophilic tendencies in homosexual subjects. This subcluster can be interpreted in the nineteenth-century *clinical homophobia* paradigm (Rossi Barilli

1999), which, as we see from the data, continues to foster hate speech in more recent times. That paradigm considered reproduction of the species to be the cornerstone of normal human psychophysical health. In the name of Darwinism, homosexuality was considered a psychopathology detectable at an objective scientific level. The striking success of Freudian psychoanalysis reignited the homophobic current in scientific thought by shifting the representation of homosexuality from a physical pathology to simple deviance in the process of individual sexual development. The extensive list of works that sought to medicalise homosexuality came to an end only in 1973, when the American Psychiatric Association, and subsequently also the World Health Organization removed homosexuality from the manual of psychopathologies.⁹

It is therefore unsurprising that biological and clinical words appear with a remarkable frequency in our study, not only as a reference strategy to qualify the queer target but also as a profitable setting for homotransphobic speech.

Finally, we highlight the dangerous potential of words such as *pulizia* (cleaning) and (*fare*) *igiene* (to clean up/sweep away), which have been resemanticised with a hateful connotation and have historically been employed in genocide storytelling, such as in the case of the Rwandan genocide in 1994 (Gagliardone, Patel, and Pohjonen 2014), see (20).

(20) Le famiglie gay non esistono. Bene, cominciamo a fare un po' di igiene. #Fontana.

'Gay families don't exist. Great, let's start a clean-up. #Fontana.'

9 Due to the World Health Organization's clinical distinction between ego-syntonic and ego-dystonic homosexuality, we see the definitive removal in 1990 and the actual application in 1994. Ego-syntonic homosexuality refers to the condition whereby the person lives and accepts their homosexuality with serenity. In contrast, ego-dystonic homosexuality is the homosexual's feelings of rejection and suffering toward their own condition. (Istituto A. T. Beck, n.d.).

These words could be more harmful than the others in the subcluster since terms connected to cleaning serve as euphemistic reformulations of words for extermination, and in addition cannot be detected by automatic systems. Nevertheless, the non-explicit harmful connotation of those meanings allows the hateful content to spread freely, and contributes to what has been called the ‘trivialization of evil’, a process that can lead to a

hypertrophy of the insensibility to evil, which means there appears a systematic substitution of the good for the worse and the worse for the bad ... Such a substitution is so easy because of the assistance of language which trivializes evil, for example, the word ‘to kill’ is replaced by the expression ‘to cause death out of compassion.’ (Drożdż 2016: 7)

Table 4.3: List of words in the Nature cluster.

Frequency	Semantic family
Human beings and animals	
4	uomo (man)
3	donna (woman); vivere* (to live)
2	razza (race)
1	*pescie [pesce] (fish); animali (animals); diventare (to become); esistenza (existence); mandrillo (mandrill); nasce (to be born); ragazzo (kid, boy); squali (sharks); umani (humans); zecche (ticks)
Health and disease	
4	pedofilo* (paedophiles)
2	farmaco (drug, medication)
1	aborto (abortion); anoressia (anorexia); ansia (anxiety); autismo (autism); autolesionismo (self-harm); disturbi (disorders); fare igiene (to clean up/sweep away); impazzire (to go crazy); patologie (pathologies); pazzi (crazy people); psicosi (psychosis); pubertà (puberty); pulizia (cleaning); ritardi mentali (mental retardation); sangue (blood); scalpo (scalp); suicidio (suicide); vizietto (bad habit) patologie (pathologies)

Third cluster: Values and Customs

Although clinical and anthropological homophobia paradigms are still relevant for the production of hate speech, current homotransphobia mostly relies on what Queer Theory (Arfini and Lo Iacono 2012) has called the *heteronormative paradigm*. The third cluster (10 per cent of total semantic families) is hence informed by Queer Studies, according to which heteronormative everyday discursive practices contribute to setting the perception of non-heterosexual and non-cisgender identities in an undesirable position of ‘otherness’. We therefore identify three subclusters, namely Morality, Family, and Religion (Table 4.4), which function as regulatory devices to distinguish ethical values from unworthy and tribal values, as we see in (21).

- (21) Razza malvagia e depravata, priva di valori umani e cristiani. Scompaia pure in fretta. Sono peggio delle lobby lgbt e pro-gender. Disgustosamente amorali e immorali. Veramente vil razza dannata che lotta per distruggere l'uomo e i suoi figli in nome di una falsa libertà. Pazzi.

‘Evil and depraved race, devoid of human and Christian values. Go away quickly. They are worse than the LGBT and pro-gender lobbies. Disgustingly amoral and immoral. Truly vile damned race struggling to destroy man and his children in the name of false freedom. Crazy people.’

The ‘family’ unit of meaning is the most frequent in the third cluster, with five references, and is thus also one of the most frequent in the entire corpus. It is important to note that the term *famiglia** (family, families) frequently co-occurs with the concept of ‘non-existence’, textually realised by single words like *finte* (fake) or *inesistenti* (non-existent), by verb phrases, such as *non esistono* (they do not exist), or by implicit meanings inherent to heterosexual and traditional families, depicted as the ‘only true families’, as in (22):

- (22) @USERNAME basta con i diritti e agevolazioni a gay e finte famiglie gay...W LA NATURA E STOP AL FINITO PERBENISMO DEI CONTRONATURA...LEGGI A FAVORE DELLA VERA E UNICA FAMIGLIA(uomo e donna).....STOP AI CONTRONATURA....NO assoluto 'll'adozione di bambini a gay e cop[p]ie gay.....

'@USERNAME Enough with the rights and benefits for gay people and fake gay families...LONG LIVE NATURE AND STOP THE FALSE SELF-RIGHTEOUSNESS OF PEOPLE WHO GO AGAINST NATURE...[we want] LAWS IN FAVOUR OF THE TRUE AND ONLY FAMILY (man and woman).....NO MORE PEOPLE WHO GO AGAINST NATURE....Absolute NO to the adoption of children by gays and gay couples.....'

Unlike other types of hate speech, a large number of homotransphobic offensive words come from the religious semantic sphere, as we see from the terms *dannata* (damned), *maledetta* (cursed), and *abominio* (abomination), confirming the reference system of values in homotransphobic prejudice.

Table 4.4: List of words in the Values and Customs cluster.

Fre- quency	Semantic family
Morality	
2	etici (ethical); immorali (immoral)
1	indegne (unworthy); servire (to serve); tribali (tribal); valori (values)
Family	
7	bambini (children)
5	famiglia* (families)
2	madre (mother)
1	adozione (adoption); casa (home); figli (children); sposati (married); tetto (roof)

Fre- quency	Semantic family
Religion	
2	depravata* (depraved) [collocate for 'race']
1	abominio (abomination); dannata (damned); glorificato (glorified); maledetta (cursed); malvagia (wicked); orrore (horror)

Fourth cluster: Affectivity and Sexuality

This cluster represents the last queer-related topic found in hate speech expression in our corpus. As can be seen from the data in [Table 4.5](#), the cluster contains only ten semantic families, representing 5 per cent of the total. In addition, four semantic families include expressions that functioned as keywords in the TWEER construction phase. If we exclude those keywords, affectivity and sexuality words actually account for only 3 per cent of the entire corpus. This is hence not a cluster that naturally emerged from the data, but is the result of our specific search that was designed to detect how many and which words were chosen by haters to talk about the primary LGBTQIA+-related topic. Once again, in accordance with the results found by Motschenbacher (2018) and Locatelli et al. (2023), a very narrow list of words concerns sexual identity in Italian hate speech, specifically *amare* (to love), *co[p]pie* (couples), *effusioni* (displays of affection), *emotività* (emotionality), and *eros** (eros, erotically). Moreover, we found no occurrences of the lexeme *lesbic** (lesbian, lesbians), nor of *trans** (transgender, transgenderism, transexual). The previous quantitative analysis of the TWEER corpus confirms this data: we found 2815 occurrences of *gay*, but only 486 of *lesbic** and 364 of *trans**. Along with frequency scores, even NLP studies reported a clear prevalence of the gay target, rather than lesbian, as the main recipient of derogatory language (Locatelli, Damo, and Nozza 2023). The significant difference in the frequency of occurrence of the word *gay* versus the other two target words is in part due to the use of the word 'gay' as generic, unmarked masculine (Thornton 2016), including male and female referents. There remains,

however, an important difference in frequency between references to non-heterosexual orientations (e.g. 'gay', 'lesbian', 'bisexual') and non-cisgender identities (e.g. 'transgender', 'transsexual', 'genderfluid', 'bi-gender').

Although the narrow dimension of our subcorpus could limit the significance of this result, it seems appropriate to interpret this data in the light of the notion of 'corpus notable absences' (Partington 2014), namely infrequent or absent usage types in the corpus: 'for example, certain grammatical constructions or lexical combinations that are in principle possible but do not or only infrequently occur in a data set [which] may instantiate discourses that are perceived to be marked or non-normative' (Motschenbacher 2018: 11). In our case, the problem-oriented qualitative analysis allowed us to note the importance of 'what gets left out' (Kulick 2005) of the discourse, such as hateful references to lesbians and trans* people.

According to Borrillo (2009), the absence of a term in the corpus may not be due to an intentional selection of hated targets, but to the heterosexist matrix that affects even hateful discursive practices. Borrillo discussed a *regime or hierarchy of sexuality* where the homo/hetero dichotomy represents an exacerbation of the male/female dichotomy, which not only determines the direction of normative sexual desire and practices based on biological predisposition (sex) but also regulates socially expected behaviours in the masculine/feminine binary (gender). According to this view, lesbian identities have become subordinate to gay (male) identities, and trans* identities have been marginalised still further, respecting a precise hierarchical order, even in hate speech production.

Table 4.5: List of words in the Affectivity and Sexuality cluster.

Fre- quency	Semantic family
25	gay* (keyword in TWEER)
8	lgbt (keyword in TWEER)
5	Sesso (sex) (keyword in TWEER)
4	gender (keyword in TWEER)
2	eros*
1	amare (to love); *coppie [coppie] (couples); desiderarli (to desire them); effusioni (displays of affection); emotività (emotionality)

Offensive and hate words

Finally, we identify two further clusters: offensive words and hate words. These clusters (20 per cent of total semantic families) were intended to represent some typical linguistic features of hate speech, such as slurs or foul language, but also to provide an account of which negative qualities or images are generally associated with the target of hate. Starting with offensive words ([Table 4.6](#)), we firstly see slurs—that is, specific words to harm a certain target, such as ‘bottana’ (whore), and more terms indicating the meaning ‘faggot(s)’: *checche*, *rottoinculo*, *froci**, *ricchione*.

Table 4.6: List of offensive words.

Fre- quency	Semantic family
Slurs	
6	froci* (faggot[s])
2	checche (faggots); rottoinculo ([literal] broken ass, wreck); ricchione (faggot)
1	bottana (whore)

Fre- quency	Semantic family
Foul language	
2	coglione* (asshole); merda (shit); palle (balls)
1	cacca (poop); incazzare (to piss off); rutto (burp); sculettare (to shake one's ass)
Negative qualities	
7	inesistenti* (non-existent)
2	inutile (useless)
1	antipatici (obnoxious); arroganza (arrogance); ciecato ([vulgar] blind); cocainomani (cocaine addicts); fastidiosi (annoying); ingordo (greedy); insostenibile (unsustainable); perditempo (time waster); pistolino ([literal] little pistol – penis); rifiuti (waste); schifo (disgust); sporca (dirty); taci (shut up); vil (coward)

It is important to note that Italian dictionaries very often provide numerous alternative terms to designate homosexuality, such as *androfilia* (androphilia), *omofilia* (homophilia), *androgamia* (androgamy), *lesbismo* (lesbianism), *saffismo* (sapphism), *uranismo* (uranism). It is much rarer to find alternatives for the term 'heterosexuality'. This imbalance is clear in the Italian dictionary Treccani (online version): while *eterosessuale* (heterosexual) shows only one alternative, namely the short form *etero*, several lexical alternatives are listed for *omosessuale*. The dictionary first defines 'gay', 'lesbian', and 'bisexual', after which we find a great degree of variation by register, such as the more informal and offensive *frocio* (faggot) or *finocchio* ([literal] fennel, for faggot), and the more formal and obsolete *sodomita* (sodomite) or *pederasta* (pederast), not to mention the countless regional variants (see also [Chapter 5](#) in this volume).¹⁰

In her discussion of the representation of minorities in language use, the German scholar, activist, and writer Kübra Gümüsay

¹⁰ Lexical entry 'eterosessuale' (heterosexual) in the online Italian dictionary 'Sinonimi e contrari' by Treccani, accessed 31 March 2022, https://www.treccani.it/vocabolario/eterosessuale_%28Sinonimi-e-Contrari%29/.

(2021) perfectly describes this specific linguistic imbalance by portraying the language system as a museum. She divides people into two opposing camps: the museum visitors, also called ‘named people’, who in our case equate to heterosexual or cisgender individuals, and people on display, or ‘unnamed people’, here the non-heterosexual and non-cisgender individuals. Given this dichotomy, visitors are not referred to by multiple labels for two reasons: first, because they embody the social standard, and second, because nobody would be particularly interested in their (sexual) behaviour or external characteristics. In contrast, the unnamed people on display represent all those *weird identities* that are for some reasons ‘deviant’ from the social norm, whose behaviour is constantly in the spotlight and needs to be classified by others.

Along with this long-standing social and linguistic division, a wide variety of labels have been used to describe sexual minorities, all of which are informed by the dominant group’s heteronormative perspective. The significant number of synonyms for ‘homosexual’ with a negative connotation range from those indicating an association with child abuse (e.g. ‘pederast’), to those denoting pseudoscientific beliefs, such as the reference to the theory of sexual inversion (e.g. *invertito*, ‘sexual invert’), alongside reduction to the synecdoche of a sexualised body part (e.g. *rottoinculo*, [literally] ‘broken ass, wreck’) and connection with recurring sins in biblical tradition (e.g. *sodomita*, ‘sodomite’).¹¹

The rich and varied set of offensive words also includes generic foul language, such as *cacca* (poop), *incazzare* (to piss off), *coglione*

11 The theory of sexual inversion gained success between the 19th and 20th centuries. The theory made no distinction between biological sex and gender identity. In fact, sex inversion was first applied in biology to describe a particular behaviour of fish and amphibians. In sexology, the phenomenon concerned both transsexual and homosexual people, as people who developed ‘inverse’ sexual desires and behaviours compared to those normally expected based on their biological sex. The term, first proposed in German as ‘konträre Sexualempfindung’, was also successful in Italian translated as ‘invertito’, and remained in use as an insult to designate homosexuals (see Wikipedia, n.d.).

(asshole, jerk), and some words describing negative qualities, such as *antipatici* (obnoxious), *perditempo* (time waster), *vil* (coward). Once again, we can observe the process of assimilation of minorities (see [section on the first cluster](#)) within the cluster: queer individuals are associated with *cocainomani* (cokeheads) and *ciecati* ([vulgar] blind people) by the juxtaposition of words in the tweet in (23):

- (23) Tutti della stessa linea di sangue. Tribali, pedofili, checche sfrenate e cocainomani. Tutti con il vizietto. Vi manderei in Africa compreso il Vostro inutile palazzo a Roma.

‘[You are] All of the same bloodline. Tribals, paedophiles, wild faggots, and cokeheads. All with the vice. I would send you to Africa including your useless palace in Rome.’

With regard to hate words ([Figure 4.4](#)), we found it more useful to schematise the cluster within the ‘distance’ conceptual space, instead of merely presenting a list of words as we did previously. [Figure 4.4](#) shows three separate but communicating rectangles of meanings. The first includes all the *boundary* terms, namely all those words that express an uncrossable border, such as *stop*, *basta* ([hortative] enough), *ritiriamo* (let’s withdraw), and *fine* ([the] end).

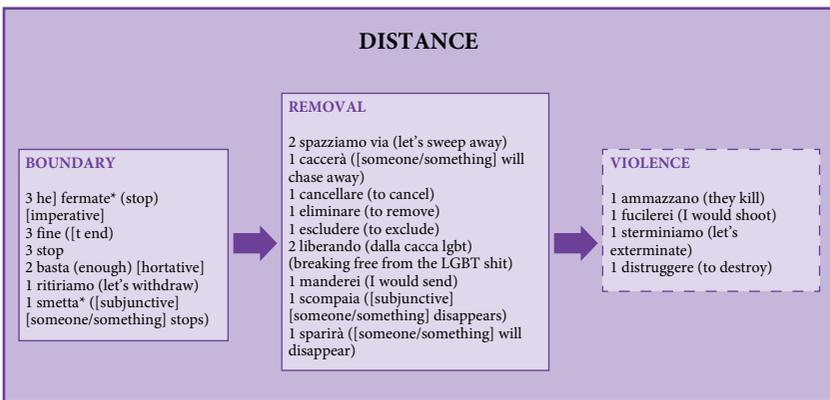


Figure 4.4: Cluster of hate words in the ‘distance’ conceptual space.

The meaning of ‘boundary’ presupposes a separation between what is included and preserved, versus what is to be left out and excluded; the same conceptual dichotomy divides in-groups from out-groups in social discursive practices. A division of this sort becomes increasingly dangerous if it has been crystallised in societies, especially in those communities that have experienced holocausts and genocide. Indeed, hate speech plays a significant and instrumental role in the perpetuation of discrimination of minorities (Gagliardone, Patel, and Pohjonen 2014): by making use of a single discursive practice, the in-group succeeds both in breaking down the minority and in reinforcing a sense of identity and community among the perpetrators. The perpetrators therefore reverse the direction of the actual hate speech: by producing fake accounts of an explicit attack waged by the minority against the in-group, haters wield hatred as the last dutiful defence available to them, as a way of defending the borders, the traditions, and even the safety of their own people, thus, transforming the ordinary social opinion gap into a dramatic ‘us vs them’ polarisation.

The immediate consequence of this dangerous polarisation can be found in the second rectangle of meanings, which is populated by *removal* terms, such as *liberando da* (breaking free from), *cancellare* (to cancel), *spazziamo via* (let’s sweep away), *caccerà* ([someone/something] will chase away), *sparirà* ([someone/something] will disappear). As the data in [Figure 4.4](#) show, the semantic space hosts all the terms that refer to a forced distancing of the out-group by the in-group; in fact, with the exception of the intransitive verbs *sparirà* ([someone/something] will disappear) and *scompaia* ([subjunctive] [someone/something] disappear), all the forms are transitive and imply significant intentionality of the part of the agent of the verb.

The last rectangle of meanings includes terms relating to *true violence*, which could relate to murder or have been historically used in context of genocide, such as *ammazzano* (they kill), *fucilerei* (I would shoot), *sterminiamo* (let’s exterminate), *distruggere* (to destroy).

Starting from the left side of [Figure 4.4](#), the first two rectangles are drawn with continuous lines, which is intended to highlight the significant connection between the groups of meanings. The third rectangle, on the other hand, has a dotted line because, as we are still dealing with speech acts, we wanted to emphasise the non-deterministic consequentiality between hate speech and hate crimes (Article 19 2015). In fact, both the social polarisation and the desire for estrangement are highly pervasive traits in hate narratives, whereas explicit incitement to extermination and explicitly violent references, especially those expressed in the first person, are fortunately very rare elements in the TWEER corpus (see [Section 4.2.3](#)). If we take the second rectangle of meanings (removal) to be an effective consequence of the first (boundary), the third rectangle (violence) can be understood as a serious and dangerous degeneration of the previous two. With this in mind, we suggest that these linguistic acts should be interpreted in light of the data regarding the distribution of stereotype labels, which accounted for 82 per cent of tweets labelled as hate speech in TWEER (see [Section 4.2.3](#)). Hence, although we recognise the significant distance between the cluster of hate words and all the previous clusters in terms of intensity of hatred, the violence rectangle can only be understood as a superficial and explicit manifestation of a dense underlying network of stereotypical narratives.

4.4 Evaluation of methodology and results

The first aim of this research was to quantify the prevalence of homotransphobic speech in an Italian Twitter corpus using the opinion mining and sentiment analysis methodology. This analysis revealed that hate speech was present in 13 per cent of the corpus, mostly conveyed by stereotypes and moderate language.

The analysis of computer-mediated texts proved to be highly effective in this context, as it allowed the analysis of specific groups of meanings in a narrow text space (see [Chapter 5](#) in this volume for similar considerations). Indeed, the binding length of tweets (280 characters in 2019) forced perpetrators of hate to be

more direct and to lay out clear argumentation strategies. At the same time, however, the impossibility of providing arguments in support of an opinion on the social network could have encouraged the production of a high number of stereotypes.

The fine-grained annotation scheme allowed us to describe hate speech behaviour on Twitter in detail, but it did have some drawbacks. First, the fact that the descriptive categories are independent of the actual presence of hate speech has caused considerable data dispersion in this study; to be precise, 20 per cent of the labels for aggressiveness, offensiveness, irony, and stereotype have been attached to harmless tweets. This occurred because the annotators had to assign six different labels to every tweet instead of analysing only harmful messages along with these descriptive features. In addition, two descriptive labels have three available values while the other two have only binary values. This value shift had the potential to cause confusion and misunderstandings during the annotation of individual tweets. The independence of the labels could be useful for general linguistic studies on computer-mediated communication, or for studies looking to explore phenomena adjacent to hate speech, such as verbal aggression (Rösner and Krämer 2016), linguistic triviality, or abusive language (Waseem et al. 2017). Yet, this kind of complexity was challenging for the goal of improving automatic hate speech detection systems. In order to achieve better performance in various detection tasks, the manual annotation should be validated by other expert or non-expert annotators by calculating the score of inter-annotator agreement. This measurement serves to establish whether the subjective opinion of the first annotator can be widely shared, therefore considered objective enough to train a machine. It is usually difficult to get a high agreement score with complex tagsets such as the one described in this chapter. In our study, we measured the inter-annotator agreement by submitting a questionnaire to non-expert annotators on Facebook. The results were far from an acceptable level of agreement: only the hate speech and irony categories showed a substantial level of agreement, while there was significant disagreement with regard

to the other four labels with two or more values; this is especially true for the intensity label, for which we observed a high level of misunderstanding. Even if expert annotators are used, opinion mining and sentiment analysis tasks always carry some degree of uncertainty and ambiguity because of the subjective interpretation and the cultural or linguistic background of the annotator.

It is also important to discuss the other two disambiguation issues involved in the hate speech detection task. Firstly, the scheme considers cases of reported speech as in (24), such as newspaper titles or statements by third parties, inevitably leading to ambiguity in the annotation. In these cases, we chose to ignore the reported hateful content by labelling the tweets as harmless.

- (24) *Mi fa ridere perché secondo mia mamma i bisessuali sono i peggiori loro vanno con tutti.*

‘It makes me laugh that according to my mum bisexuals are the worst, they screw ([literal] go with) anyone.’

We then dealt with the problem of offensive tweets that contain homotransphobic slurs but do not display any semantic link with the actual LGBTQIA+ target, as in (25). It means that the author uses a homotransphobic lexicon, but does not address actual homosexual referents.

- (25) *MACRON finirà di ammazzare il suo popolo, sto gay infame...*

‘Macron will finish killing his people, this infamous gay...’

This use of derogative words could lead to negative consequences in two ways: firstly, from a poststructuralist point of view (Motschenbacher and Stegu 2013), everyday discursive practices, and particularly idiomatic offensive language (Pinker 2007), could contribute to the social construction of the public image of the LGBTQIA+ target. Even if we do not share the desire to hurt a gay person, we still have to make the inference ‘being gay is something that negatively qualifies people’ to understand the

meaning of the sentence, thus we share, even in a passive way, the stereotype (Pistolessi 2007).

Secondly, racist and homotransphobic lexical items often co-occur in trivial messages without specific reference to any minority target, leading to false positives in automatic detection systems, as these are not sensitive to the pragmatic meaning of the sentence (Davidson et al. 2017; Malmasi and Zampieri 2017; Pamungkas, Cignarella, and Basile 2018). More broadly, negative sentiments do not necessarily coincide with hate speech: concerning homotransphobia, false positive detections may be due to users' indignation against posted discriminatory episodes or haters' comments. In these cases, it is likely that the system incorrectly recognises hate speech relying on the co-occurrence of users' anger and the queer references in the text (Locatelli, Damo, and Nozza 2023).

Looking at the findings of the lexical analysis ([Section 4.3](#)), the Politics, Nature, and Values and Customs clusters may serve as useful starting points to build full glossaries that could enrich the training phase in automatic homotransphobic detection tasks, moving beyond the detection of simple slurs and offensive language.

Specifically, we propose that attention should be paid in particular to those cases of co-occurrences between target words and apparently 'harmless words', such as *igiene* (hygiene) or *pulizia* (cleaning) (see [section on the second cluster](#)), especially in the construction '*fare* (to do) + Noun'. These constructions cannot be blocked by automatic systems because of their denotative meanings but still spread highly dangerous messages when they co-occur with a reference to a hated target.

With regard to the limits of the lexical and semantic analysis, we identified a problem in the subcorpus dimension as a result of the original goal of detecting homotransphobia. Because the group of explicit hate speech tweets was small, it was impossible to measure keyness and frequency correlations, such as collocates, which could have helped to create a better picture of the reference target; we thus believe that it would be useful to replicate the lexical and semantic analysis on a larger corpus of tweets.

Based on methodology used in corpus analysis and queer linguistics (Baker et al. 2008; Baker 2016; Motschenbacher 2018, 2019), it could be interesting to distinguish between lexical and semantic representations of single target words (e.g. ‘gay’, ‘bisexual’, ‘lesbian’, ‘genderfluid’, ‘transgender’), with the aim of exploring different collocates and co-occurrences to define sexual orientations or gender identities. It could additionally be useful to apply the same methodology to the study of queer-related noun phrases and hashtags, such as ‘rainbow families’, ‘LGBTQIA+ rights’, ‘surrogacy’, and ‘love is love’, in order to understand how the hateful storytelling changes in accordance with specific aspects of queer life.

From a sociocultural perspective, our work demonstrated that hate speech production is dependent on the heteronormative social matrix. Indeed, a large portion of the subcorpus referred to the maintenance of the status quo, by denying access to family rights for LGBTQIA+ people, by contesting a possible introduction of a sort of queer literacy in the educational system, or by simply rejecting *unnatural* non-heterosexual behaviours.

We consequently consider a further analysis of homotransphobic speech compared to the *neutral* heteronormative discursive model to a valuable addition. First, implicit hate speech is more widespread than its explicit counterpart, and it often refers to a superior natural or normative order that is presented as the only possible and right one. Second, those prejudices could also be spread within the LGBTQIA+ community through the reproduction of a dangerous top-down hierarchy of sexualities (see [section on the fourth cluster](#)).

Queer Studies have widely discussed the concept of ‘homonormativity’ (Duggan 2002; Hermann-Wilmarth and Ryan 2016; Motschenbacher 2020) by describing the homosexual adoption of heteronormative sociocultural categories in referring to or interacting with other members of the community. That is why studies of this type seek to deconstruct the sexual and gender binary categories in order to reduce discrimination against sexual minorities. Some of the consequences of ‘binary thinking’

and homonormativity could be found in the underrepresentation of lesbian and trans* individuals in activism, as mentioned previously, but also in biphobic and transphobic discrimination, which are still common within LGBTQIA+ communities. Both bisexual and trans* people, especially those who identify as non-binary instead of trans men or trans women, are often perceived as a non-existent third way, outside of the binary distinctions of homosexual/straight or male/female.

In line with the frequent references to ‘non-existent gay love and families’ in our hate speech corpus, the same kind of rejection of ‘minorities among minority’ can also be found in many other online spaces (e.g. as shown in [Table 4.7](#)); further work is therefore needed on abusive language relating to LGBTQIA+ individuals, particularly in the Italian context.

It is important to analyse homonormative speech in order to detect analogies and differences with homophobic speech produced by heterosexual and cisgender people. It would also be interesting, however, to investigate the out-group’s response to homophobic speech, both for reproduction of hate speech, such as the public call to violent actions against people who commit homotransphobic acts, and in the reappropriation of slurs by the out-group (see [Chapter 2](#) in this volume). For example, the Italian slur *frocio* is currently used as a pride device by the LGBTQIA+ community, as also happened in late 1980s with the English ‘queer’ (Perlman 2019), and it has gradually become linguistically productive with more inflected and derivational forms (e.g. *frocia*, *froce*, *frociarola*, *frociaggine*).¹²

12 The slur *frocio* commonly refers to male individuals, while in the out-group usage it could often designate female referents. Moreover, some Italian transfeminist groups often use the feminine form *froce* in a wider sense, to refer to all those individuals perceived as deviant by the social norm, regardless of their sexual orientation, thus regardless the original meaning of ‘gay man’.

Table 4.7: Biphobic examples on Facebook page Gay.it

Post on Gay.it public Facebook page	Author's translation
Il bisex è fondamentalmente uno omosex mascherato per esigenze sociali.	The bisexual is basically a homosexual disguised for social needs.
Ma esistono davvero i bisessuali?	But do bisexuals really exist?
Mi piacerebbe esserlo. L'avrei vissuto con grazia: matrimonio con erede e, consensualmente, un bel maschio come amante. Sarebbe stata una vita molto meno incasinata. Però adoro i maschi...	I would like to be. I would have lived it with grace: marriage with an heir and, by consensus, a handsome male as a lover. It would have been a much less messed-up life. But I love males...
Io non credo all'esistenza del bisex.	I don't believe in the existence of bisexuals.
I bisex esistono come patologia... anch'io ero bisex.....ma poi sono guarito e ho scelto il c...o	Bisexuals exist as a pathology...I was also bisexual.....but then I recovered, and I chose the d...k
Praticamente dovremmo festeggiare un omosessuale che si nasconde per ragioni sociali? Ah ok! (Non sono leghista prima che cominciate a sparare cavolate).	Should we basically celebrate a homosexual who is hiding for social reasons? Ah OK! (I'm not a Lega Nord voter, before you start shooting bullshit).
ah ecco la coppia aperta.....certo. Caro Davide io vivo in Inghilterra, dove a confronto l'Italia è medioevo. E sarà magari per l'opportunità che abbiamo qui di poter esprimere noi stessi senza che nulla accada, ma di bisessuali, ce ne sono ben pochi eh. Siamo tutti piuttosto convinti qui	ah here is the open couple.....sure. Dear Davide, I live in England, which makes Italy seem medieval in comparison. And maybe it will be because of the opportunity we have here to be able to express ourselves without anything happening, but there are very few bisexuals, eh. We are all pretty convinced here

4.5 Conclusion and future developments

In conclusion, our work may represent a starting point for new guidelines against online homotransphobia, and to inform future campaigns for protection laws that encompass sexual minorities, which have not yet taken hold in Italy, after the rejection of the Zan bill by the Italian Parliament in 2021 (De Carli 2021). Indeed, the

amount of homotransphobic texts circulating online is greater in religious and conservative cultural contexts such as Italy or France compared to Germany or Norway (Locatelli, Damo, and Nozza 2023: 20), and the use of overtly abusive language is highly more likely in countries with less LGBTQIA+ safety legislation. Furthermore, the clusters analysis may help to increase awareness in educational, healthcare, and corporate contexts on the heteronormative arguments underlying explicit hate, such as the denial of the existence and marginalisation of queer people (Leonard et al. 2022).

From a philosophical point of view, one of the most significant aspects of anti-homotransphobic discourse is in the treatment of hateful speech acts as inherently dangerous for collectivity and immediately harmful to the out-group. According to Jonathan Seglow (2016: 7), hate speech acts cause direct damage to the out-group in terms of agency and entitlement to self-respect. The first term refers to the ability of an individual to pursue their own goals and to affirm a personal belief within society. These goals are more respectable the more they are endorsed by other members of the social community; in this regard, Seglow stresses how self-respect is configured as a collective construct and not a strictly individual status. The concept of entitlement to self-respect refers to the respectability that individuals perceive based on their abilities, merits, titles, and rights acquired from birth and throughout their lives.

By dehumanising the out-group, hate speech flattens the individual features of the group and diminishes its merits, projecting the image of a homogeneous social category whose goals are no longer perceived as rights but as privileges. From this perspective, not only does the LGBTQIA+ community have less entitlement regarding equal marriage and parenthood, but some studies have demonstrated that the demand for these rights and further access to legal protection from hate crimes as a social minority will be perceived as an undeserved extra (Leonard et al. 2022). Further studies must therefore focus on implicit and moderate hate speech, and particularly on the concepts of political priority and privilege while analysing discrimination against social minorities.

References

- Arduini, Stefano, and Roberto Fabbri. 2013. *Che cos'è la linguistica cognitiva*. Rome: Carocci.
- Arfini, Elia A. G., and Cristian Lo Iacono. 2012. *Canone inverso: antologia di teoria queer*. Pisa: Edizioni ETS.
- Article 19. 2015. *'Hate Speech' Explained – A Toolkit*. London: Article 19. <https://www.article19.org/data/files/medialibrary/38231/'Hate-Speech'-Explained--A-Toolkit-%282015-Edition%29.pdf>.
- Assimakopoulos, Stavros, Fabienne H. Baider, and Sharon Millar, eds. 2017. *Online Hate Speech in the European Union: A Discourse-Analytic Perspective*. Berlin: Springer. <https://doi.org/10.1007/978-3-319-72604-5>.
- Attanasio, Giuseppe, and Eliana Pastor. 2020. 'PoliTeam @ AMI: improving sentence embedding similarity with misogyny lexicons for automatic misogyny identification in Italian tweets'. In *EVALITA Evaluation of NLP and Speech Tools for Italian – 17 December 2020*, edited by Valerio Basile, Danilo Croce, Maria Maro, and Lucia C. Passaro, 48–54. Turin: Accademia University Press. <https://doi.org/10.4000/books.aacademia.6807>.
- Baker, Paul. 2016. 'Gendered discourses'. In *Triangulating Methodological Approaches in Corpus Linguistic Research*, edited by Paul Baker and Jesse Egbert, 138–151. London: Routledge.
- Baker, Paul, Costas Gabrielatos, Majid Khosravinik, Michał Krzyżanowski, Tony McEnery, and Ruth Wodak. 2008. 'A useful methodological synergy? Combining Critical Discourse Analysis and corpus linguistics to examine discourses of refugees and asylum seekers in the UK press'. *Discourse and Society* 19 (3): 273–306. <https://doi.org/10.1177/0957926508088962>.
- Basile, Valerio, Danilo Croce, Maria Di Maro, and Lucia C. Passaro, eds. 2020. *Proceedings of the Seventh Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2020)*, 2765. Turin: Accademia University Press. <http://ceur-ws.org/Vol-2765/>.
- Basile, Valerio, Mirko Lai, and Manuela Sanguinetti. 2018. 'Long-term social media data collection at the University of Turin'. In *Proceedings of the Fifth Italian Conference on Computational Linguistics CLiC-It 2018*, edited by Elena Cabrio, Alessandro Mazzei, and Fabio Tamburini, 40–45. Collana Dell'Associazione Italiana Di Linguistica Computazionale. Turin: Accademia University Press. <http://books.openedition.org/aacademia/3075>.
- Bianchi, Claudia. 2017. 'Linguaggio d'odio, autorità e ingiustizia discorsiva'. *Rivista di estetica* 64: 18–34. <https://doi.org/10.4000/estetica.2059>.
- Borrillo, Daniel Ángel. 2009. *Omofobia: storia e critica di un pregiudizio*. Bari: Edizioni Dedalo.

- Bucholtz, Mary, and Kira Hall. 2004. 'Theorizing identity in language and sexuality research.' *Language in Society* 33 (October): 469–515. <https://doi.org/10.1017/S0047404504334020>.
- Caselli, Tommaso, Valerio Basile, Jelena Mitrović, and Michael Granitzer. 2021. 'HateBERT: retraining BERT for abusive language detection in English.' In *Proceedings of the Fifth Workshop on Online Abuse and Harms (WOAH 2021)*, edited by Aida Mostafazadeh Davani, Douwe Kiela, Mathias Lambert, Bertie Vidgen, Vinodkumar Prabhakaran, and Zeerak Waseem, 17–25. Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.woah-1.3>.
- Chakravarthi, Bharathi Raja, Ruba Priyadharshini, Thenmozhi Durairaj, John McCrae, Paul Buitelaar, Prasanna Kumaresan, and Rahul Ponnusamy. 2022. 'Overview of the shared task on homophobia and transphobia detection in social media comments.' In *Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion*, edited by Bharathi R. Chakravarthi, B Bharathi, John P McCrae, Manel Zarrouk, Kalika Bali, and Paul Buitelaar, 369–377. Dublin: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.ltedi-1.57>.
- Coates, Jennifer. 2013. 'The discursive production of everyday heterosexualities.' *Discourse and Society* 24 (5): 536–552. <https://doi.org/10.1177/0957926513486070>.
- Davidson, Thomas, Dana Warmley, Michael Macy, and Ingmar Weber. 2017. 'Automated hate speech detection and the problem of offensive language.' In *Proceedings of the Eleventh International AAAI Conference on Web and Social Media (ICWSM 2017)*, 512–515. Montreal: AAAI Press.
- De Carli, Mirko. 2021. 'Ddl Zan, non è prioritario I manifesti lo dimostrano.' *Il Resto del Carlino*, 13 August 2021. <https://www.ilrestodelcarlino.it/ravenna/cronaca/ddl-zan-non-e-prioritario-i-manifesti-lo-dimostrano-1.6691834>.
- De Mauro, Tullio. 2016. 'Le parole per ferire.' *Internazionale*, 27 September 2016. <https://www.internazionale.it/opinione/tullio-de-mauro/2016/09/27/razzismo-parole-ferire>.
- De Smedt, Tom, Sylvia Jaki, Eduan Kotzé, Leïla Saoud, and Maja Gwózdź, Guy De Pauw, and Walter Daelemans. 2018. 'Multilingual cross-domain perspectives on online hate speech.' In *CLiPS Technical Report 8*: 1–24. <https://doi.org/10.48550/arXiv.1809.03944>.
- Drożdż, Michał. 2016. 'Hate speech in media discourse.' *Acta Universitatis Lodziensis: Folia Litteraria Polonica* 35 (5): 19–30. <https://doi.org/10.18778/1505-9057.35.02>.
- Duggan, Lisa. 2002. 'The new homonormativity: the sexual politics of neoliberalism.' In *Materializing Democracy: Toward a Revitalized Cultural*

- Politics*, edited by Russ Castronovo and Dana D. Nelson, 175–194. Durham, NC: Duke University Press.
- Fersini, Elisabetta, Debora Nozza, and Paolo Rosso. 2020. ‘AMI @ EVALITA2020: automatic misogyny identification.’ In *EVALITA Evaluation of NLP and Speech Tools for Italian – 17 December 2020*, edited by Valerio Basile, Danilo Croce, Maria Maro, and Lucia C. Passaro, 21–28. Turin: Accademia University Press. <https://doi.org/10.4000/books.aaccademia.6764>.
- Fillmore, Charles J. 1985. ‘Frames and the semantic of understanding.’ *Quaderni di Semantica: Rivista Internazionale id Semantica Teorica e Applicata* 6: 222–254.
- Gagliardone, Iginio, Alisha Patel, and Matti Pohjonen. 2014. ‘Mapping and analysing hate speech online: opportunities and challenges for Ethiopia.’ *SSRN Electronic Journal*, 1 May 2014. <https://dx.doi.org/10.2139/ssrn.2601792>.
- Gümüşay, Kübra. 2021. *Lingua e essere*, translated by Lavinia Azzone. Rome: Fandango Libri.
- Hermann-Wilmarth, Jill M., and Caitlin L. Ryan. 2016. ‘Queering chapter books with LGBT characters for young readers: recognizing and complicating representations of homonormativity.’ *Discourse: Studies in the Cultural Politics of Education* 37 (6): 846–66. <https://doi.org/10.1080/01596306.2014.940234>.
- Istituto A. T. Beck, Centro Studi Omosessualità. n.d. ‘Lomosessualità è un disturbo mentale?’. Accessed 31 March 2022. <https://www.istitutobeck.com/omosessualita>.
- Kulick, Don. 2005. ‘The importance of what gets left out.’ *Discourse Studies* 7 (4/5): 615–624. <https://doi.org/10.1177/1461445605054408>.
- Leonard, William, Debra Parkinson, Alyssa Duncan, Frank Archer, and Catherine Orian Weiss. 2022. ‘Under pressure: developing lesbian, gay, bisexual, transgender and intersex (LGBTI) inclusive emergency services.’ *Australian Journal of Emergency Management* 37 (1): 52–58. <https://doi.org/10.47389/37.1.52>.
- Locatelli, Davide, Greta Damo, and Debora Nozza. 2023. ‘A cross-lingual study of homotransphobia on Twitter.’ In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, edited by Sunipa Dev, Vinodkumar Prabhakaran, David Adelani, Dirk Hovy, and Luciana Benotti, 16–24. Dubrovnik: Association for Computational Linguistics. <https://aclanthology.org/2023.c3nlp-1.3>.
- Liddicoat, Anthony J. 2009. ‘Sexual identity as linguistic failure: trajectories of interaction in the heteronormative language classroom.’ *Journal of Language, Identity and Education* 8 (2–3): 191–202. <https://doi.org/10.1080/15348450902848825>.
- Malmasi, Shervin, and Marcos Zampieri. 2017. ‘Detecting hate speech in social media.’ In *Proceedings of the International Conference Recent*

- Advances in Natural Language Processing, RANLP 2017*, edited by Ruslan Mitkov and Galia Angelova, 467–472. Varna: INCOMA Ltd. https://doi.org/10.26615/978-954-452-049-6_062.
- McEnery, Tony, and Andrew Hardie. 2012. *Corpus Linguistics: Method, Theory and Practice*. Cambridge: Cambridge University Press.
- Mishra, Pushkar, and Marco Del Tredici. 2018. 'Author profiling for hate speech detection.' In *Proceedings of the 27th International Conference on Computational Linguistics (COLING 2018)*, edited by Emily M. Bender, Leon Derczynski, and Pierre Isabelle. Santa Fe, NM: Association of Computational Linguistics.
- Mongie, Lauren. 2016. 'Critical Discourse Analysis as queer linguistics: religious pro- and anti-LGBT equality framing and counterframing in two letters to the editor in the City Press.' *Stellenbosch Papers in Linguistics Plus* 49: 23–43. <https://doi.org/10.5842/49-0-664>.
- Motschenbacher, Heiko. 2020. 'Walking on Wilton Drive: a linguistic landscape analysis of a homonormative space.' *Language and Communication* 72: 25–43. <https://doi.org/10.1016/j.langcom.2020.02.002>.
- Motschenbacher, Heiko. 2019. 'Methods in language, gender and sexuality studies: an overview.' *Wiener Slawistischer Almanach* 84: 43–70.
- Motschenbacher, Heiko. 2018. 'Corpus linguistics in language and sexuality studies: taking stock and future directions.' *Journal of Language and Sexuality* 7 (2): 145–174. <https://doi.org/10.1075/jls.17019.mot>.
- Motschenbacher, Heiko, and Martin Stegu. 2013. 'Queer Linguistic approaches to discourse.' *Discourse and Society* 24 (5): 519–535. <https://doi.org/10.1177/0957926513486069>.
- Musto, Cataldo, Angelo P. Sansonetti, Marco Polignano, Giovanni Semeraro, and Marco Stranisci. 2019. 'HATECHECKER: a tool to automatically detect hater users in online social networks.' In *Proceedings of the Sixth Italian Conference on Computational Linguistics (CLiC-it 2019)*, edited by Raffaella Bernardi, Roberto Navigli, and Giovanni Semeraro. Bari: CEUR-WS.
- Muti, Arianna, and Alberto Barrón-Cedeño. 2020. 'UniBO @ AMI: a multi-class approach to misogyny and aggressiveness identification on Twitter posts using ALBERTo.' In *EVALITA Evaluation of NLP and Speech Tools for Italian – 17 December 2020*, edited by Valerio Basile, Danilo Croce, Maria Maro, and Lucia C. Passaro, 29–34. Turin: Accademia University Press. <https://doi.org/10.4000/books.aaccademia.6769>.
- Norocel, Ov Cristian. 2011. 'Heteronormative constructions of Romanians: a genealogy of gendered metaphors in Romanian radical-right populism 2000–2009.' *Debatte: Journal of Contemporary Central and Eastern Europe* 19 (1–2): 453–70. <https://doi.org/10.1080/0965156x.2011.626121>.
- Nozza, Debora. 2022. 'Nozza@LT-EDI-ACL2022: ensemble modeling for homophobia and transphobia detection.' In *Proceedings of the Second*

Workshop on Language Technology for Equality, Diversity and Inclusion, edited by Bharathi R. Chakravarthi, B Bharathi, John P McCrae, Manel Zarrouk, Kalika Bali, and Paul Buitelaar, 258–264. Dublin: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.ltedi-1.37>.

- Orrù, Paolo. 2017. *Il discorso sulle migrazioni nell'Italia contemporanea*. Milan: Edizioni Franco Angeli.
- Pamungkas, Endang Wahyu, Alessandra Teresa Cignarella, and Valerio Basile. 2018. 'Automatic identification of misogyny in English and Italian tweets at EVALITA 2018 with a multilingual hate lexicon'. In *EVALITA Evaluation of NLP and Speech Tools for Italian: Proceedings of the Final Workshop 12–13 December 2018, Naples*, edited by Tommaso Caselli, Nicole Novielli, Viviana Patti, and Paolo Rosso. Turin: Accademia University Press. <https://doi.org/10.4000/books.aaccademia.4421>.
- Partington, Alan. 2014. 'Mind the gaps: the role of corpus linguistics in researching absences'. *International Journal of Corpus Linguistics* 19 (1): 118–146. <https://doi.org/10.1075/ijcl.19.1.05par>.
- Patti, Viviana, Cristina Bosco, and Rossana Damiano. 2017. 'Implicazioni etiche dell'analisi di contenuti, emozioni e interazioni nei social media'. *Sistemi intelligenti* 3: 579–600. <https://doi.org/10.1422/88510>.
- Perlman, Merrill. 2019. 'How the word 'queer' was adopted by the LGBTQ community'. *Columbia Journalism Review*, 22 January 2019. https://www.cjr.org/language_corner/queer.php.
- Peterson, Cassie. 2013. 'The lies that bind: heteronormative constructions of "family" in social work discourse'. *Journal of Gay and Lesbian Social Services* 25 (4): 486–508. <https://doi.org/10.1080/10538720.2013.829394>.
- Pinker, Steven. 2007. *The Stuff of Thought: Language as a Window into Human Nature*. London: Penguin.
- Pistolesi, Elena. 2007. 'Identità e stereotipi nel discorso conflittuale'. In *Vicini/lontani: identità e alterità nella/della lingua*, edited by Elena Pistolesi and Sabine Schwarze, 115–130. Berlin: Peter Lang Verlag.
- Poletto, Fabio, Marco Stranisci, Manuela Sanguinetti, Viviana Patti, and Cristina Bosco. 2017. 'Hate speech annotation: analysis of an Italian Twitter corpus'. In *Proceedings of the Fourth Italian Conference on Computational Linguistics (CLiC-It 2017)*, edited by Roberto Basili, Malvina Nissim, and Giorgio Satta, 263–268. Aachen: CEUR-WS. <https://doi.org/10.4000/books.aaccademia.2448>.
- Reddy, Vasu. 2002. 'Perverts and sodomites: homophobia as hate speech in Africa'. *Southern African Linguistics and Applied Language Studies* 20 (3): 163–175. <https://doi.org/10.2989/16073610209486308>.
- Rösner, Leonie, and Nicole C. Krämer. 2016. 'Verbal venting in the social web: effects of anonymity and group norms on aggressive language use in online comments'. *Social Media + Society* 2 (3). <https://doi.org/10.1177/2056305116664220>.

- Rossi Barilli, Gianni. 1999. *Il movimento gay in Italia*. Milan: Universale Economica Feltrinelli.
- Sanguinetti, Manuela, Fabio Poletto, Cristina Bosco, Viviana Patti, and 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)*, 2798–2805. Miyazaki: European Language Resources Association.
- Seglow, Jonathan. 2016. 'Hate speech, dignity and self-respect'. *Ethical Theory and Moral Practice* 19 (5): 1103–1116. <https://doi.org/10.1007/s10677-016-9744-3>.
- Silva, Leandro, Mainack Mondal, Denzil Correa, Fabricio Benevenuto, and Ingmar Weber. 2016. 'Analyzing the targets of hate in online social media'. In *Proceedings of the Tenth International Conference on Web and Social Media (ICWSM 2016)*, edited by Markus Strohmaier and Krishna Gummadi, 687–690. Cologne: AAAI Press.
- Socciarelli, Matteo. 2019. 'The "synonymy" of gay and omosessuale in Italian: a corpus linguistic analysis of two news websites'. In *TESOL Working Paper Series 17*, edited by Hanh Thi Nguyen, 46–68. Honolulu, HI: Hawaii Pacific University TESOL.
- Thornton, Anna Maria. 2016. 'Designare le donne: preferenze, raccomandazioni e grammatica'. In *Genere e linguaggio: i segni dell'uguaglianza e della diversità*, edited by Fabio Corbisero, Pietro Maturi, and Emanuela Ruspini, 15–33. Milan: Franco Angeli.
- Van Dijk, Teun A. 2004. *Ideologie, discorso e costruzione sociale del pregiudizio*, translated by Paola Villano. Rome: Carocci.
- Warner, William, and Julia Hirschberg. 2012. 'Detecting hate speech on the World Wide Web'. In *Proceedings of the Second Workshop on Language in Social Media (LSM 2012)*, edited by Sara Owsley Sood, Meena Nagarajan, and Michael Gamon, 19–26. Montreal: Association for Computational Linguistics.
- Waseem, Zeerak, Thomas Davidson, Dana Warmusley, and Ingmar Weber. 2017. 'Understanding abuse: a typology of abusive language detection subtasks'. In *Proceedings of the First Workshop on Abusive Language Online (ACL 2017)*, edited by Tommaso Caselli, Nicole Novielli, and Viviana Patti, 78–84. Vancouver: Association for Computational Linguistics.
- WCF (World Congress of Families XIII). 2019. 'Il congresso'. Accessed 2 March 2022. <https://wcfverona.org/it/about-the-congress/>.
- Wikipedia. n.d. 'Sexual inversion (sexology)'. Accessed 3 January 2023. [https://en.wikipedia.org/wiki/Sexual_inversion_\(sexology\)](https://en.wikipedia.org/wiki/Sexual_inversion_(sexology)).