

Cleansing and expanding the HURTLEX(EL) with a multidimensional categorization of offensive words

Vivian Stamou* Iakovi Alexiou† Antigone Klimi† Eleftheria Molou†
Alexandra Saivanidou† Stella Markantonatou*

*Institute for Language and Speech Processing, Athena R.C.

†Faculty of Philology, University of Athens

{vivianstamou, iakovi.alexiou, antyklimi, moloueleftheria,
alexansaivan, stiliani.markantonatou}@gmail.com

Abstract

We present a cleansed version of the Modern Greek branch of the multilingual lexicon HURTLEX.¹ The new version contains 737 offensive words. We worked bottom-up in two annotation rounds and developed detailed diagnostics of "offensiveness" by cross-classifying words on three dimensions: context, reference, and thematic domain. Our work reveals a wider spectrum of thematic domains concerning the study of offensive language than those identified in the Greek lexicographic literature as well as social and cultural aspects that are not included in the original HURTLEX categories.

1 Introduction

The term *offensive language* (OL) is used to describe "hurtful, derogatory or obscene comments made by one person to another person" and the term *hate speech* (HS) to describe speech that is possibly harmful to disadvantaged social groups.² Although both legal and ethical aspects have been considered in an effort to differentiate between HS and OL, the line between the two terms is difficult to be drawn (Davidson et al. 2017; Waseem et al. 2017) and they are often used interchangeably (Jacobs and Potter, 1998). In this work, terms in the domains of OL and HS are considered together.

Many of the studies referring to OL detection use vocabularies (Chen et al. 2012; Colla et al. 2020; Njagi et al. 2015; Pedersen 2019; Razavi et al. 2010) or patterns as a starting point and depend heavily on the selection of "seed words". Keyword-based approaches might be more effective in the case of explicit abuse according to the typology provided in Waseem et al. (2017). Also, there are strong indications that key-word and lexicon-based

approaches score better when there is a shortage of annotated corpora (Sazzed, 2021); Modern Greek (MG) is an underresourced language in terms of corpora annotated for OL.

Resource development for OL detection is an issue in itself. Firstly, "offense" is a subjective notion and as a result, the social (in general) and personal characteristics of the annotators as well as the annotation method may put bias on the resources for OL detection (lists of offensive words, corpora). The so-called "descriptive" approaches to resource development try to represent various stances in the same resource while the so-called "prescriptive" approaches try to represent few or even only one stance. High interrater scores seem to correlate with the prescriptive approach (Röttger et al., 2022). Furthermore, Schmidt and Wiegand (2017) point out that little is known about the creation process and the theoretical concepts underlying collections of offensive words. The context in which words occur also affects their offensive nature; for instance, Pelosi et al. (2017) observe that words collected in vulgar lexicons, sometimes may be considered neutral or even positive.

Our group represents female native speakers of MG with middle to high education aged 20-60; none belongs to marginal social groups. Our work is of the prescriptive persuasion. We did not make use of a pre-existing list of guidelines for recognising offensive words; instead we developed our own list of diagnostics with an iterative bottom-up procedure. We offer a cleansed version of the HURTLEX-(EL) lexicon containing 737 words after removing the wrong words and the words that were not considered offensive by all the annotators. Explanations whether the OL value of the words is context-dependent or not are offered as well as descriptions of certain contexts that trigger the offensive meanings.

¹The lexicon is available here:

https://osf.io/t5jey/?view_only=e910e28ea21e4895905aff2d0c0ac162
(archived under: DOI 10.17605/OSF.IO/T5JEY).

²<https://thelawdictionary.org/offensive-language/>.

2 OL identification studies and resources for Modern Greek

Pitenis et al. (2020) presented the first annotated MG dataset, the Offensive Greek Tweet Dataset (OGTD) that was extracted with a yet unpublished list of profane or obscene keywords (e.g., *μαλάκας* ‘asshole’, *πουτάνα* ‘whore’). Tweets were marked as “offensive”, “not offensive” or “spam”. As “offensive” were labelled tweets that contained profane or obscene language or when they could be considered offensive on the basis of the context (Pitenis 2019:32-33). These general annotation guidelines were meant for texts. Lekea and Karampelas (2018) has investigated HS in the context of terrorist argument drawing on an also unpublished list of 1265 words. Perifanos and Goutsos (2021) have combined visual and textual cues in a multimodal approach for HS detection on Twitter. 4004 tweets with the hashtag *#απέλαση* ‘deportation’ and the term *λάθρο* ‘illegal’ were annotated manually as hateful, xenophobic and racist by 3 annotators with the majority vote.

Overall, the literature on Modern Greek OL detection does not provide annotated corpora representing a wide range of registers, sizeable OL lexica or annotation guidelines. In this context, and given that lexical resources are crucial for OL identification when few or no labelled corpora exist (Sazed 2021), the Greek (EL) branch of HURTLEX (Bassignana et al., 2018) seemed a promising starting point.

HURTLEX is a domain-independent lexicon of 53 languages with offensive, aggressive and hateful words. Its kernel consists of ~1000 manually selected words corresponding to 17 fine-grained thematic categories that were enriched in a semi-automatic manner by drawing on the MultiWordnet synsets and Babelnet.^{3,4} In HURTLEX each lemma-sense pair is classified as “non-offensive” or “neutral” or “offensive”. The neutral cases were further divided into “not literally pejorative” and “negative connotation” (not a directly derogatory use). An agreement of 61% between two annotators was reported. The senses judged as non offensive were removed and two versions of the lexicon were received: one containing the translations of offensive senses and one with the additional distinction concerning the neutral cases.

³<https://multiwordnet.fbk.eu/english/home.php>.

⁴<https://babelnet.org/>.

Notably, HURTLEX aims to support the development of resources for underrepresented languages (Bassignana et al. 2018:5).

OL has been discussed in the context of MG lexicography. Efthymiou et al. (2014) show that the classification of the negative terms as derogatory, offensive, slang and taboo words in two celebrated dictionaries of MG, the LNEG2 (Babinotis, 2002) and the LKN (Triantafyllidis, 2007) do not converge. In Table 1 a tick in the sixth column denotes an overlap between the categories of OL words identified by Efthymiou et al. (2014) and our classification. Christopoulou (2012) and Xydopoulos (2012) discuss extensively experiments on the measuring of word offensiveness but do not expand on how native speakers offer the relevant evaluation.

3 Working with HURTLEX-(EL)

Although filtering has been applied to prevent noise propagation in the semi-automatically enriched HURTLEX, its EL branch still includes synsets with no offensive meaning and incorrect terms. First, we manually removed clearly incorrect terms. Two linguists agreed that these included: (i) foreign words (384 words; either in English or French), (ii) combinations of Greek and foreign words (33 words), i.e., *ευρασίας griffon*, Lit. eurasia’s griffon, (iii) about 194 meaningless phrases, i.e., *πουτίγκα κεφάλι*, Lit. pudding head, (iv) terms with morphological errors (23 words) i.e., *φυσιογνωμονική* ‘physiognomic’ instead of *φυσιογνωμική* ‘physiognomic’ (v) agreement errors (46 words), i.e., *σεξουαλικά επίθεση*, instead of *σεξουαλική επίθεση* ‘sexual assault’ (vi) different inflectional forms of the same lemma (298 words); MG makes heavy use of inflectional morphology and HURTLEX seemed unable to filter out types in the same inflectional paradigm, and (vii) archaic words (37 words), i.e., *αιχμαλωτίζων* ‘capturer’ which is an active present participle of a verb still used in MG but these particular participles belong to older forms of the language. At this stage, annotators also removed words that they all considered “unoffensive” in MG, i.e., *μοτσαρέλα* ‘mozzarella’. 2143 words (about 69% of the original HURTLEX-(EL) contents) were retained out of the 3114 original entries of HURTLEX -(EL).

Given the growing body of literature (Chakrabarty et al. 2019; Naseem et al. 2019; Ashraf et al. 2021) emphasizing the role of context in characterising a word as offensive, we adopted

an annotation schema with three categories, namely *offensive (context-independent)*, *context (context-dependent)*, following the distinction introduced in Vargas et al. (2021), and *non-offensive* entries. Representative examples were provided for terms assigned the label “context-dependent”.

Next, four independent annotators, all undergraduate linguists who offered volunteer work, assigned one of the three labels: context-independent, context-dependent, non-offensive. General diagnostics of offensiveness mainly about profane and obscene language were offered as suggestions at this stage. The interannotator agreement score in this first step was 0.77 (Fleiss kappa), which indicates an already substantial agreement.

In the final step, a somewhat different annotation procedure was adopted (see Poletto et al. 2017 for a similar approach). The four annotators were provided with a set of more detailed diagnostics of offensiveness, e.g.: “Names of animals that are stereotypically related with negative properties in the Greek culture, such as ugliness, e.g., φώκια ‘seal’ or dirt, e.g., γουρούνι ‘pig’, are offensively used when they target individuals.” These diagnostics were not developed on the basis of the classification of offensive words in the original HURTLEX or in the MG lexica (Section 2); instead, we preferred to work bottom-up and develop our own diagnostics. The motivation for this decision was that the rich material in HURTLEX-(EL) would present more classification challenges than the material in Greek printed lexica and that a Greek group’s idea of offensiveness might not be identical to that of HURTLEX, a possibility that is recognised by the HURTLEX developers (Bassignana et al. 2018:5). The annotators were asked to consult these diagnostics when classifying the terms as un/offensive but (i) they might propose changes such as deletions, additions and redefinitions of categories (ii) a term might fit to more than one category. The annotators would meet with the group leaders to discuss the diagnostics. There were three rounds of this procedure and eventually the system of thematic categories was developed as a set of diagnostics for recognising offensive words in Modern Greek; this system is presented in Section 4.

Lastly, the labels context-independent, context-dependent and non-offensive were reassigned independently by the annotators and an interannotator agreement Fleiss kappa score of 0.96 was received. We did not resort to majority vote so only 737 terms

that were shared by all the four annotators were included in the final lexicon; of them, as “context independent” were marked 448 words and as “context dependent” 289 words.

4 Annotation Diagnostics

Prose in this Section should be read with constant reference to Table 1. The final annotation diagnostics scheme comprises:

1. 17 thematic categories of offensive words
2. Tripartite distinction: offensive context-dependent, offensive context-independent and non-offensive words (Section 3). The role of the context is illustrated with the following examples: (i) the word φυτό ‘plant’ acquires derogatory meaning when it is attributed to a person (‘nerd’), (ii) the word μαλάκας ‘asshole’ loses its offensive connotation when it is used to address someone in a friendly social context (Christopoulou, 2012; Xydopoulos, 2012).
3. A subtler specification of context where words are classified by the entities that are the targets of the offensive meaning: individuals (indv.), groups, non-humans and events / properties / states (ESP). This is helpful, for instance, when individuals are assigned stereotypically negative characteristics of animals.

Below are given indicative terms and clarifications regarding the identified 17 thematic categories listed in Table 1:

1. **Social class and hierarchy:** Words implying stereotypical negative characteristics of the members of the respective social communities, e.g., χωριάτης ‘peasant’, νεόπλουτος ‘nouveau riche’, φτωχός ‘poor’, βαρώνος ‘baron’.
2. **Historical and social context:** Historical events, movements or acts are assigned a negative characterization that is absent in their historical context but it may have occurred because of their contemporary obsolete nature (Hamilton et al., 2016), e.g., σχολαστικισμός ‘scholasticism’, ηθικολόγος ‘moralist’, ακαδημαϊσμός ‘academicism’, μεσαιωνικός ‘medieval’.
3. **Crime and immoral behavior & respective agents,** e.g., δολοφονία ‘murder’ and δολοφόνος ‘murderer’, τρομοκρατία ‘terrorism’ and τρομοκράτης ‘terrorist’, ληστεία ‘robbery’, συκοφαντία ‘slander’ and σούφρωμα ‘puckering’.

4. **Religion** is viewed as a behavior not congruent with the beliefs of the Greek population and its duly constituted religion (Moon, 2018), e.g., ειδωλολατρία ‘idololatry’, μασόνος ‘mason’.

5. **Nationality/ethnicity**: Negative stereotypical ethnic characteristics are assigned to individuals of other nationalities and minorities, e.g., Εβραίος ‘Jew’, γύφτος ‘gypsy’ (Razavi et al. 2010; Warner and Hirschberg 2012). These words might be acceptable in a casual conversation if the speaker and the recipient belong to the same cultural group (Warner and Hirschberg, 2012).

6. **Politics**: In the context of democratic and liberal societies especially (Razavi et al., 2010), extreme political regimes or acts receive negative political evaluation, e.g., φασισμός ‘fascism’, χούντα ‘junta’, αποστάτης ‘renegade’.

7. **Professions of low prestige and sexual occupations**, e.g., σκαφτιάς ‘digger’, παπαράτσι ‘paparazzi’, ιερόδουλη ‘prostitute’, ζιγκολό ‘gigolo’.

8. **Animals**: Transfer of animal characteristics to humans, e.g., γουρούνι ‘pig’, γάιδαρος ‘donkey’, πρόβατα ‘cattle’, φίδι ‘snake’, τσιμπούρι ‘tick’ (Efthymiou et al., 2014).

9. **Plants**: Stereotypical negative attributes are assigned to humans regarding their cognitive skills and physical appearance, e.g., αγγούρι ‘cucumber’, πατάτες ‘potatoes’, φάβα ‘fava bean’, φυτό ‘nerd’.

10. **Characteristics of inanimates** are transferred to humans e.g., σκουπίδι ‘trash’, βαρίδι ‘sinker’.

11. **Sentiments/psychological states**: e.g., τρελός ‘crazy’, δυστυχισμένος ‘miserable’, θυμωμένος ‘mad’, μανιασμένος ‘raging’.

12. **Behavior**: People tend to criticize other people’s manner based on social norms and their own way of perceiving reality, e.g., κακότροπος ‘snappy’, λεχρίτης ‘asswipe’, εξυπνάκιας ‘smartass’, κλόουν ‘clown’.

13. **Physical and cognitive disabilities / appearance**: Assignment of specific physical or cognitive disabilities to humans (καμπούρης ‘hunchback’, τυφλός ‘blind’, χωλός ‘lame’, βλάκας ‘idiot’, κουτορνίθι ‘dumb’).

14. **Sexuality / gender identity**: Some are official terms, e.g., ομοφυλόφιλος ‘homosexual’, λεσβία ‘lesbian’, τραβεστί ‘tranny’ (Narváez et al., 2009).

15. **Taboo body parts** are context-independent offensive, e.g., αρχίδια ‘balls’, κώλος ‘ass’, παπάρι

‘whatchamacallit’, ψωλή ‘dick’. Scientific terms, e.g., χολή ‘spleen’, οπίσθια ‘buttock’ may be used offensively or as formal / scientific terminology (Crespo-Fernández, 2018).

16. **Scientific or medical terms**, e.g., ναρκισσισμός ‘narcissism’, μικρόβιο ‘germ’.

17. **Places** related to offensive occupations, e.g., μπουρδέλο ‘brothel’.

Figure 1 presents the distribution of words per diagnostic. Behavior is the most populated diagnostic followed by Crime & immoral behavior and Animals.

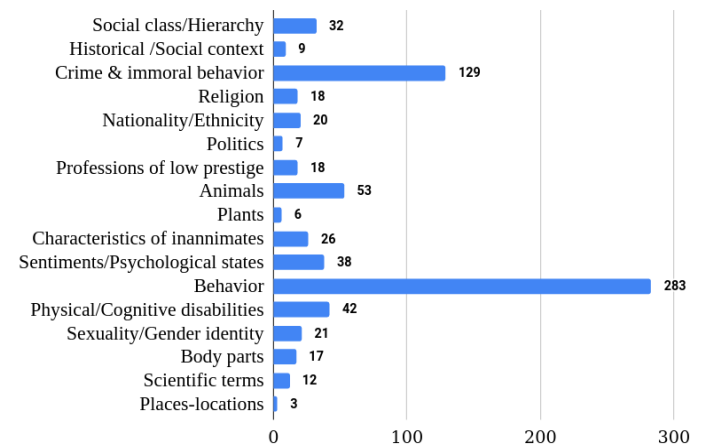


Figure 1: Word distribution per diagnostic.

5 Comparison to HURTLEX-(EL)

HURTLEX relies on a classification of OL words in 17 categories (Bassignana et al., 2018). We have defined our own diagnostics in a bottom-up iterative fashion (Section 3). The comparison of these diagnostics against the OL categories in the MG literature (sixth column of Table 1) justifies our expectations that HURTLEX would provide access to more thematic categories of offensive/derogatory words (note that all the OL categories defined in the MG literature feature among our diagnostics).

Our 17 diagnostics are equal in number with the original HURTLEX categories, but they present, probably expected, similarities and differences.

Similarities were expected because we worked on the expansion of the original 17 HURTLEX categories. However, this similarity of our independently derived diagnostics -also with the lexicographic OL categories of Greek- indicates a certain stability of OL diagnostics across different social settings, namely those of HURTLEX, of Greek lexicography which refers to the Greek society of at

	Classes	OL Target	Cont. Ind.	Cont. Dep.	Efthymiou (2014)
1.	Social class/ hierarchy	indv., groups		+	
2.	Historical/ social context	indv., groups, ESP		+	
3.	Crime immoral behavior	indv., groups, ESP	+	+	
4.	Religion	indv., groups, ESP		+	✓
5.	Nationality ethnicity	indv., groups	+	+	✓
6.	Politics	indv., groups, ESP	+	+	✓
7.	Professions of low prestige/ sexual occup.	indv., groups, ESP	+	+	
8.	Animals	indv., groups, non-human		+	
9.	Plants	indv., groups, non human		+	
10.	Characteristics of inanimates	indv., groups, non-human		+	
11.	Sentiments, psychological states	indv., ESP	+	+	
12.	Behavior	indv., groups, ESP	+	+	✓
13.	Physical/ cognitive disabilities, appearance	indv., groups, non humans	+	+	✓
14.	Sexuality gender identity	indv., groups, ESP	+	+	✓
15.	Body parts	indv., groups, ESP, non-human	+	+	✓
16.	Scientific terms	indv., groups, ESP, non-human		+	
17.	Places-locations	indv., groups, ESP, non human	+		

Table 1: Presentation of the OL diagnostics & comparison to the study by Efthymiou et al. (2014).

least 20 years ago and the contemporary Greek social settings that our group represents.

The deviation was expected because OL phe-

nomena are influenced by regional and cultural patterns (Bassignana et al. 2018). As a fact, mainly historically and culturally marked diagnostics deviate from the HURTLEX categories. The differences between HURTLEX’s categories and our diagnostics are: (i) HURTLEX’s category “SVP—words related to the seven deadly sins of the Christian tradition”: Our diagnostic 4 reflects tendencies of Greek society and contains words referring to different religions or religious states (ii) HURTLEX’s “IS—social class/ hierarchy”: Our diagnostic 1 also comprises terms denoting social and economic (dis)advantages, e.g., νεόπλουτος ‘nouveau riche’ and βαρώνος ‘baron’ (iii) We included the new diagnostic 2 “Historical / social context”, which contains contemporary terms particular to Greek history, e.g., κλέφτες ‘armatole / militiamen’ (Greek armed groups of the Ottoman occupation era); HURTLEX distributes these words in the categories “Potential negative connotations (QAS)”, “Derogatory words (CDS)” and, “Felonies and words related to crime and immoral behavior (RE)” (iv) We added the new diagnostic 5 containing terms about nationalities/minorities within the Greek ethnicity and words reflecting social and cultural differentiation, e.g., ‘Jew’, ‘gypsy’ (vi) We included the words related to sexual orientation (HURTLEX’s OM) in the single diagnostic 16 “Sexuality / gender identity”.

6 Conclusions and future work

We have discussed our experience regarding the development of an openly available, cleansed version of the Greek branch of HURTLEX; in doing so, we have defined diagnostics of offensiveness that will be useful in future offensive word and text categorisation tasks.

This was the first step in a longer-term effort that aims to offer reasonable MG lexica and corpora for the task of OL detection. On the lexicon development front we plan to study the effect of evaluative morphology on OL (Christopoulou, 2012; Stavrianaki, 2009), enlarge the lexicon semi-automatically drawing on corpora (Wiegand et al., 2018) and test its coverage and contribution to OL identification tasks using texts from a variety of registers. On the corpora development front, we intend to use the lexicon in order to leverage corpora for OL detection and for a variety of registers.

References

- Noman Ashraf, Arkaitz Zubiaga, and Alexander Gelbukh. 2021. [Abusive language detection in youtube comments leveraging replies as conversational context](#). *PeerJ Computer Science*, 7:e742.
- George Babiniotis. 2002. *Dictionary of Modern Greek Language*. Center Lexicology.
- Elisa Bassignana, Valerio Basile, and Viviana Patti. 2018. Hurltlex: A multilingual lexicon of words to hurt. In *Proceedings of the Fifth Italian Conference on Computational Linguistics (CLiC-it)*.
- Tuhin Chakrabarty, Kilol Gupta, and Smaranda Muresan. 2019. [Pay “attention” to your context when classifying abusive language](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 70–79, Florence, Italy. Association for Computational Linguistics.
- Ying Chen, Yilu Zhou, Sencun Zhu, and Heng Xu. 2012. Detecting offensive language in social media to protect adolescent online safety. *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*, pages 71–80.
- Aikaterini Christopoulou. 2012. *A lexicological analysis of slang vocabulary of Modern Greek*. PhD dissertation, University of Patras.
- Davide Colla, Tommaso Caselli, Valerio Basile, Jelena Mitrović, and Michael Granitzer. 2020. [GruPaTo at SemEval-2020 task 12: Retraining mBERT on social media and fine-tuned offensive language models](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1546–1554, Barcelona (online). International Committee for Computational Linguistics.
- Eliecer Crespo-Fernández. 2018. [Taboos in speaking of sex and sexuality](#). *The Oxford Handbook of Taboo Words and Language*, pages 41–60.
- Thomas Davidson, Dana Warmsley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Proceedings of the 11th International AAAI Conference on Web and Social Media, ICWSM ’17*, pages 512–515.
- Angeliki Efthymiou, Zoe Gavriilidou, and Eleni Papadopoulou. 2014. [Labeling of Derogatory Words in Modern Greek Dictionaries](#), pages 27–40. De Gruyter Open Poland.
- William L. Hamilton, Kevin Clark, Jure Leskovec, and Dan Jurafsky. 2016. [Inducing domain-specific sentiment lexicons from unlabeled corpora](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 595–605, Austin, Texas. Association for Computational Linguistics.
- James B. Jacobs and Kimberly Potter. 1998. *Hate Crimes: Criminal Law and Identity Politics*. Oxford University Press USA.
- Ioanna K. Lekea and Panagiotis Karampelas. 2018. Detecting hate speech within the terrorist argument: A greek case. In *Proceedings of the 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM ’18*, page 1084–1091. IEEE Press.
- Richard Moon. 2018. *Putting Faith in Hate: When Religion Is the Source or Target of Hate Speech*. Cambridge University Press.
- Dr. Rafael F. Narváez, Ilan H. Meyer, Robert M. Kertzner, Suzanne C. Ouellette, and Allegra R. Gordon. 2009. [A qualitative approach to the intersection of sexual, ethnic, and gender identities](#). *Identity*, 9(1):63–86. PMID: 27683200.
- Usman Naseem, Imran Razzak, and Ibrahim A. Hameed. 2019. Deep context-aware embedding for abusive and hate speech detection on twitter. *Aust. J. Intell. Inf. Process. Syst.*, 15:69–76.
- Dennis Njagi, Z. Zuping, Damien Hanyurwimfura, and Jun Long. 2015. [A lexicon-based approach for hate speech detection](#). *International Journal of Multimedia and Ubiquitous Engineering*, 10:215–230.
- Ted Pedersen. 2019. [Duluth at SemEval-2019 task 6: Lexical approaches to identify and categorize offensive tweets](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 593–599, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Serena Pelosi, Alessandro Maisto, Pierluigi Vitale, and Simonetta Vietri. 2017. Mining offensive language on social media. In *CLiC-it*.
- Konstantinos Perifanos and Dionysis Goutsos. 2021. [Multimodal hate speech detection in greek social media](#). *Multimodal Technologies and Interaction*, 5(7).
- Zeses Pitenis. 2019. *Detecting Offensive Posts in Greek Social Media*. Master thesis, University Of Wolverhampton, School Of Humanities.
- Zesis Pitenis, Marcos Zampieri, and Tharindu Ranasinghe. 2020. [Offensive language identification in Greek](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5113–5119, Marseille, France. European Language Resources Association.
- Fabio Poletto, Marco Antonio Stranisci, Manuela Sanguinetti, Viviana Patti, and Cristina Bosco. 2017. Hate speech annotation: Analysis of an italian twitter corpus. In *CLiC-it*.
- Amir H. Razavi, Diana Inkpen, Sasha Uritsky, and Stan Matwin. 2010. Offensive language detection using

- multi-level classification. In *Advances in Artificial Intelligence*, pages 16–27, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Paul Röttger, Bertie Vidgen, Dirk Hovy, and Janet B. Pierrehumbert. 2022. Two contrasting data annotation paradigms for subjective nlp tasks. *ArXiv* 2112.07475.
- Salim Sazed. 2021. [A lexicon for profane and obscene text identification in Bengali](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 1289–1296, Held Online. INCOMA Ltd.
- Anna Schmidt and Michael Wiegand. 2017. [A survey on hate speech detection using natural language processing](#). In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, pages 1–10, Valencia, Spain. Association for Computational Linguistics.
- Aikaterini Stavrianaki. 2009. [–άα και –ette: diminutive suffixes of Modern Greek and French. A comparative analysis and extensions regarding language learning and teaching](#). In *Proceedings of International Conference 2008, European year of intercultural dialogue: Talking with languages-cultures. Thessaloniki: Department of French Language, AUTH*, pages 759–769.
- Manolis Triantafyllidis. 2007. *Dictionary of Standard Modern Greek*. Institute of Modern Greek Studies.[Manolis Triandaphyllidis Foundation].
- Francielle Alves Vargas, Isabelle Carvalho, and Fabiana Rodrigues de Goes. 2021. Identifying offensive expressions of opinion in context. *ArXiv*, abs/2104.12227.
- William Warner and Julia Hirschberg. 2012. [Detecting hate speech on the world wide web](#). In *Proceedings of the Second Workshop on Language in Social Media*, pages 19–26, Montréal, Canada. Association for Computational Linguistics.
- Zeerak Waseem, 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*, pages 78–84, Vancouver, BC, Canada. Association for Computational Linguistics.
- Michael Wiegand, Josef Ruppenhofer, Anna Schmidt, and Clayton Greenberg. 2018. [Inducing a lexicon of abusive words – a feature-based approach](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1046–1056, New Orleans, Louisiana. Association for Computational Linguistics.
- George Xydopoulos. 2012. *Lexicology: An Introduction to the Analysis of Words and Dictionaries*. Patakis Publishers, Athens.