

Taboo Wordnet

Merrick Choo Yeu Heng and **Francis Bond**

School of Humanities

Nanyang Technological University

bond@ieee.org, merrick.emrys@gmail.com

Abstract

This paper describes the development of an online lexical resource to help detection systems regulate and curb the use of offensive words online. With the growing prevalence of social media platforms, many conversations are now conducted online. The increase of online conversations for leisure, work and socializing has led to an increase in harassment. In particular, we create a specialized sense-based vocabulary of Japanese offensive words for the Open Multilingual Wordnet. This vocabulary expands on an existing list of Japanese offensive words and provides categorization and proper linking to synsets within the multilingual wordnet. This paper then discusses the evaluation of the vocabulary as a resource for representing and classifying offensive words and as a possible resource for offensive word use detection in social media.

Content Warning: this paper deals with obscene words and contains many examples of them.

1 Introduction

The aim of this paper is to create a sense-based lexicon of offensive and potentially inappropriate terms linked to the Open Multilingual Wordnet (Bond and Foster, 2013). As well as adding new terms, we will categorize existing terms. The categorization is designed to be useful for both human and machine users. We distinguish between offensive terms, where the word itself has a negative connotation, and inappropriate terms, which may fine in some contexts, but not in others.

Real-life communication and socializing are rapidly being replaced by their online counterparts due to the overwhelming popularity and exponential growth of the use of social media platforms. Social

media platforms allow people to express their opinions and feelings on various topics, including social, cultural and political issues, mirroring tensions that are relevant in real-life conversations. While social media connects people instantly on a global scale, it also enables a wide-reaching and viral dissemination of harassing messages filled with inappropriate words. With more online conversations, the use of inappropriate words to express hostility and harass others also increases accordingly, and this is further amplified by the globalization of the internet where inappropriate words in different languages and cultures are utilized and manipulated to attack and offend. Anonymity on social media platforms enables people to be crueler and less restrained by social conventions when they use inappropriate words in these conversations. As such, the development of new linguistic resources and computational techniques for the detection, analysis and categorization of large amounts of inappropriate words online becomes increasingly important.

Machine learning from annotated corpora is a very successful approach but does not provide a general solution that can be used across domains. Our goal is to make multilingual online lexicons of inappropriate word meanings that can then be utilized in hate speech detection systems. Due to an inadequate representation and classification of inappropriate words in physical dictionaries, these online resources and systems are essential in empowering the regulation and moderation of the use of inappropriate words online. Additionally, online lexicons can be updated and edited at any time while being easily shared across the world.

This paper presents the development of a lexicon of Japanese inappropriate words that will be added to the Open Multilingual Wordnet (OMW) with links to existing synsets and the creation of new synsets to define new inappropriate words. These words in the OMW can then be used as an online lexical resource to build awareness while analyzing and identifying inappropriate words in a multilin-

gual context.

2 Background

While there is much relevant work on the detection of offensive language (Zampieri et al., 2019, 2020) ¹ lexicons of abusive words receive little attention in literature, especially for lexicons in languages other than English. Lexicons of abusive words are often manually compiled and processed specifically for a task and are not reusable in other contexts or tasks and are rarely updated once the task it was created for has been completed.

One exception is Hurltlex (Bassignana et al., 2018). It was created as a multilingual lexicon that is reusable and not task-specific or limited by context of its development. The creators of Hurltlex expanded on the lexicon of *Le parole per ferire* “words that hurt” (Mauro, 2016) and linked the words to lexical resources such as MultiWordNet and BabelNet, while translating Hurltlex into a multilingual lexicon with a combination of semi-automatic translation and expert annotation. Unfortunately, the Hurltlex release² does not include the links to the wordnets, only lists of words. It has been updated twice, with versions 1.0, 1.1 and 1.2.

There are certain criteria that a lexicon needs to fulfil in order to be effective in classifying and representing inappropriate words, let alone function as a resource for hate speech detection. The lexicon needs to be accessible, tractable, comprehensive and relevant. While Hurltlex is both tractable and accessible, it cannot be considered as truly comprehensive as much of the lexicon was translated from the original French resource. This may result in a loss of semantics and nuance, especially with regards to euphemisms and cultural expressions and has no way of adding culturally specific terms. We discuss Hurltlex more in the evaluation.

Vulgar words in dictionaries for humans are normally marked with a usage note. The naming of the note varies considerably from dictionary to dictionary, and even from edition to edition of the same dictionary (Uchida, 1997). Typically dictionaries have two or three levels, a sample is shown in Table 1.

The definitions for the terms for the American Heritage Dictionary (AHD) are shown below (cited in Uchida, 1997, p41). Note that fewer than ten

¹See also the OffensEval series of shared tasks: <https://sites.google.com/site/offensevalsharedtask/home>.

²<https://github.com/valeriobasile/hurltlex>

words are in the classes «obscene» or «usually considered vulgar».

Vulgar The label «vulgar» warns of social taboos attached to a word; the label may appear alone or in combination as «vulgar slang».

Obscene A term that is considered to violate accepted standards of decency is labeled «obscene»

Offensive This label is reserved for terms such as racial slurs that are not only insulting and derogatory, but a discredit to the user as well.

The lack of consistency in terminology, definitions of the terminology and which word is in which class show the inherent subjectivity of the decisions.

Interestingly, generally **insulting** words, such as **idiot** or **slacker**, are not normally marked in the lexicons: the only indication that a word denigrates its referent is through understanding the definition. For work on cyberbullying, such words are perhaps even more important than vulgar words.

3 Resources

To extend the wordnets, we looked at a couple of resources.

3.1 Princeton WordNet

Princeton WordNet version 3.0 (Fellbaum, 1998) has 29 different usage categories for synsets, of which we consider three to be relevant, shown in Table 2. Irrelevant categories include «synecdoche», «plural» and «trope».

The categorization is fairly hit and miss: **jap** is in **disparagement** but not **ethnic slur**, **cock** is in **obscenity** but **cunt** is not, and so forth. The same variation also occurs in the definitions. Sometimes the fact that a word is obscene is marked in other ways in the lexicon: the synset for **cock** has the definition “obscene terms for penis”. However **bugger all** is just “little or nothing at all” and while **cunt** “obscene terms for female genitals” is marked as obscene in the definition it does not have the usage category. One of the goals of this research is to create a more comprehensive list of potentially offensive words and mark them more consistently.

Pullum (2018) suggests that the fact that a slur is offensive should only be encoded in the metadata,

AHD	Examples	MWCD	Examples
vulgar	ass, dick	sometimes considered vulgar	piss, turd
obscene	fuck, cunt, shit, motherfucker	often considered vulgar	ass, balls
offensive	Polack, Jap	usually considered vulgar	fuck, cunt, dick

Table 1: Usage Notes from Dictionaries

AHD is the American Heritage Dictionary (Ed 3); MWCD is Merriam Webster’s Collegiate Dictionary (Ed 8)

Category	Frequency	Example
«disparagement»	40	suit, tree-hugger, coolie
«ethnic slur»	12	coolie, paddy, darky
«obscenity»	9	cock, bullshit, bugger all

Table 2: Princeton Wordnet Offensive Categories

not the definition itself. So something like *cock* should just be “a penis” «vulgar».³

In addition to the explicit marking, there is implicit marking in the hypernymy hierarchy: PWN has a synset *unwelcome person*, and most all of its hyponyms are insults, for example *ingrate*, *pawer*, *cad* and *sneak*.

The criteria for separating synsets is not always transparent. Maks and Vossen (2010) note that in the Dutch Wordnet, words with different connotation often appeared in the same synset. In the English wordnet, words with a different connotation are generally split into a different synset, so we have *Kraut*, *Boche*, *Jerry*, *Hun* “offensive term for a person of German descent” as a hyponym of *German* “a person of German nationality”. They noted that this structure, while allowing one to mark sentiment/connotation, is unintuitive and suggest a solution using roles. We suggest a similar solution using **Inter Register Synonymy** in Section 6.

3.2 J-lex: a list of Japanese inappropriate words

We had available a dictionary of Japanese words produced by researchers in Japan. Because they were unable to release the data themselves, they offered it to us so that we could incorporate it into the Japanese wordnet. Interestingly, the main reason they could not release the data was that they did not want their organization to be associated with a dictionary of abusive terms. However, they did want to make their lexicon available to help other researchers work on cyber-bullying. They therefore offered it to the Japanese wordnet project (Isahara

et al., 2008), a richly structured open source lexicon which is linked to wordnets in many other languages. This makes it a good means of distributing the data. The data in J-Lex was originally taken from words marked as X, “rude or X-rated term (not displayed in educational software)” by the WWW-JDICT project⁴ Breen (1995); Breen et al. (2020) with some additions by researchers on cyberbullying including Ptaszynski et al. (2010).

The list includes more than 1,600 Japanese words that are prevalent in both formal and informal speech and the words were categorized into 4 macro-categories: words related to sex, words related to bodily fluids and excrement, insulting words (used to attack and hurt) and words related to controversial topics. These macro-categories and the following sub-categories under them are not exclusive and a single word can be under multiple sub-categories. Of these 1,688 words, 1,207 are related to sex, 117 are related to bodily fluids and excrement, 468 are general insulting words while 220 are words related to controversial topics. The list of words is then further divided into 53 more specific and fine-grained sub-categories.

Looking at them, it becomes clear that not all the words are necessarily offensive: neutral words in the domain of «sex» and «excrement», like *nipple* or *urine* are fine in context but may be inappropriate out of context.

3.3 Hurtlex

Looking at the English and Japanese versions of Hurtlex,⁵ we were impressed by its size. There was a big improvement in quality from version 1.1 to 1.2,

³Metadata such as usage information will be shown encased in double angle brackets: «vulgar» while domains will be shown in single angle brackets: (linguistics).

⁴http://www.edr.org/wwwjdic/wwwjdicinf.html#code_tag

⁵<https://github.com/valeriobasile/hurtlex>

but especially in Japanese, we found many entries we considered to be errors: words that were left as English, words that were clearly mistranslated and so on.

3.3.1 Comparing J-lex and Hurltex

While Hurltex is a lexicon of words used in hate speech to attack and harass and J-lex is a general lexicon of taboo and offensive Japanese words, we expected that these lexicons about Japanese taboo and offensive words would have considerable overlap. However, despite the Japanese version of Hurltex having 5,428 unique items and J-lex having 1,688 unique items, only 154 unique words appear in both of them.

One major reason is that Hurltex is a translated lexicon, and thus is missing many native Japanese expressions. In addition, there are 758 non-translated items in the Japanese version of Hurltex. These 758 words are presented in the ISO basic Latin alphabet and examples include *animalia*, *arsehole* and *ballock*. Some of these words are neither English nor Japanese. Furthermore, translating words used in hate speech that are often lexicalized and require nuance causes words like 鳥肉 *toriniku* “chicken meat” and 小鳥 *kotori* “small bird” to be classified as hate speech despite their neutral connotations in the Japanese’s language and culture: we guess the first is a mistranslation of *chicken* “coward”, we have no idea why the second one is there.

3.4 Vulgar words

We also accessed a list of vulgar words curated by Cachola et al. (2018) from <https://www.noswearing.com/>.⁶ This had 267 English swear words, divided into «general», «homosexual» and «slur».

4 Building the Taboo Wordnet

4.1 Linking J-lex to the Japanese wordnet

In order to make the J-lex data available to the wordnets we first had to link words to senses. The first step of development consisted of linking the 1,688 unique items extracted from the Japanese lexicon J-lex to existing synsets in the Open Multilingual Wordnet. First we did this through looking up words in the Japanese wordnet, and were able to link

397 (23%) of all unique words. An example is given in (1).⁷

(1)	[lem: ja	し尿, 屎尿]
		pron: ja	しにょう <i>shinyou</i>	
		class	shit01	
		synset	OMW 14855635-n	
		def: ja	人間の体内からの排出物	
		def: en	the body wastes of human beings	

This matching was done even on low confidence Japanese entries: that is those that were automatically created but not hand-checked (Bond et al., 2008). If we got a match then we raised the confidence. Having the automatically generated low confidence entries proved to save some time. In Princeton Wordnet 3.0 this is not marked as a taboo word in any way, and the word does not appear in Hurltex.

As a second step, the remaining 1,291 words were analyzed manually with the help of other Japanese online dictionaries such as WWWJDIC. A further 421 unique words could be linked to existing synsets for a total of 818 (46%). We give some examples in (2), (3) and (4).

(2)	[lem: ja	手こき]
		pron: ja	てこき <i>tekoki</i>	
		class	sex02	
		synset	OMW 00856193-n	
		def: ja	マスターベーションを意味する俗語	
		def: en	slang for masturbation	

(3)	[lem: ja	けばい]
		pron: ja	けばい <i>kebai</i>	
		class	insult09	
		synset	OMW 02393791-a	
		def: ja	目を引く趣味悪さの	
		def: en	tastelessly showy	

(4)	[lem: ja	支那人]
		pron: ja	しなじん <i>shinajin</i>	
		class	insult15	
		synset	OMW 09698337-n	
		def: ja	中国系の人にとっては不快な言葉	
		def: en	offensive term for a person of Chinese descent	

Of the remaining 870 unique words, 71 were judged to be compositional and did not need to be included into the Japanese Wordnet. 226 words were considered as genuinely inappropriate and still relevant and thus require new synsets to encompass

⁶Taken from <https://github.com/ericholgate/vulgartwitter>.

⁷lem is the lemma, pron is the pronunciation (in hiragana, we also give the transliteration here), synset is the ID in the Japanese wordnet (normally the same as the synset offset in PWN 3.0, def is the definition.

their lemmas. The final 573 words need to be reviewed by native Japanese speakers as to whether they are truly inappropriate and whether they are used widely enough to be entered into the lexicon. A majority of the 226 new synsets are related to the domain of sexual activity rather than disparagement while there are no synsets related to the usage note of vulgar.

Examples of words deemed compositional include 変態オヤジ *hentai oyaji* “pervert old man” and 性格わるい *seikaku warui* “personality bad”. The former expression is used to describe a perverted old man and is a combination of the word 変態 *hentai* “perverted” and オヤジ *oyaji* “middle aged/old man” or “one’s own father”. Similarly, 性格わるい *seikaku warui* “personality bad” is used to describe someone who has a bad character or personality and is a combination of the word 性格 *seikaku* “character/personality” and わるい *warui* “bad”.

On the other hand, it is not as clear cut to differentiate whether words are genuinely inappropriate. For example 巨乳 *kyonyu* “huge breasts” is generally used positively but would be inappropriate in a work place. 短足 *tansoku* “short-legs” is generally insulting, but does not absolutely have to be. Words like ブヨブヨ (*buyobuyo*) meaning soft and flabby is generally offensive while 同和地区 (*dowa chiku*) meaning “untouchable area, slums” is always offensive.

We made a new Japanese extension of wordnet, that adds the new lexical entries from J-Lex to the appropriate synsets. It will be made available at <https://github.com/bond-lab/tabown> and shared with the Japanese Wordnet Project (Isahara et al., 2008).

4.2 Re-Labeling Wordnet

We decided to mark words in two different ways. First, we use the general domain category link to link words into topic domains, that are not necessarily taboo, but may be of interest to research into taboo terms. Existing topic domains include things like ⟨law⟩, ⟨music⟩ or ⟨terrorism⟩. To these we will add: ⟨*sexual activity*⟩, ⟨*excrement*⟩ and ⟨LGBTQ+⟩ (a new synset). At least in English and Japanese, these domains are potentially inappropriate without being necessarily offensive. In general, anything marked in J-Lex with *sex** will be put into the ⟨*sexual activity*⟩ topic, anything marked *shit** into ⟨*excrement*⟩ and *controversial07* will be marked as

Tag	Number	Example
⟨excrement⟩	33	shit, toilet bowel
⟨LGBTQ+⟩	9	gay, lesbian
⟨sexual⟩	274	promiscuous, arouse
⟨disparagement⟩	630	lunatic, bimbo
⟨ethnic slur⟩	14	gringo, redneck
⟨obscenity⟩	48	chickenshit, butch

Table 3: Usage and Domain tags in the Taboo Wordnet

⟨LGBTQ+⟩.⁸

For usage notes, there is very little agreement as to what should be marked in dictionaries (Sakwa, 2012). So we decided to keep the three broad categories already used by wordnet: ⟨disparagement⟩ for words which are basically insulting; ⟨obscenity⟩ for words that are considered inappropriate in many circumstances, typically because of their association with a taboo subject and ⟨ethnic slur⟩ for ethnic slurs.

We took advantage of both J-Lex and the structure of wordnet to mark offensive words. We marked with ⟨disparagement⟩ everything labelled in J-Lex as *insult** or *controversial99* as well as all hyponyms of *criminal*, *unpleasant person* or *bad person* as ⟨disparagement⟩.

Finally, we were still were missing information about which terms should be marked as ⟨obscenity⟩. To increase our coverage, we manually checked the English wordnet against the terms in <https://www.noswearing.com/> which gave another 38 vulgar synsets.

We end up with a total of 912 synsets marked in some way, with many being marked with multiple tags. The breakdown per tag is given in Table 3.

These categorizations will be made available as a stand-alone file at <https://github.com/bond-lab/tabown>, along with links to the original J-Lex categories. In addition, we will share them with the English Wordnet Project (McCrae et al., 2020). We will also release scripts to use the Open Multilingual Wordnet to generate offensive word lists for any language in the OMW using the Wn Python Library (Goodman and Bond, 2021).

⁸We also added some new words to this domain as part of a separate project by our annotators to improve the coverage of LGBTQ+ terms.

Lexicon	# Words	Comment
Wordnik	1,282	
Wordnet Original	50	
Wordnet Offensive	1,512	645 synsets
Wordnet Extended	2,095	912 synsets
HurtLex Conservative	2,228	
HurtLex Inclusive	5,965	

Table 4: Offensive Word Lists

5 Evaluation

We used a curated list to test the coverage of our enhanced wordnet. It is the list of 1,285 offensive words from the Wordnik online English dictionary.⁹ Note that, while we had originally worked on adding Japanese senses, we are using the multilingual wordnet links to produce English data here.

We compare three versions of the wordnet: **Original** which just has those entries marked as «disparagement», «ethnic slur» or «obscenity» in the original PWN 3.0; **Offensive** which has those marked as such in the Taboo Wordnet; and **Extended**, which also includes everything from the domains of «excrement», «sexuality» and «LGBTQ+». We also compared the results to the Hurlex (1.2), using both the Conservative and Inclusive lists. The resources are summarized in Table 4.

The comparison in Table 5 shows a surprisingly small overlap. The original wordnet does very badly. The Taboo wordnet matches with roughly the same accuracy as Hurtlex, with far less noise. The resources differ in their use of number. When we did error analysis, we noticed that both Wordnik and Hurtlex often had both singular and plural forms of terms (e.g. *gringo* and *gringos*) or sometimes only plural forms. We automatically lemmatized everything to singular (results shown in the final column). This improves wordnet’s coverage but degrades hurtlex (as plural forms are not longer being counted separately).

An analysis of errors shows three main causes for the lack of cover. The first is that wordnik’s entries are not necessarily lemmatized in the same way as wordnet: *beating your meat* rather than *beat one’s meat* “masturbate”, respectively. The second is that many existing synsets do not have all the possible variants, especially colloquial entries: e.g. *nut sack* for “scrotum”. Finally, there is still a long tail of new

⁹<https://www.wordnik.com/lists/offensive>, we removed a couple of non-offensive entries.

Lexicon	in Wordnik	singular
Wordnet Original	18	18
Wordnet Offensive	82	86
Wordnet Extended	163	172
HurtLex Conservative	98	95
HurtLex Inclusive	139	137

Table 5: Comparison with Wordnik

synsets: e.g., *happy ending* “A handjob, especially one after a massage”. We estimate that around 80% of the missing entries could go into existing synsets, and around 20% would need new ones (approximately 200).

6 Discussion and Future Work

The Taboo Wordnet is based on synsets, not words. This is important as words such as *しゃくはち* which could either mean *しゃくはち shakuhachi* “a Japanese and ancient Chinese longitudinal, end-blown bamboo-flute” or *しゃくはち shakuhachi* “fellatio or oral stimulation of the penis”. A word based lexicon such as HurtLex will confuse such uses.

In future work, we intend to:

- Help upstream wordnets integrate this data
- Add missing senses for existing synsets in English (as we have done for Japanese): the coverage of colloquial expressions is still low
- Create new synsets for missing concepts from J-Lex and Wordnik
- Reorganize the wordnet structures: currently a synset with offensive terms is normally linked as a hyponym of the neutral term. However, it shares the same denotation, with a different connotation. This is the relation of Inter Register Synonymy, described in Maziarz et al. (2015), and now added in the Global Wordnet Association format (McCrae et al., 2021). We will replace hyponym with inter register synonymy where appropriate.
- Add exclamation, like *fuck off* “go away”, using the extension of Morgado da Costa and Bond (2016). In Japanese, a typical equivalent is *消えろ go away* “disappear”.
- Decide how to deal with reclaimed slurs, those historically derogatory names or term that are

used or reinterpreted in a positive way, as in pride for one's social group. As slurs are generally only partly reclaimed it is important to make sure that this is clear to the dictionary user.

- Decide how to deal with expressions that were not inappropriate in the past but are currently inappropriate due to new contexts that have emerged in the recent years. As these expressions may not be unacceptable for all, it is important for the dictionary user to understand that there are two separate but related versions of the expressions.

As the Taboo Wordnet is an online resource, it can be updated frequently, and this allows for the Taboo Wordnet to reflect the ever changing status of words and expressions. Such changes in status may be due to reclamation of offensive words for in-group social pride or due to new contexts that have emerged in recent years causing some neutral words to take on a new meaning. We welcome contributions.

7 Conclusion

The main contribution of this paper is the categorization of offensive and potentially inappropriate synsets. We have also added many Japanese for these words. Through the collaborative interlingual index, they can be used for future categorization and analysis of words in other languages. These additions will be made available in the OMW as a resource for future studies on inappropriate words and may function as a sense-tagging tool for these words as well.

Acknowledgements

We would like to thank the anonymous providers of the Japanese lexical data. This research was mainly carried out during NTU's Undergraduate Research Experience on Campus Programme (URECA).

References

Elisa Bassignana, Valerio Basile, and Viviana Patti. 2018. Hurtlex: A multilingual lexicon of words to hurt. In *5th Italian Conference on Computational Linguistics, CLiC-it 2018*, volume 2253, pages 1–6. CEUR-WS. URL <http://ceur-ws.org/Vol-2253/paper49.pdf>.

Francis Bond and Ryan Foster. 2013. Linking and extending an open multilingual wordnet. In *51st Annual Meeting of the Association for Computational Linguistics: ACL-2013*, pages 1352–1362. Sofia. URL <http://aclweb.org/anthology/P13-1133>.

Francis Bond, Hitoshi Isahara, Kyoko Kanzaki, and Kiyotaka Uchimoto. 2008. Boot-strapping a WordNet using multiple existing WordNets. In *Sixth International Conference on Language Resources and Evaluation (LREC 2008)*. Marrakech.

James Breen, Timothy Baldwin, and Francis Bond. 2020. The Japanese dictionary entry: Lexicographic issues and terminology. In *Globallex Workshop on Lexicography and Neologism at Euralex 2020 (GWLN 2020)*. To appear.

James W. Breen. 1995. Building an electronic Japanese-English dictionary. Japanese Studies Association of Australia Conference (http://www.csse.monash.edu.au/~jwb/jsaa_paper/hpaper.html).

Isabel Cachola, Eric Holgate, Daniel Preoțiuc-Pietro, and Junyi Jessy Li. 2018. Expressively vulgar: The socio-dynamics of vulgarity and its effects on sentiment analysis in social media. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2927–2938. URL <http://aclweb.org/anthology/C18-1248>.

Christine Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database*. MIT Press.

Michael Wayne Goodman and Francis Bond. 2021. Intrinsically interlingual: The Wn Python library for wordnets. In *11th International Global Wordnet Conference (GWC2021)*. (this volume).

Hitoshi Isahara, Francis Bond, Kiyotaka Uchimoto, Masao Utiyama, and Kyoko Kanzaki. 2008. Development of the Japanese WordNet. In *Sixth International conference on Language Resources and Evaluation (LREC 2008)*. Marrakech.

Isa Maks and Piek Vossen. 2010. Modeling attitude, polarity and subjectivity in wordnet. In *In Proceedings of Fifth Global Wordnet Conference, Mumbai, India*.

Tullio De Mauro, editor. 2016. *Le parole per ferire*. 'Joe Cox' Committee on intolerance, xenophobia, racism and hate phenomena, of the Italian Chamber of Deputies.

- Marek Maziarz, Maciej Piasecki, Stan Szpakowicz, and Joanna Rabiega-Wisniewska. 2015. Semantic relations among nouns in polish wordnet grounded in lexicographic and semantic tradition. *Cognitive Studies*, 11:161–182.
- John P. McCrae, Michael Wayne Goodman, Francis Bond, Alexandre Rademaker, Ewa Rudnicka, and Luis Morgado da Costa. 2021. The global wordnet formats: Updates for 2020. In *11th International Global Wordnet Conference (GWC2021)*. (this volume).
- John P. McCrae, Ewa Rudnicka, and Francis Bond. 2020. English wordnet: A new open-source wordnet for English. Technical report, K Lexical News. URL <https://kln.lexicala.com/kln28/mccrae-rudnicka-bond-english-wordnet/>.
- Luís Morgado da Costa and Francis Bond. 2016. Wow! what a useful extension to wordnet! In *10th International Conference on Language Resources and Evaluation (LREC 2016)*. Portorož.
- Michał Ptaszynski, Paweł Dybala, Tatsuaki Matsuba, Fumito Masui, Rafał Rzepka, and Kenji Araki. 2010. Machine learning and affect analysis against cyber-bullying. In *Proceedings of the Linguistic And Cognitive Approaches To Dialog Agents Symposium, Rafał Rzepka (Ed.), at the AISB 2010 convention*. AISB 2010 Organizing Committee, De Montfort University, Leicester, UK. URL https://www.researchgate.net/publication/228791020_Machine_Learning_and_Affect_Analysis_Against_Cyber-Bullying.
- Geoffrey K. Pullum. 2018. *Bad Words: Philosophical Perspectives on Slurs*, chapter Slurs and Obscenities: Lexicography, Semantics, and Philosophy. OUP.
- Lydia Sakwa. 2012. Problems of usage labelling in English lexicography. *Lexikos*, 21(1). URL <https://lexikos.journals.ac.za/pub/article/view/47>.
- Hiroaki Uchida. 1997. The treatment of vulgar words in major English dictionaries. *Lexicon*, 27:152–168.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. Semeval-2019 task 6: Identifying and categorizing offensive language in social media (offenseval). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 75–86.
- Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. SemEval-2020 Task 12: Multilingual Offensive Language Identification in Social Media (OffenseEval 2020). In *Proceedings of SemEval*.