BullyFrame: Cyberbullying Meets FrameNet

Silvia Brambilla[‡], **Alessio Palmero Aprosio**[†], **Stefano Menini**[†]
[†]FBK (Trento), [‡]University of Bologna
silvia.brambilla2@unibo.it
{aprosio,menini}@fbk.eu

Abstract

English. This paper presents BullyFrame, a dataset of cyberbulling interactions collected from WhatsApp conversations in Italian and annotated with FrameNet semantic frames. We will describe the creation of the dataset discussing the problematic aspects found in the annotation process, such as the lack of coverage of FrameNet for the annotation of texts extracted from social media. Finally, we present a preliminary study that describes the relations between the frames and the cyberbullying-related annotation of the original dataset. [1]

Italiano. *Questo studio presenta BullyFrame, un dataset di conversazioni WhatsApp in italiano contenenti episodi di cyberbullismo e annotate secondo i frame semantici di FrameNet. Verrà descritta la creazione del dataset discutendo gli aspetti problematici incontrati nel processo di annotazione, come ad esempio i limiti di copertura di FrameNet per l'annotazione di testi estratti da social media. Infine, presentiamo uno studio preliminare che descrive le relazioni tra l'annotazione di FrameNet e quella del dataset originale, relativa al cyberbullismo.*

1 Introduction

The semantic analysis of a text involves the classification of predicates into a set of events, for which it is important to determine who did what, when and where. For example, in the sentence "*In 1912, the Titanic hit an iceberg on its first trip across the Atlantic*", the verb "*hit*" represents the event, "*Titanic*" is the main actor of that event, "*1912*" and "*Atlantic*" indicate when and where it took place, and so on. The process of extracting the semantic roles and relations in a sentence is called Semantic Role Labeling (SRL), and, in the last years, both resources listing possible events and corpora have been annotated with this kind of information. Examples of such datasets are FrameNet (Ruppenhofer et al., 2006) and PropBank (Palmer et al., 2005). Given the availability of these resources, over the years SRL has gained more attention and has become an important task in computational linguistics, with a growing number of works and evaluations (QasemiZadeh et al., 2019; Basili et al., 2012).

Unfortunately, the vast majority of annotated datasets relies mainly on newswire and narrative texts, and their coverage turns out to be inadequate when it comes to annotate more specific domains, such as, for instance, football domain (Torrent et al., 2014) or medicine domain (Tan et al., 2011).

Aside from that, over the last decades, ICT technologies and communication habits underwent profound changes, with the greatest part of text production in the world coming from social networks and being usually written in non-standard language.[2] This kind of communication is of fundamental importance, in particular for teenagers' social life. For instance, according to the last report by the Italian Statistical Institute (ISTAT, 2014) in Italy 82.6 of children aged 11-17 use the mobile phone every day. The use of these new technologies, however, leads also to some undesirable side effects, as the proliferation of hate speech and the digitization of traditional forms of harassment, also known as cyberbullying.

Many studies (O'Moore and Kirkham, 2001; Fekkes et al., 2006; Farag et al., 2019) have high-

[2]https://www.domo.com/learn/data-never-sleeps-6

lighted that cyberbullying can have a negative impact on the victims' psychological and emotional well-being and that, in extreme cases, it can lead to self-harm and suicidal thoughts. For this reason, some strategies have been implemented to detect and contrast this phenomenon (Van Hee et al., 2018; Zhao et al., 2016; Menini et al., 2019), but none of them makes use of SRL, and no resources on this topic based on frame semantics have been developed yet. We therefore developed BullyFrame, a dataset annotated with frame semantic annotation, where the messages are taken from a corpus of data on cyberbullying interaction in Italian, gathered through a WhatsApp experimentation with lower secondary school students (Sprugnoli et al., 2018). Our work leads to the release of the annotated corpus (see Section 3), and constitutes a feasibility study, that investigates the potential lacks of FrameNet - resource that does not claim to be exhaustive in its coverage - for the annotation of online chats that, in addition to their non-standard nature, contain offensive language and informal expressions. We show, for instance, that some frames are completely missing, such as those regarding sexual orientation, as discussed in Section 4. In other cases, FrameNet provides frames whose purpose is similar to the needed one, but cannot fit perfectly the meaning of the sentence. For example, the frame "Offenses" refers to acts that violate a legal code, but it is not used for marking offenses (or bad words) between two users, e.g. "idiota" ("idiot" - currently tagged as `Mental_property`), "stronzetta" ("asshole" - left currently with no annotation). Similarly, a sentence like "Ti ricordo che io ho ballato con Kledi" ("I remind you that I danced with Kledi") cannot be correctly annotated, as neither `Evoking`, nor `Reminder` or `Remembering_*` frames are able to capture the meaning of someone who reminds something to another person.

In Section 5, we also provide a comparison study to highlight relations between the newly-released frame annotation and the existing one regarding the type of cyberbullying expression. Results show that some of them are strictly connected (even when it is not immediate to understand).

Finally, in Section 6 we present Framy, a frame annotation tool that works as a web server and that has been used for annotating BullyFrame.

2 Related Work

The work presented in this paper spans topics from different research areas. As for the methodology, we deal with issues related to the annotation of Italian texts with FrameNet and frame annotation on social media texts. Then, as case study, we focus on the cyberbullying domain, where we witness a growing interest and a large number of novel works over the last few years.

The FrameNet database is a resource originally developed for the English language that has proven to be largely portable over different languages. This because its frames appear to be mostly language independent, as pointed out by Gilardi and Baker (2018). Nevertheless, some language specific differences can arise both at the level of frames themselves (coarse-grained level) and at the level of frame elements (FEs) (fine-grained level) (Lönneker-Rodman, 2007). As an example it is possible to recall the works of Candito et al. (2014) on French, of Ohara (2012) on Japanese and of Subirats and Sato (2004) on Spanish. In all the three languages the creation of a FrameNet-like resource required to add new frames or FEs or modify already existing ones, for instance in French some frames needed to be merged, while others needed to be split into two subframes.

For the Italian language, we rely on previous researches, carried out at the Universities of Bologna and Roma Tor Vergata (Basili et al., 2017; Vanzo et al., 2017), Fondazione Bruno Kessler in Trento (Tonelli et al., 2009; Tonelli and Pianta, 2009; Tonelli, 2010) and Pisa (Johnson and Lenci, 2011), that investigated the creation of an Italian FrameNet and first annotated Italian texts with frames.

Gerrard et al. (2017) outline how frame annotation of texts extracted from social media could be challenging because of the differences between social media data and the kind of data on which FrameNet is built, i.e. edited and well-formed sentences. For this reason as for today only few studies annotated social media texts with frame information (Kim and Hovy, 2006; Gerrard et al., 2017; ElSherief et al., 2018) even if it proved to be useful for example in identifying opinions with their holder and topic (Kim and Hovy, 2006) or in deepening the analysis of Directed and Generalized hate speech (ElSherief et al., 2018).

Works on cyberbullying try to detect and pre-

vent the phenomenon exploiting different methodologies and techniques. In particular, a dataset extracting data from Facebook has been developed at University of Pisa (Del Vigna et al., 2017), while at the University of Turin a similar corpus has been created from Twitter (Sanguinetti et al., 2018). Dinakar et al. (2011) build individual topic-sensitive binary classifiers, Van Hee et al. (2018) perform classification based on n-grams and specific features as the presence of aggressive and subjective language, while Zhao et al. (2016) apply different weights to pre-defined insulting words using them as bullying features combined with bag-of-words and latent semantic features for their classifier.

As for today, at the best of our knowledge, there are not research works that studied the possible interconnections between cyberbullying and frames.

3 Dataset Description

For the annotation of the frames related to cyberbullying we use as starting point the dataset from Sprugnoli et al. (2018). The dataset presents a collection of WhatsApp chats written by 12-13 years old students simulating instances of cyberbullying in specific scenarios.

The text of the chats is provided with annotations about *i)* the *role* of who is writing (i.e. Victim, Bully, or supporter of one of the two sides) and *ii)* labels with the *type* of offense that can be found on each message (in particular, the labels include: Threat or blackmail, General Insult, Body Shame, Sexism, Racism, Curse or Exclusion, Insult Attacking Relatives, Harmless Sexual Talk, Defamation, Sexual Harassment, Defense, Encouragement to the Harassment, and Other).

The dataset consists of 10 chats, for a total of 2192 messages (14,600 tokens) and includes 1,203 cyberbullying expressions, corresponding to 6,000 tokens.

Starting from this, we fully annotated the sentences referring to FrameNet 1.7: the resulting annotation is available for download from the resource website.[3] It is released under the Creative Commons Attribution-ShareAlike 4.0 International license.[4]

A total of 2,458 frames and 2,769 frame element have been annotated on 1,558 sentences. The remaining 1,211 sentences cannot be annotated, mainly because no corresponding frames can be found (1,180 sentences), or because there was a picture instead (19 sentences), or finally because the messages have been deleted by the user (12 sentences). Table 1 (a) shows statistics on how many frames have been annotated for each sentence. Regarding the coverage, a total of 268 unique frames and 696 unique frame elements have been found in the dataset. Table 1 (b) shows the most frequent frames that have been annotated. Finally, Table 1 (c) shows statistics on how many frame elements are annotated for each frame.

[3] https://github.com/dhfbk/bullyframe

[4] https://creativecommons.org/licenses/by-sa/4.0/

4 Frame Annotation

In order to investigate possible connections between frames and cyberbullying we annotated all the sentences of the dataset with frames and frame elements referring to the 1.7 version of FrameNet. In each sentence we tried to annotate all the possible evoked frames alongside with their frame elements.

When annotating the sentences we have to face some problems that, due to the nature of this dataset, to the differences between English and Italian, and to the nature of FrameNet itself, is not complete but that is constantly updated and enlarged.

Problematic aspects can be found on three different levels: Frames layer, Frame Elements layer and Frame Evoking Elements layer.

Frames layer: We found that some of the concepts that were evoked by lexical units (LUs) were not present in FrameNet. The missing frames could be:

a) Concepts that are new to FrameNet and that are linked to the particular nature of the text. This is the case for instance of frames that occur often in conversations or in oral communication. These concepts are often not present in FrameNet, but frequent in our dataset since it includes interactions between participants and is close to oral communication. For example we found that FrameNet does not have a frame that covers "greetings", evoked in sentences such as:

> *"Ciao ci sentiamo domani" (Bye, we'll talk tomorrow)*
>
> *"Hahahah esatto ciao e buon allenamento" (Hahahah, exactly bye and have a good training)*

Frames	Sentences
8	2
7	2
6	7
5	8
4	46
3	132
2	406
1	955
0	603
Pic	19
Del	12

Frequency	Frame
167	Silencing
138	Desirability
109	Statement
108	Correctness
107	Cause_emotion
97	Desiring
87	Awareness
83	Opinion
73	Capability
69	Intentionally_act

Frame elements	Frames
4	7
3	118
2	633
1	1121
0	332

Table 1: These three tables show: (a) the number of sentences with the corresponding amount of frame found in them; (b) the frequencies of the top 10 frames; (c) the frequencies of frame elements for each frame annotation.

"Buongiorno a tutti!" (Have a good day, everybody!)

b) Concepts that are new to FrameNet and that are linked to abusive language and cyberbullying. For example we found that bullies often refer to people's sexual orientation as an insult such as in:

"Crede di essere figo facendo il gay a danza" (He thinks he looks cool acting like a gay when he dances)

"Manco fossi gay 😂😂😂" (What am I, gay? 😂😂😂)

"Sei così effemminato che intorno a te ci sono più finocchi che in un orto" (You are so effeminate that around you there are more pansies than in a garden)

However, a frame that covers this concept is missing in FrameNet.

c) Concepts that are new to FrameNet, but that are not specifically linked to the nature of the text nor to abusive language or cyberbullying. For example in FrameNet are missing frames related with "sports" and similar activities:

"Anche tu fai calcio" (You play football as well)

"Sì e tu vai a giocare a rugby" (Yes, and you go play rugby)

"Lui non fa danza classica" (He does not do ballet)

d) Concepts that are not new to FrameNet corresponding to holes in the FrameNet hierarchy. For example FrameNet has a frame for `Silencing`, a frame for `Becoming_silent` but it does not have a frame for `Being_silent`.

Frame Elements layer: We found that not only frames were missing but that it was also possible to find missing FEs.

For example it appears to be missing the FE `Reason` for the frame `Statement`, useful for annotating sentences such as:

"Lo diciamo per il tuo bene" (We say that for your own sake)

here *"Per il tuo bene" (For your own sake)* expresses the motivation for which the speaker makes his statement and could be labeled as `Reason`.

Another example can be the frame `Ingestion` for which a FE `Quantity`, for annotating the quantity of the ingestibles eaten, appears to be missing. For example, in the sentence:

> *"Non mangiare tanto o diventi ancora più obeso" (Do not eat a lot or you will get even fatter)*

the FE label `Quantity` would be perfectly fitting for annotating the adverb *"tanto" (a lot)*.

Frame-Evoking Elements layer: Problems linked to the fact that in the sentences we tagged we find that not only words or multiword expressions (MWEs) evoke frames but that also other elements. In particular we found that frames can be evoked also by:

a) **Constructions:** For example in the sentences *"Di sicuro un cane è più bravo di lui"(A dog is better than him for sure)* or *"Noi siamo più forti di te"(We are stronger than you)* the frame `Surpassing` is evoked by the construction *"essere più X di Y"(To be Xer than Y)*" rather than by a word or a multiword expression.

b) **Emoji:** For example, in the sentence

> *"Ma tu sei già una 💩" (But you are already a 💩)*

the *"Pile of Poo"* emoji evokes the frame `Desirability`.

Aside from these three problematic layers, we found that for a considerable amount of messages it was not possible to add any frame annotation because of problems of different nature. More specifically we found that:

a) Some messages are only made of punctuation marks, mostly ellipsis, exclamation points and question marks.

b) Some messages are made of interjections or discourse markers and it is, thus, not possible to identify any frame evoking element:

> *"Ooooooooooooooooo 😂"*

> *"Ahahahahahahahahahahahah"*

c) In some other cases there are sentences that have been split into two or more messages. In these cases it is often possible to find messages in which no frame is evoked, but that constitute a FE of a frame evoked in the bigger sentence that has been split.

For example, the sentence:

> *"Ma noi verremmo con i nostri bei cori" (But we would come with our nice chant)*

has been split into two different messages *"Ma noi verremmo" (But we would come)* and *"Con i nostri bei cori" (With our nice chants)*. The first message can be annotated with the frame `Arriving` while the second message could only be annotated as the `Arriving` frame element `Depictive`.

The sentence:

> *"Neanche hai capito che è una citazione di Battiato " (You didn't even understand that this is a quote from Battiato)*

have been split into *"Neanche hai capito che è una citazione"(You didn't even understand that it is a quote)* and *"Di Battiato" (From Battiato)*. In the first message, the LU *"capire.v"(understand.v)* evokes the frame `Awareness`, and *"Che è una citazione" (That it is a quote)* instantiates its frame element `Content`, whereas the second message can only be considered as a part of it.

d) Some messages contain only affermative and negative expressions, i.e *"Yes"* or *"No"*.

e) Other messages only repeat a word or a group of words of the previous message or anticipate one word or a group of words that will be part of the subsequent message:

> *"Tu", "Tu che sei un maschio" (You, You that are a boy)*

f) Finally there are messages that only aim to correct a word or a letter previously misspelled:

> *"Ai scritto", "*Hai" (You wrote)*

> *"Bravo Bul", "*Bullo" (Good bully)*

A field that is particularly relevant is the semantic field of emotions. We found that FrameNet frames referring to this field have sometimes fuzzy boundaries and that it is sometimes hard to choose a frame over another. Moreover there are also some frames that seem to be missing: for example in FrameNet there is no frame that covers the concept of *"Expressing emotions"*

evoked by LUs such as *"weep.v"* or *"cry.v"* or *"laugh.v"*. Indeed, the first is completely missing in FN, the second is present as evoking `Make_noise`, `Communication_noise` and `Vocalization`, the third in present only as evoking `Make_noise`.

5 Annotations comparison

In order to highlight significant relations between frames and cyberbullying, we compared the frame annotation with the already existing annotation regarding the type of cyberbullying expression (see Section 3). In particular we computed their correlation using the weighted mutual information. This kind of evaluation can be useful, for instance, to predict cyberbullyng conversations using tools that automatically extract semantic information with respect to frames, such as SEMAFOR (Das et al., 2014).

The results, reported in Table 2, show some interesting outcomes. Most of them are in line with what we could have expected, but some others instead reflect the limitations of FrameNet in the annotation of this kind of interactions. For example we can see that *"General_insult"* is related with frames such as `Mental_property` or `Desirability`, this well matches with the intuitions that those frames capture respectively expressions which denigrates the interlocutor by referring to his/her lower intelligence, e.g. *"Idiota"* or *"Stupida"* (*"Idiot"*, *"Stupid"*), or to his/her scarce desirability, e.g. *"Sfigato"* (*"Loser/Lame"*). The same can be said for the pairs *"Treat_or_Blackmail"* - `Cause_harm` and *"Insult-BodyShame"* - `Aesthetics`, where the connection between the frame and the cyberbullying type appears to be straightforward. Nevertheless there are also pairs if which the connection is hard to understand. For example *"Encouragment to the Harasser"* shows a strong relation with the frame `Correctness`. This is due, once again, to the limitations of FrameNet that lacks of some frames, in this particular case it lacks of a frame for the expressions that indicate a reinforcement of what one of the interlocutors just said such as *"Esatto"* (*"Exactly"*) or *"Hai ragione"* (*"You are right"*) that are now listed under the frame `Correctness`.

Bullying annotation	Frame	wMI
Curse_or_Exclusion	Silencing	0.0672
General_Insult	Desirability	0.0304
General_Insult	Mental_property	0.0227
Encourage_Harasser	Correctness	0.0177
Curse_or_Exclusion	Desiring	0.0135
Threat_or_Blackmail	Cause_harm	0.0127
Discrimination-Sexism	Suitability	0.0083
Curse_or_Exclusion	Required_event	0.0080
General_Insult	Silencing	0.0065
Insult-BodyShame	Aesthetics	0.0046

Table 2: Correlation between the new annotations of frames and the previous ones of cyberbullying types using weighted mutual information (wMI).

6 The annotation interface

The annotation on FrameNet has been performed using a tool called Framy, developed at Fondazione Bruno Kessler and freely available on Github[5] under the Apache 2.0 license. It is written in php and needs a MySQL database to work.

The application is optimized for frame semantics annotation, and can be configured to work with every version of FrameNet. After loading the already tokenized text data using the included scripts, a human annotator can select both the lexical unit that evokes the frame and the frame elements relative to the selected words.

7 Conclusions and Future Work

In this paper, we present and release BullyFrame, an Italian resource consisting in a set of WhatsApp chats with full-text FrameNet annotations. The data, freely accessible on GitHub, increases the availability of resources in Italian. We also discuss how FrameNet lacks certain frames, as it cannot cover some expressions used mainly in the social media language. Finally, we describe Framy, a free tool that supports the manual annotation of texts w.r.t. FrameNet.

In the future, we want to extend this dataset by including other text resources, and extend FrameNet coverage for the social media domain, to deal with informal expressions and emojis.

Acknowledgments

This work has been supported by the European Commission project Hatemeter (REC-DISC-AG-

[5] `https://github.com/dhfbk/framy`

2016, action grants 2016: European citizenship rights, anti-discrimination, preventing and combating intolerance).

References

Roberto Basili, Diego De Cao, Alessandro Lenci, Alessandro Moschitti, and Giulia Venturi. 2012. EvalIta 2011: The Frame Labeling over Italian Texts Task. In *International Workshop on Evaluation of Natural Language and Speech Tool for Italian*, pages 195–204. Springer.

Roberto Basili, Silvia Brambilla, Danilo Croce, and Fabio Tamburini. 2017. Developing a large scale FrameNet for Italian: the IFrameNet experience. *CLiC-it 2017 11-12 December 2017, Rome*, page 59.

Marie Candito, Pascal Amsili, Lucie Barque, Farah Benamara Zitoune, Gaël De Chalendar, Marianne Djemaa, Pauline Haas, Richard Huyghe, Yvette Yannick Mathieu, Philippe Muller, et al. 2014. Developing a French Framenet: Methodology and first results.

Dipanjan Das, Desai Chen, André FT Martins, Nathan Schneider, and Noah A Smith. 2014. Frame-semantic parsing. *Computational linguistics*, 40(1):9–56.

Fabio Del Vigna, Andrea Cimino, Felice Dell'Orletta, Marinella Petrocchi, and Maurizio Tesconi. 2017. Hate me, hate me not: Hate speech detection on Facebook.

Karthik Dinakar, Roi Reichart, and Henry Lieberman. 2011. Modeling the detection of textual cyberbullying. In *fifth international AAAI conference on weblogs and social media*.

Mai ElSherief, Vivek Kulkarni, Dana Nguyen, William Yang Wang, and Elizabeth Belding. 2018. Hate lingo: A target-based linguistic analysis of hate speech in social media. In *Twelfth International AAAI Conference on Web and Social Media*.

Nadine Farag, Samir Abou El-Seoud, Gerard McKee, and Ghada Hassan. 2019. Bullying hurts: A survey on non-supervised techniques for cyber-bullying detection. In *Proceedings of the 2019 8th International Conference on Software and Information Engineering*, pages 85–90. ACM.

Minne Fekkes, Frans IM Pijpers, A Miranda Fredriks, Ton Vogels, and S Pauline Verloove-Vanhorick. 2006. Do bullied children get ill, or do ill children get bullied? a prospective cohort study on the relationship between bullying and health-related symptoms. *Pediatrics*, 117(5):1568–1574.

David Gerrard, Martin Sykora, and Thomas Jackson. 2017. Social media analytics in museums: extracting expressions of inspiration. *Museum Management and Curatorship*, 32(3):232–250.

Luca Gilardi and C Baker. 2018. Learning to Align across Languages: Toward Multilingual FrameNet. In *Proceedings of the International FrameNet Workshop*, pages 13–22.

Martina Johnson and Alessandro Lenci. 2011. Verbs of visual perception in Italian FrameNet. *Constructions and Frames*, 3(1):9–45.

Soo-Min Kim and Eduard Hovy. 2006. Extracting opinions, opinion holders, and topics expressed in online news media text. In *Proceedings of the Workshop on Sentiment and Subjectivity in Text*, pages 1–8. Association for Computational Linguistics.

Birte Lönneker-Rodman. 2007. Multilinguality and FrameNet. *International Computer Science Institute Technical Report*.

Stefano Menini, Giovanni Moretti, Michele Corazza, Elena Cabrio, Sara Tonelli, and Serena Villata. 2019. A System to Monitor Cyberbullying based on Message Classification and Social Network Analysis. In *Proceedings of the third Workshop on Abusive Language Online*, Florence, Italy.

Kyoko Ohara. 2012. Semantic Annotations in Japanese FrameNet: Comparing Frames in Japanese and English. In *LREC*, pages 1559–1562. Citeseer.

Mona O'Moore and Colin Kirkham. 2001. Self-esteem and its relationship to bullying behaviour. *Aggressive behavior*, 27(4):269–283.

Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational linguistics*, 31(1):71–106.

Behrang QasemiZadeh, Miriam RL Petruck, Regina Stodden, Laura Kallmeyer, and Marie Candito. 2019. SemEval-2019 Task 2: Unsupervised Lexical Frame Induction. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 16–30.

Josef Ruppenhofer, Michael Ellsworth, Myriam Schwarzer-Petruck, Christopher R Johnson, and Jan Scheffczyk. 2006. FrameNet II: Extended theory and practice.

Manuela Sanguinetti, 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)*, Miyazaki, Japan, May. European Languages Resources Association (ELRA).

Rachele Sprugnoli, Stefano Menini, Sara Tonelli, Filippo Oncini, and Enrico Piras. 2018. Creating a WhatsApp Dataset to Study Pre-teen Cyberbullying. In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, pages 51–59.

Carlos Subirats and Hiroaki Sato. 2004. Spanish framenet and framesql. In *4th International Conference on Language Resources and Evaluation. Workshop on Building Lexical Resources from Semantically Annotated Corpora. Lisbon (Portugal)*. Citeseer.

He Tan, Rajaram Kaliyaperumal, and Nirupama Benis. 2011. Building frame-based corpus on the basis of ontological domain knowledge. In *Proceedings of BioNLP 2011 Workshop*, pages 74–82. Association for Computational Linguistics.

Sara Tonelli and Emanuele Pianta. 2009. Three issues in cross-language frame information transfer. In *Proceedings of the International Conference RANLP-2009*, pages 441–448.

Sara Tonelli, Daniele Pighin, Claudio Giuliano, and Emanuele Pianta. 2009. Semi-automatic development of FrameNet for Italian. In *Proceedings of the FrameNet Workshop and Masterclass, Milano, Italy*.

Sara Tonelli. 2010. Semi-automatic techniques for extending the FrameNet lexical database to new languages.

Tiago Torrent, Maria Margarida Salomão, Fernanda Campos, Regina Braga, Ely Matos, Maucha Gamonal, Julia Gonçalves, Bruno Souza, Daniela Gomes, and Simone Peron. 2014. Copa 2014 framenet brasil: a frame-based trilingual electronic dictionary for the football world cup. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: System Demonstrations*, pages 10–14.

Cynthia Van Hee, Gilles Jacobs, Chris Emmery, Bart Desmet, Els Lefever, Ben Verhoeven, Guy De Pauw, Walter Daelemans, and Véronique Hoste. 2018. Automatic detection of cyberbullying in social media text. *PloS one*, 13(10):e0203794.

Andrea Vanzo, Danilo Croce, Roberto Basili, and Daniele Nardi. 2017. Structured Learning for Context-aware Spoken Language Understanding of Robotic Commands. In *Proceedings of the First Workshop on Language Grounding for Robotics*, pages 25–34.

Rui Zhao, Anna Zhou, and Kezhi Mao. 2016. Automatic detection of cyberbullying on social networks based on bullying features. In *Proceedings of the 17th international conference on distributed computing and networking*, page 43. ACM.