Relationships and Sentiment Analysis of Fictional or Real Characters

Paul Diac[1], Cătălina Mărănduc[1,2], and Mihaela Colhon[3]

[1]Faculty of Computer Science, Al. I. Cuza University, Iasi, Romania
[2]Academic Linguistics Institute I. Iordan – Al. Rosetti, Bucharest, Romania
[3]Department of Computer Science, University of Craiova, Craiova, Romania
paul.diac@info.uaic.ro, catalinamaranduc@gmail.com,
mcolhon@gmail.com

Abstract. In a previous work, we developed a tool that automatically extrapolated triggers, i.e. diagnostic words for sentiments and relationships, from a manually annotated corpus, the Romanian version of the novel *"Quo Vadis"* by Henryk Sinkiewicz. The NodeXL program can draw graphs of character relationships, to analyse relationships both in the fictional and the real-world. In this research, we describe how we have refined our tool, which becomes both a detector and a semiautomatic (interactive, assisted) annotator of relationships in any previously morphological annotated real or fictional story. We will also show how we improved and restructured the list of triggers manually annotated in the novel *"Quo Vadis"*. Finally, the tool will annotate the triggers in the Chat corpus, having 2,575 sentences, part of the UAIC Romanian Dependency Treebank, a balanced corpus that contains especially non-standard Romanian language. Finally, we have made graphs to analyse the relations and sentiments of communicators from the Chat corpus.

Keywords: fictional characters · graph of relations · interactive framework · real characters · sentiment analysis · semantic annotation · social-media communication · trigger

1 Introduction

Nowadays, text mining applications have to implement deep and meaningful representation of texts which usually implies discovering the entities described in texts and the relations between them. In 2012 our NLP group started a project called *"Quo Vadis"* dedicated to the semantic relations described in texts. It used the Romanian translation of Henryk Sienkiewicz's novel *"Quo Vadis"*, morphologically annotated previously. The aim of this project was to design a manually annotated corpus, with semantic data, and then, to build an automatic recognizer based on the annotations of the corpus.

The semantic annotations in the Quo Vadis project mark more types of relations mentioned in the text and the entities linked by these relations. The annotations of the relations are defined by the two boundaries of the relations, the type and subtype of the relations, the two arguments (all relations being binary), and the trigger - a word or an expression which signals the relation.

During the Quo Vadis project, the following resources and programs have been created:

- The *"Quo Vadis"* corpus[1], which is publically available on the NLP-Group@UAIC-FII site. It consists of 7,281 sentences manually annotated with semantic data. They consist of marking the textual realization of entities (persons and gods), and also marking four types of semantic relations: *referential*, *affect*, *kinship* and *social* relations. Each type has more subtypes, see [5, 6, 7];
- A Web Interface[2] for visualizing a unique type of annotation: *the co-reference relations between entities*;
- A recognizer for the semantic relations that occur between nested entities [2], i.e. entities that can include one or more other entities. Example: *"the sister of my mother"* is an entity which includes another entity, i.e. *"my mother"*.

Based on the semantic relations described in text, a summary can be automatically generated as illustrated in [4].

In 2014 a trigger detector was built, which memorized the triggers annotated by human annotators. The program generated a list of suggested triggers in certain contexts, in the entire novel. The 5,136 suggestions were validated or invalidated by human annotators. They validated 305 kinship relations, 2,315 social, and 1,219 affect relations, a total of 3,839 that includes also the 757 manually annotated triggers. The percent of validated triggers was 74% from the suggestions [3].

An improved version of this tool will be described in the presented paper. Next, we aim to verify whether the same annotated relationships in the Quo Vadis project (a fictional world) also work in the real world, between chat communicators. The first experiments have been made by selecting only one type of the relationships annotated in the Quo Vadis corpus, namely the detection of affective relationships, which simultaneously leads to a way of sentiment analysis.

We have worked on a Chat corpus, containing 39,391 words and punctuation elements, with the average 15.25 items per sentence. The corpus is morphologically and syntactically annotated, and entirely manually checked. The length of sentences, unusual in chatting, is explained by the high level of education of communicators. A POS-tagger and a syntactic parser were trained on chats and the morphological, lexical, syntactic, semantic, discursive particularities of this type of communication were analysed in [16, 17].

The main contributions of this work can be summarized as follows:

- We propose a more balanced list of semantic relations taking into account the various sentiments and feelings that can be described in various texts;
- We analyse the realization of semantic relations in the non-standardized social media language, a conversation resulting in a less-structured text. The chat style is informal, does not obey any rules;
- We propose a tool which brings together the old trigger detector with a framework for the assisted semi-automatic annotation of the relations in the fictional and non-fictional stories;

[1] http://nlptools.info.uaic.ro/Resources.jsp
[2] http://nlptools.infoiasi.ro

- Using graphs made with the NodeXL program, as in another work [3], the similarity of the relationships structure in the fictional and real worlds were shown, as in [11] [3].

The paper is organized as follows: the first section introduces the actual experiment as a further instalment and sequel of our previous work. The Related Work section summarizes the existing studies conducted in the domain of semantic data. The next section describes the tool that we have designed in order to annotate the chat corpus in a similar way to the annotations in the Quo Vadis corpus. Some statistics given in graphical form are presented in the following section. The article ends with the final conclusions and proposes future research directions.

2 Related Work

Lately, many text analysing applications have implemented semantics in their processing. A continuing growing domain that greatly exploits the semantic data extracted from text is *sentiment analysis* or *opinion mining*; they aim to identify the emotion expressed in texts. The basic goal of sentiment analysis is to identify the overall polarity of a document: positive, negative, or neutral [15].

The semantic oriented approaches usually exploit the relationship between words. State-of-the-Art studies mainly exploit term extraction methods to obtain concepts from texts [18]. In this paper, the semantic relationships between words are identified by a dependency parsing process. Paper [12] explored a new direction in the concept mining field by means of lexicon-syntactic patterns.

Sentiment analysis considers only a special kind of text, namely *affective text,* with the intended aim of analysing the emotional content of texts. The *affective text* analysis has been a popular topic of research in Natural Language Processing (NLP) and Semantic Web communities in recent years [14]. This is an open research problem, relevant for numerous NLP studies such as news stories, public blogs or forums or product reviews [13, 20].

Sentiment analysis tasks are usually designed around an already existing lexicon making use of the WordNet [9], WordNet Affect [19] or ANEW (Affective Norms for English Words) [1]. However, there are still limitations, e.g., WordNet based efforts cannot produce ratings for words not included in WordNet, including multi-word terms and proper nouns [14].

The most important features of the sentiment analysis programs are greatly determined by the quality of the used sentiment lexicons. Other important features included bag-of-word features, hash-tags, handling of negation, word shape and punctuation features, elongated words, etc. [10]. In the lexicon-based approaches, the coverage of the affective lexicon has a great impact on the accuracy scores; consequently, there is a need for methods for automatically updating the lexicon based on the already included elements.

[3] Hansen analyses the social media network and the relationships in Victor Hugo's novel "Les Miserables".

Emotion analysis emerged as a somewhat more specific task than opinion analysis, since it looks at fine-grained types of emotion [10]. Classification of sentences by emotions is done in accordance with some classes of emotions. Here it is worth recalling Ekman's (1992) six classes of emotions: *happiness*, *anger*, *sadness*, *fear*, *disgust*, and *surprise*. These emotion classes are the most frequently used ones, being associated with the facial expressions [8].

One of the problems in opinion mining systems is that sarcastic or ironic statements could easily trick these systems [10]. For an automatic system, it is not so important to distinguish between the two of them as it is to eliminate these cases in order to identify the real sentiments that are hidden by these kind of expressions.

3 Bipolar System of Annotation

The interaction with the tool obliged us to adopt a more symmetrical tagset of annotation. We tried not to get away from the annotation of the Quo Vadis corpus (which also underwent changes along the way). But since all relationships are polarized, it is useless to add prepositions such as ``of'' or ``by''. *rec-love* and *rec-hate* relationships result from the summation of two relationships in which the arguments change their place. We have added the following tags for marking sentiments: LIKE, FEARLESS, OFFEND as the negations of the tags UPSET, FEAR, WORSHIP, because the Quo Vadis system was not symmetrical. The trigger detector has been programmed to suggest us both a sentiment and its negation, in order to select one of them. The problem that remains to be solved is how to formulate rules such that an automatic trigger recognizer can detect irony and sarcasm without being assisted by a human annotator.

Example:

In the sentence *"John messed up everything. He is a very intelligent person."* the trigger detector will suggest the human assistant the annotation of *"intelligent"* as an AFFECT.WORSHIP trigger. The human will choose the button NO and immediately the trigger detector will offer the assistant the opposite trigger: AFFECT.OFFEND. The human will choose the button YES, because it is an ironical statement. The first argument of the trigger is the emitter of the statement (marked at the beginning of the sentence and annotated with the id 0), and the second pole of the relation is John.

If the negation of a feeling does not result in its opposite (a situation in fact), the human annotator will select NO in both alternatives. If a person does not love a particular person, this does not necessarily mean that she/he hates the respective person.

If, by the negation of an AFFECT it results that a person has no AFFECT, we will not annotate anything, because in this project we only deal with the AFFECT annotations.

In fact, only three positive AFFECT relationships and three negative ones are annotated in the Quo Vadis project. By renouncing the targeted variants of the six relationships, we have added 4 other AFFECT relationships, so that the palette of sentiments

becomes more comprehensive and the system can be applied to other stories than the one in the novel Quo Vadis, really dominated by the listed feelings (**Figure 1a**).

QUOVADIS RELATIONSHIPS TAGSET		OUR RELATIONSHIPS TAGSET	
POSITIVE	**NEGATIVE**	**POSITIVE**	**NEGATIVE**
LOVE	HATE	LOVE	HATE
LOVED BY	HATED BY		
REC-LOVE	REC-HATE		
	FEAR	FEARLESS	FEAR
	FEAR BY		
	UPSET	LIKE	UPSET
FRIEND OF		FRIENDLINESS	ENMITY
WORSHIP		WORSHIP	OFFEND
WORSHIP BY			

Figure 1a. Comparison between Quo Vadis and our affect relations tagset.

adoration **LOVE** happiness	anger **HATE** contempt
naughtiness **FEARLESS** heroism	horror terror **FEAR** worry
amazement **LIKE** joy cheerfulness	stress pain **UPSET** desolation
solidarity **FRIENDLINESS** care devotion	envy **ENMITY** malice
admiration **WORSHIP** enthusiasm	humiliation **OFFEND** blasphemy

Figure 1b. The real AFFECTs and the grid of our tags.

Another preliminary statement is that using this tagset, a limited number of feelings can be annotated. As in reality their range is very wide, each of the ten feelings is understood here in a very broad sense, so that it can include a multitude of feelings. In fact, we have a grid placed over a continuous of real AFFECTs, and when we make the annotation, we choose the tag that comes closest to it. Our tags form a grid, and we hope every real AFFECT can be placed in the perimeter of one of our tags. (**Figure 1.b**).

Therefore, the aim of the project is to annotate AFFECTs that are considered generic, positive or negative, and to see their proportion and direction in a certain group of real or fictitious characters. We can also make a positive and negative grading as follows:

Positive: worship > love > like > fearless

Negative: offend > hate > upset > fear

4 Trigger Detector and Assisted Annotator

The tool used in 2014 has been transformed in a multi-functional one. It has a directory called *"configurations"* that contains multiple trigger lists with their type and subtype (Social, Kinship, Affect). In another folder, called *"resources"*, we introduce the XML documents that will be annotated. In this way, linguists can use the interface

without requiring programmer assistance in order to annotate any narrative text, short or long, having the following characteristics: novel, chat, bible, blogs, comments, etc., provided they have a previous basic annotation i.e., they are segmented in sentences, and each word has an id, lemma, and the morphological analysis; and they adhere to a specific format.

To make the program run, we have to choose a configuration, and then to choose one or more XML files in which the tool will search for all the triggers listed in that configuration. The detector will display the proposals in the order they appear textually, and will ask questions to which the answer is YES or NO, e.g.:

- "Is trigger (**râde**, AFFECT.LIKE) valid for sentence 10_chat2_1500U?"
  YES / NO

The next step of the assisted annotation, is the proposal of the opposite trigger, if the option NO has been selected. In ironic or sarcastic uses, the opposite is suggested and can be validated:

- "Is trigger (**râde**, AFFECT.HATE) valid for sentence 10_chat2_1500U?"
  YES / NO

The first variant of the trigger detector was searching in the text form (the list including MWEs, negations, reflexive pronouns, etc. that have been manually annotated in Quo Vadis), which resulted in a high accuracy of detection because some polysemous words are triggers if they have certain neighbourhoods. Therefore, many triggers can escape our detection if they have small formal differences from those found in Quo Vadis. Cases will be numerous, especially if we annotate non-standard texts, such as social media or old Romanian. So, we decided that the new variant of the trigger detector should make searches for trigger **lemmas** resulting in less accuracy, that will lead to a large number of rejected proposals.

By programming the interface for trigger suggestions with both a tag and with its opposite, we eliminated the need of including the verb + negation in the trigger list. We have created only two types of conditions:

1. if the next word is *xxx*
2. if one of three words above has lemma "sine".

The first condition led us to detect the words which are triggers only if followed by a certain preposition.

Example:

- The word **ține** is a trigger for AFFECT.LOVE if it is followed by the word **la** as in *"**ține la** cineva"* (in English, *love somebody*).
- As opposite, *"**ține** în mână"* (in English, *holds in his hand*) is not a trigger for AFFECT.LOVE.

The second condition helped as detect verbs which are triggers only if preceded by a reflexive pronoun.

If we chose YES option at one of the two first steps, then in the third step the interface displays drop-down lists of the words in the sentence, from which the human annotator choses: the first and the second poles of the trigger. The relation is a vector from the first to second. Once the human annotator saves his choice, the program will

add the annotation in the XML, after the line of the word validated as a trigger, having the following form, similar to those in the Quo Vadis project:

- <relation from="15" to="17" trigger="AFFECT.love"/>

Were *"from=15"* indicates the id of the first relation argument or the source, and *"to=17"* represents the id of the second argument or the target.

The *"trigger= AFFECT.love"* gives the information about the type and subtype of the validated trigger.

By applying similar annotations in multiple resources, we will create a training corpus for the future automatic trigger and argument recognizer.

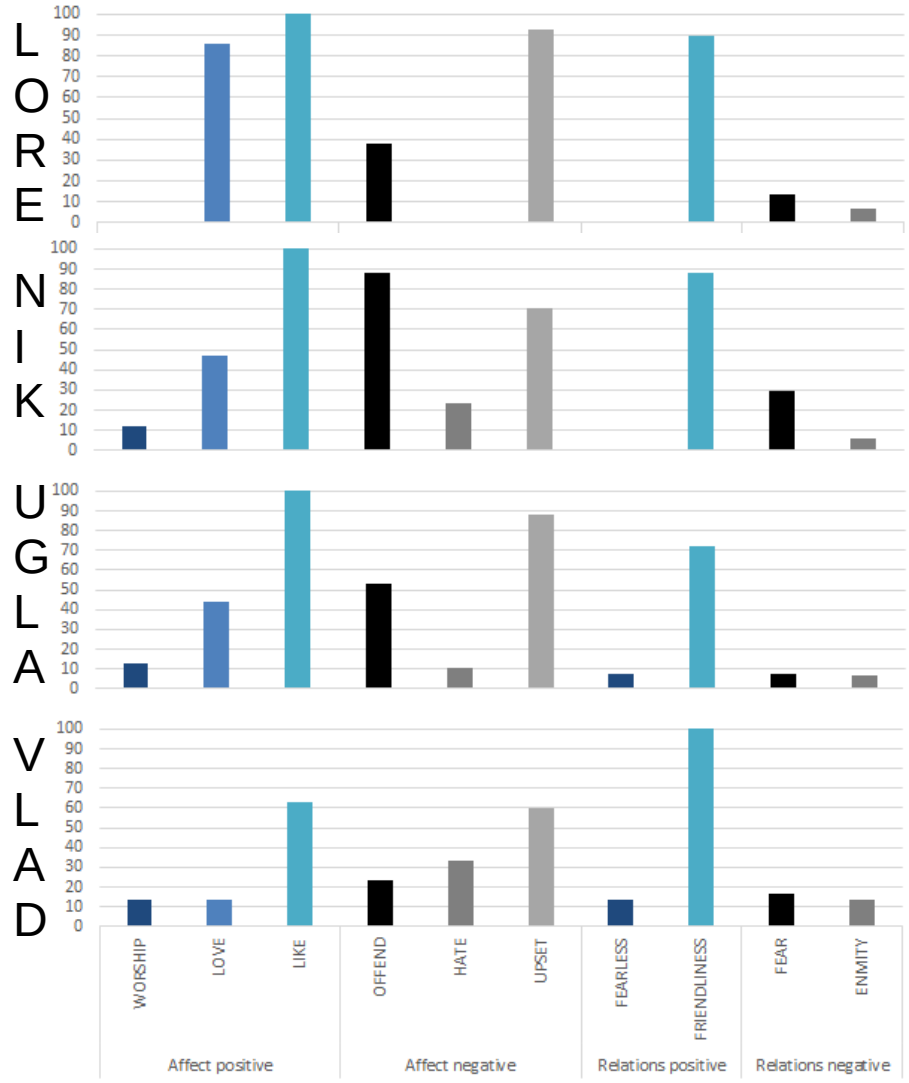

Figure 2. The frequency of the 10 analyzed affect relations for the 4 communicators.

5 Semantic Relation Graphs

Once we have annotated chats and obtained a consistent training corpus for future triggers and argument recognizer, we have also extracted data from the annotated chat corpus in order to interpret them.

The data extraction consists of the automatic transformation of ids into character names as well as the unification of multiple names for a single entity. The difficulties arise from the fact that the chat corpus is not yet annotated with the names of entities and with their co-references, such as Quo Vadis. The entities that appear as first and second person pronouns must equate with the name of communicators: the phrase issuer with the first person pronouns, and the recipient with the second person pronouns. The communicator's names appear in the id of the sentence, and are annotated with id 0, 1, 2, 3.

We computed the frequency of the 10 feelings for each of the four communicators in our chat corpus: Ugla, 62, researcher, Nik, 61, writer, Lore, 28, psychologist and Vlad, 28, economist.

For Ugla, the predominant feelings are *like* and *upset*, but she also has *fearless* and *offend* relations, she communicates without any reticence, and without bad opinions about anyone. *Fearless* and *offend* also characterize Nik. Lore *hate*s nobody, the greatest number of *hate* relations are directed towards Vlad. He has the record number of *friendliness* relations, and of *upset* sentiments, but surprisingly few AFFECT of the subtype *love*, that are dominant to Lore. (**Figure 2**).

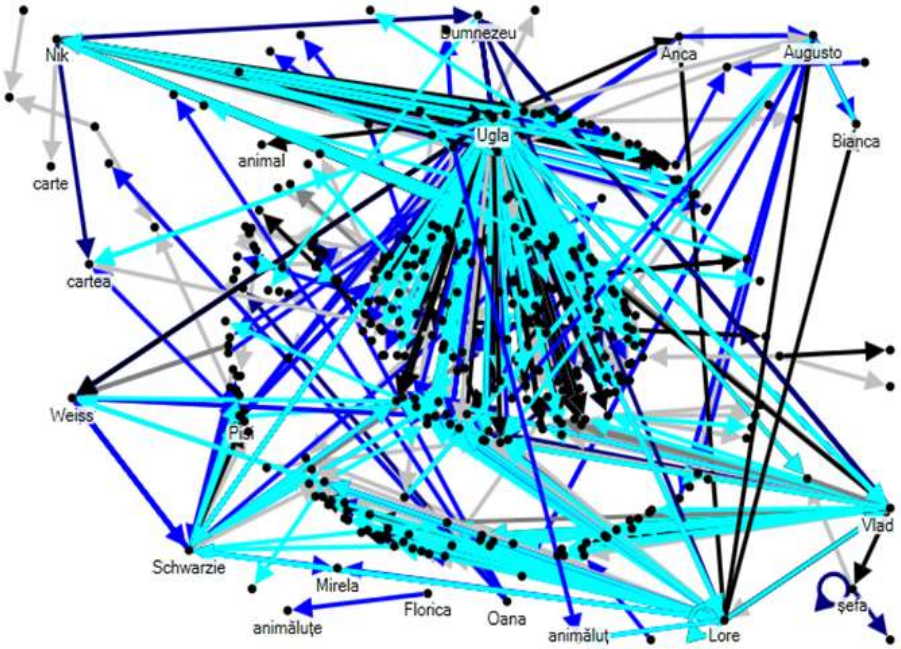

Figure 3: Graph of AFFECT relations

In order to obtain more suggestive graphs, we decided that *worship*, *love*, *like* (and their negations) are a gradient of positive or negative sentiments, while *friendliness* and *fearless* with their negatives are names of relationships between characters.

The tool annotates the triggered relation in the XML, and saves the data in an XLX table. On the first column, the first selected argument, the source of affect relationship (from…), on the second column, the second argument selected by the user, the target of the relationship (to…): the trigger, their type and subtype are saved on the 3-5 columns and the id of the sentence is saved on the last column. The two poles of a relation can be selected by the user from the words that have the morphological category annotated as noun, pronoun, numeral that does not determine a noun, or possessive adjective.

In this XLX, the trigger lemmas are replaced by the type and subtype, and the argument ids are replaced with the names of communicators or of other persons and objects they refer to. The XLX is introduced in the NodeXL program in order to draw out graphs of character relationships [11].

In the chart of Figure 3, illustrating more relations, we see a conglomerate of AFFECTs such as *like* and *offend* around Ugla, and a large fascicle of common AFFECTs of subtype *like* with Nik. At the bottom of the chart there are two reciprocal *love* relationships between two couples, Schwarzie-Weiss (black and white cats) and Lore-Vlad. On the other extreme of the chart, Augusto is characterized by *love* relationships in several directions. Oana has a friendly relationship with Ugla and Vlad, a relation that Lore rejects. All characters have AFEECTs *like* or *love* directed to animals. AFFECTs of subtype *hate* and *offend* are directed to hierarchical superiors.

It can be observed that all the characters have multiple *friendship* relationships, but they are particular orientations, and are unified only by Ugla. Most *fear* relationships start at Schwarzie and Vlad. Reflexive relationships, from a character to him or herself appear in both graphs as circles.

The common affective relationships between Ugla and Lore have in the center Schwarzie, the black cat. A strong positive mutual relationship, framed as a bow with arrows at both ends is established between Ugla and Vlad, the two characters who quarrel throughout their dialogue. (See **Figure 3**).

If we analyze the characters in the Quo Vadis novel as compared to real world communicators, we see a greater variance and a chaotic orientation of directions in expressed relationships at the latter, which are not controlled by an omniscient author.

The sentiments are various and are not focused; in the chat they can result of a concrete real event. The contradictories sentiments and their ironical expression are frequent. The characters have their own circle of relations and feelings, without connection of the feeling of the other characters.

In Quo Vadis, the target of the affective relationships of the characters can be: gods, Vinicius, Ligia, Petronius, Nero, senate, death, Christ, people, family, Seneca, Venus, Acteea, etc; as resulting from our previous research in [3].

The targets of emotional relationships of real communicators are more varied and there are not only characters, but also objects or abstractions among them: plant, installer, stove, trains, conference, editions, wedding, retirement, managers, sponsors, raven, book, foods, academy, ox, pictures, infarction, math, physics, recreation, etc. See table **1**.

Table 1. The table which generates the graph in the figure 3 and the chart in the figure 2 (small excerpt).

Argument 1	Argument 2	Trigger	Type	Subtype	Sent. id
cat	doctors	run	AFFECT	FEAR	929U
she(cat)	in darkness	bravery	AFFECT	FEARLESS	932U
me(U)	answer	not receive	AFFECT	UPSET	933U
they(dogs)	meat	seems	AFFECT	LIKE	934U
give(L)	rice with meat	delights	AFFECT	LOVE	934U
me(V)	pictures	not want	AFFECT	HATE	1062V
managers	me(V)	alert	AFFECT	FEARLESS	1062V
U	that	devil	AFFECT	OFFEND	1063U
me(V)	pictures	scared	AFFECT	FEARLESS	1067V
me(U)	well	fell	AFFECT	UPSET	1069U
me(U)	you(V)	not miss	AFFECT	HATE	1069U
they(fellows)	me(V)	condemn	AFFECT	ENMITY	1070V
managers	me(V)	say	AFFECT	FEAR	1072V
managers	pictures	indecent	AFFECT	OFFEND	1073U
me(V)	association	personal	AFFECT	FRIENDLINESS	1076V
me(V)	personal	not allowed	AFFECT	FEAR	1077V

6 Conclusions and Future Work

As an interpretation of our analysis, we observed that the real world is characterized by the variety and divergence of relations. Eventually, we could possibly use this observation to distinguish between a real story and a fictional one. The tool could also be used to extract the orientation of the feelings of real characters towards a particular political line or the preference for the consumption of certain brands.

The tool, applied here on Romanian texts, is language-independent, because in the *resources* and *configurations* folders, documents can be added in any language.

In future, the corpus presented in this paper would be diversified by introducing other communicators and then other types of texts, like social media (twitter messages, product reviews, or political comments). The corpus will also be annotated with entities and coreferences, increasing the precision and recall of suggestions. We also intend to annotate with social and kinship relations, entities, and coreferences, the Mateiu Caragiale's novel *"The Old Courtyard Princess"* published in 1915 that is already morphologically and syntactically annotated entirely supervised. By adding it to Quo Vadis and the Chat corpora, we will form a large training corpus for the trigger and arguments recognizer.

As a first step, the trigger list extracted from our corpora will be very flexible and we the framework will permit us to add triggers. The number of occurrences will be

kept in the memory and the triggers which will not have utility will be eliminated. In this way, by retaining only the productive ones, we accuracy of the tool will be increase.

The final plan is to fully automate the learning system, trained on partially supervised annotations of these corpora. This will enable annotating with considerable less effort.

References

1. Bradley, Margaret M., and Peter J. Lang. 1999. Affective norms for English words (anew): instruction manual and affective ratings. In Technical report c-1, University of Florida. The Center for Research in Psychophysiology.
2. Colhon, Mihaela, Dan Cristea, and Daniela Gîfu. 2016. Discovering Semantic Relations within Nominals. In Diana Trandabăț and Daniela Gîfu, editors, Linguistic Linked Open Data: 12th EUROLAN 2015 Summer School and RUMOUR 2015 Workshop, Sibiu, Romania, July 13-25, 2015, Revised Selected Papers. Springer International Publishing, pages 85–100. https://doi.org/10.1007/978-3-319-32942-0_6.
3. Colhon, Mihaela, Paul Diac, Cătălina Mărănduc, and Cenel-Augusto Perez. 2014. Quo vadis research areas – text analysis. In Proceedings of the 10th International Conference Linguistic Resources and Tools for Processing the Romanian Language. Alexandru Ioan Cuza University Publishing House, pages 45–56.
4. Colhon, Mihaela, Daniela Gîfu, and Dan Cristea. 2015. The Quo Vadis Story Telling. In Proceedings of the the 11th International Conference Linguistic Resources and Tools for Processing The Romanian Language (ConsILR 2015). Alexandru Ioan Cuza University Publishing House, pages 93–108.
5. Cristea, Dan, Gabriela E. Dima, Oana-Diana Postolache, and Ruslan Mitkov. 2002. Handling complex anaphora resolution cases. In Proceedings of the Discourse Anaphora and Anaphor Resolution Colloquium (DAARC 2002).
6. Cristea, Dan, Daniela Gîfu, Mihaela Colhon, Paul Diac, Anca-Diana Bibiri, Cătălina Mărănduc, and Liviu-Andrei Scutelnicu. 2015. Quo vadis: A corpus of entities and relations. In Núria Gala, Reinhard Rapp, and Gemma Bel-Enguix, eds, Language Production, Cognition, and the Lexicon. Springer International Publishing, pp. 505–543.
7. Cristea, Dan, and Eugen Ignat. 2013. Linking book characters. toward a corpus encoding relations between entities. In Proceedings of the 7th International Conference on Speech Technology and Human-Computer Dialogue (SpeD 2013). pages 1–8.
8. Ekman, Paul. 1992. An argument for basic emotions. Cognition and Emotion 6
9. Esuli, Andrea, and Fabrizio Sebastiani. 2006. Sentiwordnet: A publicly available lexical resource for opinion mining. In Proceedings of the 5th Conference on Language Resources and Evaluation (LREC'06). pp. 417–422.
10. Farzindar, Atefeh, and Diana Inkpen. 2015. Natural Language Processing for Social Media. Morgan & Claypool Publishers.
11. Hansen, Derek L., Ben Shneiderman, and Marc A. Smith. 2011. Analyzing Social Media Networks with NodeXL. Morgan Kaufmann.
12. Hearst, Marti A. 1992. Automatic acquisition of hyponyms from large text corpora. – Association for Computational Linguistics, volume 2, pp. 539–545.

13. Hu, Minqing, and Bing Liu. 2004. Mining and summarizing customer reviews. In Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining. Association for Computational Linguistics, pp. 168–77.
14. Malandrakis, N., A. Potamianos, E. Iosif, and S. Narayanan. 2007. Distributional semantic models for affective text analysis. IEEE Transactions on Audio, Speech, and Language Processing.
15. Pang, Bo, and Lillian Lee. 2008. Opinion mining and sentiment analysis. Foundations and trends in information retrieval.
16. Perez, Cenel-Augusto, Cătălina Mărănduc, and Radu Simionescu. 2016a. Including social media – a very dynamic style – in the corpora for processing Romanian language. Linguistic Linked Open Data: 12th EUROLAN 2015 Summer School and RUMOUR 2015 Workshop, Sibiu, Romania, July 13-25, 2015, Revised Selected Papers pages 139–153.
17. Perez, Cenel-Augusto, Cătălina Mărănduc, and Radu Simionescu. 2016b. Social media – processing romanian chats and discourse analysis. Computación y Sistemas 20(3):404–414. https://doi.org/10.13053/CyS-20-3-2453.
18. Poria, Soujanya, Basant Agarwal, Alexander Gelbukh, Amir Hussain, and Newton Howard. 2014. Dependency-based semantic parsing for concept level text analysis. Computational Linguistics and Intelligent Text Processing. CICLing 2014.
19. Valitutti, R. 2004. Wordnet-affect: an affective extension of wordnet. In Proceedings of the 4th International Conference on Language Resources and Evaluation. pp 1083–1086.
20. Wiebe, Janyce, and Rada Mihalcea. 2006. Word sense and subjectivity. In Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics. Association for Computational Linguistics, pp. 1065–1072.