Opinion Mining on Educational Resources at the Open University of Catalonia

Isabel Guitart, Jordi Conesa, Luis Villarejo, Àgata Lapedriza, David Masip, Antoni Pérez, Elena Planas
Open University of Catalonia, IT, Multimedia and Telecommunications Department
Rambla del Poblenou, 156,
E-08018 Barcelona, Spain
{iguitarth, jconesac, lvillarejo, alapedriza, dmasipr, aperezn, eplanash}@uoc.edu

Abstract— **In order to make improvements to teaching, it is vital to know what students think of the way they are taught. With that purpose in mind, exhaustively analyzing the forums associated with the subjects taught at the Universitat Oberta de Cataluya (UOC) would be extremely helpful, as the university's students often post comments on their learning experiences in them. Exploiting the content of such forums is not a simple undertaking. The volume of data involved is very large, and performing the task manually would require a great deal of effort from lecturers. As a first step to solve this problem, we propose a tool to automatically analyze the posts in forums of communities of UOC students and teachers, with a view to systematically mining the opinions they contain. This article defines the architecture of such tool and explains how lexical-semantic and language technology resources can be used to that end. For pilot testing purposes, the tool has been used to identify students' opinions on the UOC's Business Intelligence master's degree course during the last two years. The paper discusses the results of such test. The contribution of this paper is twofold. Firstly, it demonstrates the feasibility of using natural language parsing techniques to help teachers to make decisions. Secondly, it introduces a simple tool that can be refined and adapted to a virtual environment for the purpose in question.**

Keywords— Opinion mining; forums; students; language technology

I. INTRODUCTION

Constantly improving the quality of teaching is a general, implicit goal of every educational institution. To gauge the quality of an institution's teaching and identify scope for improvement, it is essential to know how satisfied students are with the way they are taught.

While satisfaction levels can be ascertained by classic means, such as interviews or surveys, in e-learning environments it is possible to use Natural Language Processing (NLP) techniques to the same end. That option exists because, in such environments, students communicate with one another and with lecturers by writing in a specific facility called a virtual classroom. Studying the messages posted in a virtual classroom can give an idea of the extent of student satisfaction with the teaching that takes place there.

The basic aim of the work described in this article consists of using NLP techniques to help lecturers gauge students' satisfaction with study programmes they are following or have followed. We have thus created a tool for identifying text containing opinions, objective judgments, appraisals, emotions, beliefs or speculations within the messages generated as part of teaching activity in an e-learning environment. For pilot testing purposes, the tool has been used to identify students' opinions on the UOC's Business Intelligence master's degree course in last two years.

Our work is significant in two ways. Firstly, it demonstrates the feasibility of using NLP techniques to help lecturers make decisions to improve the learning process. Secondly, it introduces a straightforward tool that can be refined and adapted to a virtual environment for the purpose in question.

The paper is structured as follows: next section, entitled "Project context", presents the academic environment where the work has been carried out. The following section explains some of the previous work undertaken in the field of opinion mining and introduces the lexico-semantic resources used in our project. That leads into a section entitled "A NLP proposal for supporting teaching decisions", which describes the computational approach developed to find out what students think about how they have learned, an approach based on the expansion of an affective dictionary. We subsequently explain the structure results of our pilot test, before ending by setting out our conclusions and future lines of work.

II. PROJECT CONTEXT

The Universitat Oberta de Catalunya (UOC) is the main virtual university in Spain, with over 45,000 students. The UOC's Virtual Campus has been the medium through which students and lecturers have communicated ever since the university was founded in 1994. The Virtual Campus has many functionalities, notably including virtual classrooms, where the university's subjects are taught. Each classroom has a communication facility (enabling lecturers to guide learning activities and students to ask questions) and other features, such as an area for learning resources, another for submitting work, etc. The communication facility uses forums and bulletin boards to record the messages that classroom users (lecturers and students) exchange with one another.

It is in a classroom's communication facility that the use of NLP techniques to help lecturers in their day-to-day activity is particularly relevant. In particular, NLP techniques could be used to identify underlying opinions in the messages posted in a classroom's communication facility. We believe that doing so in classrooms will be problematic to our purpose, which is to create a system that helps teachers to take more and better learning decisions. Using NLP techniques to classrooms will help to detect particular problems within the teaching that should be solved in a short time, such as some web resource not accessible, errors in any solution, problems of understanding of one topic, etc. However, if the focus is to detect possible improvements that can be made in mid/long term then the classrooms alone cannot provide all the information needed. The main reasons are that classroom is only being used by a fairly small number of students (80 at most in computer science degrees), the teaching activity carried out in a classroom has a very limited duration (a single semester), and most of the messages posted in a classroom deal with studying the corresponding subject rather than opinions on teaching. The factors in question led us to the idea of analyzing opinions in tutorial classrooms, which offer greater duration and a broader range of message content.

For some time, the UOC offers a personalized tutoring service to each student. Each student has a counselor who guides the student in his/her academic career. In the context of the students of postgraduate courses, counselors perform their work within a tutorial classroom, an environment similar to a classroom but which is used for longer (one year for postgraduate courses and two years for master's degrees). Our proposal has been developed in postgraduate tutorial classrooms, an environment we considered to be ideal for our pilot testing because the students use it to discuss aspects of the study programmes, job opportunities, scope for improvement and operational queries, to ask about deeper information in any of the taught topic, as well as to express opinions in other ways, such as in the form of complaints or messages of thanks. Specifically, our project has applied opinion mining techniques in the tutorial classrooms of the UOC's Business Intelligence master's degree course of the last two years.

III. COMPUTATIONAL BACKGROUND

Automatic opinion mining represents a mean of fulfilling our project's goals. It is a computational discipline within the field of NLP, which consists of detecting fragments of text in which an opinion on a given matter is expressed. There are various methodological approaches to carry out such detection, and the most common revolve around information recovery techniques, text classification techniques or dictionary-based techniques. Once opinions have been mined, they are classified. In the project described in this article, we have focused on detecting opinions expressed in tutorial forums using an affective dictionary created by scratch and expanded on the basis of semantic relationships between words, such as synonymy or superordination (generalization).

Automatically mining the opinions expressed in a text is a complicated NLP task, one for which a range of strategies rooted in different approaches have been used. One type of strategy is based on information recovery techniques as in [1] and [2], which firstly identify a text's polarity and then its affective content. Such techniques are used to recover texts discerningly, based on their polarity. A second strategy type uses supervised learning and classification techniques, such as support vector machines [3] or latent semantic analysis [4], to develop statistical models for classifying texts according to emotions [5]. The drawback of supervised learning is that relatively large quantities of tagged samples are required for model development. In our context, a supervised technique would have required texts whose sentences containing emotions had been manually tagged. A third type of strategy is based on the use of an affective dictionary, which contains words with a significant element of emotion in the language under analysis. These words may act as *triggers* for expressions of emotion. The recovery of expressions of emotion can be enhanced using lexico-semantic resources. In our case, we have used two language resources, WordNet [6] and FreeLing [7], to extend the coverage of the triggers in our affective dictionary.

A. WordNet

WordNet [6] is a lexico-semantic database organised hierarchically on the basis of synsets. A synset is a grouping of synonyms linked to other synsets through semantic relationships (synonymy, superordination, hyponymy, the "is-a" relationship, etc.). The dictionary defines a superordinate as a word that has a generic meaning in relation to one or more words with a specific meaning. For example, the word *seat* is the superordinate of *chair*, *armchair* and *stool*. We have used such semantic relationships to extend the range of triggers in our affective dictionary. A hyponym, meanwhile, is the opposite of a superordinate, i.e. a word with a more specific meaning. The word *armchair* is a hyponym of *seat*, for example. The "is-a" relationship links specific instances with abstractions that represent the common characteristics of a grouping of instances. On the basis of the "is-a" relationship, for example, it can be said that *Obama* is a person, where *Obama* is an instance and *person* is an abstraction representing the common attributes of people.

B. FreeLing

FreeLing [7] is a set of language tools that aid morphological and syntactic analyses of written text. FreeLing's form dictionary generates all the possible forms of a lemma, and we have used it to enable our tool to detect not only the lemmas in our affective dictionary but also every form of those lemmas appearing in tutorial classroom messages. For example, the forms *siento, sientes, siente, sentimos, sentís, sienten, sentid, sentida, sentidas, sentimos, sentiremos, sentirá, sentirán, sentirás, sentiré, sentiréis, sentiré,* etc., are generated on the basis of the entry *sentir#v* (denoting the Spanish verb meaning "to feel") in our affective dictionary.

IV. A NLP PROPOSAL FOR SUPPORTING TEACHING DECISIONS

As aforesaid, the aim is to provide useful information to lecturers in order to reduce the time they spend analysing students' contributions to tutorial forums at the end of every semester. The proposed system uses opinion mining techniques to automatically provide to lecturers a graphical representation of the opinions of students during last semester in order to help its analysis. In our system, that graphical result is provided by a tag cloud that shows the opinions expressed by students using different size and distribution of tags in order to represent the relevance and the number of apparitions of each opinion. Clicking on any of the tags, the teacher can navigate through the message fragments containing the corresponding opinion.

A lack of appropriate Spanish and Catalan sample texts with manually tagged expressions of emotion ruled out the development of models based on supervised learning. We thus opted to use an affective dictionary for opinion mining purposes. Since we were unable to find a Spanish affective dictionary that met our needs, we had to create a new one. Next subsection describes the architecture of the proposed system, followed by the description of how the affective dictionary was created. Thereafter, the next section explains how the system works and the different ways of augmenting the affective dictionary and the pros and cons of all of them.

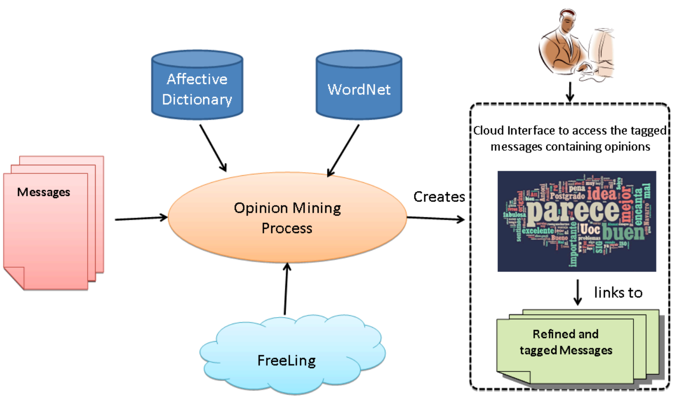

Figure 1. Opinion Mining System Architecture.

A. *Opinion Mining System Architecture*

The proposed system uses a set of language resources (WordNet, FreeLing and the affective dictionary we have created) to identify the opinions expressed in a volume of messages written by students, as shown in Figure 1. The affective dictionary consists of a list of lexical triggers manually compiled for our project. The triggers in question are pairings with the format *lemma#grammatical_category* (English examples would include *feel#verb, better#adjective* and *opinion#noun*) and are characterised by being commonly used to express an opinion in at least one of their semantic senses. WordNet provides the information on synonymy and superordination necessary to extend the affective dictionary's coverage (i.e. it increases the number of expressions of opinion the system is able to detect). Lastly, FreeLing's form dictionary is used to ensure the system detects not only the lemmas in the affective dictionary but also every form of those lemmas appearing in messages. For example, as aforesaid, different forms (*siento*, *sientes*, etc.) are generated on the basis of the entry *sentir#v* in the affective dictionary. Therefore, including just a lemma in the dictionary makes it possible to automatically generate all the forms associated with it for the purpose of detecting opinions in texts.

B. *Affective Dictionary Creation*

For the purposes of our project, we have manually compiled a dictionary of lexical triggers of expressions of emotion, i.e. a list of lemmas based on verbs, nouns and adjectives commonly used when expressing opinions, such as *creo* ("I believe" in Spanish) or *mejor* ("better"). The resultant dictionary contains 75 terms.

Precision and coverage are factors that must be taken into consideration in relation to manually compiling such a resource. In the case of the former, our method offers a very high level of precision, as the list only features words with a substantial element of opinion. Nonetheless, lexical ambiguity could result in the recovery of instances of the words in question being used in ways other than to express opinions. In that regard, semantic disambiguation techniques could be used to select contexts in which the words are used to express opinions [8]. At the time of writing, however, the state of the art in techniques of the kind in question for Spanish and Catalan is such that they are not accurate enough to make their use worthwhile [9]. Where the latter of the aforementioned factors is concerned, an affective dictionary such as ours has the drawback of limited coverage due to it having been compiled manually. One way of broadening the coverage offered by the method we have used whilst maintaining similar levels of precision consists of dictionary expansion through a lexico-semantic hierarchy, such as WordNet.

C. *System Operation*

This section explains how the proposed process works: the system basically searches within messages for the terms featured in the affective dictionary. A message that contains such a term is identified as expressing an opinion and the relevant part of the message, which is considered the triggering words in this work, is subsequently tagged to indicate the opinion triggered by the term. Once all the messages to be analysed have been examined and tagged, the extracted, marked-up data are generated in XML format with an accompanying XSL document for greater visual appeal. As Figure 2 shows, the opinions detected are displayed as a tag cloud in which a term's size is proportionate to the frequency with which it features in the messages analysed, helping the user to quickly determine the general tendency of the opinions expressed in them. When the user clicks on any of the triggers in the tag cloud, an XML file containing the forum message and the context in which the corresponding opinion is expressed is displayed in a web browser, enabling the lecturer to analyse the opinion conveyed in the message in question. To facilitate the identification of the triggers within the texts, those are highlighted.

The system can operate in four modes, which are explained in detail in subsection D, each of which uses a different set of terms to identify opinions: terms in the affective dictionary only, terms in the affective dictionary plus their synonyms, terms in the affective dictionary plus their superordinates, and terms in the affective dictionary plus their synonyms and superordinates.

The system searches for and marks up triggers according to the operational mode selected by the user, and groups them together on the basis of the original trigger.

Figure 2. Extracted opinion triggers arranged in a tag cloud.

If, for example, the option of searching for terms from the affective dictionary and their synonyms were chosen, the system would group together and mark up fragments of texts containing emotions associated with the Spanish verb *sentir* and with its synonyms (the Spanish equivalents of *consider*, *find*, *believe*, etc.) alike under the entry *sentir#v*.

D. Dictionary Expansion Based on Synonymy and Superordination

Our methods of expanding our affective dictionary on the basis of synonymy and superordination are explained below.

<u>Expansion based on synonymy</u>: if the option of extending the affective dictionary's coverage through synonymy is chosen, the system uses each entry's lemma and grammatical category to generate a list of associated synonyms on the basis of WordNet. As mentioned previously, each *lemma#grammatical_category* pairing is present in a group of synsets, and each synset contains a certain number of synonyms (also called variants). In the case of the trigger *sentir#v*, for example, there are 12 synsets in which the Spanish verb *sentir* appears. Using just WordNet's synonyms, the following list of triggers is automatically generated:

sentir#v | *saber#v* | *sentirse#v* | *considerar#v* | *encontrar#v* | *creer#v* | *deplorar#v* | *lamentar#v* | *experimentar#v* | *arrepentirse#v* | *percibir#v* | *notar#v* (etc.)

<u>Expansion based on superordination</u>: in this case, the system extends the list to enhance the method's coverage by detecting the synset to which each term in the dictionary belongs. Each synset has one or more parent (or superordinate) synsets, which also contain their respective variants or synonyms. The list of superordinates is thus also generated on the basis of the list of synsets corresponding to the trigger. For example, in the case of the trigger *sentir#v*, and using just WordNet's superordinates, the following list of triggers is automatically generated:

considerar#v | *creer#v* | *pensar#v* | *juzgar#v* | *opinar#v* | *concluir#v* | *lamentarse#v* | *quejarse#v* | *declinar#v* | *rehusar#v* | *conseguir#v* | *obtener#v* (etc.)

The system performs neither, either or both of the processes explained above according to the operational mode selected at the outset. Expanding the affective dictionary's content provides a set of root terms that identify triggers of expressions of opinion. In the case of a verb, for example, one such term might be *creer* (the Spanish verb "to believe"). Evidently, students will not have used the verb's infinitive form in their messages, but rather a variant thereof, such as *yo creo* ("I believe") or *creemos* ("we believe"). To be able to recognise the entire conjugation of the verbs and all the possible representations of the nouns in the affective dictionary, the forms associated with the *lemma#grammatical_category* pairings are generated through expansion based on FreeLing's form dictionary. All the possible forms originating from the lemmas are thus obtained, so that every way in which a lemma included in a trigger might be mentioned in a text can be detected.

If WordNet and FreeLing were not used to expand the affective dictionary, it would only be possible to detect terms explicitly listed in the resource. The use of WordNet and FreeLing makes the process more efficient, with a single entry in the affective dictionary resulting in the potential detection of over 300 items in the cases in which the method is most successful.

V. PILOT TEST

As aforesaid, the proposed method has been used to identify the opinions expressed by students in messages posted in the tutorial classrooms of the UOC's Business Intelligence master's degree. The course in consideration has two tutorial classrooms, one for first-year students and the other for second-year students. In total, 41 messages were obtained from the two classrooms (22 for 1st year and 19 for 2nd year classroom).

TABLE I. THE MESSAGES WERE EXTRACTED IN PLAIN TEXT FORMAT AND PROCESSED BY OUR TOOL. A TEACHER HAS CREATED A GOLD STANDARD IDENTIFYING BY HAND THE OPINIONS PRESENT IN EACH OF THE MESSAGES. THIS GOLD STANDARD CONSISTS OF 28 AND 36 UNIQUE OPINIONS FOR FIRST AND SECOND-YEAR CLASSROOMS RESPECTIVELY. WE APPLIED OUR TOOL TO IDENTIFY OPINIONS IN THE CLASSROOM'S MESSAGES FOLLOWING THE FOUR PRESENTED METHODOLOGIES: EXCLUSIVELY USING THE ITEMS THAT EXIST IN THE AFFECTIVE DICTIONARY, USING THE SYNONYM EXPANSION ON THE AFFECTIVE DICTIONARY, USING SUPERORDINATE EXPANSION ON THE AFFECTIVE DICTIONARY, AND USING BOTH EXPANSIONS. FROM THE NUMBERS OF TABLE I, WHICH SUMMARIZES THE NUMBER OF OPINIONS AUTOMATICALLY DETECTED, IT CAN BE CONCLUDED THAT AFFECTIVE DICTIONARY'S COVERAGE ON THE BASIS OF SUPERORDINATION AND DOING LIKEWISE ON THE BASIS OF SYNONYMY ARE BOTH EFFECTIVE METHODS.NUMBER OF UNDUPLICATED OPINIONS DETECTED PER CLASSROOM. EACH ROW CORRESPONDS TO A DIFFERENT METHOD OF EXPANDING THE AFFECTIVE DICTIONARY (AD IN THE TABLE).

	2nd-year classroom	Coverage Increase	1st-year Classroom	Coverage Increase
Affective Dictionary(AD)	31		36	

AD + superordinates	53	70.96%	50	38.88%
AD + synonyms	54	74.19%	54	50.00%
AD + synonyms + superordinates	74	138.70%	64	77.77%

In order to evaluate the precision of the different methods, we need to evaluate whether the detected opinions are relevant or not. To achieve this goal, a teacher has evaluated these items according to two binary variables. The first one, which is called *exact detection*, takes a value of one when the marked item constitutes the core of the opinion, and zero otherwise. The second one, *proximity detection*, takes a value of one when the marked item is not the opinion's core but helps to retrieve a useful opinion in the very same sentence where it is present. Results of this evaluation are shown in Table II. An example of exact detection is the opinion "agree" within the message "... I agree with the clarification about this topic...". In the message the verb "agree" represents the agreement opinion of the writer. On the other hand, the term "know" appears many times as a proximity detection since it does not represent knowledge within the text "... I would like to know ....", but it helps to identify a doubt represented by "would like to know".

TABLE II. NUMBER OF UNDUPLICATED OPINIONS MARKED AS EXACT DETECTION (ED), PROXIMITY DETECTION (PD), TOTAL NUMBER OF UNDUPLICATED OPINIONS MARKED (TOTAL) AND NUMBER OF UNDUPLICATED OPINIONS IN THE GOLD STANDARD (GE). ROWS SHOW WHETHER A EXPANSION METHOD HAS BEEN USED

	2nd-year classroom				1st-year classroom			
	ED	PD	Total	GE	ED	PD	Total	GE
Affective Dictionary	13	19	31	36	11	12	36	28
AD + Superord.	**23**	**27**	53	36	13	16	50	28
DA + Synonyms	16	23	54	36	14	16	54	28
DA + Syns + superord.	**23**	**27**	74	36	**15**	**17**	64	28

Tables III and IV show the performance of each one of the four methodologies applied by our tool in both classrooms. Regarding proximity detection, the best methodology across classrooms is the superordination expansion. Regarding the exact detection, the superordination expansion achieves the best results in the 2nd year classroom. But, in the 1st year classroom, the methodology exclusively based on the affective dictionary performs slightly better. In this case, the high precision balances the lack of coverage and generates a slightly better F1 value.

TABLE III. PRECISION (P), RECALL (R) AND BALANCED F1 SCORE (F1) VALUES FOR EXACT DETECTION.

Exact detection	2nd-year classroom			1st-year classroom		
	P	R	F1	P	R	F1
Affective dict (AD)	41.9	36.1	38.8	**30.5**	39.3	**34.3**
AD+ superordinates	**43.4**	**63.9**	**51.7**	26.0	46.4	33.3
AD + synonyms	29.6	44.4	35.5	25.9	50.0	34.1
AD + synonyms + superordinates	31.1	**63.9**	41.8	23.4	**53.6**	32.6

TABLE IV. PRECISION (P), RECALL (R) AND BALANCED F1 SCORE (F1) FOR PROXIMITY DETECTION.

Proximity detection	2nd-year classroom			1st-year classroom		
	P	R	F1	P	R	F1
Affective dict (AD)	**61.2**	52.8	56.7	**33.3**	42.8	37.4
AD+ superordinates	50.9	**75.0**	**60.6**	32.0	57.1	**41.0**
AD + synonyms	42.6	63.9	51.1	29.6	57.1	38.9
AD + synonyms + superordinates	36.5	**75.0**	49.1	29.6	**60.7**	36.9

VI. CONCLUSIONS AND FURTHER WORK

The new virtual systems of eLearning provide new opportunities to create advanced services that behave intelligently and improve learning processes. So is possible due to the fact that most of the interactions that occur within the learning experiences are stored in a digital format that can be analyzed by using techniques of data mining. Following such idea this paper presents a simple approach that uses opinion mining techniques to provide relevant information to teachers in order to support their decisions. The article also shows the results of applying the developed tool for mining opinions on in the context of the tutorial classrooms of one of the UOC's master's degree programmes during two years.

The developed architecture involves the use of an affective dictionary and different language resources, such as WordNet and FreeLing, to extend the method's coverage. Different methods for extending the affective dictionary are proposed. The results obtained in the tests show that dictionary expansion based on superordination and synonymy substantially improves the method's coverage. The differences between synonym-based expansion and superordinate-based expansion are not significant, and we are thus unable to state that one method is better than the other. The two methods have the advantage of being compatible, as the spectacular increase in coverage produced by using them together demonstrates. It should be highlighted that both expansion methods are useful for exact and proximity detection. Disaggregating results on a classroom basis, performance is better on the 2nd-year classroom. This fact could be attributed to the higher level of expertise of the students reflected on the way they express their opinions in the messages. It is more common to find clarification messages regarding communication skills coming from the 1st-year classroom than from the 2nd-year.

Our work is significant in two ways. Firstly, it has shown that using language technology techniques for opinion mining to support teaching activity is both feasible and effective. Secondly, it has generated a tool that teaching staff may find useful for assessing how a semester has gone and detecting potential scopes for improvement.

The tool is a finished product. Nonetheless, its potential is such that many improvements could be made to enhance its functionalities or make it easier to adopt. Such improvements can be divided into three broad areas, namely integration into the UOC's campus, improving efficiency in terms of opinion detection and improving usability.

The proposed tool is designed to be used after the teaching in order to help in the analysis of what happened in the last course. We feel that integrating the tool into the UOC's campus in such a way as to allow for the online analysis of forum messages would be beneficial. Leaving aside more technical considerations, such as the tool's inclusion in classrooms, certain operational aspects ought to be refined, possibly involving, for example, the creation of new functionalities for filtering by date and topic, which would help lecturers identify the general view on a teaching programme at a given point in time or on a particular topic.

There are various ways in which efficiency in terms of opinion detection could be enhanced. Firstly, the list of triggers could be generated using other techniques, and the results compared to that of the current method. For example, one option would consist of compiling large quantities of text and manually marking up opinions (positive and negative), resulting in a training corpus for opinion detection purposes. Such a corpus would provide a foundation for the development of a statistical model based on automatic learning, which would make it possible to automatically detect opinions contained in texts with a degree of coverage likely to be greater than that offered by manually drawing up a list of triggers. Another possibility, in relation to which some preliminary work has been undertaken, would be the introduction of resources such as WordNet-Affect [10]. An extension of WordNet's domains, WordNet-Affect marks up certain synsets with affective tags to indicate that they have an element of emotion, as well as marking up synsets as having positive, negative, neutral or ambiguous valence. Where coverage is concerned, FreeLing could be used to lemmatise text instead of to generate the different forms of triggers. It would thus be possible to search for triggers in a lemmatised text, with the benefits that the morphological disambiguation performed by FreeLing would entail. If, for example, the Spanish word *crea* appeared in a text, its lemma could be *creer#v* or *crear#v*, as it is part of the conjugation of both the verbs in question (the first of which means "to believe", the second "to create"). Lemmatising with FreeLing would give the corresponding lemma with a high level of accuracy, thus filtering out instances of *crea* corresponding to *crear#v*, which, in principle, is not used in expressions of opinion. Finally, the opinions can be refined by detecting also their polarity (whether an opinion is good or bad) and the topic they are dealing with (what the opinion is about).

While the information (tag clouds and xml-tagged documents) generated by the tool is comprehensible and easy to use, we feel that there ought to be more ways of viewing it. The tag cloud generated by the tool only provides a quick vision of the most common opinions, which can be of use only in few circumstances. One option could be using the communication facilities' directory structure or threads to group opinions together on the basis of the topics they cover. Another interesting representation would be in a form of a balanced scorecard that provides information about the main indicators of the learning process to the teachers in one integrated graphical view and facilitates the navigation thru the information.

Acknowledgment

This work has been partially supported by department of Innovation of the Universitat Oberta de Catalunya (APLICA/2010).

References

[1] P. Turney, "Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews", Proceedings of the 40th annual meeting of the Association for Computational Linguistics, 2002, pp. 417–424.

[2] J. Read, "Recognising affect in text using pointwise-mutual information", Master Thesis, University of Sussex, 2004.

[3] C. Corinna and V. Vapnik, "*Support-Vector Networks"* Machine Learning, 20, 1994.

[4] S. Deerwester, S. Dumais, G. Furnas, T. Landauer, R. Harshman, "Indexing by Latent Semantic Analysis". *Journal of the American Society for Information Science* 41 (6): 391-407, 1990.

[5] G. Leshed, J. Kaye, "Understanding how bloggers feel: recognizing affect in blog posts". *CHI extended abstracts* (pp. 1019-1024), ACM 2006.

[6] G. A. Miller, R. Beckwith, C. D. Fellbaum, D. Gross, K. Miller, "WordNet: An online lexical database". *Int. J. Lexicograph.*, 3, 4, 1990 pp. 235-244.

[7] X. Carreras, I. Chao, L. Padró, M. Padró, "FreeLing: An Open-Source Suite of Language Analyzers". *Proceedings of the 4th International Conference on Language Resources and Evaluation* (LREC'04).

[8] D. García, F. Alías, "Identificación de emociones a partir de texto usando desambiguación semantic". *Procesamiento del lenguaje natural* (no. 40, April 2008, pp. 75-82). ISSN 1135-5948.

[9] M. Taulé, M.A. Martí, "SENSEVAL, una aproximació computacional al significant". *Digithum*, no. 5, UOC. ISSN 1575-2275, 2003.

[10] R. Valitutti, "WordNet-Affect: an Affective Extension of WordNet". *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC'04).*