OpeNER: Open Tools to Perform Natural Language Processing on Accommodation Reviews

Aitor García-Pablos[a], Montse Cuadros[b] and Maria Teresa Linaza[a]

[a]Department of eTourism and Cultural Heritage
[b]Department of Human Speech and Language Technologies
Vicomtech-IK4, Spain
{agarciap, mcuadros, mtlinaza}@vicomtech.org

Abstract

Opinion mining is crucial for hoteliers and other tourism industries in order to improve their service from the analysis of services failures and recovery. The extensive use of the Internet and social networks has shifted the way tourism information is shared and spread. Travel agencies, hotels, restaurants, tourist destinations and other actors require the aid of new technologies to get an insight of the vast amount of customer generated reviews. Develop and integrate text analysis technologies is usually difficult and expensive, because it involves the use of Natural Language Processing techniques. This paper introduces the OpeNER European project, a set of free Open Source and ready-to-use text analysis tools to perform text processing tasks like Named Entity Recognition and Opinion detection. The paper also provides an example of a possible application of the OpeNER results in the geolocation of hotel reviews..

Keywords: user generated reviews, language processing, sentiment analysis

1 Introduction

Opinion mining is crucial for hoteliers and other tourism industries in order to improve their service from the analysis of failures and recovery. The extensive use of the so-called Web2.0 and social networks has issued a big change on the way tourism information is shared and spread (Liu *et al.*, 2013). Travel agencies, hotels, restaurants, tourist destinations and any other service providers cannot control the immense data flow generated via thousands of online reviews, comments and interactions between past and potential customers. The classic and monolithic self-promotion techniques are losing strength against the social media and the word-of-mouth enabled by the information era.

In order to keep track of what is going on over the Web about their brands (and their competitors), companies require tools to cope with the vast amount of content generated every day. This kind of tools should allow them to gain control over what is being said; to tackle negative opinions; to detect trends in user behaviour and, in general, to take smarter decisions.

Currently, there are many companies offering solutions to this problem. Most of them involve the automatic analysis of text coming from different online sources (e.g. websites like TripAdvisor, social networks as Twitter or Facebook, etc.). The automatic analysis of text requires Natural Language Processing (NLP) tools and techniques. Some of the existing tools and software libraries are Open Source and free, but the heterogeneity and diversity of technologies, output formats and system requirements make it difficult to integrate them to build a customized analysis system.

To develop and maintain such a system requires both expertise and an investment of time and resources that may not be affordable by the tourism value chain.

This paper introduces the results of the OpeNER[1] European project, which aims at providing a set of Open Source and ready-to-use tools to perform NLP analysis in six languages, including English, Spanish, Italian, Dutch, German and French. The application of the results of OpeNER to the customer reviews in the tourism sector should enable the automatic extraction of textual feedback on the basis of NLP technologies specially focused on opinion mining. The remaining of this paper is structured as follows. Section 2 presents a brief state of the art describing some common NLP tasks, approaches and existing Open Source tools in each case. Section 3 describes the OpeNER project, explaining the motivations and a general overview of the objectives of the project. Finally the Section 4 shows the conclusions.

2 State of the art

Natural Language Processing (NLP) is a field of Computer Science that studies the use of automatic ways to process natural language. As it has been mentioned before, automatic processing of text is becoming more and more important in the tourism sector due to the large amount of content generated by users every minute. Thus, (semi) automatic ways of processing is needed to extract valuable information. NLP is a very wide research field, with many subfields addressing specific tasks, from breaking a text into basic units to ease further processing (i.e. sentence splitting, phrase chunking, tokenizing, etc.) to more complex ones like semantic analysis.

2.1 Processing text

In order to process a text, it is first necessary to determine its language. There are currently many Open Source language identification tools that implement state-of-the-art algorithms, achieving a precision over 99% for tens of languages. The most popular approaches are based on statistical distributions and probabilities of character level *n*-grams (Rehurek and Kolkus, 2009), which are sequences of *n* characters. It is proven that every language has its own particular distribution of such *n*-grams.

Once the language has been identified, tokenization is commonly the following step of any text processing pipeline (Webster and Kit, 1992). It is the process of breaking a text into its fundamental pieces (i.e. tokens), which are likely to be a word, a number, a punctuation mark, or a particular combination of them.

Part-of-Speech tagging (PoS-tagging) is the next step that assigns grammatical categories to words in a text. Basically, it states that a word in a particular context is a noun, a verb, an adjective, an adverb, etc. It can also provide more information, like the gender and number of a word, or the person in case of verbs. PoS-taggers are usually based on stochastic methods like Hidden Markov Models or Maximum Entropy, trained on sets of pre-annotated data (Brants, 2000; Collins, 2002). The accuracy achieved by state-of-the-art taggers varies from one language to another and relies heavily on available training datasets (Giesbrecht and Evert, 2009).

[1] http://www.opener-project.eu

Furthermore, Named Entity Recognition and Classification (also known as NERC) locate and classify rigid entity designators in text such as proper names (Nadeau and Sekine, 2007). The concept of "entity" varies from one system to another. In the tourism field, the main entities are names of people, organizations and location names (countries, cities, or any other kind of geographical location). In other contexts, also dates, numeric expressions and/or currencies are detected.

The previously detected entities are disambiguated in order to distinguish the entities referred from a set of potential candidates using Named Entity Disambiguation and Linking techniques. When possible, detected named entities are linked to well-known ontologies or knowledge-bases (Sil *et al*, 2012) like the Wikipedia's page of that entity. This allows uniquely identifying that entity according to a certain namespace or vocabulary (Rao, McNamee and Dredze, 2013), and aggregating or manipulating more precisely all the mentions to the same entity in order to avoid confusions with other entities with similar names.

On the other hand, two different mentions in a text may refer to the same real-world entity. For example, in the following comment, "I stayed in NH in Brussels and Zurich and I really liked *them* because of *their* modern and stylish design and big rooms", the word *them* refers to "NH in Brussels and Zurich", and so does the word *their*. Detecting which mentions co-refer to the same entity is known as co-reference resolution (Bagga and Baldwin, 1999). To solve co-referent expressions, both linguistic and domain knowledge are required. One of the best performing systems is a multi-pass sieve co-reference resolution system (Lee *et al.*, 2011).

Finally, sentiment analysis and opinion mining are closely related fields which refer to the application of NLP techniques to extract subjective information about how someone expresses a feeling (negative, positive or neutral) about something (Pang and Lee, 2008). These tasks are increasingly important for determining the opinion about products and services, and brand reputation on the Internet. Usually, this information is the sentiment of the so-called "opinion holder" towards a particular "opinion target" (a topic, an entity or some part or feature of it) (Liu, 2010). Ideally, this task is about retrieving "who" is opining "what" about "which entity" in each given piece of text. The time can be also important, especially when the opinions and sentiments change very quickly.

There are plenty of different approaches to perform sentiment analysis and opinion mining. Not all the available systems and techniques aim to extract the same type of information or with the same granularity. Some are oriented to just finding the overall polarity of a full sentence, paragraph or document, while others aim at finding the polarity on a product/service feature basis (e.g. distinguishing whether a particular opinion is about the rooms of a hotel or about the breakfast).

Furthermore, most of them involve machine learning techniques combined with specific language resources. Usually, those tools are language and domain dependent (i.e. they work better for the language and domain they were developed for and require minor or major adaptions to work in other languages or application domains).

2.2 Application to the tourism sector

The increasing growth and popularity of user-generated contents on the Web has led to a new area of research in the application of text mining techniques. Applications of sentiment analysis and opinion mining based on text reviews have grown very quickly during the last decade in the tourism sector.

The earliest approaches focused on sentiment analysis of product reviews, which were clustered as positive or negative on the basis of specific sentiment structures (Hu and Liu, 2004; Lau, Lee and Hoo, 2005; Popescu and Etzioni, 2005). Four steps were defined for online text mining: definition of mining context and concepts; data collection; dictionary construction; and data analysis. Several analysis have been done related to the profile of the hotel or the price of the room.

More recently, sentiment classification of consumer reviews is addressing bigger challenges, since the opinion mining systems try to deal with more complex tasks and results, as customers may provide a mixed review, combining positive and negative aspects of the same product or service. Ghose, Ipeirotis and Li (2009) used a 4-grams Dynamic Language Model classifier to acquire a subjectivity confidence score for each sentence in a hotel review and derive the mean and standard deviation of this score. The analysis of the content focused on polarity classification, sentiment classification of customer reviews, or the automated extraction of product attributes. They have further used text-mining techniques to incorporate textual information from hotel reviews in demand estimation models on the basis of the user-generated hotel reviews from Travelocity and TripAdvisor.

Ye, Zhang and Law (2009) presented a study to analyse the existing approaches to perform automatic classifications based on the sentiment analysis of online reviews related to travel destinations. Furthermore, the study analyses different supervised machine learning algorithms and their effect on the different amount of training corpus to various performance measurements in terms of accuracy, precision, and recall in the sentiment classification of online reviews about tourist destinations. The algorithms evaluate the reviews about seven popular travel destinations in Europe and North America.

On the other hand, Xiang and Gretzel (2010) have applied text analysis to understand the queries extracted from a number of transaction logs from search engines. Although generally speaking accommodation and transportation were the most searched information, there were differences depending on the size of the destination and its tourist level. Furthermore, there were strong associations between keywords used and specific destinations, reflecting the knowledge about them.

Moreover, Lee, Singh, & Chan (2011) used text mining techniques to extract keywords from descriptive comments from hotel customers in order to identify areas of service failures and recovery actions. CATPAC software was used to classify algorithms and identify main topics based on the frequency of key terms. Furthermore, Kasper and Vela (2011) have implemented a service for hotel managers that collects customer reviews from various sites on the web; analyzes and classifies the textural content of the review; and presents the results in a precise way. Its main disadvantage is that it is only available in German.

Finally, Gräbner *et al.* (2012) have proposed a system that classifies customer reviews of hotels on the basis of sentiment analysis techniques. The study includes building a lexicon with a semantic orientation; the application of sentiment analysis to generate a classification of customer reviews; and the evaluation of the results with quantitative ratings.

3 The OpeNER project

OpeNER is a European project which aims at providing a set of Open Source tools to perform text processing tasks like Named Entity Recognition, sentiment analysis and opinion detection. The objective is to offer a set of ready-to-use tools and software modules to process texts in six different languages, plus the capabilities to easily extend them to new languages and application domains. The Open Source nature of the project (i.e. the source code is open and freely available) should enable the potential community of users to take the existing OpeNER tools as a starting point, and extend and integrate them to build their custom text analysis systems.

During the OpeNER project, different text processing modules have been developed for six major European languages (English, Spanish, French, Italian, Dutch and German). These modules include the following functionalities: language detection; sentence splitting and tokenisation; Part-of-Speech tagging; Named Entity Detection and Classification; Named Entity Linking; co-reference resolution; and sentiment analysis and opinion detection. OpeNER also provides some tools to perform domain adaptation of the existing resources (e.g. to adapt sentiment lexicons to a new domain, to train new models for opinion detection, etc.). Some of the provided tools are based on already available third-party tools, like Apache OpenNLP framework or DBpedia Spotlight that have been adapted and conveniently wrapped to achieve the versatility and modularity desired for the OpeNER modules.

One of the main features of the tools is the modularity of each component (i.e. understanding a component as the piece of software in charge of a particular NLP task). This modularity is achieved using a single yet expressive data representation format called KAF (Bosma, Vossen and Soroa, 2009).

The OpeNER project has been evaluated in the tourism domain. During the project, a manual annotation campaign allowed annotating a hotel review dataset for each of the six languages officially handled by OpeNER. These datasets were then used to train specific models to analyse hotel customer reviews, and also to evaluate the performance of the resulting system. Additionally a set of reference applications was built in order to analysis potential added value services in the tourism domain.

3.1 OpeNER general architecture

OpeNER is built on an individual module basis. Each module receives a single input; performs a single text processing task; and returns a single output. Both the input and the output are documents in KAF format, which allows a very easy integration and chaining between different modules to build a full analysis pipeline.

Fig. 1~~Fig. 1~~ shows a possible way of chaining OpeNER modules to perform different analysis. The output is always a document in KAF format that can act as the input to another module. KAF documents include all the information obtained in each analysis separated in individual layers. Each module works only on a single KAF layer (creating it from scratch or completing the information of an existing layer). OpeNER provides tools to parse and work with KAF documents and tools to convert them to a more human readable format like JSON.

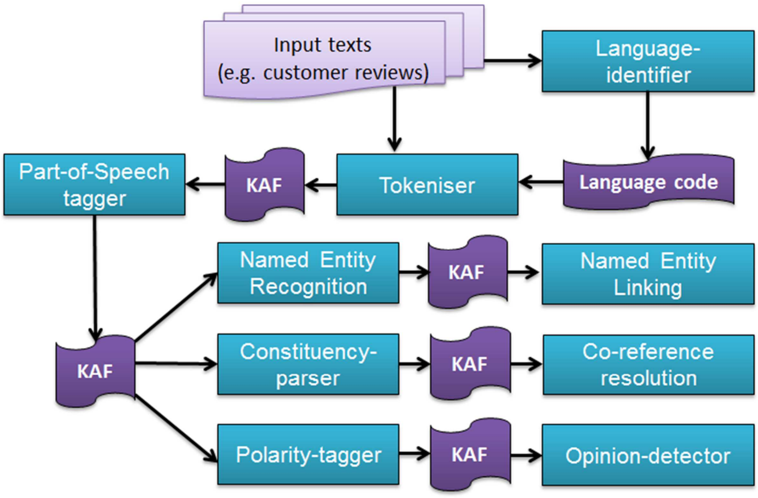

Fig. 1. A possible text analysis pipeline chaining OpeNER modules

The following text from a hotel review will be taken as an example.

"I have been at Albergo Acquarello hotel at Lugano and I liked the beautiful decoration. The rooms were very comfortable. On the other hand, the restaurant was really expensive."

First, the text to be analysed is sent to the language identifier which returns the language code corresponding to the language detected in the text. Secondly, the tokeniser module receives the text and the language code, and performs the tokenisation of the words outputting the result as a KAF document. Such document is the input for the Part-of-Speech tagger module, which outputs the same KAF document with additional information coming from the Part-of-Speech tagging process.

The language identifier correctly detects the language as English; the tokeniser breaks the text into individual sentences and tokens (i.e. separating words and punctuation marks); and the Part-of-Speech tagger annotates each word as being a noun, a verb, an adjective, etc. An illustrated representation of the result can be found at Fig. 2~~Fig. 2~~.

All this information is represented in KAF[2], which is sent to the Named Entity Recognition module to detect entities. The analysis detects two entities in the text: *Albergo Acquarello* and *Lugano*. The former has been classified as an "organisation"

[2] KAF documents are XML files too verbose to be represented in this paper. More information about KAF format and examples can be found at the OpeNER website

(the *Albergo Acquarello hotel*), while the latter has been defined as a geo-spatial location (Lugano, Switzerland). After sending the result to the Named Entity Linking module, the mention to *Lugano* has been linked to its entry in DBpedia. This allows determining which "*Lugano*" entity is the text about (in case there is more than one possible "*Lugano*" in the world) and obtaining additional metadata about the entity if available (e.g. the geo-coordinates, the population, the country, etc.).

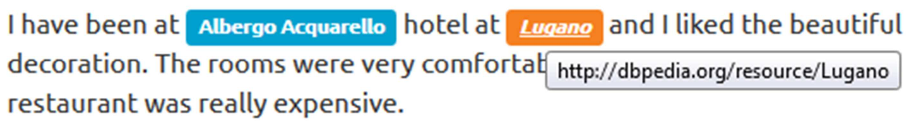

Fig. 2. Coloured representation of the Named Entity Recognition result (Albergo Acquarello and Lugano as "organisation" and "location" respectively)

If the Polarity-tagger module is invoked, the analysis of the sentiment and opinion-related information are obtained. The result is illustrated in the Fig. 3~~Fig. 3~~. The module assigns a polarity (positive, negative) to the words in the text according to a sentiment lexicon (i.e. a dictionary that states the most probable polarity for a word inside the given domain). The detected positive and negative words have been highlighted with different colours, as well as the intensifiers (i.e. the words that intensify the polarity of the surrounding words).

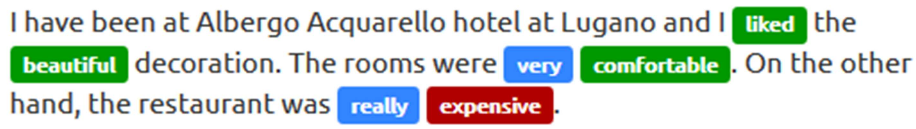

Fig. 3. Detected polarity of the words highlighted with different colours

The polarity information is a first step to get an insight about the sentiment of the review. The Opinion detector module goes further and detects whole expressions; classifies them as being positive or negative (e.g. some expressions may contain words of a certain polarity but the overall expression might not inherit it); and tries to find the target of that expression (i.e. the particular object or feature which the opinion is about).

For example, Fig. 4~~Fig. 4~~ shows the possible representation of the triplet of information the OpeNER opinion detector tries to fulfil. One is the "opinion holder" (i.e. the author of the opinion itself). In a standard hotel review, the opinion holder of all the opinions in that review is the author of the review implicitly. When there is an explicit opinion holder, it appears as "Somebody" in the example. The second part of the triplet is the opinion expression itself, which is the word or group of words that comprise an opinion or a particular sentiment towards something. An opinion expression can be positive, negative or neutral.

Finally, the opinion target is the object/feature being reviewed (i.e. the object being assessed in the corresponding opinion expression). The opinion target (also called aspect term, feature term, etc.) is very important to obtain a fine grained sentiment score. It is crucial to be able to aggregate the opinions on a per-feature basis to assess the strengths and weaknesses of a product or service (e.g. hotel rooms are positively perceived while the breakfast service is negatively evaluated).

Somebody said **"liked"** about **decoration**

Somebody said **"very comfortable"** about **The rooms**

Somebody said **"really expensive"** about **the restaurant**

Fig. 4. An inline representation of the information obtained by the Opinion detector

3.2 Evaluating OpeNER in the tourism sector

One of the OpeNER project evaluation scenarios has been the tourism sector, more precisely, the hotel domain. During the customization of the platform to the tourism sector, a set of hotel reviews has been manually annotated with sentiment and opinion related information. The reviews were extracted from online customer review websites like Zoover[3]. Further factors were taken into account to avoid bias in the extracted content, apart from choosing reviews for the six languages involved in OpeNER (English, Spanish, French, Italian, Dutch and German). For example, the chosen reviews were equally distributed among variables like the home country of the reviewer, the motivation for the stay at that hotel (work or leisure), etc. Such data is usually available as metadata annexed to the reviews. The final set of hotel reviews included about 200 reviews per language.

The annotation campaign consisted on two or more people (native speakers or with a deep knowledge of the language they were annotating) tagging the reviews according certain annotation guidelines with the help of a customized annotation tool. Per each review, the annotations consisted on tagging the opinion expressions and when possible, the corresponding opinion holders and opinion targets. Also, other valuable information was manually tagged, like the polarity of the words or the general category of the opinion target (e.g. both “coffee” and “orange juice” belong to the “breakfast” category, while “towel” and “shower” belong to the “bathroom” category).

These annotated reviews were then used to train the models that enable the work of the Opinion detector module. It is based on different machine learning techniques like Conditional Random Fields (CRF) (Sutton and McCallum, 2012) and Support Vector Machines (SVM) (Brereton and Lloyd, 2010) that must be trained over a previously annotated dataset. A certain amount of the annotated hotel reviews was used for the training while the remaining subsets were employed to perform a formal evaluation of the resulting opinion detection models. The results of this evaluation are shown in

Table 1~~Table 1~~. The results vary for each language due to the impact of the different number of annotated reviews among them. Additionally, not all languages have the same complexity; issues like morphology and sparse vocabulary affect performance.

[3] http://www.zoover.com

Table 1. OpeNER opinion detector evaluation results

Tool	Language	Precision	Recall	F-Score	Method	Dataset
Opinion detector	en	85,52%	58,45%	69,44%	CRF + SVM	OpeNER manual hotel annotations
Opinion detector	nl	82,8%	51,77%	63,71%	CRF + SVM	OpeNER manual hotel annotations
Opinion detector	de	75,64%	48,88%	59,38%	CRF + SVM	OpeNER manual hotel annotations
Opinion detector	es	74,41%	46,55%	57,27%	CRF + SVM	OpeNER manual hotel annotations
Opinion detector	it	65,47%	40,39%	49,96%	CRF + SVM	OpeNER manual hotel annotations
Opinion detector	fr	70,94%	46,28%	56,02%	CRF + SVM	OpeNER manual hotel annotations

OpeNER already provides these pre-trained models and also tools to perform further adaptation and annotation of more reviews. It should be noted that the amount of annotated reviews and the annotation quality have a direct impact on the performance of the Opinion detector module. There are other linguistic resources, like the opinion lexicon (i.e. the dictionary that holds the polarity of the words), that can be improved and tuned to better fit the target domain (e.g. a word may denote different sentiments in different domains).

OpeNER has also been tested by building different reference applications to serve as an example of potential best-practices of the OpeNER tools and technologies. One relevant example is shown in Fig. 5~~Fig. 5~~. Tour-pedia (Marchetti *et al.*, 2014) is an application that geolocates the sentiment analysis of hotel reviews using emoticons to provide a quick overview of the positive or negative feedback provided by customers in their reviews on the social media. Reviews and other metadata (e.g. location metadata on a map) from customers have been extracted from different sources like Google Places or FourSquare. The content of the reviews has been processed with OpeNER tools to obtain a measure of the polarity and draw the appropriate emoticon. Tour-pedia is an illustrative example of how to build an added-value service on top of the text processing capabilities provided by OpeNER.

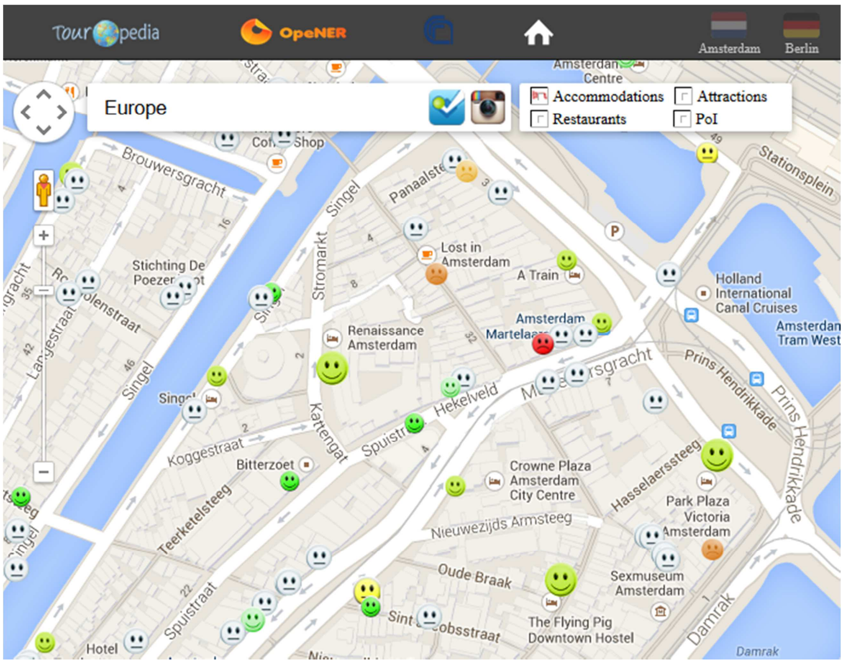

Fig. 5. A screenshot of Tour-pedia based on some of the OpeNER technologies

4 Conclusions

The large amount of text content generated everyday over the Internet is both a big opportunity and a challenge. There are many ways in which customers can provide their opinion and feedback about products and services. This also applies to tourist destinations, hotels, restaurants and other services. Currently, there are many specialized websites to write reviews and provide feedback, plus the omnipresence of the social networks to exchange information publicly. At the same time, there are many companies offering services to monitor this content and give an insight about what is being said about a particular service or brand.

This paper describes some of the outcomes of the OpeNER project which aims at bringing text processing technologies a step closer to SMEs and other kind of end-users interested in analysing textual content. OpeNER is an Open Source project which provides ready-to-use tools and modules to create a custom analysis pipeline with Named Entity Recognition, Sentiment Analysis and Opinion Mining capabilities. OpeNER is based on a single data representation format (KAF) to enable a simple integration between the different modules and ease the extension and development of new modules and components.

The evaluation scenario of the OpeNER tools was the tourism sector, more precisely hotel reviews written by customers. During the development and domain customization of the platform to the tourism sector, a set of hotel reviews has been manually annotated with sentiment and opinion related information.

These annotated reviews were then used to train the machine learning models that enable the work of the Opinion detector module. A certain amount of the annotated hotel reviews were used for the training while the remaining subsets were employed to perform a formal evaluation of the resulting opinion detection models.

OpeNER also provides tools to improve or further customise the system for the tourism sector or to extend some of the existing modules to new domains. This adaptation to a new domain requires the generation of some specific resources, like sentiment lexicons and opinion detection models trained on pre-annotated content of the target domain. Additionally, a set of reference applications was built in order to foresee potential uses of OpeNER technology. The Open Source nature of the project provides a good entry point to the language processing technologies and enables SMEs to extend the provided software and build their own analysers and products upon it.

5 Acknowledgements

OpeNER has been funded by the European Commission under the FP7-ICT-2011-SME- DCL-296451.

6 References

Agerri R., Cuadros M., Gaines S., & Rigau G. (2013) *OpeNER: Open Polarity Enhanced Named Entity Recognition*, Proceedings of the 29th Annual Meeting of Sociedad Española para el Procesamiento del Lenguaje Natural, SEPLN'13. Madrid, España. Procesamiento del Lenguaje Natural. Vol. 51 pp. pp 215-218.

Bagga, A., & Baldwin, B. (1999). Cross-document event coreference: Annotations, experiments, and observations. *Proceedings of the Workshop on Coreference and its Applications*.

Bacciu, C., Lo Duca, A., Marchetti, A., & Tesconi., M (2014). Accommodations in Tuscany as Linked Data. Proceedings of the 9th edition of the Language Resources and Evaluation Conference.

Bosma, W., Vossen, P., & Soroa, A. (2009). KAF: a generic semantic annotation format. *Proceedings of the GL2009 Workshop on Semantic Annotation*.

Brants, T. (2000). TnT: a statistical part-of-speech tagger. *Proceedings of the sixth conference on Applied natural language processing*, (1).

Brereton, R. G., & Lloyd, G. R. (2010). Support vector machines for classification and regression. *The Analyst*, *135*, 230–267.

Collins, M. (2002). Discriminative training methods for hidden Markov models: theory and experiments with perceptron algorithms. *Proceedings of the ACL-02 conference on Empirical methods in natural language processing*, 1–8.

Ghose, A., Ipeirotis, P., & Li, B. (2009). The economic impact of user-generated content on the Internet: Combining text mining with demand estimation in the hotel industry, (December).

Giesbrecht, E., & Evert, S. (2009). Is Part-of-Speech Tagging a Solved Task? An Evaluation of POS Taggers for the German Web as Corpus. *Web as Corpus Workshop WAC5*, 27.

Gräbner, D., Zanker, M., Fliedl, G., & Fuchs, M. (2012). *Classification of customer reviews based on sentiment analysis*.

Hu, M., & Liu, B. (2004). Mining opinion features in customer reviews. *AAAI*.

Kasper, W., & Vela, M. (2011). Sentiment analysis for hotel reviews. *Computational Linguistics-Applications Conference*, *231527*, 45–52.

Lau, K., Lee, K., & Ho, Y. (2005). Text mining for the hotel industry. *Cornell Hotel and Restaurant Administration Quarterly*, *46*(3), 344–362.

Lee, H., Peirsman, Y., Chang, A., Chambers, N., Surdeanu, M., & Jurafsky, D. (2011). Stanford ' s Multi-Pass Sieve Coreference Resolution System at the CoNLL-2011 Shared Task. *Proceedings of the Fifteenth Conference on Computational Natural Language Learning: Shared Task. Association for Computational Linguistics* (pp. 28–34).

Lee, M. J., Singh, N., & Chan, E. S. W. (2011). Service failures and recovery actions in the hotel industry: A text-mining approach. *Journal of Vacation Marketing*, *17*(3), 197–207.

Liu, B. (2010). Sentiment Analysis and Subjectivity. *Handbook of natural language processing*, 1–38.

Liu, S., Law, R., Rong, J., Li, G., & Hall, J. (2013). Analyzing changes in hotel customers' expectations by trip mode. International Journal of Hospitality Management, 34, 359–371.

Nadeau, D., & Sekine, S. (2007). A survey of named entity recognition and classification. *Lingvisticae Investigationes*, *30*(1), 3–26.

Pang, B., & Lee, L. (2008). Opinion Mining and Sentiment Analysis. *Foundations and Trends® in Information Retrieval*, *2*(1–2), 1–135.

Popescu, A., & Etzioni, O. (2005). Extracting product features and opinions from reviews. *Natural language processing and text mining*, (October), 339–346.

Rao, D., McNamee, P., & Dredze, M. (2013). Entity Linking : Finding Extracted Entities in a Knowledge Base. *Multi-source, Multilingual Information Extraction and Summarization* (pp. 93–115).

Řehůřek, R., & Kolkus, M. (2009). Language identification on the web: Extending the dictionary method. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (Vol. 5449 LNCS, pp. 357–368).

Sil, A., Cronin, E., Nie, P., Yang, Y., Popescu, A.-M., & Yates, A. (2012). Linking named entities to any database. *EMNLP-CoNLL 2012* (pp. 116–127).

Sutton, C., & McCallum, A. (2012). An Introduction to Conditional Random Fields. *Foundations and Trends in Machine Learning*, *4*, 267–373.

Webster, J. J., & Kit, C. (1992). Tokenization as the Initial Phase in NLP. *Proc. of COLING-92* (pp. 1106–1110).

Xiang, Z., & Gretzel, U. (2010). Role of social media in online travel information search. *Tourism Management*, *31*(2), 179–188. Elsevier Ltd.

Ye, Q., Zhang, Z., & Law, R. (2009). Sentiment classification of online reviews to travel destinations by supervised machine learning approaches. *Expert Systems with Applications*, *36*(3), 6527–6535. Elsevier Ltd.

Corrections regarding the comments from the reviewers:

One of the reviewers made no other suggestion but to include an additional reference, to
http://dx.doi.org/10.1016/j.ijhm.2012.11.011
That reference has been included.

On the other hand the second reviewer had many suggestions, and we have addressed most of them. These are:
Suggestion 1: change the title in order to focus on the application of Opener (eg, adding “opinion” before text); as it is now the focus is on the project and on NLP in general. Even better would be: hotel reviews analysis.
Changes made: the title has been slightly changed according to the suggestions

Suggestion 2: opinion mining and sentiment analysis are not the same thing: do not use them as synonymous (eg in the abstract, keywords)
Changes made: after revising the literature (e.g. Pang & Lee 2002) and even the Wikipedia entry for “sentiment analysis”, we think that the terms are reasonably well used together along the article. Anyway we have revised the text and made some small edits to further ensure this point.

Suggestion 3: The following sentence should also be revised: “Applying OpeNER to the tourism sector will lead to the use of a linguistics-based text mining model to extract detailed information about customer experiences from textual feedback. “ Tourism sector is not only textual opinions
Changes made: The sentence has been revised and edited accordingly

Suggestion 4: Explain the meaning of granularity in the following sentence: “with the same granularity. “
Changes made: It is true that the meaning of the sentence is not clear enough. The paragraph has been reworked to add and explanation and clarify the meaning of the “same granularity” expression.

Suggestion 5: Are the analysis in “with the same granularity. “

automatic? semi-automatic? manual? The same remark applies to the following approaches

Changes made: This entails with the previous suggestion made and the performed changes.

Suggestion 6: Change: “More recently, sentiment classification of consumer reviews is getting very challenging, since customers provide a mixed review, combining positive and negative aspects of the same product or service. “: It is not a recent problem, it has been addressed more recently

Changes made: We completely agree, the sentence has been corrected.

Suggestion 7: “This paper describes the results of the OpeNER project which aims at bringing text processing technologies a step closer to SMEs and other kind of end-users interested in analysing textual content. “: are all the results reported in the paper?!

Changes made: As the reviewer has noticed, not all the results of the OpeNER project are reported, so the sentence has been reworked to clarify this point

Suggestion 8: add an explanation of the reasons why results for recall and precision are different in different languages

Changes made: an explanation has been added