Weakly-supervised Aspect Based Sentiment Analysis*

Análisis de Sentimiento Basado en Aspectos Débilmente Supervisado

Aitor García-Pablos, Montse Cuadros
Vicomtech-IK4
P/ Mikeletegi 57
San Sebastián, Spain
{agarciap,mcuadros}@vicomtech.org

German Rigau Claramunt
IXA Group, Basque Country University
P/ Manuel Lardizabal 1
San Sebastián, Spain
german.rigau@ehu.eus

Resumen: Presentamos un sistema débilmente supervisado basado en *topic modelling* y representación continua de palabras para realizar análisis de sentimiento basado en aspectos en diferentes idiomas y dominios usando sólo unas pocas palabras semilla.
Palabras clave: ABSA, topic modelling, continuous word embeddings

Abstract: We present a weakly-supervised system based on topic modelling and continuous word embedding to perform aspect based sentiment analysis for different languages and domains just by adapting a few seed words.
Keywords: ABSA, topic modelling, continuous word embeddings

1 Introduction

Sentiment Analysis is a relevant area of Natural Language Processing that deals with the automatic analysis of subjective opinions (Pang y Lee, 2008; Liu, 2012). In few words, sentiment analysis tries to discover if an opinion contains a positive or a negative sentiment. In addition, the so-called Aspect Based Sentiment Analysis (ABSA) goes further, trying to discover to which domain aspect (i.e. particular feature or facet of the evaluated entity) corresponds each analysed opinion.

2 Proposed system

The proposed system is a combination of several unsupervised methods inside a topic modelling approach. The only supervised information injected consists of a few seed words that are used to define the desired domain aspects and the sentiment polarity.

The core of the system is an extension over the well-known Latent Dirichlech Allocation (LDA) (Blei, Ng, y Jordan, 2003) with additional variables to model domain aspects, aspect-term and opinion-word separation and sentiment polarity. Continuous word embeddings, like Word2Vec (Mikolov et al., 2013), are used to calculate word similarity among words in each document and the configured seed words for domains aspects and polarities. The calculated similarity values are used to bias the topic modelling hyper-priors, guiding the process towards the configured domain aspects. In addition, a simple Maximum Entropy (MaxEnt) classifier helps separating aspect terms from opinion words during the topic modelling process. The training instances for the MaxEnt classifier are bootstrapped from the same seed words that define the domain aspects and polarity, avoiding additional supervision.

The result is W2VLDA, a weakly-supervised topic model that models each corpus sentences according to the specified domain aspects and polarities. Figure 1 shows the model in plate notation.

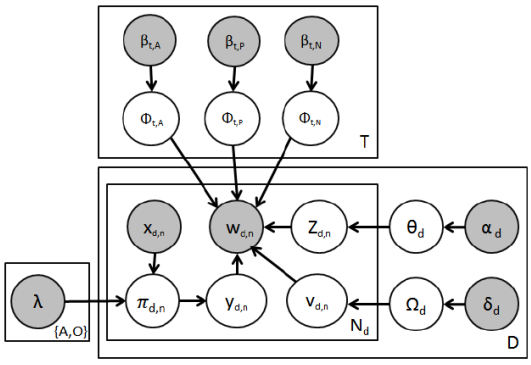

Figure 1: W2VLDA model in plate notation.

* This work was supported by Vicomtech-IK4 and by the project TUNER - TIN2015-65308-C5-1-R (MINECO/FEDER, UE)

3 Results and evaluation

We evaluate W2VLDA on both domain aspect classification and sentiment polarity classification using the SemEval 2016 task 5 datasets (Pontiki et al., 2016). These datasets consist of customer reviews about restaurants in several languages annotated with a category (domain aspect) and sentiment polarity. We only use the three main domains aspects (*food*, *service*, *ambience*), because the other domain aspects, *drinks* and *location*, represent less than the 5% of the labelled instances. We evaluate a binary polarity classification (*positive*, *negative*), discarding *neutral* examples. We perform experiments for four languages: English, Spanish, French and Dutch.

Table 1 shows an example of domain aspect configuration for restaurants reviews domain, for English and Spanish, used for the experiments. These seed words (only one per domain aspect), and their equivalents for other languages, are the only language resource required by the system in order to work[1]. Table 2 shows top words of some modelled domain aspects. The system admits more than one seed word per aspect or polarity, so we have performed additional experiments with different combinations and number of seed words, and the results remain quite stable. Using more than one seed word per aspect or polarity leads to better results.

Aspect/Polarity	Seeds (EN)	Seeds (ES)
food	chicken	pollo
service	service	servicio
ambience	ambience	ambiente
drinks	drinks	bebidas
location	location	ubicación
positives	excellent	excelente
negatives	horrible	horrible

Table 1: Example of seed words (one per topic) used to monitor certain aspects of restaurant reviews in several languages, including the general polarity seeds

Table 3 shows the result of the domain aspect classification experiments, compared to several baselines. *Naive Bayes* (NB) and *Multilayer Perceptron* (MLP) are supervised baselines trained using the manually labelled data from the datasets and evaluated using 10-fold cross validation. The *Majority* baseline uses the most frequent class. The *W2VLDA_NO* is the same topic model but without the hyperprior bias calculation using word embeddings. *W2VLDA* is the proposed method. Table 4 shows the results for analogous experiments for the sentiment polarity classification, compared against the same baselines.

[1] A list of stopwords is also required to obtain better results.

Dom. aspect	Aspect terms	Positives	Negatives
Food	chicken	moist	undercooked
	beef	goat	dry
	pork	smoked	drenched
	tuna	seared	soggy
Service	staff	helpful	inattentive
	workers	polite	rude
	employees	efficient	unfriendly
	chefs	prompt	wearing
Ambiance	lighting	modern	bad
	wall	beautiful	loud
	interior	chic	noisy
	vibe	nice	dark

Table 2: Top words for several domains aspects from customer restaurants reviews. Each domain aspect is split into three word distributions: aspect terms, positive words and negative words.

Domain aspect classification				
	EN	ES	FR	NL
NB	0.492	0.497	0.472	0.457
MLP	0.554	0.564	0.496	0.464
Majority baseline	0.333	0.333	0.333	0.333
W2VLDA_NO	0.313	0.374	0.356	0.315
W2VLDA	**0.781**	**0.633**	**0.586**	**0.473**

Table 3: Comparison of domain aspect classification results for several languages against several baselines.

Sentiment polarity classification				
	EN	ES	FR	NL
NB	0.672	0.577	0.587	0.563
MLP	0.711	0.602	0.583	0.577
Majority baseline	0.500	0.500	0.500	0.500
W2VLDA_NO	0.531	0.552	0.534	0.523
W2VLDA	**0.773**	**0.723**	**0.628**	**0.623**

Table 4: Comparison of sentiment polarity classification results for several languages against several baselines.

4 Conclusions

We propose W2VLDA, a weakly-supervised ABSA system that combines topic modelling and continuous word embeddings using only a few seed words, one per desired domain aspect and polarity. We evaluate it for different languages using the SemEval2016 dataset beating several supervised baselines.

References

Blei, D. M., A. Y. Ng, y M. I. Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.

Liu, B. 2012. Sentiment analysis and opinion mining. *Synthesis Lectures on Human Language Technologies*, 5(1):1–167.

Mikolov, T., K. Chen, G. Corrado, y J. Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *arXiv preprint arXiv:1301.3781*, páginas 1–12, Enero.

Pang, B. y L. Lee. 2008. Opinion mining and sentiment analysis. *Foundations and trends in information retrieval*, 2(1-2):1–135.

Pontiki, M., D. Galanis, H. Papageorgiou, I. Androutsopoulos, S. Manandhar, M. AL-Smadi, M. Al-Ayyoub, Y. Zhao, B. Qin, O. De Clercq, V. Hoste, M. Apidianaki, X. Tannier, N. Loukachevitch, E. Kotelnikov, N. Bel, S. M. Jiménez-Zafra, y G. Eryiğit. 2016. SemEval-2016 Task 5: Aspect Based Sentiment Analysis. En *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, páginas 19–30.