



HAL
open science

A comparative study of target-based and entity-based opinion extraction

Joseph Lark, Emmanuel Morin, Sebastián Peña Saldarriaga

► **To cite this version:**

Joseph Lark, Emmanuel Morin, Sebastián Peña Saldarriaga. A comparative study of target-based and entity-based opinion extraction. 18th International Conference on Computational Linguistics and Intelligent Text Processing (CICLing), Apr 2017, Budapest, Hungary. 10.1007/978-3-319-77116-8_16 . hal-01537896

HAL Id: hal-01537896

<https://hal.science/hal-01537896v1>

Submitted on 13 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A comparative study of target-based and entity-based opinion extraction

Joseph Lark and Emmanuel Morin and Sebastián Peña Saldarriaga

University of Nantes

LS2N CNRS 6004

Nantes, France

joseph.lark@univ-nantes.fr, emmanuel.morin@univ-nantes.fr,

sebastian@dictanova.com

Abstract. Opinion target extraction is a crucial task of opinion mining, aiming to extract occurrences of the different entities of a corpus that are subjects of an opinion. In order to produce a readable and comprehensible opinion summary, these occurrences are aggregated under higher order labels, or entities, in a second task. In this paper we argue that combining the two tasks, *i.e.* extracting opinion targets using entities as labels instead of binary labels, yields better results for opinion extraction. We compare the binary and the multi-class approaches on available datasets in English and French, and conduct several investigation experiments to explain the promising results. Our experiments show that an entity-based labelling not only improves opinion extraction in a single domain setting, but also let us combine training data from different domains to improve the extraction, a result that has never been observed on target-based training data.

1 Introduction

The field of sentiment analysis has attracted much interest over the recent years, and several frameworks have been proposed to tackle the challenges it presents. One of the main framework for sentiment analysis is aspect-based sentiment analysis (ABSA), which is particularly suited for the analysis of consumer reviews. This framework has been designed for summarizing points of interest (or *entities*) and causes of interest (or *aspects*) from every occurrence of opinion expression in a corpus. Consequently, the main subtasks in this framework include finding these occurrences of opinion (or *targets*) on the many subjects in the corpus and associating targets to an entity and an aspect. Initial works [1–3] on formalisation of the problem led to a binary annotation of the target extraction task, labelling as *target* a continuous span of text in a sentence representing an occurrence of a subject in an opinion expression. For instance, in the sentence “*The waiter is unfriendly but the menu is delicious*”, *waiter* and *menu* are opinion targets. Entity extraction and aspect extraction are then treated as additional classification tasks on the opinion targets. In the given example, *waiter* is an occurrence of the SERVICE entity and *menu* an occurrence of FOOD.

While this formulation of the opinion target extraction problem has helped tremendously on designing well performing systems, the binary annotation of targets can seem suboptimal for the quite complex language phenomenon of opinion. A known limitation of this formulation is that opinion target extraction is very sensible to the topic of the corpus, which is often referred to as domain specificity [1, 2, 4, 5]. We suggest that this limitation is in fact correlated to entity specificity and experiment a multi-class representation of the opinion target extraction task. To this end we use entities as target labels, in a manner similar to named entity recognition, where entities are labelled differently according to the concept they represent, but are consistent across domains.

After a brief review of related work (Section 2), we argue in this work in favour of a multi-class representation of opinion extraction over the current binary representation (Section 3). We compare extraction results on the SemEval ABSA datasets (Section 4), and first observe that an entity-based model improves the performance in a single domain setting. Moreover we find that in a cross-domain setting, where target-based opinion extraction learning has shown to be disadvantageous, an entity-based model improves the extraction (+1.68 points on F1 on average in English and +2.78 in French). Finally, we analyse opinion entities occurrences in the annotated data to explain these results (Section 5), and put forth that coherence of opinion words towards entities is critical for opinion extraction.

2 Related work

This work is related to the formalisation of the ABSA problem, and especially to the representation of opinion target occurrences and their associated semantic category.

Definitions for ABSA tasks are fairly recent. While basic definitions such as sentiment polarity, opinion words or opinion targets remained stable since initial work on the subject [6–10], the definition of an opinion entity has been regularly revisited. The core idea of opinion entity extraction is to consolidate all occurrences of opinion targets that refer to the same object (*e.g.* a phone), object feature (*e.g.* a phone screen size), or abstract notion (*e.g.* the price or practicality of a phone) under a unique label.

Hu & Liu, 2004 [7] define this task as the last step of opinion summarization, following the prior steps of entity (or *feature*) extraction, opinion word extraction and opinion orientation prediction. In their work, only explicit entities are extracted (*i.e.* opinion targets occurrences matching the entity term). Liu *et al.* 2005 [11] find implicit occurrences of entities by building a dictionary of variants from key entity terms, such as *weight* from *heavy* or *price* from *cost*. Kim & Hovy, 2006 [12] introduce the definition of an opinion *topic* as “an object an opinion is about”, which very much corresponds to what is most known in recent work as an opinion entity [2]. The approach for entity extraction presented in their work relates to ours as they first identify semantic roles, using semantic frames, to find opinion entities. However, these semantic roles are not specific to opinion

targets as in our approach. In a similar manner, Mukherjee & Liu, 2012 [13] suggest a topic modelling method to infer an opinion target entity (or *category*) from manually generated seed terms. Kobayashi et al. 2007 [14] formalise the opinion extraction task using a two-level hierarchy of opinion targeted objects, as in the current ABSA framework. Ding et al. 2008 [15] introduce the definition of an opinion entity as an identifiable concept (an object, person, event, etc.) in a taxonomy of components and for which a set of attributes can be defined. SemEval ABSA workshops [3, 16, 17] have largely contributed to the definition of the aspect-based sentiment analysis tasks and have led to the production of labelled data. In ABSA2014, entities are associated to sentences only and not opinion targets. In ABSA2015 and ABSA2016, one of the subtasks is to find associations of two types of semantic classes for each opinion target: opinion entities and opinion categories. In this definition of the opinion extraction, a word or multi-word expression can be labelled as *target*. A target is then an occurrence of the targeted entity, which may not share the same textual form as this particular occurrence. Finally, a category of opinion describes the precise aspect that is being criticised. Intuitively, an *entity* is therefore a concept that can be the subject of an opinion or not, while an *aspect* is a subjective attribute of this concept that calls for an opinion.

3 Entity-based opinion extraction

The task of opinion target extraction is to find occurrences of subjects towards which one expresses an opinion. To this end, a widely adopted approach is to consider that a subject can either be an opinion target or not. In particular, sentences such as the following are to be disambiguated :

- *We went to this restaurant based on prior internet comments.*
- *I was very disappointed with this **restaurant**.*

While both contain an occurrence of *restaurant*, only the occurrence of the second sentence is an opinion target. In order to summarize opinions of a corpus, existing works suggest to infer the opinion entity as an additional piece of information associated with the opinion target [3, 7, 11, 14, 16, 17]. We differ from this approach by directly extracting entities occurrences that are subject to an opinion. We question the need for binary target extraction and argue that, in a manner similar to named entity recognition, entity labels improve opinion extraction, assuming that these are coherently defined. Using existing concepts from previous formalisation works, we suggest the labelling of targets as entities rather than as a binary information. Despite being present in the literature this formalisation has, to the best of our knowledge, not been much studied from a opinion target extraction point of view. In particular, it has never been exploited to improve opinion target extraction or to tackle the domain adaptation problem, two applications we cover in this paper.

3.1 Coherence of opinion entities

In addition to the coherence of targets towards opinion entities, which usually share a hyperonymy dependency, the concept of entity in opinion target extraction is strongly related to the use of opinion words. Besides some very generic adjectives such as *good* and *bad*, opinion words are associated with specific types of opinion targets, which very often coincide with opinion entities [4, 5]. In the example shown in Section 1, *waiter* is associated with *unfriendly* and *menu* with *delicious*; the opposite associations seem highly improbable (*delicious waiter* or *unfriendly menu*). This linguistic coherence in opinion expressions towards each entity is to us a motivation to investigate entity-based opinion extraction.

3.2 Domain adaptation through opinion entities

In addition to a finer-grained extraction, we see in the annotation of opinion targets using entity labels an opportunity to tackle the domain adaptation problem in the context of opinion target extraction with a novel approach. Indeed, a domain can be defined as a set of entities, each entity being a label for opinion target extraction. Using this formalisation let us use the fact that different domains can share some entities, thus possibly sharing training data. This approach differs from existing works on domain adaptation in the sense that we do not adapt a closed and well-defined first model (specific or general) to another domain. Our hypothesis is that each domain in the context of opinion mining in user reviews is composed of several entities that can be shared across domains. When building an opinion target extraction model for a new domain, the domain adaptation task could thus be shifted to an identification task of the entities that compose the new domain. The new model would benefit in training from previously annotated data in a modular manner, as pictured on Figure 1.




Fig. 1: Illustration of entity modularity : as several entities are shared across domains (in this example, Restaurant, Hotel and Museum), training data could be mutualised.

4 Experiments and results

In these experiments, we compare a target-based annotation and an entity-based annotation for the task of opinion target extraction on the English and French SemEval ABSA datasets. We first describe the datasets and the extraction method we use in our experiments, and provide results for single domain and cross-domain settings.

4.1 Datasets

We conducted our study on customer reviews datasets from the SemEval ABSA workshops¹ [16, 17]. These include restaurant and hotel reviews in English, and restaurant and museum reviews in French. Each dataset was annotated by a native linguist, who indicated for each sentence offsets of opinion targets and associated entities. Additional information on the datasets is shown in Table 1.

Corpus	Reviews	Sentences	Targets
Restaurants (English)	440	2,676	2,529
Hotels (English)	30	266	264
Restaurants (French)	455	2,427	2,484
Museums (French)	162	687	582

Table 1: Number of reviews, sentences and targets for each corpus.

The datasets are very relevant for our study as these cover different domains sharing some of their entities, as shown in Table 2. Indeed, this configuration let us demonstrate the usefulness of cross-domain entities in opinion target learning. Besides SemEval ABSA workshops, the datasets have been used as evaluation material for several works. However, to the best of our knowledge, there is no existing work on comparing annotations for opinion extraction.

Restaurants (English and French)	Hotels (English only)	Museums (French only)
AMBIENCE, DRINKS*, FOOD*, LOCATION , RESTAURANT, SERVICE	FACILITIES, FOOD AND DRINKS*, HOTEL, LOCATION , ROOMS, ROOMS AMENITIES, SERVICE	COLLECTIONS, FACILITIES, LOCATION , MUSEUM, SERVICE , TOUR GUIDING

Table 2: Opinion entities for each corpus. Shared entities are indicated in bold, and starred entities are related but don't share the same label.

¹ <http://alt.qcri.org/semEval2016/task5/>

4.2 Extraction method

Following Jakob & Gurevych, 2010 [1] and similar works that have performed well in the SemEval ABSA workshops, we train a Conditional Random Fields [18] model for target and entity extraction. Both extractions are formulated as sequence labelling tasks, and only differ from one another by the nature of the annotation: while target labels are binary, entities are annotated following a multi-class labelling, as shown on Figure 2. We use the CRF++² toolkit for our experiments, with a segmentation on sentences. Features for each word entry include the word unigram, word bigram, part-of-speech tag and lemma of the preceding, current and following word.

Word	Lemma	POS tag	Target label	Entity label
Excellent	excellent	ADJ	0	0
atmosphere	atmosphere	NOUN	Target	AMBIENCE
,	,	.	0	0
delicious	delicious	ADJ	0	0
dishes	dish	NOUN	Target	FOOD
good	good	ADJ	0	0
and	and	CON	0	0
friendly	friendly	ADJ	0	0
service	service	NOUN	Target	SERVICE
.	.	.	0	0

Fig. 2: Annotation example on a sentence from the English train corpus. The 0 label represents the Outside class in both target-based and entity-based annotations.

We use a simple, class descriptive annotation (Target/Outside for target-based extraction and *Entity name*/Outside for entity-based extraction) instead of the often used BIO format, as early results indicated a better performance using simpler labels. Breck *et al*, 2007 [19] made a similar observation for opinion expression extraction, and pointed that the absence of contiguous annotations, as in our case, could explain the fact that the BIO format does not shape best the labelled data. Finally, we resolve ambiguous cases for the multi-class scenario, *i.e.* when probability of outside class is less than 0.5, by selecting the most probable entity class.

4.3 Single domain opinion extraction

In order to compare a target-based and an entity-based opinion extraction, we train two distinct CRF models on the same train corpus, namely the Restaurants reviews corpus, and use the same features; only labels were replaced to compare the results from the extractions. We conduct this experiment on the two languages for which this type of annotation is available, English and French.

² <https://taku910.github.io/crfpp/>

Train corpus	Precision†	Recall†	F1
Targets	67.60	62.23	64.81
Entities	68.43	62.88	65.54

Table 3: Target-based and entity-based opinion extraction on English Restaurants reviews. The p -value for precision is $2.48e-3$ and $3.91e-3$ for recall.

Train corpus	Precision	Recall	F1
Targets	74.27	59.25	65.92
Entities	74.64	60.03	66.55

Table 4: Target-based and entity-based opinion extraction on French Restaurants reviews. The p -value for precision is $3.31e-1$ and $3.29e-1$ for recall.

Comparison of the results, reported in Tables 3 and 4, shows that the entity annotation enhance both precision (+0.83 percentage points in English, +0.37 pp in French) and recall (+0.65 pp in English, +0.78 pp in French) for opinion target extraction. Significance testing using a t-test showed extraction in English to be significant, but less so in French. Nonetheless, the closeness of the results between the two extractions questions the need for a target extraction step, as the end goal. Intuitively, this result supports the hypothesis of a linguistic coherence of entities over opinion targets, in other words that entity labels help disambiguate the target extraction more than they add ambiguity. The fact that this behaviour can be observed on both languages also favours this idea.

4.4 Cross-domain opinion target learning

Using an identical framework, we now want to compare target-based and entity-based opinion extraction in a cross-domain setting. To this end, we train both models using additional out-of-domain data from the SemEval ABSA dataset. As described in Section 4.1, such out-of-domain data include hotel reviews for the English corpus and museum reviews for the French corpus. Results shown in Tables 5 and 6 demonstrate best the usefulness of an entity-based annotation over a target-based annotation. On one hand, we can see on line 1 that when adding target-based training data from another domain, the extraction is generally less performing.

Train corpus	Precision	Recall	F1
Targets (R)	67.60	62.23	64.81
Targets (R+H)	67.25	61.91	64.47
Entities (R)	73.84	51.70	60.81
Entities (R+H)	74.37	52.67	61.66

Table 5: Target and entity cross-domain opinion extraction on English Restaurants (R) and Hotels (H) datasets.

Train corpus	Precision	Recall	F1
Targets (R)	74.27	59.25	65.92
Targets (R+M)	72.69	58.33	64.72
Entities (R)	74.64	60.03	66.55
Entities (R+M)	73.34	61.57	66.94

Table 6: Target and entity cross-domain opinion extraction on French Restaurants (R) and Museums (M) datasets.

On the other hand, adding entity-based out-of-domain training data yields opposite results, as we can see on line 2. This tends to confirm that training data on shared entities improve the extraction, or that differentiating exclusive entities help disambiguate non relevant contexts. Only precision in the case of the French corpus is lower in the cross-domain setting, which may be due to the fact that museum reviews are less similar to restaurants ones than hotels reviews. To further investigate these results, we run an entity-by-entity opinion extraction on single and cross-domain datasets. When analysing results of this extraction, as it can be seen on Tables 7 and 8, we can observe that F1 for these entities is significantly better (+6.5 pp for LOCATION and +3.18 pp for SERVICE in English datasets, +5.36 pp for LOCATION and +1.97 pp for SERVICE in French datasets).

Train corpus	Restaurants			Restaurants + Hotels			Gain (pp)
Entity	Precision	Recall	F1	Precision	Recall	F1	
AMBIENCE	76.6	61.02	67.92	80.43	62.71	70.48	+2.56
DRINKS	78.95	41.67	54.55	82.35	38.89	52.83	-1.72
FOOD	67.10	47.69	55.76	68.42	48.00	56.42	+0.66
LOCATION	100.00	40.00	57.14	58.33	70.00	63.64	+6.50
RESTAURANT	58.97	28.05	38.02	59.46	26.83	36.97	-1.05
SERVICE	78.95	69.44	73.89	81.44	73.15	77.07	+3.18

Table 7: Single domain and cross-domain entity learning on the English Restaurants and Hotels reviews dataset.

Evolution of entities that are exclusive to the restaurant domain is less trivial. Not only results from the cross-domain model can be better or worse than those of the single domain model, but in this case differences are not consistent across languages. For instance, F1 for the RESTAURANT entity in the English dataset has decreased (-1.05 pp), while it displays a consistent improvement (+2.6 pp) in the French dataset.

Train corpus	Restaurants			Restaurants + Museums			Gain (pp)
Entity	Precision	Recall	F1	Precision	Recall	F1	
AMBIENCE	86.21	66.67	75.19	92.00	61.33	73.60	-1.59
DRINKS	73.33	32.35	44.90	72.22	38.23	50.00	+5.10
FOOD	65.37	53.00	58.54	72.96	53.62	61.81	+3.27
LOCATION	60.00	13.64	22.22	57.14	18.18	27.58	+5.36
RESTAURANT	56.45	51.47	53.85	62.50	51.47	56.45	+2.60
SERVICE	87.90	80.15	83.85	92.37	80.14	85.82	+1.97

Table 8: Single domain and cross-domain entity learning on the French Restaurants and Museums reviews dataset.

5 Opinion entity analysis

In this section we conduct a study on coherence of opinion entities in the restaurants reviews datasets to best explain our results.

5.1 Target terms and opinion words

We analyse the coherence of entity labels for opinion target extraction by measuring the number of target terms and opinion words by entity and across entities, as these two types of lexical elements are characteristic of the expressed opinion [2]. Coherence for each entity is here represented by the uniqueness of these elements (columns 3 and 6 of Tables 9 and 10) and coherence across entities is measured by their exclusiveness for a given entity (columns 4 and 7 of Tables 9 and 10).

Entity	#Target terms	#Unique targets	#Exclusive targets	#OW	#Unique OW	#Exclusive OW
AMBIENCE	287	115 (40.07%)	99 (86.09%)	165	85 (51.52%)	33 (38.82%)
DRINKS	133	59 (44.36%)	58 (98.31%)	48	23 (47.92%)	4 (17.39%)
FOOD	1,311	541 (41.27%)	538 (99.45%)	776	193 (24.87%)	105 (54.40%)
LOCATION	32	16 (50.00%)	10 (62.50%)	18	15 (83.33%)	3 (20.00%)
RESTAURANT	343	119 (34.69%)	99 (83.19%)	214	90 (42.06%)	32 (35.56%)
SERVICE	432	62 (14.35%)	56 (90.32%)	250	122 (48.80%)	59 (48.36%)

Table 9: Target terms and opinion words (OW) by entity in the English dataset.

Entity	#Target terms	#Unique targets	#Exclusive targets	#OW	#Unique OW	#Exclusive OW
AMBIENCE	253	42 (16.60%)	32 (76.19%)	124	69 (55.65%)	29 (42.03%)
DRINKS	123	47 (38.21%)	46 (97.87%)	43	33 (76.74%)	4 (12.12%)
FOOD	1,248	400 (32.05%)	392 (98.00%)	540	231 (42.78%)	136 (58.87%)
LOCATION	56	28 (50.00%)	20 (71.43%)	25	19 (76.00%)	6 (31.58%)
RESTAURANT	247	42 (17.00%)	25 (59.52%)	127	75 (59.06%)	30 (40.00%)
SERVICE	536	49 (9.14%)	46 (93.88%)	321	150 (46.73%)	78 (52.00%)

Table 10: Target terms and opinion words (OW) by entity in the French dataset.

The main observation here seems to be the fact that measures are consistent on both languages. Indeed, relative order of entities with regards to uniqueness and exclusivity of target terms as well as opinion words is very similar in English and French, despite the fact that entities were annotated on different datasets and by different experts. This is a strong argument towards an entity-based representation of opinion target as it tends to show a coherent conceptual coherence in entities in addition to the sheer homogeneity of target terms for each language. However, understanding the relation between coherence and extraction

performance is not trivial as uniqueness and exclusivity are not correlated. For instance, the entity FOOD is represented by a large number of target terms – mainly descriptions of the different dishes – that are highly exclusive to this entity, while LOCATION is represented by a small number of target terms, including *restaurant* or *place* that are shared by other entities such as RESTAURANT or AMBIENCE. When crossing these measures with the per-entity evaluation of opinion extraction (Tables 7 and 8), we can see that entities that are best recognised, namely AMBIENCE and SERVICE, are those combining a high rate of exclusive opinion words and small number of unique target terms. In next experiments, we investigate how this coherence affects opinion extraction learning by conducting entity-by-entity active learning iterations.

5.2 Opinion entity learning

We analyse learning iterations from target-based and entity-based training datasets through evaluations at the target level and at the entity level. Batches of 50 sentences are sampled from the training part of the restaurants reviews corpus in English and French. We used the uncertainty sampling strategy [20], *i.e.* we added on each iteration the 50 sentences with the lowest global output sequence probability, using annotations provided in the SemEval datasets.

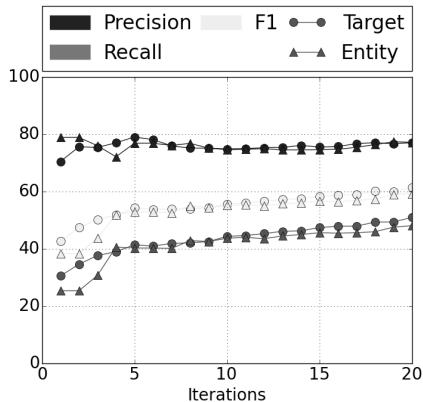


Fig. 3: Learning curves for target extraction in the English dataset using target and entity-based labelling.

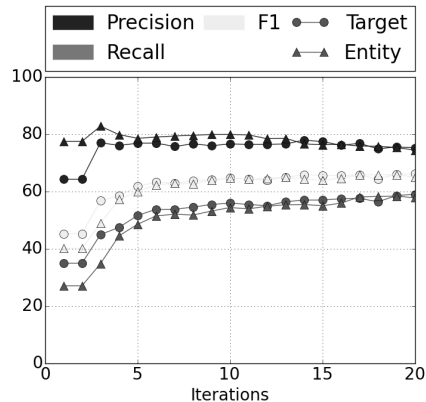


Fig. 4: Learning curves for target extraction in the French dataset using target and entity-based labelling.

Although a possible drawback of a multi-class framework is an increased need for training examples, evaluation at the target level shows that the entity-based model quickly converge to a learning curve identical to the one of the target-based model. As it can be seen on Figures 3 and 4, the entity-based model starts from a lower recall and a higher precision than the target-based model on the initial batch of reviews, and stabilises before the first five iterations.

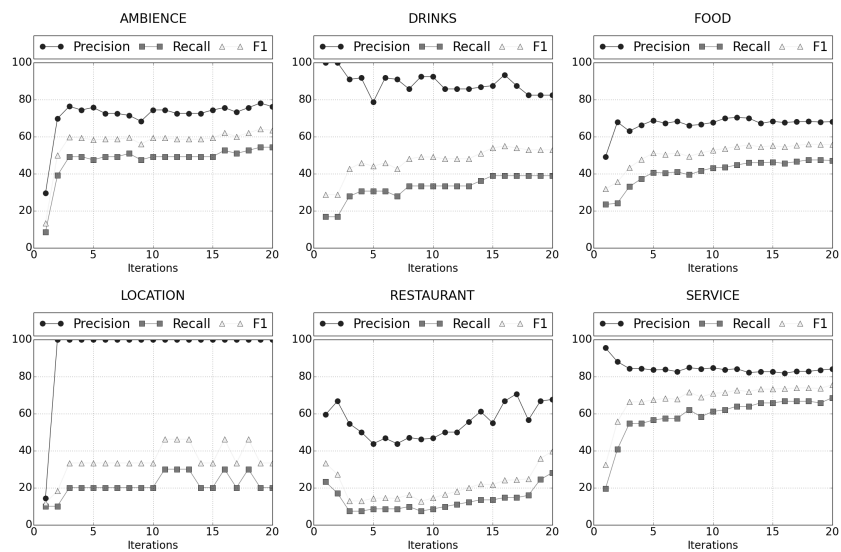


Fig. 5: Learning curves for entity extraction in the English Restaurants dataset.




Fig. 6: Learning curves for entity extraction in the French Restaurants dataset.

Evaluation at the entity level, displayed on Figures 5 and 6, shows that learning can be even faster for entities such as AMBIENCE and SERVICE, which we previously highlighted for their high coherence. Learning curves for other entities, LOCATION, DRINKS and RESTAURANT, are very chaotic. Again these results are surprisingly similar for both languages. From the observation of lexical elements associated with opinion entities shown on Tables 9 and 10, it seems that opinion word coherence impacts the most opinion entity learning. Indeed the common factor of the three entities for which learning is fast and steady (AMBIENCE, FOOD and SERVICE) is a high rate of exclusive opinion words, whereas other metrics are less conclusive. Specifically, the rate of exclusive target terms does not appear as important as we assumed. An example of this observation is the fact that the AMBIENCE entity present a stable learning curve in spite of a relatively low rate of exclusive target terms.

6 Conclusion

In this paper we compare a target-based and an entity-based annotation for opinion extraction. From an initial intuition that the complex nature of opinion expression in language requires a fine-grained labelling, we investigate how this is depicted on real data. We use available customer reviews datasets in English and French, labelled on opinion targets and their associated entities. Our experiments show that an entity-based labelling not only improves opinion extraction in a single domain setting, but also let us combine training data from different domains to improve the extraction, a result that has never been achieved on target-based training data.

Elements as to why entity annotation improves opinion extraction are strongly related to the coherence of elements in the lexical context of opinion targets. We found that the exclusivity of opinion words towards an entity is directly linked to the capacity of the model to correctly learn to recognise its occurrences. The exclusivity of target terms representing an entity contribute to a lesser extent to the quality of the learning process. In our observations, this metric is only correlated to the convergence rate of the learning curve.

In our sense, these observations are particular signs of a need for a larger framework. In a manner similar to named entity recognition, where relevant items are defined by distinct categories (person, location, company, *etc.*), we see in entity-based opinion the opportunity to build a multi-class model for opinion extraction based on the opinion linguistic context rather than on the domain of the analysed corpus. A great advantage of a model of this kind would be to ease the domain adaptation problem. While target-based opinion extraction is very sensitive to the domain of the dataset it is trained on, we demonstrate in our experiments that an entity-based opinion extraction model could benefit from training data of multiple domains. In future works, we will investigate how multiple domains can be covered using this framework.

References

1. Jakob, N., Gurevych, I.: Extracting Opinion Targets in a Single- and Cross-Domain Setting with Conditional Random Fields. In: Proceedings of EMNLP 2010, Cambridge, MA, USA (2010) 1035–1045
2. Liu, B.: Sentiment Analysis and Opinion Mining. Synthesis Lectures on Human Language Technologies (2012)
3. Pontiki, M., Galanis, D., Pavlopoulos, J., Papageorgiou, H., Androutsopoulos, I., Manandhar, S.: Semeval-2014 task 4: Aspect based sentiment analysis. In: Proceedings of SemEval 2014, Dublin, Ireland (2014)
4. Li, F., Pan, S., Jin, O., Yang, Q., Zhu, X.: Cross-domain co-extraction of sentiment and topic lexicons. In: Proceedings of ACL 2010, Uppsala, Sweden (2012) 410–419
5. Hamilton, W.L., Clark, K., Leskovec, J., Jurafsky, D.: Inducing domain-specific sentiment lexicons from unlabeled corpora. arXiv preprint arXiv:1606.02820 (2016)
6. Hatzivassiloglou, V., McKeown, K.R.: Predicting the semantic orientation of adjectives. In: Proceedings of ACL 1997, Madrid, Spain (1997) 174–181
7. Hu, M., Liu, B.: Mining and summarizing customer reviews. In: Proceedings of KDD 2004, Seattle, WA, USA (2004) 168–177
8. Kim, S.M., Hovy, E.: Determining the sentiment of opinions. In: Proceedings of COLING 2004, Geneva, Switzerland (2004) 1367
9. Bethard, S., Yu, H., Thornton, A., Hatzivassiloglou, V., Jurafsky, D.: Automatic extraction of opinion propositions and their holders. In: Proceedings of AAAI 2004, Palo Alto, CA, USA (2004)
10. Popescu, A.M., Etzioni, O.: Extracting Product Features and Opinion from Reviews. In: Proceedings of HLT-EMNLP 2005, Vancouver, Canada (2005) 339–346
11. Liu, B., Hu, M., Cheng, J.: Opinion observer: analyzing and comparing opinions on the web. In: Proceedings of WWW 2005, Chiba, Japan (2005) 342–351
12. Kim, S.M., Hovy, E.: Extracting opinions, opinion holders, and topics expressed in online news media text. In: Proceedings of SST 2006, Sydney, Australia (2006)
13. Mukherjee, A., Liu, B.: Aspect extraction through semi-supervised modeling. In: Proceedings of ACL 2012, Jeju Island, Korea (2012) 339–348
14. Kobayashi, N., Inui, K., Matsumoto, Y.: Opinion mining from web documents: Extraction and structurization. In: Transactions of the Japanese Society for Artificial Intelligence. (2007) 227–238
15. Ding, X., Liu, B., Yu, P.S.: A holistic lexicon-based approach to opinion mining. In: Proceedings of WSDM 2008, Palo Alto, CA, USA (2008) 231–240
16. Pontiki, M., Galanis, D., Papageorgiou, H., Manandhar, S., Androutsopoulos, I.: Semeval-2015 task 12: Aspect based sentiment analysis. In: Proceedings of SemEval 2015, Denver, CO, USA (2015) 486–495
17. Pontiki, M., Galanis, D., Papageorgiou, H., Androutsopoulos, I., Manandhar, S., Al-Smadi, M., Al-Ayyoub, M., Zhao, Y., Qin, B., De Clercq, O., Hoste, V., Apidianaki, M., Tannier, X., Loukachevitch, N., Kotelnikov, E., Bel, N., Jiménez-Zafra, S.M., Eryigit., G.: Semeval-2016 task 5: Aspect based sentiment analysis, San Diego, CA, USA (2016)
18. Lafferty, J.D., McCallum, A., Pereira, F.C.N.: Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: Proceedings of ICML 2001, San Francisco, CA, USA (2001) 282–289
19. Breck, E., Choi, Y., Cardie, C.: Identifying expressions of opinion in context. In: IJCAI. Volume 7. (2007) 2683–2688
20. Lewis, D.D., Gale, W.A.: A sequential algorithm for training text classifiers. In: Proceedings of SIGIR 1994, Dublin, Ireland (1994) 3–12