Lechuzo: Weakly-Supervised System for Fine-Grained Sentiment Analysis

Nir Ofek, Lior Rokach

Ben-Gurion University of the Negev, Israel

{nirofek,liorrk}@bgu.ac.il

Abstract. Traditional sentiment analysis approaches suffer from two major drawbacks: coarse granularity - polar opinions can co-occur even within the same sentence, and ambiguity - opinion-bearing terms can convey polar sentiment in different contexts. Consider the following laptop review: "the big plus was a large screen but having a large battery made me change my mind," where polar opinions co-occur in the same sentence, and the opinion term that describes the opinion targets ("large") encodes polar sentiments: a positive for *screen*, and a negative for *battery*. To parse these differences, our approach is to identify opinions with respect to the specific opinion targets, while taking to take the context into account. Moreover, considering that there is a problem of obtaining an annotated training set in each context, our approach trains unsupervised model and cascading to a weakly-supervised model where the ground truth opinion for target aspects is never given.

Keywords: fine-grained sentiment analysis, opinion mining.

1 Introduction

With the proliferation of user-generated-content (UGC) on websites, subjective information in the form of reviews, blogs, and bulletin boards becomes widely available and accessible. Such information constitutes fertile ground for sentiment analysis which can be highly useful for decision-making.

A traditional sentiment analysis approaches aim to extract sentiment at the document or sentence level. However, this approach has two major drawbacks:

1. Granularity - polar opinions can co-occur even within the same sentence. Consider the following statement: the screen is ideal, the battery is great, but the price is unaffordable. A positive sentiment is conveyed for both screen and battery; however a negative sentiment is conveyed for price.
2. Ambiguity - in lexical or concept-based approaches, opinion-bearing terms can convey polar sentiment in different contexts. For example, the adjective small may convey positive sentiment when is modifying the noun device, whereas negative sentiment is conveyed when the noun meal is modified.

adfa, p. 1, 2011.

The proposed system focuses on calculating sentiment at the aspect level, rather than by extracting the overall sentiment of a document or a sentence. An aspect refers to a set of terms that relate to a certain topic in a certain context. The system alleviates the first drawback mentioned above, since it is provided with the set of terms that represent the governor (or target) aspect – on which the sentiment will be calculated.

Adjectives, which are words that describe or modify other elements in a sentence, and are frequently used to directly convey facts and opinions about the nouns they modify. As such, they are the backbones of our system; therefore, this paper elaborates mainly on disambiguating the polarity of adjectives across aspects. The proposed system involves calculating for every adjective a sentiment score with respect to each aspect, to constitute an aspect-specific lexicon. Moreover, we use WordNet to further expand the lexicon after the more frequent (and hopefully important) adjectives have been detected.

Since sentiment is often conveyed in a latent or more complex way, the system is able to identify concepts and expand the identified set by using SenticNet 3, and to disambiguate their polarity in the relevant context, mainly by using the adjective lexicons. For example, our system can successfully predict the sentiment in the sentence "*the pool looks large*" or "*I just love the view*" and to associate it to the relevant aspect.

In order to address the issue of obtaining an annotated training set for each aspect, the suggested learning approach for learning adjectives' polarity for every aspect is unsupervised. This approach ensures that the method is not bounded to any specific domain or website.

To summarize, our method has the following properties: (1) it can be trained with unsupervised data, (2) it can determine an adjective's polarity with respect to the target aspect, (3) for each sentiment score of an aspect, it is able to display the adjectives for which the sentiment was computed, to explain the result and provide pros and cons, and (4) the predicted sentiment is in the range of [0:1] to reflect how positive (or negative) the sentiment is, in contrast to deterministic decision which usually provides the sentiment score as either negative or positive.

In addition, in case that the overall rating is available for training, the system can utilize this information to learn a fallback model.

We evaluate our method using the reviews dataset of Wang et al. (2010). A demonstration of the system, only as a proof-of-concept, is given on www.iofek.com; however the full capacity of the system is not yet demonstrated online.

2 Description

The system starts with identifying important aspects in the text, defined by repeating nouns that often bearing-opinion adjectives are related to. In the next step aspects are clustered to a single topical aspect, i.e., each topical aspect will be represented by a set of aspects. For example, the sentiment of the topical aspect *room* can be calculated as the average sentiment of the aspects: room, bed, bathroom and view.

Once identified, the system aims to learn the polarity score of adjectives according to each aspect. That is due to the propensity of people to convey sentiment through adjectives and since this information (adjectives and their polarity per aspect) is used to derive the sentiment of more complex concepts that the system identified or extracted by using SenticNet 3. We further elaborate mainly on learning the polarity of adjectives.

The process of generating aspect-specific lexicon is presented in Algorithm 1. The algorithm is an iterative process that starts with a seed lexicon and expands it to construct an aspect-specific sentiment lexicon. In each iteration we start with the current aspect-specific lexicon, and by processing a set of relevant reviews, we search for new adjectives that are not in the lexicon and modify (i.e., are connected to) the aspect. For the last constraint, we suggest a classification approach to connect adjectives to the nouns they modify. A new adjective is added to the lexicon only if it is connected with a conjunction to anther adjective which is in the lexicon. In this case, the polarity of the new adjective is derived by considering the conjunction pattern and the polarity of the known adjective. The input to the algorithm includes the following:

- A seed lexicon (*SL*) - a set of adjectives paired with their corresponding polarity to reflect how positive/negative each adjective is (1 for positive and 0 for negative). The polarity paired with each adjective pertaining to this lexicon should not be dependent on the context (domain and aspect). Therefore, the polarity of these adjectives is set as a-prior convention. For example, the polarity of the adjectives excellent and amazing should always be positive. Two classes of adjectives must be excluded from the seed lexicon: ambiguous adjectives (such as great which may be very good or big) and adjectives that are used to express polar sentiment in different contexts (such as big which can be negative to describe a device or positive in the context of the description of a meal).
- Aspect (*A*) - the noun entity for which we create the lexicon. Examples of aspects include: the laptop battery, the food in the hotel, or the size of a camera. In fact, an aspect can be represented by a set of topical words.
- Reviews (*R*) - a set of reviews and opinions that are relevant to aspect A.
- Conjunction patterns (*C*) - a set of conjunctions (for example, *and*) and their polarity patterns. The polarity pattern is in the form of $<p_1, c, p_2>$ where p_1 is the polarity of the first adjective, c is the conjunction and p_2 is the polarity of the second adjective. For example <positive, and, positive> indicate that for the conjunction *and*, given that the polarity of the first adjective is positive (i.e., p_1 is positive), the polarity of the second adjective is positive as well. Note that this represents an aspect-specific set of conjunction patterns.

The output of Algorithm 1 is an extended set of aspect-specific lexicons which includes the seed lexicon as well.

The process of creating the aspect-dependent lexicon is as follows. First, the extended lexicon of aspect A (EL_A) is initialized with the seed lexicon (*SL*) (line 1). Then, the following steps are repeated n times (n is a configurable parameters) (line 2). For each review $r_i \in R$ (line 4) we identify all adjectives (line 5). Then, for each identified adjective a we check if a is modifying aspect A (i.e., a is related to A); if true, a is added to the modifying aspect set (*ModAdj*) (lines 7-9). Then, for each pair of modifying aspects a_1 and a_2 in *ModAdj*, we check whether the two aspects are connected with a conjunction in review r_i (lines 10-11).

If a_1 and a_2 are connected with a conjunction c, then, if one of the two conjunctions (let's assume a_1 – without loss of generality) is in the current extended lexicon (EL_A), and the second adjective (a_2) is not in EL_A, we add a_2 to a temporary set of adjectives (*NewAdj*). *NewAdj* is set as an empty set in each iteration and holds all new adjectives that were identified in the current iteration (i.e., after iterating over all reviews in R). The polarity of a_2 is set according to the conjunction pattern of $c \in C$ (lines 12-14). At the end of each iteration all new adjectives (*NewAdj*) are added to the extended lexicon of A (EL_A) (line 15) with their corresponding scores. For each new adjective a2, its synonyms are added to the lexicon with the same score, whereas antonyms are added to the lexicon with polar score. The set is retrieved from the WordNet graph, which is a publicly available and lexical rich resource.

ALGORITHM 1. Constructing Lexicon

Input: $SL = \{<adj_k, pol_k> \mid \forall k \; pol_k = \begin{cases} 1, if\ positive \\ 0, if\ negative \end{cases}\}$

$R = \{r_1, \ldots, r_n\}$

A – the entity for which we create the lexicon

C – a list of known conjunctions

Output: EL_A – the extended lexicon

1. $EL_A = SL$
2. **for each** $iter = 1 \ldots n$
3. $NewAdj = \emptyset$
4.   **for each** $r_i \in R$
5. $Adj \leftarrow$ findAdjectives(r_i)
6. $ModAdj = \emptyset$
7.   **for each** $a \in Adj$
8.   **if** isModifyingAspect(A,a) **then**
9. $ModAdj \leftarrow ModAdj \cup \{a\}$
10.   **for each** $<a_1,a_2> \in ModAdj$ x $ModAdj$
11. $conj \leftarrow$ extractConjunction(a_1,a_2,r_i)
12.   **if** $a_1 \in EL_A$ **and** $a_2 \notin EL_A$ **and** $conj \in C$ **then**
13. $p_2 \leftarrow$ getPolarity(A, a_1,$conj$)
14.   updatePolarity($NewAdj$, a_2, p_2)
15. $EL_A \leftarrow EL_A \cup NewAdj$
16. **return** EL_A

After the complete adjective expansion process is done we employ a statistical test. The test is conducted for each adjective a_2, and takes into account co-occurring adjectives from the previous iteration, to dispute the hypothesis that the polarity score computed for a_2 is drawn from a set of adjectives whose polarity scores distribute randomly. Hence, we withdraw some non-informative adjectives from the lexicons.

To this end, the polarity of the adjectives that are modifying the target aspect can be used to calculate its sentiment. This approach can obtain a relatively high precision rate. As a result, in some cases still the target aspect does not have any modifying adjectives, or the modifying adjective does not include in the aspect's lexicon. Aiming to increase recall, we employ a cascade approach to utilize concepts that are mainly extracted by using SenticNet 3, as fallbacks. To disambiguate the polarity of the concepts, we are using the polarity of the participating adjectives. In case of the concept adjectives do not exist in the lexicon or there are no adjectives, the polarity of concepts will be determined based on the average overall rating of reviews that contain the concept, in our training set.

The system can output an overall sentiment for a given sentence, based on the calculated sentiment for each aspect in the sentence.

3 Evaluation

The evaluation was conducted on data that was collected from TripAdvisor.com – one of the major hotel rating and reviews websites. Since this website is very popular, it contains a large number of reviews, which is important for the training phase. Moreover, in addition to the textual content TripAdvisor provides multiple user ratings for each review. Each review can have an overall rating, as well as ratings for seven aspects: value, room, location, cleanliness, check-in/front desk, service, and business services, in the range of 1 to 5 stars. This feature of TripAdvisor is important to evaluate the results. Since our method predicts the sentiment per aspect, user rating, once provided, will be regarded as ground-truth.

The training set is obtained from Wang et al. (2010) and contains nearly 240,000 hotel reviews collected in a one month period (from February 14, 2009 to March 15, 2009). We compute the MSE (Mean Square Error), MAE (Mean Absolute Error) and MAP@10 to measure the performance. Since it is use to summarize sentiment per aspects, we evaluate our method per hotel and aspect; i.e., the predicted scores are aggregated per hotel by averaging the aspect's scores across all of the hotel reviews.

4 Discussion

The main models of the system are trained with unsupervised data. Only as a fallback method we utilize the overall score for a weakly-supervised learning. This approach was taken since the overall score is coarse, while we wish to associates sentiment with every aspect. Consider the following review taken from Tripadvisor.com, rated as '*terrible*' (1 of 5 points): *"Nice kitchenette, good location next to Museum station. Aircon unit is standalone and controls fully adjustable*". No doubt that the overall rating is not in accordance with the text. A conclusive overall score cannot take into consideration divergent opinions.