Panicheva, P.; Cardiff, J.; Rosso, P. (2013). Identifying subjective statements in news titles using a personal sense annotation framework. Journal of the American Society for Information Science and Technology. 64(7):1411-1422. doi:10.1002/asi.22841.

Identifying Subjective Statements in News Titles Using a Personal Sense Annotation Framework

Polina Panicheva, John Cardiff
Department of Computing
Institute of Technology Tallaght, Dublin, Ireland
email: polina.panicheva, john.cardiff@ittdublin.ie

Paolo Rosso
Natural Language Engineering Lab. - ELiRF
Universidad Politecnica de Valencia, Spain
email: prosso@dsic.upv.es

Corresponding author: John Cardiff, john.cardiff@ittdublin.ie

Abstract

Subjective language contains information about private states. The goal of subjective language identification is to identify that a private state is expressed, without considering its polarity or specific emotion. A component of word meaning, "Personal Sense", has clear potential in the field of subjective language identification, as it reflects a meaning of words in terms of unique personal experience and carries the personal characteristics.

In this paper, we investigate how Personal Sense can be harnessed for the purpose of identifying subjectivity in news titles. In the process, we develop a new Personal Sense annotation framework, for annotating and classifying subjectivity, polarity and emotion. The Personal Sense framework yields a high performance in a fine-grained sub-sentence subjectivity classification. Our experiments demonstrate lexico-syntactic features to be useful for the identification of subjectivity indicators and the targets which receive the subjective Personal Sense.

Introduction

Subjective language is language containing information about private states, i.e. opinions and emotions (Wiebe et al, 2004). The goal of subjective language identification is to identify that a private state is expressed, without going into detail about its polarity or its specific emotion. On one hand, it is a preliminary stage in opinion mining: before identifying an opinion as positive or negative, it is necessary to identify it as an opinion, as opposed to a fact description etc. Furthermore, it may serve as a technique for separating facts from points of view, classifying opinionated text and identifying ideological perspective of the author.

Subjectivity has been defined in (Wiebe, 1994) as private states, i.e. "states … not open to objective observation or verification" (Pang & Lee, 2008). This includes opinions, emotions, moods etc. Research in subjectivity analysis has increased significantly in recent years, due largely to the vast growth and availability of personal texts in the blogosphere. In polarity classification, authorship attribution, background characteristics identification of authors, and in the basic subjectivity identification, various features have been proposed and high results achieved. However, it is widely overlooked that a personal component of meaning - Personal Sense (Leontev, 1978) - is a very important feature of any subjective text. Moreover, it forms unique idiolect features and reflects personal preferences in text; although in theory its role in subjective language research is obvious, as it is defined as a former of subjective consciousness.

Leontev (Leontev, 1978) stated that consciousness is subjective, and defined two types of word-meaning: significance, being the meaning shared by the speakers of a language and representing a part of the objective reality, and Personal Sense, representing subjective characteristics in consciousness, in terms of unique experience of a person. Thus, Personal Sense, serving as a building block for the subjective consciousness, can be harnessed from the writings of bloggers, in order to more accurately deduce information about their opinions, private states and sentiments.

By way of illustration, consider the following examples, taken from a debate (“The Green Line”, sourced from bitterlemons.org) about establishing a border between Israel and Palestine

The "green line" is invisible, undocumented and unfounded in international law[...] it sets a precedent of substituting principles of international law with agreements signed under duress. (Example 1)

Despite these trans-boundary movements, the line remained an important point of separation between the two territories.[...] the green line-with some minor deviations-has the greatest likelihood of constituting the formal international boundary between two independent states. (Example 2)

Both pieces contain information about the green line, not serving as a border between two independent states yet - and this is where opinions begin - the first author believes it to be illegitimate and gives a negative assessment of the possibility of it becoming a formal and legal object. The second author, on the other hand, assesses it positively as one of the formers of two independent states. Both authors describe the same phenomena, but use different words relating to it. The words '*undocumented, invisible, duress*' in the first passage and '*important, independent*' in the second are the clues that help us detect subjectivity expressed. An automatic subjectivity identification tool uses broadly the same technique: it captures subjective clues in text and relates them to certain objects or topics of discussion.

In this paper, we set out to investigate how Personal Sense can be harnessed for the purpose of identifying subjectivity in news titles. We provide an annotation schema for the Personal Sense ‘target’ and ‘indicator’ constructions covering emotion, polarity and subjectivity-objectivity in terms of the Personal Sense. We proceed to analyze the subjectivity-objectivity issue. Assuming that the subjective Personal Sense patterns are constructed in text on a regular basis using lexical and syntactic elements, we perform an experiment on the automatic detection of subjectivity in text, as opposed to objective expressions not containing any subjective emotion. First, we demonstrate that subjective expressions are more accurately described using a combination of lexical and syntactic information than by using lexical means only. Next, we select a number of lexical and syntactic features for the identification of the subjective patterns in text. We apply the Personal Sense technique to pairs of words, at least one of which is a noun: thereby identifying the Personal Sense of the noun in the pair. We argue that the suggested features, including the syntactic path between the words in the pair and lexical information about the Personal Sense indicator word, are useful for the identification of the subjective Personal Sense. We use the resulting subjective and objective word-pairs for the subjectivity classification applying the suggested feature set. Thus we learn to identify automatically the word-pairs, connected by a certain syntactic path, bearing an emotion, as opposed to the pairs that do not bear any emotional content. The results confirm our expectations and demonstrate the lexico-syntactic features to be useful for the identification of subjectivity indicators and the targets which receive the subjective

Personal Sense.

The rest of the paper is organized as follows. In the next section, we present the background for our research, including its theoretical foundations - the discussion of Personal Sense, and existing practical approaches to sentiment analysis, especially the ones related closely to our research direction. In Section 3 we present the framework for our research - the Personal Sense annotation schema. Our experiments and analysis of the results are presented in Section 4 and finally in Section 5 the conclusions of our experiments are given, and directions for future work are outlined.

Background

There has been considerable work done on the identification of private states in language (Pang & Lee, 2008). However, it is important to underline the meaning of the emotion’s intentional object, which has not yet been investigated from a lexico-syntactic viewpoint. The authors of (Scherer, 1999) discuss intentionality and the intentional object as an inherent characteristic of emotions, as well as appraisal. While the focus of previous work has been on the overall emotion, mood, or more broadly, private state of a sentence or even a text, we pursue a more fine-grained goal of identifying emotions intended at specific objects, attaching a subjective polar or emotional Personal Sense to the object in the text.

Emotion is focused on an object or event relevant to a person's motivation (Scherer, 1999). This stimulus is called an *intentional object* of the emotion. A word in a text acquires a subjective emotional Personal Sense if it is the intentional object of an emotion expressed in the text. Thus, there should be the object of the emotion, presented by a word in the text: we call this the *target* word and investigate its Personal Sense; and a word or a construction that indicates the emotion (i.e., adds the emotional Personal Sense to the target word) is called the *indicator*. In practice, there is another important element that influences the polarity and emotion of the *target* together with *indicator*: the *intermediate* element that may alter or increase the *target* polarity and emotion against the *indicator* ones.

Personal Sense

Our consciousness is subjective. It is subjective in a sense that the objective world is filtered to be perceived in terms of our physical and spiritual needs. In (Leontev, 1978) word-meanings are called "the most important 'formers' of human consciousness". The dual nature of consciousness is defined by Leontev, as "a picture of the world, opening up before the subject, in which he himself, his actions, and his conditions are included" (Leontev, 1978). The author underlines the fact that consciousness is potentially unlimited to reflect objective reality, but is actually determined and thus limited by personal needs, goals and activities.

As the most important form of consciousness, word-meaning is also considered in a dual manner, combining the objective (shared between the speakers of the same language) representation of reality, and the subjective which serves as a building block and an object of individual consciousness. Thus, Leontev suggests a distinction between significance and sense as two types of meaning, exemplifying the distinction with an exam mark: the significance of an exam mark is shared, as everyone who has ever studied knows the meaning of the "exam mark" and its consequences. On the other hand, in individual consciousness an exam mark acquires a certain sense in terms of actual goals of a person, such as advancing their career, impressing those around them for a student obtaining the mark; or being a successful teacher for an examiner; or a decision on how many students stay for a repeat year and how many get a scholarship for a college official, etc. Generalizing this difference, meanings of words in language have a two-fold nature in this respect: a shared and abstract one, and a personal but more actual one. Perception and reflection of the objective reality in individual consciousness is always connected with achievement of personal concrete goals and performing actions, to satisfy their needs, regardless of whether the motives are perceived consciously by the individual or not. The needs and motives make a constant contribution to the filtering of reality in consciousness by evaluating the significance of objects for the individual, thus ascribing Personal Sense to the objects and objective circumstances, in addition to their objective meaning.

Thus, we understand Personal Sense as a component of word-meaning different for each individual, reflecting an object in word-meaning in terms of unique experience of a person. Personal Sense, as word-meaning, is not manifested explicitly in text or speech. Word-meaning in text is analyzed with latent techniques, for example, Latent Semantic Analysis (Dumais, 2004), with successful applications to various linguistic tasks such as Word Sense Disambiguation (Navigli, 2009; Agirre & Edmonds, 2006). Personal Sense as a characteristic of individual language use should be studied in texts by different individuals separately, taking into account the personality of the authors or their objective characteristics.

Related Experiments on Subjectivity Identification and Personal Sense

Subjectivity is identified with subjective clues. These are lexical items (words or collocations) which contain subjectivity and attach it to the analyzed sentence or document. Initially, a vocabulary of subjectivity clues is constructed (see, for example, (General Inquirer, 2000)). The vocabulary clues are in turn identified in text and used to predict subjective pieces (Wiebe, Bruce & O'Hara, 1999). Sometimes the subjective clues can also be used objectively (Wiebe et al, 2004). Finally, a sentence or document is identified as subjective or objective using the selected lexical and probably some additional features and a classification algorithm. For instance, the authors of (Yu & Hatzivassiloglou, 2003) and (Wiebe et al, 2004) apply the described subjectivity identification scheme and report results of 96.5% F-measure and 94% accuracy respectively, in a document-level classification.

Much research work is dedicated to document-level subjectivity identification, eg. (Spertus, 1997). In (Popescu & Etzioni, 2005) the authors describe the OPINE system which focuses on extracting and analyzing opinion phrases corresponding to specific features in specific sentences. In our work the analysis is focused on a fine-grained subsentence-level, with the main classification element being a pair of words.

We identify the subjectivity clues using existing affective vocabularies. The syntactic constructions containing the clues play a unique role in our work: they are used among other features for the clues disambiguation, and serve as patterns for regular subjectivity constructions extraction. Thus, especially relevant to our work is the syntactic approach to opinion mining based on extraction patterns (Riloff & Wiebe, 2003), in which a news text dataset consisting of 34,000 sentences is investigated. The authors apply an algorithm to learn syntactic patterns associated with subjectivity. The resulting patterns are bootstrapped in order to classify sentences as subjective or objective. In our experiments we use a similar syntactic pattern extraction approach, but we apply it in a new object-oriented framework to classify word pairs.

A well-known task in subjectivity analysis is the polarity classification of documents: texts that contain, for example, product reviews can be automatically divided into two groups: positive and negative reviews, and sometimes a neutral class is also introduced. A broad description of the work done in the area of subjectivity analysis and polarity classification can be found in (Pang & Lee, 2008). We have demonstrated in previous experiments that the concept of Personal Sense has the potential to be of value in the linguistic study of subjectivity, opinion and sentiment, as it was defined just as the element that implies subjectivity in word-meaning and serves as a bridge between the objective world and the subjective consciousness (Panicheva, Cardiff & Rosso, 2009). We have also applied it in the fields of authorship attribution and authors' background identification (Panicheva, Cardiff & Rosso, 2010). These fields of study have become increasingly important in recent years, as the volume of subjective information being made available on the world wide web increases exponentially.

Various linguistic features are used in opinion mining and polarity classification. The features are traditionally divided into syntactic and lexical. As we are concerned with word-meaning and Personal Sense and how it is represented in text, our work lies in the area of the lexical approach, and as a result of our experiments we provide an efficient lexico-syntactic framework for subjectivity classification using Personal Sense. In

(Panicheva, Cardiff & Rosso, 2009), we give a detailed description of how Personal Sense can be successfully applied to polarity classification of products.

In our approach, the features represent information about the position of words in relation to other words, whereas in other works (Pang & Vaithyanathan, 2002) information about the position of the words in terms of the whole text is employed, for example if the word is used closer to the middle or to the end of the document. Word unigrams in various combinations with syntactic information and higher-order n-grams are reported to yield the best results: 82 % accuracy in a two-fold classification task. Word-meanings are used usually in their polar or emotional aspects, as for example in (Snyder & Barzilay, 2007). Adjectives, along with other features, such as parts of speech, syntax constructions, the use of negation are the main classes of features used and compared in polarity classification, see (Pang & Lee, 2008) for numerous examples. In (Ding, Liu & Yu, 2008), the authors propose a system which infers the semantic orientation of an opinion word based on the review context. This system facilitates the combination of multiple opinion words in the same sentence. In (Riloff, Patwardhan & Wiebe, 2006) a hierarchy of lexical features is presented, the information gain of different features is discussed, and the hierarchy is employed for the selection of the best features for opinion analysis.

An approach to opinion classification utilized very successfully is based on product features and their characteristics discussed in the reviews (Balahur Dobrescu & Montoyo Guijarro, 2009). The basic principle of the algorithm is that there are certain features important for a product to be successful, for example in the case of a camera they would be resolution, size and design − a camera with high resolution, tiny size and nice design is likely to be rated high or positive. However, in our approach we exploited the task which involves a greater degree of subjectivity in the sense that for other objects, like movies or news, there are no clearly defined features which can have ‘positive’ or ‘negative’ values: for example, a ‘complicated plot’ can be considered an advantage by one movie critic and a shortcoming by another.

In (Panicheva, Cardiff & Rosso, 2010), we demonstrated that Personal Sense can be used to reflect subjective concept structures specific to a certain professional background. These structures, known as Personal Sense thesauri, were used to infer an author’s perspective from the texts they write. The results confirmed that with certain restrictions, the method of representing the similarities in the personalized thesauri could be used to reflect similarities in occupation of the authors.

Personal Sense in News Headlines

The dataset that we are considering in this section has been described in (Strapparava & Mihalcea, 2007) and is widely used in research (Bhowmick, Basu & Mitra, 2009; Bao et al 2009). It consists of 1,250 news headlines from major newspapers as BBC News, and from the Google News search engine. The titles are annotated in a fine-grained manner with six basic emotions (Anger, Disgust, Fear, Joy, Sadness, Surprise) and valence (Positive/Negative). For each emotion there is a scale ranging from 0 to 100, indicating the degree of the emotion presence in the sentence. Valence is represented by a number ranging from -100 to +100, with 0 indicating a neutral headline, -100 and +100 represent a highly negative and a highly positive headline respectively. In Table 1, we show some example headlines from the dataset in order to illustrate the manner in which the emotions are annotated (a distinct dominant emotion is shown in each case).

Insert Table 1 here

In order to demonstrate the motivation for identifying Personal Sense in headlines, let us consider an example from the dataset.

Nigeria hostage feared dead is freed (Example 3)

This sentence is characterized, according to the provided gold standard, by positive (31 out of 100) valence and a considerable contribution of joy, surprise, fear, and a slight impact of sadness and anger (in descending order). Consider an artificial counterexample, not occurring in the current dataset:

The hostage supposed to be freed is dead (Example 4)

According to our intuition, the sentence should acquire negative valence, and the dominant emotion (the emotion characterized by the highest impact number) should be sadness. However, if we analyze the sentence as a whole in terms of emotional words contained in it, we will get the same pair '*freed*', '*dead*' for both sentences, and an additional '*feared*' for the first one. This does not give us insight about the opposing polarity values for the examples. It is only when we approach the syntactic level and realize that the predicates of the two sentences contain the opposite meaning in relation to the same passive subject, that we can understand why the two examples containing the same meaningful set of words acquire opposite polarity value and different dominant emotions.

The subject of the sentence also plays a very important role in the analysis. In example 3, two opposing emotions are expressed about the ‘hostage’. The joy about the hostage being free overweighs the fear about him/her being dead and becomes an overall sentence emotion; but the fear about the hostage is also present. It is only when we apply both of these emotions to their intentional object - the hostage - that we can compare them in terms of their impact on the overall sentence emotion, and conclude that the joy overweighs the fear, being expressed in a higher-level syntactic dependency, i.e. the predicate of the sentence.

For these reasons, we consider it very important to analyze sentiment and emotion in a fine-grained way, attributing emotions to the Personal Sense of their intentional objects, with syntactic paths serving as formal representation of such attribution (compare the usage of linguistic clues in (Wiebe et al, 2004)).

Of particular interest to us is the fact that the intentional object of an emotion acquires emotional Personal Sense. In Example 3, the Personal Sense of ‘hostage’ contains the emotions of fear and joy, expressed by ‘*feared (dead)*’ and ‘*is freed*’. Moreover, the fact that the author expresses fear about the hostage being dead, and happiness about them being freed, i.e. the negative Personal Sense of the word ‘hostage’ expressed by ‘feared (dead)’, and the positive one, represented by ‘is freed’, delivers some important information about the author: the author clearly opposes the actions of the kidnappers. Although this seems normally the case, it could be different if, for example, the author wrote an extremist slogan and argued for the same demands as the group of terrorists or kidnappers. Thus, Personal Sense indicates the social affiliation of the writer, and depends largely on the intended reader: the customer who buys newspapers that describe ideas in an appropriate manner. Among other characteristics, the writer expresses appropriate emotions in the Personal Sense of the objects and events described.

The Personal Sense Annotation Schema

We follow the theoretical considerations that emotions are directed at objects or events causing them, named *intentional objects* and appraised in terms of personal motivation (Scherer, 1999), and provide the new annotation framework for Personal Sense in text.

An emotion expressed in a sentence consists at least of a pair of words [*indicator, target*], where *indicator* expresses the presence of emotion itself, and *target* denotes the intentional object of the emotion. We exemplify this using the sentence from Example 3, in which the author explicitly expresses fear about the hostage being dead, and implicitly introduces happiness about him/her having been freed. The fear about the hostage is due to the possibility of his/her death, thus in the first case, the Personal Sense indicator is the word '*dead*'. '*Feared'* in this case serves as an intermediate element, which is described further. The word '*dead*' clearly indicates negative polarity, but it can ascribe different Personal Sense emotion to its target word: '*fear*' in the case of example 3, but '*sadness*' in the case of Example 4. The specific emotion does not depend always on the indicator word, but on the syntactic connection between the indicator and the target, and the nature of the target word itself in some cases. We cannot identify clearly if the word '*dead*' in an isolated position implies either '*sadness*' or '*fear*'. Accordingly we do not attribute the implied emotion to the indicator word, but rather to the target word. On the other hand, polarity is attributed to both of them, as it is usually unambiguously represented by the indicator word (clearly negative in the case of ‘*dead*’), and is clearly present in the Personal Sense of the target word (we feel bad about the hostage fearing that he/she is dead), but can acquire the opposite value in some cases.

Consider the following example from the dataset:

Rights ***group halts violent*** *Nepal strikes.* (Example 5),

where the (violent) *'strikes'* contain negative polarity, but ascribe a positive Personal Sense to the word ‘*group*’, because it ‘*halts*’ the strikes.

This example demonstrates that in some cases an additional element is necessary in order to describe the Personal Sense relation correctly. As described in the example above, it is the word ‘*halt*’. It transforms the target polarity to the opposite of the indicator one, causing the ‘*group*’ to acquire a positive Personal Sense, despite the negative polarity of the ‘*strikes*’ indicator. Because of this transformation it is useful to include this word as an ‘*intermediate*’ element. Intuitively it is a word that stands between the indicator and the target in terms of syntax, and it can transform the polarity of the relation radically, as in the example sentence above. However, in the dataset we encountered a significant number of such words that clearly occupy the same syntactic position, intermediate between the indicator and the target, but which do not transform the polarity, although they can be substituted by words that would transform it. Consider the following examples from the dataset:

Snow causes airport closures in Britain. (Example 6),

Stenson defends his title at Dubai. (Example 7).

If we substitute ‘*defends*’ by ‘*loses*’ in Example 7, we will get a negative Personal Sense for ‘*Stenson*', despite the positive ‘*title*'. The word ‘*defends*’ in this case occupies a very important intermediate position, but does not bring a transformation to the polarity. This is why we also consider such words as *intermediate* elements, and define a *modality* attribute for them, which takes the negative value when the intermediate element in question changes the polarity between the corresponding indicator and target elements, and the neutral or positive value when polarity value is preserved. The resulting annotation schemas are presented in Table 2.

The table describes and gives examples of the *indicator*, *intermediate* and *target* annotation schemas designed for our annotation. The indices in the ‘Attribute value' column are used to highlight the attributes that ought to have the same or corresponding value:

- id of the indicator, intermediate and target should be the same, in order to process the words in a single Personal Sense relation;

- polarity and emotion of the target should be consistent with each other:

  - joy should be used with positive polarity,
  - the rest of the emotions with negative polarity;
  - surprise can occur with both.

Insert Table 2 here

Table 3 contains some actual examples of the polarity and emotional Personal Sense relation from the dataset. The headlines are not annotated explicitly with subjectivity/objectivity: an element is supposed to acquire a subjective Personal Sense if and only if it is annotated with *any* emotion and *any* polarity. Any word serving as a *target* would thus acquire subjectivity, and its *indicator* would at the same time be a subjectivity indicator. Thus we get a hierarchy of classes: first, a word-pair ([*group*, *violent*] in Example 5) belongs to the *subjective* class, if it acquires any emotion and polarity, and the *objective* class, if it does not acquire any. After the manual annotation, the subjective lexical and syntactic items are determined automatically, as the ones which acquire personal sense at least once inside the word-pairs, according to the annotated dataset. Next, if the pair is *subjective*, the indicator and the target have a *positive* or a *negative* polarity. Polarity may be different for the indicator and the target (as in the case of Example 5: *positive* for *group*, *negative* for *violent*), and then an intermediate element is introduced that explains the difference (*halts* in Example 5). Moreover, the target element having *negative* polarity acquires Personal Sense containing an emotion: *anger*, *disgust*, *fear*, *sadness*, or *surprise*. The target with *positive* polarity acquires *joy* or *surprise*. (In the case of Example 5 the emotion attached to *group* is *joy*).

Insert Table 3 here

By annotating all the cases of pairs of subjective (or emotional) expressions and their intended objects, we have defined the area of "expressing subjectivity/emotion towards an object in text" extensionally. In other words, we have covered all the actual occurrences of the emotional expressions towards an object in the current dataset. The emotional expression towards an object can be considered a semantic relation, similar to *date of birth* or *headquarters* described in (Suchanek, Ifrim & Weikum, 2006). It is more complex in our case, in the sense that it covers a variety of linguistic phenomena from a formal point of view. It cannot be described by a single syntactic construction or co-occurring with items belonging to a single lexical class. In the following section we consider the emotional relation with the formal lexical and syntactic phenomena in text.

Experiments in Personal Sense Identification

Lexical and Lexico-Syntactic Approach

Lexical items are one of the traditional features used to define semantic relations in text (Gamallo, Agustini & Lopes, 2005). Our first assumption was that the emotional relation is solely characterized by the indicator lexical items, i.e., given a specific word, we assume that it does or does not indicate a subjective expression. To test this assumption, we computed the conditional probability for a lexical item to indicate a subjective expression, given that a lexical item is a subjective clue. The mean result was 83.64%, with the standard deviation of 0.26. This is a modest result, indicating the highest possible subjectivity identification accuracy that we can get, if we correctly identify a word to be subjective or objective. The main reason that we see for such a modest estimation is that most of the words can indicate subjectivity in one case, while preserving an objective meaning in another case; in other words,

"many expressions with subjective usages have objective usages as well"

(citation 6.1, (Wiebe et al, 2004)). We exemplify this consideration with the actual dataset, even with such a seemingly non-ambiguous (in terms of subjectivity and emotion) word as ‘*good*’.

Consider the following two sentences from the dataset:

PM: Havana deal a good experiment. (Example 8)

and

Bad reasons to be good. (Example 9)

In both sentences the word ‘*good*’ is automatically tagged as an adjective. In Example 8 it is annotated as a subjectivity indicator, with the target word being ‘*experiment*', but it does not indicate any subjectivity in Example 9 according to our annotation, as there is no obvious target word or intentional object present in the sentence. Thus, the lexical information alone does not define the subjectivity area accurately enough.

We assume that an additional feature useful for the subjectivity-objectivity distinction is the syntactic dependency path connecting the indicator and its target word. After we added this type of information, we computed the conditional probability of an expression being subjective, given a subjective lexical item (items which serve as subjectivity indicators at least once in the dataset) and a subjective dependency path (dependency paths which are associated at least once with subjectivity). This time the mean result was 92.50% with the standard deviation of 0.18, which proved in the paired statistical t-test to be higher than the lexical-only based result with a 99% statistical significance (df = 667, $p < 0.01$).

We conclude that emotional Personal Sense may be characterized more accurately by a combination of lexical information and syntactic information. To develop on the statement from citation 6.1, if "expression" is considered not a lexical item alone, but a pair consisting of the lexical item and a syntactic path to the potential target, then the expressions with subjective usages are more likely to have no objective usages. To classify whether or not a word indicates emotional Personal Sense in a sentence, it is useful to identify the dependency construction that contains the lexical item, because the same lexical items vary in their subjective and emotional impact depending on the dependency structures in which they occur.

Identifying Subjective Personal Sense

In the current section we are investigating the Personal Sense of nouns, although they are not the only part of speech providing the targets for the Personal Sense annotation. With the targets of the Personal Sense relation limited to nouns, 475 ‘target-indicator’ word pairs were annotated and automatically extracted.

Our assumption is that the new framework for subjectivity classification based on the Personal Sense detection yields significant results, even when applied to fine-grained subjectivity recognition. The goal of the experiment is to test the performance of the lexico-syntactic technique based on our Personal Sense framework, and to estimate the value of the suggested lexical and syntactic features for the subjectivity classification.

We have performed a classification experiment with pairs of words, the first being a noun, and the second one potentially representing the Personal Sense indicator. There were 475 word-pairs annotated with subjectivity, and 17,500 pairs containing no subjectivity. We divided the neutral items into 37 random groups, in order to use all the large volume of the non-subjective data, and performed the classification experiment with each of them, in order for the dataset to be balanced with respect to subjective and objective newstitles, i.e. there were the same number of pairs representing appraisal and emotional Personal Sense, as the number of pairs containing no Personal Sense, in each classification subset. We grouped the same 475 subjective word-pairs in turn with each of the 37 sets of neutrals, thus getting 37 result sets for every feature set and 37*2 - 1 = 73 degrees of freedom for any paired t-test between a pair of the feature sets. Each time we got 950 = 475 neutral + 475 subjective items, and in the experiment we ran a 5-fold cross-validation against each set. That is, we randomly divided 950 into two equal (and balanced) training and test sets 5 times, on each occasion changing only the "randomness" of our choice.

Lexical and Syntactic Classification Features

We investigated a number of features for the classification of pairs of words as containing or not containing subjectivity. First of all, this was the word itself, the lexical item (LEX) that was potentially a Personal Sense indicator for the noun in question; and the part of speech for the lexical item (POS). Secondly, this was the syntactic path (PATH) from the potential indicator to the potential target noun, according to the dependencies identified by the Stanford parser (We applied the Stanford parser “as is”, as our goal was to analyse the subjectivity results given the tools available, without going into their performance quality). As in real-life classification, we would often encounter new lexical items as potential Personal Sense indicators, accordingly we used some features derived from SentiWordNet (Esuli & Sebastiani, 2006) and General

Inquirer (General Inquirer, 2000) in order to represent the degree of subjectivity for lexical items themselves.

SentiWordNet contains information about a number of senses for each word (115,400 senses altogether), often more than one, sometimes more than ten or twenty. Each of the senses is characterized by a synset and a score for positivity and negativity, both ranging from 0 to 1. For every word we used the following 8 features derived from its SentiWordNet profile. The features are presented in Table 4.

Insert Table 4 here

The General Inquirer (General Inquirer, 2000) is a tool for content-analysis of textual data, based on a number of word categories, whose distributions are calculated for a text and used for different goals, with polarity classification and emotion recognition among them. We used 14 categories from the General Inquirer vocabulary, namely: Neg, Ngtv, Virtue, EMOT, Pstv, Pos, Hostile, Pleasure, Vice, EVAL, Pain, NEGAFF, POSAFF, Eval. We applied the relief feature selection measure described in (Kira & Rendell, 1992) in order to sort the features according to their a priori usefulness for the items to classify. The results of the feature relief estimation, in descending order, are presented in Table 5.

Insert Table 5 here

In Table 6, we show a sample feature vector, using as an illustration the sentences presented in Examples 3 and 4. The question marks indicate that the lexical item has not occured in the lexical resource from which the information was used, so the feature value in this case stays "undefined". The PATH feature was encoded as follows:

- If the two words were directly connected, their connection path contained four elements: the POS of the current source and the target words (u'NN' and u'NN' in the case of "mortar assault"), direction of the syntactic dependency ("down" in the case of "mortar", "up" in the case of "assault"), and the dependency path name (de Marneffe & Manning, 2008).
- If the words are interconnected via an additional word, such as "hostage" and "dead" in the example "Nigeria hostage feared dead is freed.", the path contains a sequence of all the individual paths from the target to the indicator, i.e. [u'NN', 'up', u'nsubj', u'VB'] for "hostage feared" and [u'VB', 'down', u'acomp', u'VB'] for "feared dead".

Insert Table 6 here

Results of the Subjective Personal Sense Identification

We performed a classification experiment with 475 word pairs with subjective Personal Sense and a random group of 475 objective word pairs. We used Naïve Bayes (NB) as the classification algorithm, using a machine learning tool, Weka 3.6.1 described in (Witten & Frank, 2005). NB was selected as the other classifiers, SVM and Decision Tree, performed worse in our initial experiments, probably because of the large number of features, the sparseness of some features' values, and in the case of SVM, due to the lack of sufficient training examples. We tested different feature combinations, including the ones based on the relief measure threshold. The results are presented in Table 7 using three evaluation measures: the F-measures (Van Rijsbergen, 1979) for objective and subjective groups separately, and the mean accuracy.

Insert Table 7 here

Almost all the feature sets yielded significant performance in terms of accuracy and F-measure: all of them are higher than 0.84. It is only when we leave out the LEX feature, i.e. the potential indicator word, we get a much lower result, but still a very satisfactory one: 73% accuracy, 0.70 and 0.75 F-measure for objective and subjective word pairs respectively. On one hand, it shows that the semantic features derived from the lexical sources make a useful contribution and allow the prediction of the subjective impact of a lexeme successfully. On the other hand, a considerable difference between the results shows that the most effective way to predict a word's subjective Personal Sense impact in the test set is to learn it in the training set. Evaluating it in principle using the existing lexical resources is less successful.

We performed the experiments, leaving out one of the features each time, in order to evaluate the impact of every individual feature separately. The word serving as the potential Personal Sense indicator, represented by the LEX feature, has proved to play the most important role in the classification: leaving out this feature has decreased the performance most of all. According to our expectations, the syntactic dependency path is also very important, as it also decreased the result compared to using all the features with a higher than 99% significance ($df = 73$, $p < 0.01$) in terms of accuracy and both subjective and objective F-measures. The third feature that introduced significant decrease when leaving it out is the part of speech of the indicator (POS), yielding 94-98% statistically significant difference in the t-test results for the three evaluation measures ($df = 73$, $0.02 < p < 0.06$ for the three measures).

Analysis of Results

It is apparent that there are three most useful features: the lexical item serving as a potential Personal Sense indicator (LEX), its part of speech, and the syntactic dependencies path between the indicator and the target noun (PATH). The classification results give additional evidence for this.

It is important to notice that from these three most important features, the part of speech (POS) plays an exceptional role, as in the classification it played a considerably smaller role than LEX and PATH, but it has been identified as the most important feature in terms of the relief measure. We find this result revealing, knowing that the relief measure is based on the ability of the feature in question to discriminate between the items that are similar to each other but belong to different classes. In the actual classification experiment this is not always the case, and the different classes are distinguished in addition by the features from the lexical resources. However, when such an intuition is used to evaluate the features a priori, the part of speech plays a very important role, because a different part of speech for the same word form can make a big difference in terms of subjectivity, especially when most of the other features stay the same.

Not surprisingly, leaving out the features derived from the lexical resources did not change the result dramatically: first of all, our features obviously contained redundant information bringing considerable noise to the dataset; on the other hand, most of the information was contained in the three most useful features, and for the closed dataset with a limited vocabulary the LEX feature represented all the lexical information and made the lexical resources features redundant.

The best results overall were achieved in the following two cases: using the three main features plus the General Inquirer features (which all had positive relief), *or* the three main features with the SentiWordNet features that had positive relief. Both of these feature sets, and the set consisting of the 22 positive relief features, performed better than the set containing of all the features, with a statistical significance higher than 99.99% ($df = 73$, $p < 0.001$). This shows that, first of all, the features characterized by a positive relief actually contribute positively to the classification experiment. On the other hand, it proves that both the SentiWordNet and the General Inquirer features increase the performance, but can be substituted for one another and do not significantly affect the performance when used together compared to using any one of them.

It is important to notice that the use of the syntactic path as a feature increased the performance considerably, but the absence of this feature did not completely decrease the performance. This supports our contention discussed earlier, confirming that a word often brings subjectivity to the whole sentence, not only to a specific word denoting an object. We realize that this is reasonable if we consider the following examples from the dataset: "**Bombers** *kill* **shoppers**", "Mortar **assault** leaves at least **18** *dead*", or "**Goal** *delight* for **Sheva**". The Personal Sense indicator words are in italics, and the target words are shown in bold. In each case both of the nouns (or parts of speech acting as nouns) occurring in the sentence acquire a subjective Personal Sense, in these cases with the same polarity. However, the emotion contained in the Personal Sense of the different nouns would be different: '*fear*' for '*bombers*' and '*assault*', '*sadness*' for '*shoppers*' and '*18*'. It means that the meaning of the syntactic path feature can be more clearly analyzed in experiments with a more fine-grained emotion classification, not a binary subjectivity-objectivity one.

Conclusions

There has been a vast amount of work done in subjectivity identification. We develop a new subjectivity annotation scheme based on the notion of Personal Sense. First of all, our theoretical motivation was a consideration underlined in (Scherer, 1999), that intentional object is an inherent property of emotion. Secondly, we have shown by analyzing a news headlines sentiment dataset, that a fine-grained approach to subjectivity is essential, because it allows for a more detailed emotion, subjectivity and polarity annotation, thus enabling more accurate identification of different emotions expressed towards different intentional objects in text. Thus we proposed a Personal Sense based annotation scheme and applied it to the dataset in question.

In order to classify the resulting subjective and objective expressions automatically, we added to the widely used subjectivity clues approach a syntactic feature, serving as a regular connection between expressed emotion and its intentional object. In an initial experiment we showed that the resulting feature set describes the subjectivity identifier are more accurately than the baseline lexical approach.

Finally, we performed a subjectivity-objectivity classification experiment, using the suggested fine-grained sub-sentence level annotation and a set of lexical and additional syntactic features, evaluating the impact of different features and comparing the effect of features derived from two lexical resources.

The fine-grained linguistic approach we have presented here, based on the concept of Personal Sense, is designed for annotating and analyzing automatically subjectivity, polarity and emotions in text. It yields a high performance in the subjectivity classification of word-pairs in the news titles dataset, establishing a useful background for identifying polarity and emotions based on the same annotation scheme. The features introduced in the Personal Sense approach are appropriate for the classification, mostly making a positive contribution to the result. The relative impact of the features is realistically estimated with the chosen relief measure.

Both lexical resources used to infer subjective categories for identifying subjective Personal Sense, SentiWordNet and General Inquirer, influence the classification result positively, which proves their appropriateness for the task. However, the impact of the derived categories is not high and cannot substitute for the usage of the lexical items themselves, at least for the limited dataset available. More experiments on a larger dataset are necessary to investigate the meaning of the resources more thoroughly. Moreover, in the current experiment the use of both resources yielded considerable noise, resulting in a slightly lower classification result. Equally good results were achieved when using one of the two resources separately, which indicates that they are interchangeable in the current setting.

Although the syntactic path feature was meaningful - its absence decreased the subjectivity classification results - it is obvious that the decrease was not dramatic. This is in line with the intuition described above and followed by most of the research in subjectivity, that sentiment is distributed over all the objects in the sentence and in a large number of cases is not restricted to only one noun.

The sentiment of a word-pair can have different impact on the sentence and document subjectivity. At the lowest, target-indicator level, the consideration is represented by applying the 'Modifier' element to the word-pair. In this case different types of modifiers relate to the resulting subjectivity class in different ways. In the case of larger structures, a sentence can naturally contain a number of subjective word-pairs. In this case it also appears to be a question of grammar, i.e. a complex of rules. The input might be the subjective word-pairs and syntactic relations between them, and the output would denote the subjectivity degree. Such a grammar could be established by collecting a considerable amount of data and generalizing the obvious patterns with additional information from theoretical linguistics. Another approach would again be to collect and annotate the same amount of data and infer the grammar by statistical machine learning.

It is important in our work that subjectivity cannot be calculated from words or collocations. We have argued that the "indicator" words impact "target" words in different ways, which depends on their POSs, syntactic paths, words occurring between them in syntactic paths, which would be more obvious when analyzing polarity, etc. Moreover, it is not just the syntactic pairs of words that matter, but each class of words, the largest being the POS, would "react" individually to the attachment of a subjective indicator-word. This is why we selected nouns only for our experiment in the current framework. This is our a priori assumption, which can be best described theoretically in terms of the Personal Sense framework, with the experiments we have described here representing a step towards Personal Sense implementation.

Acknowledgements

The work of Paolo Rosso has been done within the EC WIQEI IRSES project (grant no. 269180) FP 7 Marie Curie People Framework, the MICINN Text-Enterprise 2.0 project (TIN2009-13391-C04-03) Plan I+D+I, and the VLC/CAMPUS Microcluster on Multimodal Interaction in Intelligent Systems. We are grateful to the anonymous reviewers for their helpful comments.

References

Agirre E. and Edmonds P. (Eds.). Word Sense Disambiguation: Algorithms and applications. Springer, Text, Speech and Language Technology , Vol. 33, 2006

Bao, Sh., Xu, Sh., Zhang, Li, R., Su,Zh., Han, D Yan., Yu, Y., 2009. Joint Emotion-Topic Modeling for Social Affective Text Mining. ICDM 2009, pp. 699-704.

Balahur Dobrescu, A., Montoyo Guijarro, J.A., 2009. A Semantic Relatedness Approach to Classifying Opinion from Web Reviews. Procesamiento del lenguaje natural. N. 42 (2009), pp. 47-54.

Bhowmick, Pl. K., Basu, A., Mitra, P., 2009. Reader Perspective Emotion Analysis in Text through Ensemble based Multi-Label Classification Framework. Computer and Information Science, 2 (4). pp. 64-74.

de Marneffe, M., Manning, C., 2008. Stanford typed dependencies manual. Stanford

University, Technical report.

Ding, X., Liu, B., Yu, P., 2008. A Holistic Lexicon-Based Approach to Opinion Mining. Proceeding WSDM '08, Proceedings of the international conference on Web search and web data mining.

Dumais, S., 2004. Latent semantic analysis. Annual Review of Information Science and Technology (ARIST), Volume 38, Chapter 4, 189-230

Esuli, A., Sebastiani, F., 2006. SentiWordNet: A Publicly Available Lexical Resource for Opinion Mining. In: Proceedings of LREC-06, 5th Conference on Language Resources and Evaluation, May 22-28, 2006. Genova, Italy. pp. 417-422.

Gamallo, P., Agustini, A., Lopes, G., 2005. Clustering Syntactic Positions with Similar Semantic Requirements. In: Computational Linguistics 31/1. pp. 107-146.

General Inquirer, The, 2000. http://www.wjh.harvard.edu/~inquirer/homecat.htm.

Kira, K. & Rendell, L., 1992. A practical approach to feature selection. In D. Sleeman and P. Edwards, editors, Proc. 9th Int'l Conf. on Machine Learning, pages 249-256, Aberdeen, Morgan Kaufmann Publishers.

Leontev, A.N., 1978. Activity, Consciousness, and Personality. Hillsdale: Prentice-Hall.

Navigli. R. Word Sense Disambiguation: a Survey. ACM Computing Surveys, 41(2), ACM Press, 2009, pp. 1-69.

Pang, B., Lee, L., Vaithyanathan, S., 2002. Thumbs Up? Sentiment Classification Using Machine Learning Techniques. In: Proceedings of the ACL-02 conference on Empirical methods in natural language processing (EMNLP) - Volume 10. Morristown, NJ, USA: Association for Computational Linguistics, pp. 79-86.

Pang, B., Lee. L., 2008. Opinion Mining and Sentiment Analysis. N.Y.: Now Publishers Inc.

Panicheva, P., Cardiff, J., Rosso, P., 2009. A Co-occurrence Based Personal Sense Approach to Opinion Mining. In: Workshop on Opinion Mining and Sentiment Analysis (WOMSA), held at the 2009 CAEPIA-TTIA 13th Conference of the Spanish Association for Artificial Intelligence. November 13, 2009, Sevilla, Spain.

Panicheva, P., Cardiff, J., Rosso, P., 2010. Identifying Writers' Background by Comparing Personal Sense Thesauri. In: Natural Language Processing and Information Systems, 15th International Conference on Applications of Natural Language to Information Systems, NLDB 2010. June 23-25, Cardiff, UK.

Popescu, A-M., Etzioni, O., 2005. Extracting Product Features and Opinions from Reviews, In: Proceedings of the Human Language Technology Conference and the Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP)

Riloff, E., Wiebe, J., 2003. Learning Extraction Patterns for Subjective Expressions, In: Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing (EMNLP-03) - Volume 10. Morristown, NJ, USA: Association for Computational Linguistics. pp. 105-112.

Riloff, E., Patwardhan, S., Wiebe, J., 2006. Feature Subsumption for Opinion Analysis. In: Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing (EMNLP). Sydney, Australia. Morristown, NJ, USA: Association for

Computational Linguistics, pp. 440-448.

Scherer, K. R., 1999. Appraisal theories. In:. Handbook of Cognition and Emotion. Chichester: Wiley. pp. 637-663.

Snyder, B., Barzilay, R., 2007. Multiple aspect ranking using the good grief algorithm. In: Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL 007), pp. 300-307.

Spertus, E., 1997. Smokey: Automatic Recognition of Hostile Messages. In: Proceedings of the Eighth Annual Conference on Innovative Applications in Artificial Intelligence (IAAI-97). pp. 1058-1065.

Strapparava, C., Mihalcea, R., 2007. SemEval-2007 Task 14: Affective Text. In: Proceedings of the 4th International Workshop on the Semantic Evaluations (SemEval 2007), Prague, Czech Republic, June 2007.

Suchanek, F.M., Ifrim, G., Weikum, G., 2006. LEILA: Learning to Extract Information by Linguistic Analysis. In: Proceedings of the 2nd Workshop on Ontology Learning and Population: Bridging the Gap between Text and Knowledge. July 2006. Sydney, Australia. pp. 18-25.

Van Rijsbergen, C.J., 1979. Information Retrieval, 2nd edition. Dept. of Computer Science, University of Glasgow, 1979.

Wiebe, J., Bruce, R., O'Hara, T., 1999. Development and Use of a Gold Standard Data

Set for Subjectivity Classifications. In: Proc. 37th Annual Meeting of the Assoc. for Computational Linguistics (ACL-99), June 23-26, University of Maryland, pp. 246-253

Wiebe, J., Wilson, Th., Bruce, R., Bell, M., Martin, M., 2004. Learning Subjective Language. Computational Linguistics 30 (3).

Wiebe, J.M., 1994. Tracking Point of View in Narrative. In: Computational Linguistics, vol. 20, 1994. pp. 233-287.

Witten, I.H., Frank, E., 2005. Data Mining: Practical Machine Learning Tools and Techniques. 2nd Edition, San Francisco: Morgan Kaufmann.

Yu, H., Hatzivassiloglou, V., 2003. Towards Answering Opinion Questions: Separating Facts from Opinions and Identifying the Polarity of Opinion Sentences, Proceedings of the 2003 conference on Empirical methods in natural language processing, July 11, 2003. pp.129-136.

TABLE 1. Example Annotations for Emotion Indicators.

Sentence	Anger	Disgust	Fear	Joy	Sadness	Surprise
Anti-U.S. Attack Videos Spread on the Internet	**63**	45	58	0	54	0
Griffiths scorns Withnail 'play'	12	**35**	0	4	3	0
Tropical Storm Isaac forms in Atlantic	0	0	**86**	0	10	0
Game on! London exhibition celebrates the history of video games	0	0	0	**70**	0	6
Parachutist dies at bridge-jump festival	6	6	33	6	**88**	0
Pot smokers may avoid Alzheimer's, study says	0	0	4	10	0	**66**

TABLE 2. Annotation schemas for the Personal Sense annotation of polarity and emotion in Example 5.

Element name	Attribute name	Attribute type, possible values	Attribute value	Attribute value for Example 5
Indicator	Id	Integer	id_i	67
'violent'	Polarity	string: neg, pos	pol_g	Neg
Target	Id	Integer	id_i	67
'group'	Polarity	string: neg, pos	pol_j	Pos
	Emotion	string: anger, disgust, fear, joy, sadness, surprise	$emotion_j$	Joy
Intermediate	Id	Integer	id_i	67
'halts'	Modality	String: negative, neutral, positive	mod_y	negative

TABLE 3. A sample of the annotated Personal Sense relation components.

Id	Sentence	Emotional relation elements		
		Indicator	Intermediate	Target
1	Mortar assault leaves at least 18 dead.	Dead		assault
2	Mortar assault leaves at least 18 dead.	Dead		18
3	Goal delight for Sheva.	Delight		goal
4	Goal delight for Sheva.	Delight		Sheva
5	Nigeria hostage feared dead is freed.	Dead	feared	hostage
6	Nigeria hostage feared dead is freed.	Freed		hostage
7	Bombers kill shoppers.	Kill		bombers
8	Bombers kill shoppers.	Kill		shoppers
9	Vegetables, not fruit, slow brain decline.	Decline	slow	vegetables
10	Rights group halts violent Nepal strikes.	Violent		strikes
11	Rights group halts violent Nepal strikes.	Strikes	halts	group
12	Snow causes airport closures in Britain.	Closures	causes	snow
13	Stenson defends his title at Dubai.	Title	defends	Stenson

TABLE 4. Subjectivity features based on the SentiWordNet profile of a word

Name of the feature	Significance of the feature
NUMALL	number of the word's SentiWordNet senses
BIGGESTVALUE	maximum ‘positivity’ or ‘negativity’ value with the respective sign
BIGGESTSUM	the biggest sum of the positive and the negative value with the respective sign
NUMPOSSUM	the number of senses for the current word for which this sum is positive
NUMNEGSUM	the number of senses for the current word for which this sum is negative
NUMBIGPOSSUM	number of senses for which the sum was higher than 0.25 and positive
NUMBIGNEGSUM	number of senses for which the sum was higher than 0.25 and negative
BIGGESTNUMOTHER0	maximum positive or negative value in the sense, for which the other value was equal to 0

TABLE 5. The results of the feature relief estimation.

Feature	Relief measure		Feature	Relief measure
POS	0.3388		NUMBIGPOSSUM	0.0024
PATH	0.2311		Hostile	0.0024
LEX	0.1594		Pleasure	0.0012
Neg	0.0232		Vice	0.0009
Virtue	0.0229		EVAL	0.0004
NUMALL	0.0127		Pain	0.0003
NUMNEGSUM	0.0111		NEGAFF	0.0002
NUMPOSSUM	0.0105		POSAFF	0.0001
Ngtv	0.0048		Eval	0.00007
NUMBIGNEGSUM	0.0047		BIGGESTNUMOTHER0	-0.0129
EMOT	0.0043		BIGGESTVALUE	-0.0156
Pstv	0.0040		BIGGESTSUM	-0.0177
Pos	0.0040			

TABLE 6. Example feature vector based on Examples 3 and 4

word	Dead	Mortar	Assault
pos	VB	NN	NN
PATH	[[u'NN', 'up', u'nsubj', u'VB'], [u'VB', 'down', u'acomp', u'VB']]	[[u'NN', 'down', u'nn', u'NN']]	[[u'NN', 'up', u'nn', u'NN']]
numall	21	?	4
numpossum	4	?	1
numnegsum	14	?	1
numbigpossum	2	?	1
numbignegsum	14	?	1
biggestsum	-0.75	?	-0.375
biggestvalue	-0.75	?	-0.375
Biggestnumother0	-0.75	?	-0.375
Pos	0	?	0
Pstv	0	?	1
Neg	1	?	1
Ngtv	1	?	1
Hostile	0	?	0
Eval	0	?	0
EVAL	0	?	0
POSAFF	0	?	0
NEGAFF	0	?	0
Virtue	0	?	1
Vice	0	?	0
Pain	0	?	0
Pleasure	0	?	0
EMOT	0	?	0
Subjectivity	Subj	Obj	Obj

TABLE 7. The results of the subjectivity classification experiment for word pairs using Personal Sense scheme.

Features	F-measure for Objective items	F-measure for Subjective items	Mean Accuracy
All	0.85	0.85	0.86
POS+WORD+PATH, relief > 0.01	0.86	0.88	0.87
14 features with relief > 0.001	0.87	0.87	0.87
22 features with relief > 0.0	0.87	0.87	0.87
No LEX	0.70	0.75	0.73
No POS	0.85	0.85	0.85
No PATH	0.85	0.84	0.84
No NUMALL	0.86	0.85	0.86
No NUMPOSSUM	0.86	0.86	0.86
Leaving out any other feature	0.86	0.86	0.86
LEX+POS+PATH+SentiWordNet 11 features	0.86	0.86	0.86
LEX+POS+PATH+GenInq 17 features	0.87	0.88	0.87
LEX+POS+PATH+SentiWordNet with positive relief only, 8 features	0.87	0.88	0.87