

Sentiment Analysis: How to Derive Prior Polarities from SentiWordNet

Marco Guerini

Trento RISE
Via Sommarive 18
38123 Povo in Trento, Italy
m.guerini@trentorise.eu

Lorenzo Gatti

Trento RISE
Via Sommarive 18
38123 Povo in Trento, Italy
l.gatti@trentorise.eu

Marco Turchi

Fondazione Bruno Kessler
Via Sommarive 18
38123 Povo in Trento, Italy
turchi@fbk.eu

Abstract

Assigning a positive or negative score to a word out of context (i.e. a word’s prior polarity) is a challenging task for sentiment analysis. In the literature, various approaches based on SentiWordNet have been proposed. In this paper, we compare the most often used techniques together with newly proposed ones and incorporate all of them in a learning framework to see whether blending them can further improve the estimation of prior polarity scores. Using two different versions of SentiWordNet and testing regression and classification models across tasks and datasets, our learning approach consistently outperforms the single metrics, providing a new state-of-the-art approach in computing words’ prior polarity for sentiment analysis. We conclude our investigation showing interesting biases in calculated prior polarity scores when word Part of Speech and annotator gender are considered.

1 Introduction

Many approaches to sentiment analysis make use of lexical resources – i.e. lists of positive and negative words – often deployed as baselines or as features for other methods (usually machine learning based) for sentiment analysis research (Liu and Zhang, 2012). In these lexica, words are associated with their prior polarity, i.e. if that word out of context evokes something positive or something negative. For example, *wonderful* has a positive connotation – prior polarity – while *horrible* has a negative one. These approaches have the advantage of not

needing deep semantic analysis or word sense disambiguation to assign an affective score to a word and are domain independent (they are thus less precise but more portable).

SentiWordNet (henceforth SWN) is one of these resources and has been widely adopted since it provides a broad-coverage lexicon – built in a semi-automatic manner – for English (Esuli and Sebastiani, 2006). Given that SWN provides polarities scores for each word sense (also called ‘posterior polarities’), it is necessary to derive prior polarities from the posteriors. For example, the word *cold* has a posterior polarity for the meaning “having a low temperature” – like in “*cold* beer” – that is different from the one in “*cold* person” which refers to “being emotionless”. This information must be considered when reconstructing the prior polarity of *cold*.

Several formulae to compute prior polarities starting from posterior polarities scores have been used in the literature. However, their performance varies significantly depending on the adopted variant. We show that researchers have not paid sufficient attention to this *posterior-to-prior polarity* issue. Indeed, we show that some variants outperform others on different datasets and can represent a fairer state-of-the-art approach using SWN. On top of this, we attempt to outperform the state-of-the-art formula using a learning framework that combines the various formulae together.

In detail, we will address five main research questions: (i) is there any relevant difference in the posterior-to-prior polarity formulae performance (both in regression and classification tasks), (ii) is there any relevant variation in prior polarity values

if we use different releases of SWN (i.e. SWN_1 or SWN_3), (iii) can a learning framework boost performance of such formulae, (iv) considering word Part of Speech (PoS), is there any relevant difference in formulae performance, (v) considering the gender dimension of the annotators (male/female) and the sentiment dimension (positive/negative), is there any relevant difference in SWN performance.

In Section 2 we briefly describe our approach and how it differentiates from similar sentiment analysis tasks. Then, in Sections 3 and 4, we present SentiWordNet and overview various posterior-to-prior polarity formulae based on this resource that appeared in the literature (included some new ones we identified as potentially relevant). In Section 5 we describe the learning approach adopted on prior-polarity formulae. In Section 6 we introduce the ANEW and General Inquirer resources that will be used as gold standards. Finally, in the two last sections, we present a series of experiments, both in regression and classification tasks, that give an answer to the aforementioned research questions. The results support the hypothesis that using a learning framework we can improve on state-of-the-art performance and that there are some interesting phenomena connected to PoS and annotator gender.

2 Proposed Approach

In the broad field of Sentiment Analysis we will focus on the specific problem of posterior-to-prior polarity assessment, using both regression and classification experiments. A general overview on the field and possible approaches can be found in (Pang and Lee, 2008) or (Liu and Zhang, 2012).

For the regression task, we tackled the problem of assigning affective scores (along a continuum between -1 and 1) to words using the posterior-to-prior polarity formulae. For the classification task (assessing whether a word is either *positive* or *negative*) we used the same formulae, but considering just the sign of the result. In these experiments we will also use a learning framework which combines the various formulae together. The underlying hypothesis is that by blending these formulae, and looking at the same information from different perspectives (i.e. the posterior polarities provided by SWN combined in various ways), we can give a better prediction.

The regression task is harder than binary classification, since we want to assess not only that *pretty*, *beautiful* and *gorgeous* are positive words, but also to define a partial or total order so that *gorgeous* is more positive than *beautiful* which, in turn, is more positive than *pretty*. This is fundamental for tasks such as affective modification of existing texts, where words' polarity together with their score are necessary for creating multiple graded variations of the original text (Guerini et al., 2008). Some of the work that addresses the problem of sentiment strength are presented in (Wilson et al., 2004; Paltooglou et al., 2010), however, their approach is modeled as a multi-class classification problem (*neutral*, *low*, *medium* or *high* sentiment) at the sentence level, rather than a regression problem at the word level. Other works such as (Neviarouskaya et al., 2011) use a fine grained classification approach too, but they consider emotion categories (*anger*, *joy*, *fear*, etc.), rather than sentiment strength categories. On the other hand, even if approaches that go beyond pure prior polarities – e.g. using word bigram features (Wang and Manning, 2012) – are better for sentiment analysis tasks, there are tasks that are intrinsically based on the notion of words' prior polarity. Consider copywriting, where evocative names are a key element to a successful product (Özbal and Strapparava, 2012; Özbal et al., 2012). In such cases no context is given and the brand name alone, with its perceived prior polarity, is responsible for stating the area of competition and evoking semantic associations. For example *Mitsubishi* changed the name of one of its SUV for the Spanish market, since the original name *Pajero* had a very negative prior polarity, as it meant 'wanker' in Spanish (Piller, 2003).

To our knowledge, the only work trying to address the SWN posterior-to-prior polarity issue, comparing some of the approaches appeared in the literature is (Gatti and Guerini, 2012). However, in our previous study we only considered a regression framework, we did not use machine learning and we only tested SWN_1 . So, we took this work as a starting point for our analysis and expanded on it.

3 SentiWordNet

SentiWordNet (Esuli and Sebastiani, 2006) is a lexical resource in which each entry is a set of

lemma-PoS pairs sharing the same meaning, called “synset”. Each synset s is associated with the numerical scores $\text{Pos}(s)$ and $\text{Neg}(s)$, which range from 0 to 1. These scores – automatically assigned starting from a bunch of seed terms – represent the positive and negative valence (or posterior polarity) of the synset and are inherited by each lemma-PoS in the synset. According to the structure of SentiWordNet, each pair can have more than one sense and each of them takes the form of `lemma#PoS#sense-number`, where the smallest sense-number corresponds to the most frequent sense.

Obviously, different senses can have different polarities. In Table 1, the first 5 senses of `cold#a` present all possible combinations, included mixed scores (`cold#a#4`), where positive and negative valences are assigned to the same sense. Intuitively, mixed scores for the same sense are acceptable, as in “cold beer” (positive) vs. “cold pizza” (negative).

PoS	Offset	Pos(s)	Neg(s)	SynsetTerms
a	1207406	0.0	0.75	cold#a#1
a	1212558	0.0	0.75	cold#a#2
a	1024433	0.0	0.0	cold#a#3
a	2443231	0.125	0.375	cold#a#4
a	1695706	0.625	0.0	cold#a#5

Table 1: First five *SentiWordNet* entries for `cold#a`

In our experiments we use two different versions of SWN: SentiWordNet 1.0 (SWN_1), the first release of SWN, and its updated version SentiWordNet 3.0 (Baccianella et al., 2010) – SWN_3 . In SWN_3 the annotation algorithm used in SWN_1 was revised, leading to an increase in the accuracy of posterior polarities over the previous version.

4 Prior Polarities Formulae

In this section we review the main strategies for computing prior polarities used in previous studies. All the proposed approaches try to estimate the prior polarity score from the posterior polarities of all the senses for a single lemma-PoS. Given a lemma-PoS with n senses (`lemma#PoS#n`), every formula f is independently applied to all the $\text{Pos}(s)$ and $\text{Neg}(s)$. This produces two scores, $f(\text{posScore})$ and $f(\text{negScore})$, for each lemma-PoS. To obtain a unique prior polarity for each lemma-PoS, $f(\text{posScore})$ and $f(\text{negScore})$ can be

mapped according to different strategies:

$$f_m = \begin{cases} f(\text{posScore}) & \text{if } f(\text{posScore}) \geq \\ & f(\text{negScore}) \\ -f(\text{negScore}) & \text{otherwise} \end{cases}$$

$$f_d = f(\text{posScore}) - f(\text{negScore})$$

where f_m computes the absolute maximum of the two scores, while f_d computes the difference between them. It is worth noting that $f(\text{negScore})$ is always positive by construction. To obtain a final prior polarity that ranges from -1 to 1, the negative sign is imposed. So, considering the first 5 senses of `cold#a` in Table 1, $f(\text{posScore})$ will be derived from the $\text{Pos}(s)$ values $\langle 0.0, 0.0, 0.0, 0.125, 0.625 \rangle$, while $f(\text{negScore})$ from $\langle 0.750, 0.750, 0.0, 0.375, 0.0 \rangle$. Then, the final polarity strength returned will be either f_m or f_d .

The formulae (f) we tested are the following:

fs. In this formula only the first (and thus most frequent) sense is considered for the given lemma-PoS. This is equivalent to considering only the SWN score for lemma-PoS#1. Based on (Neviarouskaya et al., 2009; Agrawal and Siddiqui, 2009; Guerini et al., 2008; Chowdhury et al., 2013), this is the most basic form of prior polarities.

mean. It calculates the mean of the positive and negative scores for all the senses of the given lemma-PoS. This formula has been used in (Thet et al., 2009; Denecke, 2009; Devitt and Ahmad, 2007; Sing et al., 2012).

uni. Based on (Neviarouskaya et al., 2009), it considers only those senses that have a $\text{Pos}(s)$ greater than or equal to the corresponding $\text{Neg}(s)$, and greater than 0 (the *stronglyPos* set). In case posScore is equal to negScore , the one with the highest weight is returned, where weights are defined as the cardinality of *stronglyPos* divided by the total number of senses. The same applies for the negative senses. This is the only method, together with *rnd*, for which we cannot apply f_d , as it returns a positive or negative score according to the weight.

uniw. Like *uni* but without the weighting system.

w1. This formula weighs each sense with a geometric series of ratio 1/2. The rationale behind this choice is based on the assumption that more frequent

senses should bear more “affective weight” than rare senses when computing the prior polarity of a word. The system presented in (Chaumartin, 2007) uses a similar approach of weighted mean.

w2. Similar to the previous one, this formula weighs each lemma with a harmonic series, see for example (Denecke, 2008).

On top of these formulae, we implemented some new formulae that were relevant to our task and have not been implemented before. These formulae mimic the ones discussed previously, but they are built under a different assumption: that the saliency (Giora, 1997) of a word’s prior polarity might be more related to its posterior polarities score, rather than to sense frequencies. Thus we ordered *posScore* and *negScore* by strength, giving more relevance to ‘valenced’ senses. For instance, in Table 1, *posScore* and *negScore* for `cold#a` become $\langle 0.625, 0.125, 0.0, 0.0, 0.0 \rangle$ and $\langle 0.750, 0.750, 0.375, 0.0, 0.0 \rangle$ respectively.

w1s and **w1n.** Like *w1* and *w2*, but senses are ordered by strength (sorting $\text{Pos}(s)$ and $\text{Neg}(s)$ independently).

w1n and **w2n.** Like *w1* and *w2* respectively, but without considering senses that have a 0 score for both $\text{Pos}(s)$ and $\text{Neg}(s)$. Our motivation is that “empty” senses are mostly noise.

w1sn and **w2sn.** Like *w1s* and *w2s*, but without considering senses that have a 0 score for both $\text{Pos}(s)$ and $\text{Neg}(s)$.

median: return the median of the senses ordered by polarity score.

All these prior polarities formulae are compared against two gold standards (one for regression, one for classification) both one by one, as in the works mentioned above, and combined together in a learning framework (to see whether combining these features – that capture different aspect of prior polarities – can further improve the results).

Finally, we implemented two variants of a prior polarity random baseline to assess possible advantages of approaches using SWN:

rnd. This formula represents the basic baseline random approach. It simply returns a random number between -1 and 1 for any given `lemma#PoS`.

swnrnd. This formula represents an advanced random approach that incorporates some “knowl-

edge” from SWN. It takes the scores of a random sense for the given `lemma#PoS`. We believe this is a fairer baseline than *rnd* since SWN information can possibly constrain the values. A similar approach has been used in (Qu et al., 2008).

5 Learning Algorithms

We used two non-parametric learning approaches, Support Vector Machines (SVMs) (Shawe-Taylor and Cristianini, 2004) and Gaussian Processes (GPs) (Rasmussen and Williams, 2006), to test the performance of all the metrics in conjunction. SVMs are non-parametric deterministic algorithms that have been widely used in several fields, in particular in NLP where they are the state-of-the-art for various tasks. GPs, on the other hand, are an extremely flexible non-parametric probabilistic framework able to explicitly model uncertainty, that, despite being considered state-of-the-art in regression, have rarely been used in NLP. To our knowledge only two previous works did so (Polajnar et al., 2011; Cohn and Specia, 2013).

Both methods take advantage of the *kernel trick*, a technique used to embed the original feature space into an alternative space where data may be linearly separable. This is performed by the kernel function that transforms the input data in a new structure, called *kernel*. How it is used to produce the prediction is one of the main differences between SVMs and GPs. In classification SVMs use the geometric mean to discriminate between the positive and negative classes, while the GP model uses the posterior probability distribution over each class. Both frameworks support learning algorithms for regression and classification. An exhaustive explanation of the two methodologies can be found in (Shawe-Taylor and Cristianini, 2004) and (Rasmussen and Williams, 2006).

In the SVM experiments, we use *C*-SVM and ϵ -SVM implemented in the LIBSVM toolbox (Chang and Lin, 2011). The selection of the kernel (linear, polynomial, radial basis function and sigmoid) and the optimization of the parameters are carried out through grid search in 10-fold cross-validation.

GP regression models with Gaussian noise are a rare exception where the exact inference with likelihood functions is tractable, see §2 in (Rasmussen

and Williams, 2006). Unfortunately, this is not valid for the classification task – see §3 in (Rasmussen and Williams, 2006) – where an approximation method is required. In this work, we use the Laplace approximation method proposed in (Williams and Barber, 1998). Different kernels are tested (covariance for constant functions, linear with and without automatic relevance determination (ARD)¹, Matern, neural network, etc.²) and the linear logistic (*lll*) and probit regression (*prl*) likelihood functions are evaluated in classification. In our classification experiments we tried all possible combinations of kernels and likelihood functions, while in the regression tests we ranged only on different kernels. All the GP models were implemented using the GPML Matlab toolbox³. Unlike SVMs, the optimization of the kernel parameters can be performed without using grid search, but the optimal parameters can be obtained iteratively, by maximizing the marginal likelihood (or in classification, the Laplace approximation of the marginal likelihood). We fix at 100 the maximum number of iterations.

An interesting property of the GPs is their capability of weighting the features differently according to their importance in the data. This is referred to as the automatic variance determination kernel. As demonstrated in (Weston et al., 2000), SVMs can benefit from the application of feature selection techniques especially when there are highly redundant features. Since the prior polarities formulae tend to cluster in groups that provide similar results (Gatti and Guerini, 2012) – creating noise for the learner – we want to understand whether feature selection approaches can boost the performance of SVMs. For this reason, we also test feature selection prior to the SVM training. For that we used Randomized Lasso, or stability selection (Meinshausen and Bühlmann, 2010). Re-sampling of the training data is performed several times and a Lasso regression model is fit on each sample. Features that appear in a given number of samples are retained. Both the fraction of the data to be sampled and the threshold to select the features can be configured. In our

¹*linone* and *linard* in the result tables, respectively.

²More detailed information on the available kernels are in §4 (Rasmussen and Williams, 2006)

³<http://www.gaussianprocess.org/gpml/code/matlab/doc/>

experiments we set the sampling fraction to 75%, the selection threshold to 25% and the number of resamples to 1,000. We refer to these as *SVMfs*.

6 Gold Standards

To assess how well prior polarity formulae perform, a gold standard with word polarities provided by human annotators is needed. There are many such resources in the literature, each with different coverage and annotation characteristics. ANEW (Bradley and Lang, 1999) rates the valence score of 1,034 words, which were presented in isolation to annotators. The SO-CAL entries (Taboada et al., 2011) were collected from corpus data and then manually tagged by a small number of annotators with a multi-class label. These ratings were further validated through crowdsourcing. Other resources, such as the General Inquirer lexicon (Stone et al., 1966), provide a binomial classification (either *positive* or *negative*) of sentiment-bearing words. The resource presented in (Wilson et al., 2005) uses a similar binomial annotation for single words; another interesting resource is WordNetAffect (Strapparava and Valitutti, 2004) but it labels words senses and it cannot be used for the prior polarity validation task.

In the following we describe in detail the two resources we used for our experiments, namely ANEW for the regression experiments and the General Inquirer (GI) for the classification ones.

6.1 ANEW

ANEW (Bradley and Lang, 1999) is a resource developed to provide a set of normative emotional ratings for a large number of words (roughly 1 thousand) in the English language. It contains a set of words that have been rated in terms of pleasure (affective valence), arousal, and dominance. In particular for our task we considered the valence dimension. Since words were presented to subjects in isolation (i.e. no context was provided) this resource represents a human validation of prior polarities scores for the given words, and can be used as a gold standard. For each word ANEW provides two main metrics: $anew_{\mu}$, which correspond to the average of annotators votes, and $anew_{\sigma}$, which gives the variance in annotators scores for the given word. In the same way these metrics are also provided for

the male/female annotator groups.

6.2 General Inquirer

The Harvard General Inquirer dictionary is a widely used resource, built for automatic text analysis (Stone et al., 1966). Its latest revision⁴ contains 11789 words, tagged with 182 semantic and pragmatic labels, as well as with their part of speech. Words and their categories were initially taken from the Harvard IV-4 Psychosociological Dictionary (Dunphy et al., 1974) and the Lasswell Value Dictionary (Lasswell and Namenwirth, 1969). For this paper we consider the `Positive` and `Negative` categories (1,915 words the former, 2,291 words the latter, for a total of 4,206 affective words).

7 Experiments

In order to use the ANEW dataset to measure prior polarities formulae performance, we had to assign a PoS to all the words to obtain the `SWN lemma#PoS` format. To do so, we proceeded as follows: for each word, check if it is present among both `SWN1` and `SWN3` lemmas; if not, lemmatize the word with the TextPro tool suite (Pianta et al., 2008) and check if the lemma is present instead⁵. If it is not found (i.e., the word cannot be aligned automatically), remove the word from the list (this was the case for 30 words of the 1,034 present in ANEW). The remaining 1,004 lemmas were then associated with all the PoS present in SWN to get the final `lemma#PoS`. Note that a lemma can have more than one PoS, for example, *writer* is present only as a noun (`writer#n`), while *yellow* is present as a verb, a noun and an adjective (`yellow#v`, `yellow#n`, `yellow#a`). This gave us a list of 1,484 words in the `lemma#PoS` format.

In a similar way we pre-processed the GI words that uses the generic `modif` label to indicate either adjective or adverb (noun and verb PoS were instead consistently used). Finally, all the sense-disambiguated words in the `lemma#PoS#n` format were discarded (1,114 words out of the 4,206 words with positive or negative valence).

⁴<http://www.wjh.harvard.edu/~inquirer/>

⁵We did not lemmatize everything to avoid duplications (for example, if we lemmatize the ANEW entry *addicted*, we obtain *addict*, which is already present in ANEW).

After the two datasets were built this way, we removed the words for which the `posScore` and `negScore` contained all 0 in both `SWN1` and `SWN3` (523 `lemma#PoS` for ANEW and 484 for the GI dataset), since these words are not informative for our experiments. The final dataset included 961 entries for ANEW and 2,557 for GI. For each `lemma#PoS` in GI and ANEW, we then applied the prior polarity formulae described in Section 4, using both `SWN1` and `SWN3` and annotated the results.

According to the nature of the human labels (real numbers or -1/1), we ran several regression and classification experiments. In both cases, each dataset was randomly split into 70% for training and the remaining for test. This process was repeated 5 times to generate different splits. For each partition, optimization of the learning algorithm parameters was performed on the training data (in 10-fold cross-validation for SVMs). Training and test sets were normalized using the z-score.

To evaluate the performance of our regression experiments on ANEW we used the Mean Absolute Error (*MAE*), that averages the error over a given test set. Accuracy was used for the classification experiments on GI instead. We opted for accuracy – rather than F1 – since for us True Negatives have same importance as True Positives. For each experiments we reported the average performance and the standard deviation over the 5 random splits. In the following sections, to check if there was a statistically significant difference in the results, we used Student’s t-test for regression experiments, while an approximate randomization test (Yeh, 2000) was used for the classification experiments.

In Tables 2 and 3, the results of regression experiments over the ANEW dataset, using `SWN1` and `SWN3`, are presented. The results of the classification experiments over the GI dataset, using `SWN1` and `SWN3` are shown in Tables 4 and 5. For the sake of interpretability, results are divided according to the main approaches: randoms, posterior-to-prior formulae, learning algorithms. Note that for classification we report the generics *f* and not the *f_m* and *f_d* variants. In fact, both versions always return the same classification answer (we are classifying according to the sign of *f* result and not its strength). For the GPs, we report the two best configurations only.

	MAE μ	MAE σ
<i>rnd</i>	0.652	0.026
<i>swnrnd_m</i>	0.427	0.011
<i>swnrnd_d</i>	0.426	0.009
<i>uniw_m</i>	0.420	0.009
<i>max_m</i>	0.419	0.009
<i>fs_d</i>	0.413	0.011
<i>fs_m</i>	0.412	0.009
<i>uni</i>	0.410	0.010
<i>uniw_d</i>	0.406	0.007
<i>w1sn_m</i>	0.405	0.011
<i>max_d</i>	0.404	0.005
<i>w2sn_m</i>	0.402	0.011
<i>median_d</i>	0.401	0.014
<i>w1_d</i>	0.401	0.010
<i>w1n_d</i>	0.399	0.008
<i>mean_d</i>	0.398	0.010
<i>w2_d</i>	0.398	0.010
<i>median_m</i>	0.397	0.015
<i>w1sn_d</i>	0.397	0.008
<i>w2sn_d</i>	0.397	0.008
<i>w2n_d</i>	0.397	0.008
<i>w1s_m</i>	0.396	0.010
<i>w1_m</i>	0.396	0.010
<i>w1n_m</i>	0.394	0.009
<i>mean_m</i>	0.393	0.011
<i>w2s_d</i>	0.393	0.008
<i>w1s_d</i>	0.393	0.009
<i>w2s_m</i>	0.392	0.010
<i>w2_m</i>	0.391	0.011
<i>w2n_m</i>	0.391	0.012
<i>GP_{inard}</i>	0.398	0.014
<i>GP_{inone}</i>	0.398	0.014
<i>SVM</i>	0.367	0.010
<i>SVM_{fs}</i>	0.366	0.011
AVERAGE	0.398	0.010

Table 2: MAE results for metrics using SWN_1

	MAE μ	MAE σ
<i>rnd</i>	0.652	0.026
<i>swnrnd_d</i>	0.404	0.013
<i>swnrnd_m</i>	0.402	0.010
<i>max_m</i>	0.393	0.009
<i>fs_d</i>	0.382	0.008
<i>uniw_m</i>	0.382	0.015
<i>fs_m</i>	0.381	0.010
<i>median_m</i>	0.377	0.008
<i>uniw_d</i>	0.377	0.012
<i>median_d</i>	0.377	0.011
<i>uni</i>	0.376	0.010
<i>max_d</i>	0.372	0.011
<i>mean_d</i>	0.371	0.010
<i>w1sn_m</i>	0.371	0.011
<i>w2sn_m</i>	0.369	0.010
<i>w1_d</i>	0.368	0.010
<i>w2_d</i>	0.367	0.010
<i>mean_m</i>	0.367	0.010
<i>w1_m</i>	0.365	0.010
<i>w2sn_d</i>	0.364	0.011
<i>w1sn_d</i>	0.364	0.010
<i>w1s_m</i>	0.363	0.009
<i>w1n_d</i>	0.362	0.009
<i>w2s_d</i>	0.362	0.010
<i>w2_m</i>	0.362	0.010
<i>w1s_d</i>	0.362	0.009
<i>w1n_m</i>	0.362	0.007
<i>w2n_d</i>	0.361	0.010
<i>w2s_m</i>	0.360	0.009
<i>w2n_m</i>	0.359	0.009
<i>GP_{inone}</i>	0.356	0.008
<i>GP_{inard}</i>	0.355	0.008
<i>SVM</i>	0.333	0.004
<i>SVM_{fs}</i>	0.333	0.003
AVERAGE	0.366	0.009

Table 3: MAE results for regression using SWN_3

8 General Discussion

In this section we sum up the main results of our analysis, providing an answer to the various questions we introduced at the beginning of the paper:

SentiWordNet improves over random. One of the first things worth noting – in Tables 2, 3, 4 and 5 – is that the random approach (*rnd*), as expected, is the worst performing metric, while all other approaches, based on SWN , have statistically significant improvements both for MAE and for Accuracy ($p < 0.001$). So, using SWN for posterior-to-prior polarity computation brings benefits, since it increases the performance above the baseline in words’ prior polarity assessment.

SWN_3 is better than SWN_1 . With respect to

SWN_1 , using SWN_3 enhances performance, both in regression (MAE μ 0.398 vs. 0.366, $p < 0.001$) and classification (Accuracy μ 0.710 vs. 0.771, $p < 0.001$) tasks. Since many of the approaches described in the literature use SWN_1 their results should be revised and SWN_3 should be used as standard. This difference in performance can be partially explained by the fact that, even after pre-processing, for the ANEW dataset 137 lemma#PoS have all senses equal to 0 in SWN_1 , while in SWN_3 they are just 48. In the GI lexicon the numbers are 233 for SWN_1 and 69 for SWN_3 .

Not all formulae are created equal. The formulae described in Section 4 have very different results, along a continuum. While inspecting every differ-

	Acc. μ	Acc. σ
<i>rnd</i>	0.447	0.019
<i>swn_rnd_m</i>	0.639	0.026
<i>swn_rnd_d</i>	0.646	0.021
<i>fs_m</i>	0.659	0.020
<i>uni</i>	0.684	0.017
<i>median</i>	0.686	0.022
<i>uniw</i>	0.702	0.019
<i>max</i>	0.710	0.022
<i>w1</i>	0.712	0.021
<i>w1n</i>	0.713	0.022
<i>w2n</i>	0.714	0.023
<i>w2</i>	0.715	0.021
<i>mean</i>	0.718	0.023
<i>w2s</i>	0.719	0.023
<i>w2sn</i>	0.719	0.023
<i>w1s</i>	0.719	0.023
<i>w1sn</i>	0.719	0.023
<i>GP^{ll}_{linard}</i>	0.721	0.026
<i>GP^{prl}_{linard}</i>	0.722	0.025
<i>SVM</i>	0.733	0.021
<i>SVMfs</i>	0.743	0.021
Average	0.710	0.022

Table 4: Accuracy results for classification using SWN_1

ence in performance is out of the scope of the present paper, we can see that there is a strong difference between best and worst performing formulae both in regression (in Table 2 $w2n_m$ is better than $uniw_m$, in Table 3 $w2n_m$ is better than max_m) and classification (in Table 4 $w1sn_m$ is better than fs_m , in Table 5 $w2_m$ is better than fs_m) and these differences are all statistically significant ($p < 0.001$). Again, these results indicate that the previous experiments in the literature that use SWN as a baseline should be revised to take these results into account. Furthermore, the new formulae we introduced, based on the “posterior polarities saliency” hypothesis, proved to be among the best performing in all experiments. This entails that there is room for inspecting new formulae variants other than those already proposed in the literature.

Selecting just one sense is not a good choice.

On a side note, the approaches that rely on only one sense polarity (namely fs , $median$ and max) have similar results which do not differ significantly from $swnrnd$ (for max_m , fs_d and fs_m in Table 2, and for max_m in Table 3). These same approaches are also far from the best performing formulae: in Table 3, $median_d$ differs from $w2n_m$ ($p < 0.05$), as do max_m , max_d , fs_m and fs_d ($p < 0.001$); in Ta-

	Acc. μ	Acc. σ
<i>rnd</i>	0.447	0.019
<i>swn_rnd_d</i>	0.700	0.030
<i>swn_rnd_m</i>	0.706	0.034
<i>fs</i>	0.723	0.014
<i>median_m</i>	0.742	0.016
<i>uni</i>	0.750	0.015
<i>uniw</i>	0.762	0.023
<i>max</i>	0.769	0.019
<i>w2s</i>	0.777	0.017
<i>w2sn</i>	0.777	0.017
<i>w1s</i>	0.777	0.017
<i>w1sn</i>	0.777	0.017
<i>w1n</i>	0.780	0.021
<i>w2n</i>	0.780	0.022
<i>mean</i>	0.781	0.018
<i>w1</i>	0.781	0.021
<i>w2</i>	0.781	0.021
<i>SVM</i>	0.779	0.016
<i>GP_l</i>	0.779	0.018
<i>GP_g</i>	0.781	0.018
<i>SVMfs</i>	0.792	0.014
Average	0.771	0.018

Table 5: Accuracy results for classification using SWN_3

ble 3, fs , max and $median$ in both their f_m and f_d variants are significantly different from the best performing $w2n_m$ ($p < 0.001$). For classification, in Table 4 and 5 the difference between the corresponding best performing formula and the single senses formulae is always significant (at least $p < 0.01$). Among other things, this finding entails, surprisingly, that taking the first sense of a lemma#POS in some cases has no improvement over taking a random sense, and that in all cases it is one of the worst approaches with SWN . This is surprising since in many NLP tasks, such as word sense disambiguation, algorithms based on most frequent sense represent a very strong baseline⁶.

Learning improvements. Combining the formulae in a learning framework further improves the results over the best performing formulae, both in regression (MAE μ with SWN_1 0.366 vs. 0.391, $p < 0.001$; MAE μ with SWN_3 0.333 vs. 0.359, $p < 0.001$) and in classification (Accuracy μ for SWN_1 is 0.743 vs. 0.719, $p < 0.001$; Accuracy μ for SWN_3 is 0.792 vs. 0.781, not significant $p = 0.07$). Another thing worth noting is that, in regression, GPs are outperformed by both versions of

⁶In SemEval 2010, only 5 participants out of 29 performed better than the most frequent threshold (Agirre et al., 2010).

SVM ($p < 0.001$), see Tables 2 and 3. This is in contrast with the results presented in (Cohn and Specia, 2013), where GPs on the single task are on average better than SVMs. In classification, GPs have similar performance to SVM without feature selection, and in some cases (see Table 5) even slightly better. Analyzing the selected kernels for GPs and SVMs, we notice that in most of the splits SVMs prefer the radial based function, while the best performance with the GPs are obtained with linear kernels with and without ARD. There is no significant difference in using linear logistic and probit regression likelihoods. In all our experiments, SVM with feature selection leads to the best performance. This is not surprising due the high level of redundancy in the formulae scores. Interestingly, inspecting the most frequent selected features by *SVMfs*, we see that features from different groups are selected, and even the worst performing formulae can add information, confirming the idea that viewing the same information from different perspectives (i.e. the posterior polarities provided by SWN combined in various ways) can give better predictions.

To sum up: the new state-of-the-art performance level in prior-polarity computation is represented by the *SVMfs* approach using *SWN₃*, and this should be used as the reference from now on.

9 PoS and Gender Experiments

Next, we wanted to understand if the performance of our approach, using *SWN₃*, was consistent across word PoS. In Table 6 we report the results for the best performing formulae and learning algorithm on the GI PoS classes. In particular for ADJ there are 1,073 words, 922 for NOUN and 508 for VERB. We discarded adverbs since the class was too small to allow reliable evaluation and efficient learning (only 54 instances). The results show a greater accuracy for adjectives ($p < 0.01$), while performance for nouns and verbs are similar.

	SVMfs		<i>best_f</i>	
	Acc. μ	Acc. σ	Acc. μ	Acc. σ
ADJ	0.829	0.019	0.821	0.016
NOUN	0.784	0.021	0.765	0.023
VERBS	0.782	0.052	0.744	0.046

Table 6: Accuracy results for PoS using *SWN₃*

Finally we test against the male and female ratings provided by ANEW. As can be seen from Table 7, SWN approaches are far more precise in predicting Male judgments rather than Female ones (MAE_{μ} goes from 0.392 to 0.323 with the best formula and from 0.369 to 0.292 with *SVMfs*, both differences are significant $p < 0.001$). Instead, in Table 8 – which displays the results along gender and polarity dimensions – there is no statistically significant difference in *MAE* on positive words between male and female, while there is a strong statistical significance for negative words ($p < 0.001$).

Interestingly, there is also a large difference between positive and negative affective words (both for male and female dimensions). This difference is maximum for male scores on positive words compared to female scores on negative words (0.283 vs. 0.399, $p < 0.001$). Recent work by Warriner et al. (2013) inspected the differences in prior polarity assessment due to gender.

At this stage we can only note that prior polarities calculated with SWN are closer to ANEW male annotations than female ones. Understanding why this happens would require an accurate examination of the methods used to create WordNet and SWN (which will be the focus of our future work).

	Male		female	
	MAE μ	MAE σ	MAE μ	MAE σ
SVMfs	0.292	0.020	0.369	0.008
best_f	0.323	0.022	0.392	0.010

Table 7: MAE results for Male vs Female using *SWN₃*

	Male		female	
	MAE μ	MAE σ	MAE μ	MAE σ
Pos	0.283	0.022	0.340	0.009
Neg	0.301	0.029	0.399	0.013

Table 8: MAE for Male/Female - Pos/Neg using *SWN₃*

10 Conclusions

We have presented a study on the posterior-to-prior polarity issue, i.e. the problem of computing words’ prior polarity starting from their posterior polarities. Using two different versions of SentiWordNet and 30 different approaches that have been proposed in the literature, we have shown that researchers have not paid sufficient attention to this issue. Indeed, we

showed that the better variants outperform the others on different datasets both in regression and classification tasks, and that they can represent a fairer state-of-art baseline approach using SentiWordNet. On top of this, we also showed that these state-of-the-art formulae can be further outperformed using a learning framework that combines the various formulae together. We conclude our analysis with some experiments investigating the impact of word PoS and annotator gender in gold standards, showing interesting phenomena that requires further investigation.

Acknowledgments

The authors thanks José Camargo De Souza for his help with feature selection. This work has been partially supported by the Trento RISE PerTe project.

References

- E. Agirre, O.L. De Lacalle, C. Fellbaum, S.K. Hsieh, M. Tesconi, M. Monachini, P. Vossen, and R. Segers. 2010. Semeval-2010 task 17: All-words word sense disambiguation on a specific domain. In *Proceedings of the 5th International Workshop on Semantic Evaluation (IWSE '10)*, pages 75–80, Uppsala, Sweden.
- S. Agrawal and T.J. Siddiqui. 2009. Using syntactic and contextual information for sentiment polarity analysis. In *Proceedings of the 2nd International Conference on Interaction Sciences: Information Technology, Culture and Human (ICIS '09)*, pages 620–623, Seoul, Republic of Korea.
- S. Baccianella, A. Esuli, and F. Sebastiani. 2010. SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the 7th Conference on International Language Resources and Evaluation (LREC '10)*, pages 2200–2204, Valletta, Malta.
- M.M. Bradley and P.J. Lang. 1999. Affective norms for English words (ANEW): Instruction manual and affective ratings. Technical Report C-1, University of Florida.
- C.C. Chang and C.J. Lin. 2011. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27.
- F.R. Chaumartin. 2007. UPAR7: A knowledge-based system for headline sentiment tagging. In *Proceedings of the 4th International Workshop on Semantic Evaluations (IWSE '07)*, pages 422–425, Prague, Czech Republic.
- F.M. Chowdhury, M. Guerini, S. Tonelli, and A. Lavelli. 2013. Fbk: Sentiment analysis in twitter with tweetsted. In *Second Joint Conference on Lexical and Computational Semantics (*SEM): Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval '13)*, volume 2, pages 466–470, Atlanta, Georgia, USA, June.
- T. Cohn and L. Specia. 2013. Modelling annotator bias with multi-task gaussian processes: An application to machine translation quality estimation. In *Proceedings of the 51th Annual Meeting of the Association for Computational Linguistics (ACL '13)*, pages 32–42, Sofia, Bulgaria.
- K. Denecke. 2008. Accessing medical experiences and information. In *Proceedings of the 18th European Conference on Artificial Intelligence, Workshop on Mining Social Data (MSoDa '08)*, Patras, Greece.
- K. Denecke. 2009. Are SentiWordNet scores suited for multi-domain sentiment classification? In *Proceedings of the 4th International Conference on Digital Information Management (ICDIM '09)*, pages 32–37, Ann Arbor, MI, USA.
- A. Devitt and K. Ahmad. 2007. Sentiment polarity identification in financial news: A cohesion-based approach. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics (ACL '07)*, pages 984–991, Prague, Czech Republic.
- D.C. Dunphy, C.G. Bullard, and E.E.M. Crossing. 1974. Validation of the General Inquirer Harvard IV Dictionary. Paper presented at the Pisa Conference on Content Analysis.
- A. Esuli and F. Sebastiani. 2006. SentiWordNet: A publicly available lexical resource for opinion mining. In *Proceedings of the 5th Conference on International Language Resources and Evaluation (LREC '06)*, pages 417–422, Genova, Italy.
- L. Gatti and M. Guerini. 2012. Assessing sentiment strength in words prior polarities. In *Proceedings of the 24th International Conference on Computational Linguistics (COLING '12)*, pages 361–370, Mumbai, India.
- R. Giora. 1997. Understanding figurative and literal language: The graded salience hypothesis. *Cognitive Linguistics*, 8:183–206.
- M. Guerini, O. Stock, and C. Strapparava. 2008. Valentino: A tool for valence shifting of natural language texts. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC '08)*, pages 243–246, Marrakech, Morocco.
- H.D. Lasswell and J.Z. Namenwirth. 1969. The Lasswell value dictionary. *New Haven*.
- B. Liu and L. Zhang. 2012. A survey of opinion mining and sentiment analysis. *Mining Text Data*, pages 415–463.

- N. Meinshausen and P. Bühlmann. 2010. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473.
- A. Neviarouskaya, H. Prendinger, and M. Ishizuka. 2009. Sentiful: Generating a reliable lexicon for sentiment analysis. In *Proceedings of the 3rd Affective Computing and Intelligent Interaction (ACII '09)*, pages 363–368, Amsterdam, Netherlands.
- A. Neviarouskaya, H. Prendinger, and M. Ishizuka. 2011. Affect analysis model: novel rule-based approach to affect sensing from text. *Natural Language Engineering*, 17(1):95.
- G. Özbal and C. Strapparava. 2012. A computational approach to the automation of creative naming. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (ACL '12)*, pages 703–711, Jeju Island, Korea.
- G. Özbal, C. Strapparava, and M. Guerini. 2012. Brand Pitt: A corpus to explore the art of naming. In *Proceedings of the 8th International Conference on Language Resources and Evaluation (LREC '12)*, pages 1822–1828, Istanbul, Turkey.
- G. Paltoglou, M. Thelwall, and K. Buckley. 2010. Online textual communications annotated with grades of emotion strength. In *Proceedings of the 3rd International Workshop of Emotion: Corpora for research on Emotion and Affect (satellite of LREC '10)*, pages 25–31, Valletta, Malta.
- B. Pang and L. Lee. 2008. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1-2):1–135.
- E. Pianta, C. Girardi, and R. Zanolì. 2008. The TextPro tool suite. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC '08)*, pages 2603–2607, Marrakech, Morocco.
- I. Piller. 2003. Advertising as a site of language contact. *Annual Review of Applied Linguistics*, 23:170–183.
- T. Polajnar, S. Rogers, and M. Girolami. 2011. Protein interaction detection in sentences via gaussian processes: a preliminary evaluation. *International journal of data mining and bioinformatics*, 5(1):52–72.
- L. Qu, C. Toprak, N. Jakob, and I. Gurevych. 2008. Sentence level subjectivity and sentiment analysis experiments in NTCIR-7 MOAT challenge. In *Proceedings of the 7th NTCIR Workshop Meeting (NTCIR '08)*, pages 210–217, Tokyo, Japan.
- C.E. Rasmussen and C.K.I. Williams. 2006. *Gaussian processes for machine learning*. MIT Press.
- J. Shawe-Taylor and N. Cristianini. 2004. *Kernel methods for pattern analysis*. Cambridge university press.
- J.K. Sing, S. Sarkar, and T.K. Mitra. 2012. Development of a novel algorithm for sentiment analysis based on adverb-adjective-noun combinations. In *Proceedings of the 3rd National Conference on Emerging Trends and Applications in Computer Science (NC-ETACS '12)*, pages 38–40, Shillong, India.
- P.J. Stone, D.C. Dunphy, and M.S. Smith. 1966. *The General Inquirer: A Computer Approach to Content Analysis*. MIT press.
- C. Strapparava and A. Valitutti. 2004. WordNet-Affect: an affective extension of WordNet. In *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC '04)*, pages 1083 – 1086, Lisbon, Portugal.
- M. Taboada, J. Brooke, M. Tofiloski, K. Voll, and M. Stede. 2011. Lexicon-based methods for sentiment analysis. *Computational linguistics*, 37(2):267–307.
- T.T. Thet, J.C. Na, C.S.G. Khoo, and S. Shakthikumar. 2009. Sentiment analysis of movie reviews on discussion boards using a linguistic approach. In *Proceedings of the 1st international CIKM workshop on Topic-sentiment analysis for mass opinion (TSA '09)*, pages 81–84, Hong Kong.
- S. Wang and C.D. Manning. 2012. Baselines and bigrams: Simple, good sentiment and topic classification. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (ACL '12)*, pages 90–94, Jeju Island, Korea.
- A.B. Warriner, V. Kuperman, and M. Brysbaert. 2013. Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior research methods*, pages 1–17.
- J. Weston, S. Mukherjee, O. Chapelle, M. Pontil, T. Poggio, and V. Vapnik. 2000. Feature selection for SVMs. In *Proceedings of the 14th Conference on Neural Information Processing Systems (NIPS '00)*, pages 668–674, Denver, CO, USA.
- C.K.I. Williams and D. Barber. 1998. Bayesian classification with gaussian processes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 20(12):1342–1351.
- T. Wilson, J. Wiebe, and R. Hwa. 2004. Just how mad are you? Finding strong and weak opinion clauses. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI '04)*, pages 761–769, San Jose, CA, USA.
- T. Wilson, J. Wiebe, and P. Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing (HLT/EMNLP '05)*, pages 347–354, Vancouver, Canada.
- A. Yeh. 2000. More accurate tests for the statistical significance of result differences. In *Proceedings of the 18th International Conference on Computational Linguistics (COLING '00)*, pages 947–953, Saarbrücken, Germany.