

Generating Sentiment Lexicons for German Twitter

Uladzimir Sidarenka and Manfred Stede

Applied Computational Linguistics

UFS Cognitive Science

University of Potsdam / Germany

{sidarenk|stede}@uni-potsdam.de

Abstract

Despite substantial progress made in developing new sentiment lexicon generation (SLG) methods for English, the task of transferring these approaches to other languages and domains in a sound way still remains open. In this paper, we contribute to the solution of this problem by systematically comparing semi-automatic translations of common English polarity lists with the results of the original automatic SLG algorithms, which were applied directly to German data. We evaluate these lexicons on a corpus of 7,992 manually annotated tweets. In addition to that, we also collate the results of dictionary- and corpus-based SLG methods in order to find out which of these paradigms is better suited for the inherently noisy domain of social media. Our experiments show that semi-automatic translations notably outperform automatic systems (reaching a macro-averaged F_1 -score of 0.589), and that dictionary-based techniques produce much better polarity lists as compared to corpus-based approaches (whose best F_1 -scores run up to 0.479 and 0.419 respectively) even for the non-standard Twitter genre. All reimplementations of the compared systems and the resulting lexicons of these methods are available online at <https://github.com/WladimirSidorenko/SentiLex>.

1 Introduction

Sentiment lexicons play a crucial role in many existing and emerging opinion mining applications. Not only do they serve as a valuable source of features for supervised classifiers (Mohammad et al., 2013; Zhu et al., 2014) but they also achieve competitive results when used as the main component of a sentiment analysis system (Taboada et al., 2011). Due to this high impact and tremendous costs of building such lexicons manually, devising new algorithms for an automatic generation of polarity lists has always been an area of active research in the sentiment analysis literature (Liu, 2012, pp. 79-91). Nevertheless, despite some obvious progress in this field (Cambria et al., 2016), the applicability of these approaches to other languages and text genres still raises questions: It is, for instance, unclear whether simply translating the existing English sentiment resources would produce better results than applying the methods that were initially proposed for their creation directly to the target language. Furthermore, for automatic systems which draw their knowledge from lexical taxonomies, such as WORDNET (Miller, 1995), it remains unanswered whether these approaches would also work for languages in which such resources are much smaller in size, and, even if they would, whether the resulting lexicons would then be general enough to carry over to more colloquial texts. Finally, for methods which derive their polarity lists from text corpora, it is not clear whether these approaches would still yield an acceptable quality when operating on inherently noisy input data.

In this paper, we try to analyze these and other problems in detail, using the example of German Twitter. More precisely, given a collection of German microblogs with manually labeled polar terms and prior polarities of these expressions, we want to find an SLG method that can best predict these terms and their semantic orientation. For this purpose, we compare the existing German sentiment lexicons

This work is licenced under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

(most of which were semi-automatically translated from popular English resources) with the results of common automatic dictionary- and corpus-based SLG approaches.

We begin our study by describing the data set which will be used in our evaluation. Afterwards, in Section 3, we introduce the metrics with which we will assess the quality of various polarity lists. Then, in Section 4, we evaluate three most popular existing German sentiment lexicons—the German Polarity Clues (Waltinger, 2010), SentiWS (Remus et al., 2010), and Zurich Polarity List of Clematide and Klenner (2010), subsequently comparing them with popular automatic SLG approaches in Section 5. Finally, after estimating the impact of different seed sets on the automatic methods and performing a qualitative analysis of their entries, we draw our conclusions and outline directions for future research in the final part of this paper.

To avoid unnecessary repetitions, we deliberately omit a summary of related work, since most of the popular SLG algorithms will be referenced in the respective evaluation sections anyway. We should, however, note that, apart from the research on the automatic lexicon generation, our study is also closely related to the experiments of Andreevskaia and Bergler (2008) and the “Sentiment Analysis in Twitter” track of the SemEval competition (Nakov et al., 2013; Rosenthal et al., 2014; Rosenthal et al., 2015). In contrast to the former work, however, where the authors trained a supervised classifier on one domain and applied it to another in order to determine the polarities of the sentences, we *explicitly model a situation where no annotated training data are available*, thus looking for the most general unsupervised SLG strategy which performs best regardless of the target domain, and we also *evaluate these strategies on the level of lexical phrases only*. Furthermore, unlike in the SemEval track, where the organizers also provided participants with sufficient labeled in-domain training sets and then asked them to predict the contextual polarity of pre-annotated polar expressions in the test data, we *simultaneously try to predict polar terms and their prior polarities, learning both of them without supervision*.

2 Data

We perform our evaluation on the publicly available Potsdam Twitter Sentiment corpus (PotTS; Sidarenka, 2016).¹ This collection comprises 7,992 microblogs pertaining to the German federal elections, general political life, papal conclave 2013, as well as casual everyday conversations. Two human experts annotated these posts with polar terms and their prior polarities,² reaching a substantial agreement of 0.75 binary κ (Cohen, 1960).³ We used the complete data set labeled by one of the annotators as our test corpus, getting a total of 6,040 positive and 3,055 negative terms including multi-word expressions. However, since many of these expressions were emoticons, which, on the one hand, were a priori absent in common lexical taxonomies due to their colloquial nature and therefore not amenable to dictionary-based SLG systems but, on the other hand, could be easily captured by regular expressions, we decided to exclude non-alphabetic smileys altogether from our study. This left us with a set of 3,459 positive and 2,755 negative labeled terms (1,738 and 1,943 unique expressions respectively), whose κ -agreement run up to 0.59. Besides the test set, we selected a small subset of 400 tweets from the other annotator and used it as development data for tuning the hyper-parameters of the tested approaches.⁴

3 Evaluation Metrics

A central question to our experiments are the evaluation metrics that we should use for measuring lexicon quality. Usually, this quality is estimated either *intrinsically* (i.e., taking a lexicon in isolation and immediately assessing its accuracy) or *extrinsically* (i.e., considering the lexicon within the scope of a bigger application such as a supervised classifier which utilizes lexicon’s entries as features).

¹We use version 0.1.0 of this corpus.

²The annotators had been asked to judge the semantic orientation of a term irrespective of its possible negations. They could, however, consider the context for determining whether a particular reading of a polysemous word in the text was subjective or not.

³A detailed inter-annotator agreement study of this corpus is provided in (Sidarenka, 2016).

⁴That way, we only used the labeled corpus for evaluation or parameter optimization, other resources—GERMANET (Hamp and Feldweg, 1997) and the German Twitter Snapshot (Scheffler, 2014)—were used for training the methods.

Traditionally, intrinsic evaluation of English sentiment lexicons amounts to comparing these polarity lists with the General Inquirer (GI; Stone, 1966)—a manually compiled set of 11,895 words annotated with their semantic categories—by taking the intersection of the two resources and estimating the percentage of matches in which automatically induced polar terms have the same polarity as the GI entries. This evaluation method, however, is somewhat problematic: First of all, it is not easily transferable to other languages, since even a manual translation of the GI lexicon is not guaranteed to cover all language- and domain-specific polar expressions. Secondly, due to the intersection, this method does not penalize for a low recall so that a lexicon consisting of just two terms *good*⁺ and *bad*⁻ will have the highest possible score, often surpassing polarity lists with a greater number of entries. Finally, this comparison does not account for polysemy. As a result, an ambiguous word only one of whose (possibly rare) senses is subjective will always be ranked the same as a purely polar term.

Unfortunately, an extrinsic evaluation does not always provide a solution in this case, since, depending on the type of the extrinsic system (e.g., a document classifier), it might still presuppose a large data set for training other system components and, furthermore, might yield overly high scores, which, however, are mainly due to these extrinsic modules rather than the quality of the lexicons themselves.

Instead of using these approaches, we opt for a direct comparison of the induced polarity lists with an existing annotated corpus, since this type of evaluation allows us to solve at least three of the previously mentioned issues: It does account for the recall, it does accommodate polysemous words,⁵ and it does preclude intermediate components which might artificially boost the results. In particular, in order to check a lexicon against the PotTS data set, we construct a case-insensitive trie (Knuth, 1998, pp. 492–512) from the lexicon entries and match this trie against the contiguously running corpus text,⁶ simultaneously comparing it with the actual word forms and lemmas of corpus tokens.⁷ A match is considered correct iff the matched entry absolutely corresponds to the (possibly lemmatized) expert’s annotation and has the same polarity as the one specified by the human coder. That way, we estimate the precision, recall, and F_1 -score for each particular polarity class (positive, negative, and neutral), considering all words absent in the lexicons (not annotated in the corpus) as neutral.

4 Semi-Automatic Lexicons

We first apply the above metric to estimate the quality of the existing German resources: the German Polarity Clues (GPC; Waltinger, 2010), SentiWS (SWS; Remus, 2010), and the Zurich Polarity List (ZPL) of Clematide and Klenner (2010).

The GPC set comprises 10,141 subjective entries automatically translated from the English sentiment lexicons Subjectivity Clues (Wilson et al., 2005) and SentiSpin (Takamura et al., 2005), with a subsequent manual correction of these translations, and several synonyms and negated terms added by the authors. The SWS lexicon includes 1,818 positively and 1,650 negatively connoted terms, also providing their part-of-speech tags and inflections (resulting in a total of 32,734 word forms). Similarly to the GPC, the authors used an English sentiment resource—the GI lexicon of Stone et al. (1966)—to bootstrap their polarity list, manually revising these automatic translations afterwards. In addition to that, Remus et al. (2010) also expanded their set with words and phrases frequently co-occurring with positive and negative seed lexemes using collocation information obtained from a corpus of 10,200 customer reviews and the German Collocation Dictionary (Quasthoff, 2010). Finally, the Zurich Polarity List features 8,000 subjective entries taken from GERMANET synsets (Hamp and Feldweg, 1997). These synsets were manually annotated with their prior polarities by human experts. Since the authors, however, found the number of polar adjectives obtained that way insufficient for running further classification experiments, they automatically enriched this lexicon with more attributive terms by analyzing conjoined corpus collocations using the method of Hatzivassiloglou and McKeown (1997).

⁵Recall that the annotators of the PotTS data set were asked to annotate a polar expression iff its actual sense in the respective context was polar.

⁶In other words, we successively compare lexicon entries with the occurrences of corpus tokens in the same linear order as these occurrences appear in the text.

⁷We use the TREE TAGGER of Schmid (1995) for lemmatization.

Lexicon	Positive Expressions			Negative Expressions			Neutral Terms			Macro F_1	Micro F_1
	Precision	Recall	F_1	Precision	Recall	F_1	Precision	Recall	F_1		
GPC	0.209	0.535	0.301	0.195	0.466	0.275	0.983	0.923	0.952	0.509	0.906
SWS	0.335	0.435	0.379	0.484	0.344	0.402	0.977	0.975	0.976	0.586	0.952
ZPL	0.411	0.424	0.417	0.38	0.352	0.366	0.977	0.979	0.978	0.587	0.955
GPC \cap SWS \cap ZPL	0.527	0.372	0.436	0.618	0.244	0.35	0.973	0.99	0.982	0.589	0.964
GPC \cup SWS \cup ZPL	0.202	0.562	0.297	0.195	0.532	0.286	0.985	0.917	0.95	0.51	0.901

Table 1: Evaluation of semi-automatic German sentiment lexicons.

GPC – German Polarity Clues (Waltinger, 2010), SWS – SentiWS (Remus et al., 2010), ZPL – Zurich Polarity Lexicon (Clematide and Klenner, 2010)

For our evaluation, we tested the three lexicons in isolation and also built their union and intersection in order to check for “synergy” effects. The results are shown in Table 1. As can be seen from the statistics, with a few exceptions, the highest scores for all classes as well as the best macro- and micro-averaged F_1 -measures are achieved by the intersection of all three lexicons. On the other hand, as expected, the highest recall of polar expressions (and consequently the best precision at recognizing neutral terms) is attained by the union of these resources. The only case where individual lexicons are able to outperform these combinations is observed for the F_1 -score of the negative class, where both SentiWS and ZPL show better results than their intersection, which is mainly due to the higher recall of these two polarity lists.

5 Automatic Methods

A natural question which arises upon the evaluation of the existing semi-automatic resources is how well fully automatic methods can perform in comparison with these lexicons. Traditionally, automatic SLG algorithms have been grouped into dictionary- and corpus-based ones, with their own complementary strengths and weaknesses. Dictionary-based approaches, for instance, incorporate distilled linguistic knowledge from a typically manually labeled lexical database, but lack any domain specificity. Corpus-based methods, on the other hand, can operate directly on unannotated in-domain data, but often have to deal with an extreme noisiness of their input. Since it was unclear which of these properties would have a stronger impact on the net results, we decided to reimplement the most commonly used algorithms from both of these paradigms and evaluate them on the PotTS corpus.

5.1 Dictionary-Based Approaches

For dictionary-based methods, we adopted the systems proposed by Hu and Liu (2004), Blair-Goldensohn et al. (2008), Kim and Hovy (2004), Esuli and Sebastiani (2006), as well as the min-cut and label-propagation approaches of Rao and Ravichandran (2009), and the random-walk algorithm described by Awadallah and Radev (2010).

The first of these works (Hu and Liu, 2004) expanded a given set of seed terms with known semantic orientations by propagating polarity values of these terms to their WORDNET synonyms and passing reversed polarity scores to the antonyms of these words. Later on, this idea was further refined by Blair-Goldensohn et al. (2008), who obtained polarity labels for new terms by multiplying a score vector \vec{v} containing the orientation scores of the known seed words (-1 for negative expressions and 1 for positive ones) with an adjacency matrix A constructed for the WORDNET graph. With various modifications, the core idea of passing the polarity values through a lexical graph was adopted in almost all of the following dictionary-based works: Kim and Hovy (2004), for instance, computed the polarity class for a new word w by multiplying the prior probability of this class with the likelihood of the word w occurring among the synonyms of the seed terms with the given semantic orientation, choosing at the end the polarity which maximized this equation. Other ways of bootstrapping polarity lists were proposed by Esuli and Sebastiani (2006), who created their SENTIWORDNET resource using a committee of Rocchio and SVM classifiers trained on successively expanded sets of polar terms; Rao and Ravichandran (2009), who adopted the min-cut approach of Blum et al. (2004), also comparing it with the label-propagation algorithm of Zhu and Ghahramani (2002); and, finally, Awadallah and Radev (2010), who used a random

walk method by estimating the polarity of an unknown word as the difference between an average number of steps a random walker had to make in order to reach a term from the positive or negative set.

Since some of these approaches relied on different seed sets or pursued different objectives (two-versus three-way classification), we decided to unify their settings and interfaces for the sake of our experiments. In particular, we were using the same translated seed list of Turney and Littman (2003) for all methods, expanding this set by 10 neutral terms (“neutral” *neutral*, “sachlich” *objective*, “technisch” *technical*, “finanziell” *financial* etc.).⁸ Additionally, we enhanced all binary systems to ternary classifiers, so that each tested method could differentiate between positive, negative, and neutral terms. In the final step, we applied these methods to GERMANET (Hamp and Feldweg, 1997)—a German equivalent of the English WORDNET (Miller, 1995), which, however, is much smaller in size, having 20,792 less synsets for the three common parts of speech (nouns, adjectives, and verbs) than the Princeton resource.

Lexicon	# of Terms	Positive Expressions			Negative Expressions			Neutral Terms			Macro F_1	Micro F_1
		Precision	Recall	F_1	Precision	Recall	F_1	Precision	Recall	F_1		
SEED SET	20	0.771	0.102	0.18	0.568	0.017	0.033	0.963	0.999	0.981	0.398	0.962
HL	5,745	0.161	0.266	0.2	0.2	0.133	0.16	0.969	0.96	0.965	0.442	0.93
BG	1,895	0.503	0.232	0.318	0.285	0.093	0.14	0.968	0.991	0.979	0.479	0.959
KH	356	0.716	0.159	0.261	0.269	0.044	0.076	0.965	0.997	0.981	0.439	0.962
ES	39,181	0.042	0.564	0.078	0.033	0.255	0.059	0.981	0.689	0.81	0.315	0.644
RR _{minicut}	8,060	0.07	0.422	0.12	0.216	0.073	0.109	0.972	0.873	0.92	0.383	0.849
RR _{lbl-prop}	1,105	0.567	0.176	0.269	0.571	0.046	0.085	0.965	0.997	0.981	0.445	0.962
AR	23	0.768	0.1	0.176	0.568	0.017	0.033	0.963	0.999	0.981	0.397	0.962
HL \cap BG \cap RR _{lbl-prop}	752	0.601	0.165	0.259	0.567	0.045	0.084	0.965	0.997	0.981	0.441	0.962
HL \cup BG \cup RR _{lbl-prop}	6,258	0.166	0.288	0.21	0.191	0.146	0.165	0.97	0.958	0.964	0.446	0.929

Table 2: Evaluation of dictionary-based approaches.

HL – Hu and Liu (2004), BG – Blair-Goldensohn et al. (2008), KH – Kim and Hovy (2004), ES – Esuli and Sebastiani (2006), RR – Rao and Ravichandran (2009), AR – Awadallah and Radev (2010)

The results of this evaluation are shown in Table 2. This time, the situation is much more varied, as different systems can achieve best results on just some aspects of certain classes but can hardly attain best overall scores in all categories. This is, for instance, the case for the positive and negative polarities, where the best precision scores are reached by the seed set in the first case and the label propagation algorithm of Rao and Ravichandran (2009) in the second case. However, with respect to the recall, both of these polarity lists perform notably worse than the approach of Esuli and Sebastiani (2006). Yet other systems—the matrix-vector method of Blair-Goldensohn et al. (2008) and the union of the three overall top-scoring systems respectively—reach the highest F_1 -scores for these two classes. Nevertheless, we can still notice three main tendencies in this evaluation: *i*) the method of Esuli and Sebastiani (2006) generally gets the highest recall of polar terms and, consequently, achieves the best precision in recognizing neutral words, but suffers from a low precision for the positive and negative polarities; *ii*) simultaneously five systems attain the same best F_1 -scores on recognizing neutral terms, which, in turn, leads to the best micro-averaged F_1 -results for all polarity classes; and, finally, *iii*) the system of Blair-Goldensohn et al. (2008) shows the best macro-averaged performance. This approach, however, is extremely susceptible to its hyper-parameter settings (in particular, we considered the maximum number of times the initial vector \vec{v} was multiplied with the adjacency matrix A as such a parameter and noticed a dramatic decrease of method’s scores after the fifth iteration).

5.2 Corpus-Based Approaches

An alternative way to generate polarity lists is to use corpus-based approaches. In contrast to dictionary-based methods, these systems typically operate immediately on raw texts and are, therefore, virtually independent of any manually annotated linguistic resources. This flexibility, however, might come at the cost of a reduced accuracy due to an inherent noisiness of the unlabeled data. The most prominent representatives of this class of algorithms are the approaches proposed by Takamura et al. (2005), Velikovich et al. (2010), Kiritchenko et al. (2014), and Severyn and Moschitti (2015), which we briefly describe in this section.

⁸All translated seed sets are provided along with the source code for this paper.

Drawing on the pioneering work of Hatzivassiloglou and McKeown (1997), in which the authors expanded an initial list of polar adjectives by analyzing coordinately conjoined terms from a text corpus, Takamura et al. (2005) enhanced this algorithm, extending it to other parts of speech and also incorporating semantic links from WORDNET in addition to the co-occurrence statistics extracted from the corpus. After representing the final set of terms as an electron lattice, whose edge weights corresponded to the contextual and semantic links between words, the authors computed the most probable polarity distribution for this lattice by adopting the Ising spin model from statistical mechanics.

The approach of Velikovich et al. (2010) was mainly inspired by the label-propagation algorithm of Rao and Ravichandran (2009), with the crucial difference that, instead of taking an averaged sum of the adjacent neighbor values when propagating the label scores through the graph, the authors took the maximum of these scores in order to prune unreliable, noisy corpus links. Similarly, Kiritchenko et al. (2014) built on the method of Turney and Littman (2003) and computed polarity scores for new words by taking the difference of their PMI associations with noisy labeled positive and negative classes. Finally, Severyn and Moschitti (2015) trained a supervised SVM classifier on a distantly labeled data set and included the top-ranked unigram and bigram features in their final lexicon.

For our evaluation, we applied these methods to the German Twitter Snapshot (Scheffler, 2014)—a collection of 24 M microblogs gathered in April, 2013, constructing the collocation graph from the lemmatized word forms of this corpus and only considering words which appeared at least four times in the analyzed data. We again were using the TREETAGGER of Schmid (1995) for lemmatization and GERMANET (Hamp and Feldweg, 1997) for deriving semantic links between word vertices for the method of Takamura et al. (2005).

Lexicon	# of Terms	Positive Expressions			Negative Expressions			Neutral Terms			Macro F_1	Micro F_1
		Precision	Recall	F_1	Precision	Recall	F_1	Precision	Recall	F_1		
SEED SET	20	0.771	0.102	0.18	0.568	0.017	0.033	0.963	0.999	0.981	0.398	0.962
TKM	920	0.646	0.134	0.221	0.565	0.029	0.055	0.964	0.998	0.981	0.419	0.962
VEL	60	0.764	0.102	0.18	0.568	0.017	0.033	0.963	0.999	0.98	0.398	0.962
KIR	320	0.386	0.106	0.166	0.568	0.017	0.033	0.963	0.996	0.979	0.393	0.959
SEV	60	0.68	0.102	0.177	0.568	0.017	0.033	0.963	0.999	0.981	0.397	0.962
TKM \cap VEL \cap SEV	20	0.771	0.102	0.18	0.568	0.017	0.033	0.963	0.999	0.981	0.398	0.962
TKM \cup VEL \cup SEV	1,020	0.593	0.134	0.218	0.565	0.029	0.055	0.964	0.998	0.98	0.418	0.962

Table 3: Evaluation of corpus-based approaches.

TKM – Takamura et al. (2005), VEL – Velikovich et al. (2010), KIR – Kiritchenko et al. (2014), SEV – Severyn and Moschitti (2015)

The results of these experiments are shown in Table 3. This time, we can observe a clear superiority of Takamura et al.’s method, which not only achieves the best recall and F_1 in recognizing positive and negative items but also attains the highest micro- and macro-averaged results for all three polarity classes. The cardinality of the other induced lexicons, however, is much smaller than the size of Takamura et al.’s polarity list. Moreover, these lexicons also show absolutely identical scores for the negative expressions as the original seed set. Since these results were somewhat unexpected, we decided to investigate the reasons for possible problems. As it turned out, the macro-averaged F_1 -values of these methods were rapidly going down on the held-out development set as the number of their induced polar terms increased. Since we considered the lexicon size as one of the hyper-parameters of the tested approaches, we immediately stopped populating these lexicons when we noticed a decrease in their results. As a consequence, only the highest-ranked terms (all of which had the positive polarity) were included in the final lists.

One of the reasons for such rapid quality decrease was the surprisingly high positive bias of the initial seed set: While converting the original seed list of Turney and Littman (2003) to German, we translated the English word “correct” as “richtig”. This German word, however, also has another reading which means *real* (as in *a real fact* or *a real sports car*) and which was much more frequent in the analyzed snapshot, often appearing in an unequivocally negative context, e.g., “ein richtiger Bombenanschlag” (*a real bomb attack*) or “ein richtiger Terrorist” (*a real terrorist*). As a consequence of this, methods relying on distant supervision had to deal with an extremely unbalanced training set (the automatically labeled corpus that we distantly obtained for the approach of Kiritchenko et al. (2014) using these seeds,

for instance, had 716,210 positive versus 92,592 negative training instances).

6 Effect of Seed Sets

Since the set of the initial seed terms appeared to play an important role for at least three of the tested methods, we decided to analyze the impact of this factor in more detail by repeating our experiments with the seed lists proposed by Hu and Liu (2004), Kim and Hovy (2004), Esuli and Sebastiani (2006), and Remus et al. (2010). For this purpose, we manually translated the seed sets of Hu and Liu (2004) and Kim and Hovy (2004) into German. Since the authors, however, only provided some examples of their seeds without specifying the full lists, we filled up our translations with additional polar terms to match the original cardinalities. A different procedure was applied to obtain the seed set of Esuli and Sebastiani (2006)—since this resource comprised a vast number of neutral terms (the authors considered as neutral all words from the General Inquirer lexicon which were not marked there as either positive or negative), we automatically translated the neutral subset of these seeds with the help of a publicly available translation site (<http://www.dict.cc>), using the first suggestion returned by this service for each original English term.




Figure 1: Macro-averaged F_1 -scores of the dictionary-based approaches with different seed sets.

The updated results for the dictionary-based approaches with the alternative seed sets are shown in Figure 1. This time, we again can notice superior scores achieved by the method of Blair-Goldensohn et al. (2008), which not only performs better than the other systems on average but also seems to be less susceptible to the varying quality and size of the different seed lists. The remaining methods typically achieve their best macro-averaged results with either of the two top-scoring polarity sets—the seed list of Kim and Hovy (2004) or the seed set of Esuli and Sebastiani (2006). This is, for instance, the case for the method of Kim and Hovy (2004) and the min-cut approach of Rao and Ravichandran (2009), whose performance with the native Kim-Hovy seed set is on par with their results achieved using the Turney-Littman seeds. The label-propagation and random walk algorithms can even strongly benefit from the seeds provided by Kim and Hovy (2004). The remaining two methods—Hu and Liu (2004) and Esuli and Sebastiani (2006)—work best in combination with the initial polarity set proposed by Esuli and Sebastiani (2006).

A slightly different situation is observed for the corpus-based approaches as shown in Figure 2. Except for the method of Takamura et al. (2005), all three remaining methods—Velikovich et al. (2010), Kiritchenko et al. (2014), and Severyn and Moschitti (2015)—show very similar (though not identical) scores. Moreover, these scores are also very close to the results achieved by the respective seed sets without any expansion. The primary reasons for this were again the positive bias of the distantly labeled tweets and the consequently premature stopping of the expansion.

Following the suggestion of one of the reviewers, we additionally included two more seed sets in our evaluation: gold precision and emoticons. The former list contained just two polar terms—“gut” (*good*⁺) and “schlecht” (*bad*⁻)—which showed an almost perfect precision on the PotTS data set.⁹ The

⁹Unfortunately, we could not include more terms in this seed set due to a high lexical ambiguity of other polar words. Even in our proposed prototypical seed list, one of the terms—“gut” (*good*)—could have another rather rare reading (*manor*) when used as a noun.




Figure 2: Macro-averaged F_1 -scores of the corpus-based approaches with different seed sets.

latter seed set consisted of two regular expressions: one for capturing positive smileys and another one for matching negative emoticons. As can be seen from the figure, these lists, however, could hardly outperform any of our initially used seed sets.

7 Analysis of Entries

Besides investigating the effects of different hyper-parameters and seeds, we also decided to have a closer look at the actual results produced by the tested methods. For this purpose, we extracted ten highest scored entries (not counting the seed terms) from each automatic lexicon and summarized them in Table 4.

Rank	HL	BG	KH	ES	RR** mincut	RR _{lbl-prop}	TKM	VEL	KIR	SEV
1	perfekt <i>perfect</i>	fleißig <i>diligent</i>	anrühlich <i>indecent</i>	namenlos <i>nameless</i>	planieren <i>to plane</i>	prunkvoll <i>splendid</i>	Stockfotos <i>stock photos</i>	Wahlkampfge- schenk <i>election gift</i>	Suchmaschinen <i>search engines</i>	Scherwey <i>Scherwey</i>
2	mustergültig <i>immaculate</i>	böse <i>evil</i>	unecht <i>artificial</i>	ruhelos <i>restless</i>	Erdschicht <i>stratum</i>	sinnlich <i>sensual</i>	BMKS65 <i>BMKS65</i>	Ordensge- schichte <i>order history</i>	#gameinsight <i>#gameinsight</i>	krebsen <i>to crawl</i>
3	vorbildlich <i>commendable</i>	beispielhaft <i>exemplary</i>	irregulär <i>irregular</i>	unbewaffnet <i>unarmed</i>	gefallen <i>please</i>	pompös <i>ostentatious</i>	Ziya <i>Ziya</i>	Indologica <i>Indologica</i>	#androidgames <i>#androidgames</i>	kaschieren <i>to conceal</i>
4	beispielhaft <i>exemplary</i>	edel <i>noble</i>	drittklassig <i>third-class</i>	interesselos <i>indifferent</i>	Zeiteinheit <i>time unit</i>	unappetitlich <i>unsavory</i>	Shoafoundation <i>shoah found.</i>	Indologie <i>Indology</i>	Selamat <i>selamat</i>	Davis <i>Davis</i>
5	exzellent <i>excellent</i>	tüchtig <i>proficient</i>	sinnlich <i>sensual</i>	reizlos <i>unattractive</i>	Derivat <i>derivate</i>	befehlsgemäß <i>as ordered</i>	T1199 <i>T1199</i>	Energieverbrauch <i>energy consumption</i>	Pagi <i>Pagi</i>	#Klassiker <i>#classics</i>
6	exzeptionell <i>exceptional</i>	emsig <i>busy</i>	unprofessionell <i>unprofessional</i>	würdelos <i>undignified</i>	Oberfläche <i>surface</i>	vierschrotig <i>beefy</i>	Emilay55 <i>Emilay55</i>	Schimmelbildung <i>mold formation</i>	#Sparwelt <i>#savingsworld</i>	Nationalismus <i>nationalism</i>
7	außergewöhnlich <i>extraordinary</i>	eifrig <i>eager</i>	abgeschlagen <i>exhausted</i>	absichtslos <i>unintentional</i>	Essbesteck <i>cutlery</i>	regelmäßig <i>regularly</i>	Eneramo <i>Eneramo</i>	Hygiene <i>hygiene</i>	#Seittest <i>#Seittest</i>	Kraftstoff <i>fuel</i>
8	außerordentlich <i>exceptionally</i>	arbeitsam <i>hardworking</i>	gefällig <i>pleasing</i>	ereignislos <i>uneventful</i>	ablösen <i>to displace</i>	wahrheitsgemäß <i>true</i>	GotzeID <i>GotzeID</i>	wasserd <i>waterp</i>	Gameinsight <i>Gameinsight</i>	inaktiv <i>idle</i>
9	viertklassig <i>fourth-class</i>	mustergültig <i>exemplary</i>	mustergültig <i>exemplary</i>	regellos <i>irregular</i>	Musikveranstaltung <i>music event</i>	fettig <i>greasy</i>	BSH65 <i>BSH65</i>	heizkostensparen <i>saving heating costs</i>	#ipadgames <i>#ipadgames</i>	8DD <i>8DD</i>
10	sinnreich <i>ingenious</i>	vorbildlich <i>commendable</i>	unrecht <i>wrong</i>	fehlerfrei <i>accurate</i>	Gebrechen <i>afflictions</i>	lumpig <i>shabby</i>	Saymak. <i>Saymak.</i>	Referenzarchi- tekturen <i>reference architectures</i>	Fitnessstraining <i>fitness training</i>	Mailadresse <i>mail address</i>

Table 4: Top ten polar terms produced by the automatic methods.

** – the min-cut method of Rao and Ravichandran (2009) returns an unsorted set

As can be seen from the table, the approaches of Hu and Liu (2004), Blair-Goldensohn et al. (2008), Kim and Hovy (2004), as well as the label-propagation algorithm of Rao and Ravichandran (2009) produce almost perfect polarity lists. The SENTIWORDNET approach of Esuli and Sebastiani (2006), however, already features some spurious terms (e.g., “absichtslos” *unintentional*) among its top-scored entries. Finally, the min-cut approach of Rao and Ravichandran (2009) returns a set of mainly objective terms, which, however, is rather due to the fact that this method performs a cluster-like partitioning of the lexical graph without ranking the words assigned to a cluster.

An opposite situation is observed for the corpus-based systems: The top-scoring polarity lists returned by these approaches not only include many apparently objective terms but are also difficult to interpret in

general, as they contain a substantial number of slang and advertising terms (e.g., “BMKS65”, “#gameinsight”, “#androidgames” etc.). This again supports the hypothesis that an extreme content noisiness of the input domain might pose considerable difficulties to sentiment lexicon generation methods.

8 Conclusions and Future Work

Based on the above observations and our experiments, we can formulate the main conclusions that we come to in this paper as follows:

- semi-automatic translations of common English polarity lists notably outperform automatic SLG approaches that are applied directly to non-English data;
- despite their allegedly worse ability to accommodate new domains, dictionary-based methods are still superior to corpus-based systems (at least in terms of the proposed intrinsic evaluation), provided that a sufficiently big lexical taxonomy exists for the target language;
- a potential weakness of the dictionary-based algorithms, however, is their susceptibility to different hyper-parameter settings and the size and composition of the initial seed sets;
- nevertheless, the effect of the seed sets might be even stronger for the corpus-based approaches which rely on distant supervision, if the resulting noisy labeled training set becomes highly unbalanced.

In this respect, there appears to be a great need for a corpus-based method which can both benefit from in-domain data and be resistant to non-balanced training sets; and we are, in fact, currently working on such an algorithm. By taking advantage of the recent advances in deep learning and distributional semantics, we aim to show an efficient way of getting suitable vector representations for polar terms and generating high-quality sentiment lexicons from these automatically learned vectors.

Acknowledgments

We thank the anonymous reviewers for their suggestions and comments.

References

- Alina Andreevskaia and Sabine Bergler. 2008. When specialists and generalists work together: Overcoming domain dependence in sentiment tagging. In *ACL 2008, Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics, June 15-20, 2008, Columbus, Ohio, USA*, pages 290–298. The Association for Computer Linguistics.
- Ahmed Hassan Awadallah and Dragomir R. Radev. 2010. Identifying text polarity using random walks. In Jan Hajic, Sandra Carberry, and Stephen Clark, editors, *ACL 2010, Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, July 11-16, 2010, Uppsala, Sweden*, pages 395–403. The Association for Computer Linguistics.
- Sasha Blair-Goldensohn, Tyler Neylon, Kerry Hannan, George A. Reis, Ryan McDonald, and Jeff Reynar. 2008. Building a sentiment summarizer for local service reviews. In *In NLP in the Information Explosion Era*.
- Avrim Blum, John D. Lafferty, Mugizi Robert Rwebangira, and Rajashekar Reddy. 2004. Semi-supervised learning using randomized mincuts. In Carla E. Brodley, editor, *Machine Learning, Proceedings of the Twenty-first International Conference (ICML 2004), Banff, Alberta, Canada, July 4-8, 2004*, volume 69 of *ACM International Conference Proceeding Series*. ACM.
- Nicoletta Calzolari, Khalid Choukri, Bente Maegaard, Joseph Mariani, Jan Odijk, Stelios Piperidis, Mike Rosner, and Daniel Tapias, editors. 2010. *Proceedings of the International Conference on Language Resources and Evaluation, LREC 2010, 17-23 May 2010, Valletta, Malta*. European Language Resources Association.
- Erik Cambria, Björn W. Schuller, Yunqing Xia, and Bebo White. 2016. New avenues in knowledge bases for natural language processing. *Knowl.-Based Syst.*, 108:1–4.

- Simon Clematide and Manfred Klenner. 2010. Evaluation and extension of a polarity lexicon for German. In *Proceedings of the First Workshop on Computational Approaches to Subjectivity and Sentiment Analysis*, pages 7–13.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46.
- Andrea Esuli and Fabrizio Sebastiani. 2006. SentiWordNet: a high-coverage lexical resource for opinion mining. Technical Report ISTI-PP-002/2007, Institute of Information Science and Technologies (ISTI) of the Italian National Research Council (CNR), October.
- Birgit Hamp and Helmut Feldweg. 1997. GermaNet - a lexical-semantic net for German. In *In Proceedings of ACL workshop Automatic Information Extraction and Building of Lexical Semantic Resources for NLP Applications*, pages 9–15.
- Vasileios Hatzivassiloglou and Kathleen McKeown. 1997. Predicting the semantic orientation of adjectives. In Philip R. Cohen and Wolfgang Wahlster, editors, *35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics, Proceedings of the Conference, 7-12 July 1997, Universidad Nacional de Educación a Distancia (UNED), Madrid, Spain.*, pages 174–181. Morgan Kaufmann Publishers / ACL.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In Won Kim, Ron Kohavi, Johannes Gehrke, and William DuMouchel, editors, *KDD*, pages 168–177. ACM.
- Soo-Min Kim and Eduard H. Hovy. 2004. Determining the sentiment of opinions. In *COLING 2004, 20th International Conference on Computational Linguistics, Proceedings of the Conference, 23-27 August 2004, Geneva, Switzerland*.
- Svetlana Kiritchenko, Xiaodan Zhu, and Saif M. Mohammad. 2014. Sentiment Analysis of Short Informal Texts. *J. Artif. Intell. Res. (JAIR)*, 50:723–762.
- Donald E. Knuth. 1998. *The Art of Computer Programming, Volume 3: (2Nd Ed.) Sorting and Searching*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA.
- Bing Liu. 2012. *Sentiment Analysis and Opinion Mining*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.
- George A. Miller. 1995. WordNet: A Lexical Database for English. *Communications of ACM*, 38(11):39–41, November.
- Saif M. Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu. 2013. NRC-Canada: Building the State-of-the-Art in Sentiment Analysis of Tweets. *CoRR*, abs/1308.6242.
- Preslav Nakov, Sara Rosenthal, Zornitsa Kozareva, Veselin Stoyanov, Alan Ritter, and Theresa Wilson. 2013. SemEval-2013 Task 2: Sentiment Analysis in Twitter. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 312–320, Atlanta, Georgia, USA, June. Association for Computational Linguistics.
- Uwe Quasthoff. 2010. *Deutsches Kollokationswörterbuch*. deGruyter, Berlin, New York.
- Delip Rao and Deepak Ravichandran. 2009. Semi-supervised polarity lexicon induction. In Alex Lascarides, Claire Gardent, and Joakim Nivre, editors, *EACL 2009, 12th Conference of the European Chapter of the Association for Computational Linguistics, Proceedings of the Conference, Athens, Greece, March 30 - April 3, 2009*, pages 675–682. The Association for Computer Linguistics.
- Robert Remus, Uwe Quasthoff, and Gerhard Heyer. 2010. SentiWS - A publicly available German-language resource for sentiment analysis. In Calzolari et al. (Calzolari et al., 2010).
- Sara Rosenthal, Alan Ritter, Preslav Nakov, and Veselin Stoyanov. 2014. SemEval-2014 Task 9: Sentiment Analysis in Twitter. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 73–80, Dublin, Ireland, August. Association for Computational Linguistics and Dublin City University.
- Sara Rosenthal, Preslav Nakov, Svetlana Kiritchenko, Saif Mohammad, Alan Ritter, and Veselin Stoyanov. 2015. Semeval-2015 task 10: Sentiment analysis in twitter. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 451–463, Denver, Colorado, June. Association for Computational Linguistics.

- Tatjana Scheffler. 2014. A German Twitter Snapshot. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Hrafn Loftsson, Bente Maegaard, Joseph Mariani, Asunción Moreno, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*, Reykjavik, Iceland, May 26-31, 2014., pages 2284–2289. European Language Resources Association (ELRA).
- Helmut Schmid. 1995. Probabilistic part-of-speech tagging using decision trees. In *Proceedings of the ACL SIGDAT-Workshop*.
- Aliaksei Severyn and Alessandro Moschitti. 2015. On the automatic learning of sentiment lexicons. In Rada Mihalcea, Joyce Yue Chai, and Anoop Sarkar, editors, *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Denver, Colorado, USA, May 31 - June 5, 2015, pages 1397–1402. The Association for Computational Linguistics.
- Uladzimir Sidarenka. 2016. PotTS: The Potsdam Twitter Sentiment Corpus. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Sara Goggi, Marko Grobelnik, Bente Maegaard, Joseph Mariani, H  l  ne Mazo, Asunci  n Moreno, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portoro  , Slovenia, May 23-28, 2016*. European Language Resources Association (ELRA).
- Philip J. Stone, Dexter C. Dunphy, Marshall S. Smith, and Daniel M. Ogilvie. 1966. *The General Inquirer: A Computer Approach to Content Analysis*. MIT Press, Cambridge, MA.
- Maite Taboada, Julian Brooke, Milan Tofiloski, Kimberly D. Voll, and Manfred Stede. 2011. Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37(2):267–307.
- Hiroya Takamura, Takashi Inui, and Manabu Okumura. 2005. Extracting semantic orientations of words using spin model. In Kevin Knight, Hwee Tou Ng, and Kemal Oflazer, editors, *ACL 2005, 43rd Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, 25-30 June 2005, University of Michigan, USA*. The Association for Computer Linguistics.
- Peter D. Turney and Michael L. Littman. 2003. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Trans. Inf. Syst.*, 21(4):315–346.
- Leonid Velikovich, Sasha Blair-Goldensohn, Kerry Hannan, and Ryan T. McDonald. 2010. The viability of web-derived polarity lexicons. In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 2-4, 2010, Los Angeles, California, USA*, pages 777–785. The Association for Computational Linguistics.
- Ulli Waltinger. 2010. GermanPolarityClues: A Lexical Resource for German Sentiment Analysis. In Calzolari et al. (Calzolari et al., 2010).
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT/EMNLP 2005, Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference, 6-8 October 2005, Vancouver, British Columbia, Canada*. The Association for Computational Linguistics.
- Xiaojin Zhu and Zoubin Ghahramani. 2002. Learning from labeled and unlabeled data with label propagation, cmu-cald-02-107. Technical report, Carnegie Mellon University.
- Xiaodan Zhu, Svetlana Kiritchenko, and Saif Mohammad. 2014. Nrc-canada-2014: Recent improvements in the sentiment analysis of tweets. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 443–447, Dublin, Ireland, August. Association for Computational Linguistics and Dublin City University.