

SemEval-2014 Task 9: Sentiment Analysis in Twitter

Sara Rosenthal

Columbia University

sara@cs.columbia.edu

Preslav Nakov

Qatar Computing Research Institute

pnakov@qf.org.qa

Alan Ritter

Carnegie Mellon University

rittera@cs.cmu.edu

Veselin Stoyanov

Johns Hopkins University

ves@cs.jhu.edu

Abstract

We describe the *Sentiment Analysis in Twitter* task, ran as part of SemEval-2014. It is a continuation of the last year's task that ran successfully as part of SemEval-2013. As in 2013, this was the most popular SemEval task; a total of 46 teams contributed 27 submissions for subtask A (21 teams) and 50 submissions for subtask B (44 teams). This year, we introduced three new test sets: (i) regular tweets, (ii) sarcastic tweets, and (iii) LiveJournal sentences. We further tested on (iv) 2013 tweets, and (v) 2013 SMS messages. The highest F1-score on (i) was achieved by *NRC-Canada* at 86.63 for subtask A and by *TeamX* at 70.96 for subtask B.

1 Introduction

In the past decade, new forms of communication have emerged and have become ubiquitous through social media. Microblogs (e.g., Twitter), Weblogs (e.g., LiveJournal) and cell phone messages (SMS) are often used to share opinions and sentiments about the surrounding world, and the availability of social content generated on sites such as Twitter creates new opportunities to automatically study public opinion.

Working with these informal text genres presents new challenges for natural language processing beyond those encountered when working with more traditional text genres such as newswire. The language in social media is very informal, with creative spelling and punctuation, misspellings, slang, new words, URLs, and genre-specific terminology and abbreviations, e.g., RT for re-tweet and #hashtags¹.

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Page numbers and proceedings footer are added by the organisers. Licence details: <http://creativecommons.org/licenses/by/4.0/>

¹Hashtags are a type of tagging for Twitter messages.

Moreover, tweets and SMS messages are short: a sentence or a headline rather than a document.

How to handle such challenges so as to automatically mine and understand people's opinions and sentiments has only recently been the subject of research (Jansen et al., 2009; Barbosa and Feng, 2010; Bifet et al., 2011; Davidov et al., 2010; O'Connor et al., 2010; Pak and Paroubek, 2010; Tumasjan et al., 2010; Kouloumpis et al., 2011).

Several corpora with detailed opinion and sentiment annotation have been made freely available, e.g., the MPQA newswire corpus (Wiebe et al., 2005), the movie reviews corpus (Pang et al., 2002), or the restaurant and laptop reviews corpora that are part of this year's SemEval Task 4 (Pontiki et al., 2014). These corpora have proved very valuable as resources for learning about the language of sentiment in general, but they do not focus on tweets. While some Twitter sentiment datasets were created prior to SemEval-2013, they were either small and proprietary, such as the i-sieve corpus (Kouloumpis et al., 2011) or focused solely on message-level sentiment.

Thus, the primary goal of our SemEval task is to promote research that will lead to better understanding of how sentiment is conveyed in Social Media. Toward that goal, we created the SemEval Tweet corpus as part of our inaugural Sentiment Analysis in Twitter Task, SemEval-2013 Task 2 (Nakov et al., 2013). It contains tweets and SMS messages with sentiment expressions annotated with contextual phrase-level and message-level polarity. This year, we extended the corpus by adding new tweets and LiveJournal sentences.

Another interesting phenomenon that has been studied in Twitter is the use of the #sarcasm hashtag to indicate that a tweet should not be taken literally (González-Ibáñez et al., 2011; Liebrecht et al., 2013). In fact, sarcasm indicates that the message polarity should be flipped. With this in mind, this year, we also evaluate on sarcastic tweets.

In the remainder of this paper, we first describe the task, the dataset creation process and the evaluation methodology. We then summarize the characteristics of the approaches taken by the participating systems, and we discuss their scores.

2 Task Description

As SemEval-2013 Task 2, we included two subtasks: an expression-level subtask and a message-level subtask. Participants could choose to participate in either or both. Below we provide short descriptions of the objectives of these two subtasks.

Subtask A: Contextual Polarity Disambiguation

Given a message containing a marked instance of a word or a phrase, determine whether that instance is positive, negative or neutral in that context. The instance boundaries were provided: this was a classification task, not an entity recognition task.

Subtask B: Message Polarity Classification

Given a message, decide whether it is of positive, negative, or neutral sentiment. For messages conveying both positive and negative sentiment, the stronger one is to be chosen.

Each participating team was allowed to submit results for two different systems per subtask: one constrained, and one unconstrained. A constrained system could only use the provided data for training, but it could also use other resources such as lexicons obtained elsewhere. An unconstrained system could use any additional data as part of the training process; this could be done in a supervised, semi-supervised, or unsupervised fashion.

Note that constrained/unconstrained refers to the data used to train a classifier. For example, if other data (excluding the test data) was used to develop a sentiment lexicon, and the lexicon was used to generate features, the system would still be constrained. However, if other data (excluding the test data) was used to develop a sentiment lexicon, and this lexicon was used to automatically label additional Tweet/SMS messages and then used with the original data to train the classifier, then such a system would be considered unconstrained.

3 Datasets

In this section, we describe the process of collecting and annotating the 2014 testing tweets, including the sarcastic ones, and LiveJournal sentences.

Corpus	Positive	Negative	Objective / Neutral
Twitter2013-train	5,895	3,131	471
Twitter2013-dev	648	430	57
Twitter2013-test	2,734	1,541	160
SMS2013-test	1,071	1,104	159
Twitter2014-test	1,807	578	88
Twitter2014-sarcasm	82	37	5
LiveJournal2014-test	660	511	144

Table 1: Dataset statistics for Subtask A.

3.1 Datasets Used

For training and development, we released the Twitter train/dev/test datasets from SemEval-2013 task 2, as well as the SMS test set, which uses messages from the NUS SMS corpus (Chen and Kan, 2013), which we annotated for sentiment in 2013.

We further added a new 2014 Twitter test set, as well as a small set of tweets that contained the #sarcasm hashtag to determine how sarcasm affects the tweet polarity. Finally, we included sentences from LiveJournal in order to determine how systems trained on Twitter perform on other sources. The statistics for each dataset and for each subtask are shown in Tables 1 and 2.

Corpus	Positive	Negative	Objective / Neutral
Twitter2013-train	3,662	1,466	4,600
Twitter2013-dev	575	340	739
Twitter2013-test	1,572	601	1,640
SMS2013-test	492	394	1,207
Twitter2014-test	982	202	669
Twitter2014-sarcasm	33	40	13
LiveJournal2014-test	427	304	411

Table 2: Dataset statistics for Subtask B.

3.2 Annotation

We annotated the new tweets as in 2013: by identifying tweets from popular topics that contain sentiment-bearing words by using SentiWordNet (Baccianella et al., 2010) as a filter. We altered the annotation task for the sarcastic tweets, displaying them to the Mechanical Turk annotators without the #sarcasm hashtag; the Turkers had to determine whether the tweet is sarcastic on their own. Moreover, we asked Turkers to indicate the degree of sarcasm as (a) definitely sarcastic, (b) probably sarcastic, and (c) not sarcastic.

As in 2013, we combined the annotations using intersection, where a word had to appear in 2/3 of the annotations to be accepted. An annotated example from each source is shown in Table 3.

Source	Example	Polarity
Twitter	Why would you [still]- wear shorts when it's this cold?! I [love]+ how Britain see's a bit of sun and they're [like 'OOOH]+ LET'S STRIP!	positive
SMS	[Sorry]- I think tonight [cannot]- and I [not feeling well]- after my rest.	negative
LiveJournal	[Cool]+ posts , dude ; very [colorful]+ , and [artsy]+ .	positive
Twitter Sarcasm	[Thanks]+ manager for putting me on the schedule for Sunday	negative

Table 3: Example of polarity for each source of messages. The target phrases are marked in [...], and are followed by their polarity; the sentence-level polarity is shown in the last column.

3.3 Tweets Delivery

We did not deliver the annotated tweets to the participants directly; instead, we released annotation indexes, a list of corresponding Twitter IDs, and a download script that extracts the corresponding tweets via the Twitter API.² We provided the tweets in this manner in order to ensure that Twitter's terms of service are not violated. Unfortunately, due to this restriction, the task participants had access to different number of training tweets depending on when they did the downloading. This varied between a minimum of 5,215 tweets and the full set of 10,882 tweets. On average the teams were able to collect close to 9,000 tweets; for teams that did not participate in 2013, this was about 8,500. The difference in training data size did not seem to have had a major impact. In fact, the top two teams in subtask B (coooolll and TeamX) trained on less than 8,500 tweets.

4 Scoring

The participating systems were required to perform a three-way classification for both subtasks. A particular marked phrase (for subtask A) or an entire message (for subtask B) was to be classified as *positive*, *negative* or *objective/neutral*. We scored the systems by computing a score for predicting positive/negative phrases/messages. For instance, to compute positive precision, p_{pos} , we find the number of phrases/messages that a system correctly predicted to be positive, and we divide that number by the total number it predicted to be positive. To compute positive recall, r_{pos} , we find the number of phrases/messages correctly predicted to be positive and we divide that number by the total number of positives in the gold standard. We then calculate F1-score for the positive class as follows $F_{pos} = \frac{2(p_{pos} + r_{pos})}{p_{pos} * r_{pos}}$. We carry out a similar computation for F_{neg} , for the negative phrases/messages. The overall score is then $F = (F_{pos} + F_{neg})/2$.

²<https://dev.twitter.com>

We used the two test sets from 2013 and the three from 2014, which we combined into one test set and we shuffled to make it hard to guess which set a sentence came from. This guaranteed that participants would submit predictions for all five test sets. It also allowed us to test how well systems trained on standard tweets generalize to sarcastic tweets and to LiveJournal sentences, without the participants putting extra efforts into this. The participants were also not informed about the source the extra test sets come from.

We provided the participants with a scorer that outputs the overall score F and a confusion matrix for each of the five test sets.

5 Participants and Results

The results are shown in Tables 4 and 5, and the team affiliations are shown in Table 6. Tables 4 and 5 contain results on the two progress test sets (tweets and SMS messages), which are the official test sets from the 2013 edition of the task, and on the three new official 2014 testsets (tweets, tweets with sarcasm, and LiveJournal). The tables further show macro- and micro-averaged results over the 2014 datasets. There is an index for each result showing the relative rank of that result within the respective column. The participating systems are ranked by their score on the Twitter-2014 test-set, which is the official ranking for the task; all remaining rankings are secondary.

As we mentioned above, the participants were not told that the 2013 test sets would be included in the big 2014 test set, so that they do not over-tune their systems on them. However, the 2013 test sets were made available for development, but it was explicitly forbidden to use them for training. Still, some participants did not notice this restriction, which resulted in their unusually high scores on Twitter2013-test; we did our best to identify all such cases, and we asked the authors to submit corrected runs. The tables mark such resubmissions accordingly.

Most of the submissions were constrained, with just a few unconstrained: 7 out of 27 for subtask A, and 8 out of 50 for subtask B. In any case, the best systems were constrained. Some teams participated with both a constrained and an unconstrained system, but the unconstrained system was not always better than the constrained one: sometimes it was worse, sometimes it performed the same. Thus, we decided to produce a single ranking, including both constrained and unconstrained systems, where we mark the latter accordingly.

5.1 Subtask A

Table 4 shows the results for subtask A, which attracted 27 submissions from 21 teams. There were seven unconstrained submissions: five teams submitted both a constrained and an unconstrained run, and two teams submitted an unconstrained run only. The best systems were constrained. All participating systems outperformed the majority class baseline by a sizable margin.

5.2 Subtask B

The results for subtask B are shown in Table 5. The subtask attracted 50 submissions from 44 teams. There were eight unconstrained submissions: six teams submitted both a constrained and an unconstrained run, and two teams submitted an unconstrained run only. As for subtask A, the best systems were constrained. Again, all participating systems outperformed the majority class baseline; however, some systems were very close to it.

6 Discussion

Overall, we observed similar trends as in SemEval-2013 Task 2. Almost all systems used supervised learning. Most systems were constrained, including the best ones in all categories. As in 2013, we observed several cases of a team submitting a constrained and an unconstrained run and the constrained run performing better.

It is unclear why unconstrained systems did not outperform constrained ones. It could be because participants did not use enough external data or because the data they used was too different from Twitter or from our annotation method. Or it could be due to our definition of *unconstrained*, which labels as unconstrained systems that use additional tweets directly, but considers unconstrained those that use additional tweets to build sentiment lexicons and then use these lexicons.

As in 2013, the most popular classifiers were SVM, MaxEnt, and Naive Bayes. Moreover, two submissions used deep learning, *coooolll* (Harbin Institute of Technology) and *ThinkPositive* (IBM Research, Brazil), which were ranked second and tenth on subtask B, respectively.

The features used were quite varied, including word-based (e.g., word and character n -grams, word shapes, and lemmata), syntactic, and Twitter-specific such as emoticons and abbreviations. The participants still relied heavily on lexicons of opinion words, the most popular ones being the same as in 2013: MPQA, SentiWordNet and Bing Liu’s opinion lexicon. Popular this year was also the NRC lexicon (Mohammad et al., 2013), created by the best-performing team in 2013, which is top-performing this year as well.

Preprocessing of tweets was still a popular technique. In addition to standard NLP steps such as tokenization, stemming, lemmatization, stopword removal and POS tagging, most teams applied some kind of Twitter-specific processing such as substitution/removal of URLs, substitution of emoticons, word normalization, abbreviation lookup, and punctuation removal. Finally, several of the teams used Twitter-tuned NLP tools such as part of speech and named entity taggers (Gimpel et al., 2011; Ritter et al., 2011).

The similarity of preprocessing techniques, NLP tools, classifiers and features used in 2013 and this year is probably partially due to many teams participating in both years. As Table 6 shows, 18 out of the 46 teams are returning teams.

Comparing the results on the progress Twitter test in 2013 and 2014, we can see that *NRC-Canada*, the 2013 winner for subtask A, have now improved their F1 score from 88.93 to 90.14, which is the 2014 best score. The best score on the Progress SMS in 2014 of 89.31 belongs to *ECNU*; this is a big jump compared to their 2013 score of 76.69, but it is less compared to the 2013 best of 88.37 achieved by *GU-MLT-LT*. For subtask B, on the Twitter progress testset, the 2013 winner *NRC-Canada* improves their 2013 result from 69.02 to 70.75, which is the second best in 2014; the winner in 2014, *TeamX*, achieves 72.12. On the SMS progress test, the 2013 winner *NRC-Canada* improves its F1 score from 68.46 to 70.28. Overall, we see consistent improvements on the progress testset for both subtasks: 0-1 and 2-3 points absolute for subtasks A and B, respectively.

#	System	Uncon- strain.?	2013: Progress		2014: Official			2014: Average	
			Tweet	SMS	Tweet	Tweet sarcasm	Live- Journal	Macro	Micro
1	NRC-Canada		90.14 ₁	88.03 ₄	86.63 ₁	77.13 ₅	85.49 ₂	83.08 ₂	85.61 ₁
2	SentiKLUE		90.11 ₂	85.16 ₈	84.83 ₂	79.32 ₃	85.61 ₁	83.25 ₁	85.15 ₂
3	CMUQ-Hybrid*		88.94 ₄	87.98 ₅	84.40 ₃	76.99 ₆	84.21 ₃	81.87 ₃	84.05 ₃
4	CMU-Qatar*		89.85 ₃	88.08 ₃	83.45 ₄	78.07 ₄	83.89 ₅	81.80 ₄	83.56 ₄
5	ECNU	✓	87.29 ₆	89.26 ₂	82.93 ₅	73.71 ₈	81.69 ₇	79.44 ₇	81.85 ₆
6	ECNU		87.28 ₇	89.31 ₁	82.67 ₆	73.71 ₉	81.67 ₈	79.35 ₈	81.75 ₇
7	Think.Positive	✓	88.06 ₅	87.65 ₆	82.05 ₇	76.74 ₇	80.90 ₁₂	79.90 ₆	81.15 ₉
8	Kea*		84.83 ₁₀	84.14 ₁₀	81.22 ₈	65.94 ₁₇	81.16 ₁₁	76.11 ₁₃	80.70 ₁₀
9	Lt.3		86.28 ₈	85.26 ₇	81.02 ₉	70.76 ₁₃	80.44 ₁₃	77.41 ₁₁	80.33 ₁₃
10	senti.ue		84.05 ₁₁	78.72 ₁₆	80.54 ₁₀	82.75 ₁	81.90 ₆	81.73 ₅	81.47 ₈
11	LyS		85.69 ₉	81.44 ₁₂	79.92 ₁₁	71.67 ₁₀	83.95 ₄	78.51 ₁₀	82.21 ₅
12	UKPDIPF		80.45 ₁₅	79.05 ₁₄	79.67 ₁₂	65.63 ₁₈	81.42 ₉	75.57 ₁₄	80.33 ₁₁
13	UKPDIPF	✓	80.45 ₁₆	79.05 ₁₅	79.67 ₁₃	65.63 ₁₉	81.42 ₁₀	75.57 ₁₅	80.33 ₁₂
14	TJP		81.13 ₁₄	84.41 ₉	79.30 ₁₄	71.20 ₁₂	78.27 ₁₅	76.26 ₁₂	78.39 ₁₅
15	SAP-RI		80.32 ₁₇	80.26 ₁₃	77.26 ₁₅	70.64 ₁₄	77.68 ₁₈	75.19 ₁₇	77.32 ₁₆
16	senti.ue*	✓	83.80 ₁₂	82.93 ₁₁	77.07 ₁₆	80.02 ₂	79.70 ₁₄	78.93 ₉	78.83 ₁₄
17	SAIL		78.47 ₁₈	74.46 ₂₀	76.89 ₁₇	65.56 ₂₀	70.62 ₂₂	71.02 ₂₁	72.57 ₂₁
18	columbia.nlp [◊]		81.50 ₁₃	74.55 ₁₉	76.54 ₁₈	61.76 ₂₂	78.19 ₁₆	72.16 ₁₉	77.11 ₁₈
19	IIT-Patna		76.54 ₂₀	75.99 ₁₈	76.43 ₁₉	71.43 ₁₁	77.99 ₁₇	75.28 ₁₆	77.26 ₁₇
20	Citius	✓	76.59 ₁₉	69.31 ₂₁	75.21 ₂₀	68.40 ₁₅	75.82 ₂₀	73.14 ₁₈	75.38 ₁₉
21	Citius		74.71 ₂₁	61.44 ₂₅	73.03 ₂₁	65.18 ₂₁	71.64 ₂₁	69.95 ₂₂	71.90 ₂₂
22	IITPatna		70.91 ₂₃	77.04 ₁₇	72.25 ₂₂	66.32 ₁₆	76.03 ₁₉	71.53 ₂₀	74.45 ₂₀
23	SU-sentilab		74.34 ₂₂	62.58 ₂₄	68.26 ₂₃	53.31 ₂₅	69.53 ₂₃	63.70 ₂₄	68.59 ₂₃
24	Univ. Warwick*		62.25 ₂₆	60.12 ₂₆	67.28 ₂₄	58.08 ₂₄	64.89 ₂₅	63.42 ₂₅	65.48 ₂₅
25	Univ. Warwick*	✓	64.91 ₂₅	63.01 ₂₃	67.17 ₂₅	60.59 ₂₃	67.46 ₂₄	65.07 ₂₃	67.14 ₂₄
26	DAEDALUS		67.42 ₂₄	63.92 ₂₂	60.98 ₂₆	45.27 ₂₇	61.01 ₂₆	55.75 ₂₆	60.50 ₂₆
27	DAEDALUS	✓	61.95 ₂₇	55.97 ₂₇	58.11 ₂₇	49.19 ₂₆	58.65 ₂₇	55.32 ₂₇	58.17 ₂₇
	Majority baseline		38.1	31.5	42.2	39.8	33.4		

Table 4: **Results for subtask A.** The * indicates system resubmissions (because they initially trained on Twitter2013-test), and the [◊] indicates a system that includes a task co-organizer as a team member. The systems are sorted by their score on the Twitter2014 test dataset; the rankings on the individual datasets are indicated with a subscript. The last two columns show macro- and micro-averaged results across the three 2014 test datasets.

Finally, note that for both subtasks, the best systems on the Twitter-2014 dataset are those that performed best on the 2013 progress Twitter dataset: *NRC-Canada* for subtask A, and *TeamX* (Fuji Xerox Co., Ltd.) for subtask B.

It is interesting to note that the best results for Twitter2014-test are lower than those for Twitter2013-test for both subtask A (86.63 vs. 90.14) and subtask B (70.96 vs 72.12). This is so despite the baselines for Twitter2014-test being higher than those for Twitter2013-test: 42.2 vs. 38.1 for subtask A, and 34.6 vs. 29.2 for subtask B. Most likely, having access to Twitter2013-test at development time, teams have overfitted on it. It could be also the case that some of the sentiment dictionaries that were built in 2013 have become somewhat outdated by 2014.

Finally, note that while some teams such as *NRC-Canada* performed well across all test sets, other such as *TeamX*, which used a weighting scheme tuned specifically for class imbalances in tweets, were only strong on Twitter datasets.

7 Conclusion

We have described the data, the experimental setup and the results for SemEval-2014 Task 9. As in 2013, our task was the most popular one at SemEval-2014, attracting 46 participating teams: 21 in subtask A (27 submissions) and 44 in subtask B (50 submissions).

We introduced three new test sets for 2014: an in-domain Twitter dataset, an out-of-domain LiveJournal test set, and a dataset of tweets containing sarcastic content. While the performance on the LiveJournal test set was mostly comparable to the in-domain Twitter test set, for most teams there was a sharp drop in performance for sarcastic tweets, highlighting better handling of sarcastic language as one important direction for future work in Twitter sentiment analysis.

We plan to run the task again in 2015 with the inclusion of a new sub-evaluation on detecting sarcasm with the goal of stimulating research in this area; we further plan to add one more test domain.

#	System	Uncon- strain.?	2013: Progress		2014: Official			2014: Average	
			Tweet	SMS	Tweet	Tweet sarcasm	Live- Journal	Macro	Micro
1	TeamX		72.12 ₁	57.36 ₂₆	70.96 ₁	56.50 ₃	69.44 ₁₅	65.63 ₃	69.99 ₅
2	coooolll		70.40 ₃	67.68 ₂	70.14 ₂	46.66 ₂₄	72.90 ₅	63.23 ₁₂	70.51 ₂
3	RTRGO		69.10 ₅	67.51 ₃	69.95 ₃	47.09 ₂₃	72.20 ₆	63.08 ₁₃	70.15 ₃
4	NRC-Canada		70.75 ₂	70.28 ₁	69.85 ₄	58.16 ₁	74.84 ₁	67.62 ₁	71.37 ₁
5	TUGAS		65.64 ₁₃	62.77 ₁₁	69.00 ₅	52.87 ₁₂	69.79 ₁₃	63.89 ₆	68.84 ₈
6	CISUC_KIS*		67.56 ₈	65.90 ₆	67.95 ₆	55.49 ₅	74.46 ₂	65.97 ₂	70.02 ₄
7	SAIL		66.80 ₁₁	56.98 ₂₈	67.77 ₇	57.26 ₂	69.34 ₁₇	64.79 ₄	68.06 ₁₀
8	SWISS-CHOCOLATE		64.81 ₁₈	66.43 ₅	67.54 ₈	49.46 ₁₆	73.25 ₄	63.42 ₁₀	69.15 ₆
9	Synalp-Empathic		63.65 ₂₃	62.54 ₁₂	67.43 ₉	51.06 ₁₅	71.75 ₉	63.41 ₁₁	68.57 ₉
10	Think_Positive	✓	68.15 ₇	63.20 ₉	67.04 ₁₀	47.85 ₂₁	66.96 ₂₄	60.62 ₁₈	66.47 ₁₅
11	SentiKLUE		69.06 ₆	67.40 ₄	67.02 ₁₁	43.36 ₃₀	73.99 ₃	61.46 ₁₄	68.94 ₇
12	JOINT_FORCES	✓	66.61 ₁₂	62.20 ₁₃	66.79 ₁₂	45.40 ₂₆	70.02 ₁₂	60.74 ₁₇	67.39 ₁₂
13	AMLERIC		70.09 ₄	60.29 ₂₀	66.55 ₁₃	48.19 ₂₀	65.32 ₂₆	60.02 ₂₁	65.58 ₂₀
14	AUEB		63.92 ₂₁	64.32 ₈	66.38 ₁₄	56.16 ₄	70.75 ₁₁	64.43 ₅	67.71 ₁₁
15	CMU-Qatar*		65.11 ₁₇	62.95 ₁₀	65.53 ₁₅	40.52 ₃₈	65.63 ₂₅	57.23 ₂₇	64.87 ₂₄
16	Lt.3		65.56 ₁₄	64.78 ₇	65.47 ₁₆	47.76 ₂₂	68.56 ₂₀	60.60 ₁₉	66.12 ₁₇
17	columbia_nlp [◊]		64.60 ₁₉	59.84 ₂₁	65.42 ₁₇	40.02 ₄₀	68.79 ₁₉	58.08 ₂₅	65.96 ₁₉
18	LyS		66.92 ₁₀	60.45 ₁₉	64.92 ₁₈	42.40 ₃₃	69.79 ₁₄	59.04 ₂₂	66.10 ₁₈
19	NILC_USP		65.39 ₁₅	61.35 ₁₆	63.94 ₁₉	42.06 ₃₄	69.02 ₁₈	58.34 ₂₄	65.21 ₂₁
20	senti.ue		67.34 ₉	59.34 ₂₃	63.81 ₂₀	55.31 ₆	71.39 ₁₀	63.50 ₇	66.38 ₁₆
21	UKPDIPF		60.65 ₂₉	60.56 ₁₇	63.77 ₂₁	54.59 ₇	71.92 ₇	63.43 ₈	66.53 ₁₃
22	UKPDIPF	✓	60.65 ₃₀	60.56 ₁₈	63.77 ₂₂	54.59 ₈	71.92 ₈	63.43 ₉	66.53 ₁₄
23	SU-FMI* [◊]		60.96 ₂₈	61.67 ₁₅	63.62 ₂₃	48.34 ₁₉	68.24 ₂₁	60.07 ₂₀	64.91 ₂₃
24	ECNU		62.31 ₂₇	59.75 ₂₂	63.17 ₂₄	51.43 ₁₄	69.44 ₁₆	61.35 ₁₅	65.17 ₂₂
25	ECNU	✓	63.72 ₂₂	56.73 ₂₉	63.04 ₂₅	49.33 ₁₇	64.08 ₃₁	58.82 ₂₃	63.04 ₂₇
26	Rapanakis		58.52 ₃₂	54.02 ₃₅	63.01 ₂₆	44.69 ₂₇	59.71 ₃₇	55.80 ₃₁	61.28 ₃₂
27	Citius	✓	63.25 ₂₄	58.28 ₂₄	62.94 ₂₇	46.13 ₂₅	64.54 ₂₉	57.87 ₂₆	63.06 ₂₆
28	CMUQ-Hybrid*		63.22 ₂₅	61.75 ₁₄	62.71 ₂₈	40.95 ₃₇	65.14 ₂₇	56.27 ₃₀	63.00 ₂₈
29	Citius		62.53 ₂₆	57.69 ₂₅	61.92 ₂₉	41.00 ₃₆	62.40 ₃₃	55.11 ₃₃	61.51 ₃₁
30	KUNLPLab		58.12 ₃₃	55.89 ₃₁	61.72 ₃₀	44.60 ₂₈	63.77 ₃₂	56.70 ₂₉	62.00 ₂₉
31	senti.ue*	✓	65.21 ₁₆	56.16 ₃₀	61.47 ₃₁	54.09 ₉	68.08 ₂₂	61.21 ₁₆	63.71 ₂₅
32	UPV-ELIRF		63.97 ₂₀	55.36 ₃₃	59.33 ₃₂	37.46 ₄₂	64.11 ₃₀	53.63 ₃₇	60.49 ₃₃
33	USP_Biocom		58.05 ₃₄	53.57 ₃₆	59.21 ₃₃	43.56 ₂₉	67.80 ₂₃	56.86 ₂₈	61.96 ₃₀
34	DAEDALUS	✓	58.94 ₃₁	54.96 ₃₄	57.64 ₃₄	35.26 ₄₄	60.99 ₃₅	51.30 ₃₉	58.26 ₃₅
35	IIT-Patna		52.58 ₄₀	51.96 ₃₇	57.25 ₃₅	41.33 ₃₅	60.39 ₃₆	52.99 ₃₈	57.97 ₃₆
36	DejaVu		57.43 ₃₆	55.57 ₃₂	57.02 ₃₆	42.46 ₃₂	64.69 ₂₈	54.72 ₃₄	59.46 ₃₄
37	GPLSI		57.49 ₃₅	46.63 ₄₂	56.06 ₃₇	53.90 ₁₀	57.32 ₄₁	55.76 ₃₂	56.47 ₃₇
38	BUAP		56.85 ₃₇	44.27 ₄₄	55.76 ₃₈	51.52 ₁₃	53.94 ₄₄	53.74 ₃₆	54.97 ₃₉
39	SAP-RI		50.18 ₄₄	49.00 ₄₁	55.47 ₃₉	48.64 ₁₈	57.86 ₄₀	53.99 ₃₅	56.17 ₃₈
40	UMCC_DLSI_Sem		51.96 ₄₁	50.01 ₃₈	55.40 ₄₀	42.76 ₃₁	53.12 ₄₅	50.43 ₄₀	54.20 ₄₂
41	IBM_LEG		54.51 ₃₈	46.62 ₄₃	52.26 ₄₁	34.14 ₄₆	59.24 ₃₈	48.55 ₄₃	54.34 ₄₁
42	Alberta		53.85 ₃₉	49.05 ₄₀	52.06 ₄₂	40.40 ₃₉	52.38 ₄₆	48.28 ₄₄	51.85 ₄₄
43	lsis_lif		46.38 ₄₆	38.56 ₄₇	52.02 ₄₃	34.64 ₄₅	61.09 ₃₄	49.25 ₄₁	54.90 ₄₀
44	SU-sentilab		50.17 ₄₅	49.60 ₃₉	49.52 ₄₄	31.49 ₄₇	55.11 ₄₂	45.37 ₄₇	51.09 ₄₅
45	SINAI		50.59 ₄₂	57.34 ₂₇	49.50 ₄₅	31.15 ₄₉	58.33 ₃₉	46.33 ₄₆	52.26 ₄₃
46	IITPatna		50.32 ₄₃	40.56 ₄₆	48.22 ₄₆	36.73 ₄₃	54.68 ₄₃	46.54 ₄₅	50.29 ₄₆
47	Univ. Warwick		39.17 ₄₈	29.50 ₄₉	45.56 ₄₇	39.77 ₄₁	39.60 ₄₉	41.64 ₄₈	43.19 ₄₈
48	UMCC_DLSI_Graph		43.24 ₄₇	36.66 ₄₈	45.49 ₄₈	53.15 ₁₁	47.81 ₄₇	48.82 ₄₂	46.56 ₄₇
49	Univ. Warwick	✓	34.23 ₅₀	24.63 ₅₀	45.11 ₄₉	31.40 ₄₈	29.34 ₅₀	35.28 ₄₉	38.88 ₄₉
50	DAEDALUS		36.57 ₄₉	40.86 ₄₅	33.03 ₅₀	28.96 ₅₀	40.83 ₄₈	34.27 ₅₀	35.81 ₅₀
	Majority baseline		29.2	19.0	34.6	27.7	27.2		

Table 5: **Results for subtask B.** The * indicates system resubmissions (because they initially trained on Twitter2013-test), and the [◊] indicates a system that includes a task co-organizer as a team member. The systems are sorted by their score on the Twitter2014 test dataset; the rankings on the individual datasets are indicated with a subscript. The last two columns show macro- and micro-averaged results across the three 2014 test datasets.

In the 2015 edition of the task, we might also remove the constrained/unconstrained distinction.

Finally, as there are multiple opinions about a topic in Twitter, we would like to focus on detecting the sentiment trend towards a topic.

Acknowledgements

We would like to thank Kathleen McKeown and Smaranda Muresan for funding the 2014 Twitter test sets. We also thank the anonymous reviewers.

Subtasks	Team	Affiliation	2013?
B	Alberta	University of Alberta	
B	AMI.ERIC	AMI Software R&D and Université de Lyon (ERIC LYON 2)	yes
B	AUEB	Athens University of Economics and Business	yes
B	BUAP	Benemérita Universidad Autónoma de Puebla	
B	CISUC_KIS	University of Coimbra	
A, B	Citius	University of Santiago de Compostela	
A, B	CMU-Qatar	Carnegie Mellon University, Qatar	
A, B	CMUQ-Hybrid	Carnegie Mellon University, Qatar (different from the above)	
A, B	columbia_nlp	Columbia University	yes
B	coooll	Harbin Institute of Technology	
A, B	DAEDALUS	Daedalus	
B	DejaVu	Indian Institute of Technology, Kanpur	
A, B	ECNU	East China Normal University	yes
B	GPLSI	University of Alicante	
B	IBM.EG	IBM Egypt	
A, B	IITPatna	Indian Institute of Technology, Patna	
A, B	IIT-Patna	Indian Institute of Technology, Patna (different from the above)	
B	JOINT.FORCES	Zurich University of Applied Sciences	
A	Kea	York University, Toronto	yes
B	KUNLPLab	Koç University	
B	Isis_lif	Aix-Marseille University	yes
A, B	Lt_3	Ghent University	
A, B	LyS	Universidade da Coruña	
B	NILC_USP	University of São Paulo	yes
A, B	NRC-Canada	National Research Council Canada	yes
B	Rapanakis	Stamatis Rapanakis	
B	RTRGO	Retresco GmbH and University of Gothenburg	yes
A, B	SAIL	Signal Analysis and Interpretation Laboratory	yes
A, B	SAP-RI	SAP Research and Innovation	
A, B	senti.ue	Universidade de Évora	yes
A, B	SentiKLUE	Friedrich-Alexander-Universität Erlangen-Nürnberg	yes
B	SINAI	University of Jaén	yes
B	SU-FMI	Sofia University	
A, B	SU-sentilab	Sabancı University	yes
B	SWISS-CHOCOLATE	ETH Zurich	
B	Synalp-Empathic	University of Lorraine	
B	TeamX	Fuji Xerox Co., Ltd.	
A, B	Think_Positive	IBM Research, Brazil	
A	TJP	University of Northumbria at Newcastle Upon Tyne	yes
B	TUGAS	Instituto de Engenharia de Sistemas e Computadores, Investigação e Desenvolvimento em Lisboa	yes
A, B	UKPDIPF	Ubiquitous Knowledge Processing Lab	
B	UMCC_DLSI_Graph	Universidad de Matanzas and Univarsidad de Alicante	yes
B	UMCC_DLSI_Sem	Universidad de Matanzas and Univarsidad de Alicante (different from above)	yes
A, B	Univ. Warwick	University of Warwick	
B	UPV-ELiRF	Universitat Politècnica de València	
B	USP.Biocom	University of São Paulo and Federal University of São Carlos	

Table 6: Participating teams, their affiliations, subtasks they have taken part in, and an indication about whether the team participated in SemEval-2013 Task 2.

References

- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation*, LREC '10, Valletta, Malta.
- Luciano Barbosa and Junlan Feng. 2010. Robust sentiment detection on Twitter from biased and noisy data. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*, COLING '10, pages 36–44, Beijing, China.
- Albert Bifet, Geoffrey Holmes, Bernhard Pfahringer, and Ricard Gavaldà. 2011. Detecting sentiment change in Twitter streaming data. *Journal of Machine Learning Research, Proceedings Track*, 17:5–11.
- Tao Chen and Min-Yen Kan. 2013. Creating a live, public short message service corpus: the NUS SMS corpus. *Language Resources and Evaluation*, 47(2):299–335.
- Dmitry Davidov, Oren Tsur, and Ari Rappoport. 2010. Semi-supervised recognition of sarcasm in Twitter and Amazon. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning*, CoNLL '10, pages 107–116, Uppsala, Sweden.
- Kevin Gimpel, Nathan Schneider, Brendan O'Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A. Smith. 2011. Part-of-speech tagging for Twitter: Annotation, features, and experiments. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, ACL-HLT '11, pages 42–47, Portland, Oregon, USA.
- Roberto González-Ibáñez, Smaranda Muresan, and Nina Wacholder. 2011. Identifying sarcasm in Twitter: a closer look. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Short Papers*, ACL-HLT '11, pages 581–586, Portland, Oregon, USA.
- Bernard Jansen, Mimi Zhang, Kate Sobel, and Abdur Chowdury. 2009. Twitter power: Tweets as electronic word of mouth. *J. Am. Soc. Inf. Sci. Technol.*, 60(11):2169–2188.
- Efthymios Kouloumpis, Theresa Wilson, and Johanna Moore. 2011. Twitter sentiment analysis: The good the bad and the OMG! In *Proceedings of the Fifth International Conference on Weblogs and Social Media*, ICWSM '11, Barcelona, Catalonia, Spain.
- Christine Liebrecht, Florian Kunneman, and Antal Van den Bosch. 2013. The perfect solution for detecting sarcasm in tweets #not. In *Proceedings of the 4th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 29–37, Atlanta, Georgia, USA.
- Saif Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu. 2013. NRC-Canada: Building the state-of-the-art in sentiment analysis of tweets. In *Proceedings of the Seventh international workshop on Semantic Evaluation Exercises*, SemEval-2013, pages 321–327, Atlanta, Georgia, USA.
- 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 '13, pages 312–320, Atlanta, Georgia, USA.
- Brendan O'Connor, Ramnath Balasubramanian, Bryan Routledge, and Noah Smith. 2010. From tweets to polls: Linking text sentiment to public opinion time series. In *Proceedings of the Fourth International Conference on Weblogs and Social Media*, ICWSM '10, Washington, DC, USA.
- Alexander Pak and Patrick Paroubek. 2010. Twitter based system: Using Twitter for disambiguating sentiment ambiguous adjectives. In *Proceedings of the 5th International Workshop on Semantic Evaluation*, SemEval '10, pages 436–439, Uppsala, Sweden.
- Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up?: Sentiment classification using machine learning techniques. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing - Volume 10*, EMNLP '02, pages 79–86.
- Maria Pontiki, Harris Papageorgiou, Dimitrios Galanis, Ion Androutsopoulos, John Pavlopoulos, and Suresh Manandhar. 2014. SemEval-2014 task 4: Aspect based sentiment analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation*, SemEval '14, Dublin, Ireland.
- Alan Ritter, Sam Clark, Mausam, and Oren Etzioni. 2011. Named entity recognition in tweets: An experimental study. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP '11, pages 1524–1534, Edinburgh, Scotland, UK.
- Andranik Tumasjan, Timm Sprenger, Philipp Sandner, and Isabell Welpe. 2010. Predicting elections with Twitter: What 140 characters reveal about political sentiment. In *Proceedings of the Fourth International Conference on Weblogs and Social Media*, ICWSM '10, Washington, DC, USA.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39(2-3):165–210.