

CPH: Sentiment analysis of Figurative Language on Twitter #easypeasy #not

Sarah McGillion Héctor Martínez Alonso Barbara Plank

University of Copenhagen, Njalsgade 140, 2300 Copenhagen S, Denmark
zhg159@alumni.ku.dk, alonso@hum.ku.dk, bplank@cst.dk

Abstract

This paper describes the details of our system submitted to the SemEval 2015 shared task on sentiment analysis of figurative language on Twitter. We tackle the problem as regression task and combine several base systems using stacked generalization (Wolpert, 1992). An initial analysis revealed that the data is heavily biased, and a general sentiment analysis system (GSA) performs poorly on it. However, GSA proved helpful on the test data, which contains an estimated 25% non-figurative tweets. Our best system, a stacking system with backoff to GSA, ranked 4th on the final test data (Cosine 0.661, MSE 3.404).¹

1 Introduction

Sentiment analysis (SA) is the task of determining the sentiment of a given piece of text. The amplitude of user-generated content produced every day raises the importance of accurate automatic sentiment analysis, for applications ranging from, e.g., reputation analysis (Amigó et al., 2013) to election results prediction (Tjong Kim Sang and Bos, 2012). However, figurative language is pervasive in user-generated content, and figures of speech like irony, sarcasm and metaphors impose relevant challenges for a sentiment analysis system usually trained on literal meanings. For instance, consider the following example:² @CIA We hear you're looking for sentiment analysis to detect sarcasm in Tweets. That'll be easy! #SLA2014 #irony. Irony or sarcasm

¹After submission time we discovered a bug in ST2, which means that the results on the official website are of the GSA and not of the stacking system with backoff.

²From the training data, label: -1.24; GSA prediction: +5.

does not result always in the exact opposite sentiment and therefore it is not as simple as just inverting the scores from a general SA system. Only few studies have attempted SA on figurative language so far (Reyes and Rosso, 2012; Reyes et al., 2013).

The prediction of a fine-grained sentiment score (between -5 and 5) for a tweet poses a series of challenges. First of all, accurate language technology on tweets is hard due to *sample bias*, i.e., collections of tweets are inherently biased towards the particular time (or way, cf. §2) they were collected (Eisenstein, 2013; Hovy et al., 2014). Secondly, the notion of figurativeness (or its complementary notion of literality) does not have a strong definition, let alone do irony, sarcasm, or satire. As pointed out by Reyes and Rosso (2012), “there is not a clear distinction about the boundaries among these terms”. Yet alone attaching a fine-grained score is far from straightforward. In fact, the gold standard consists of the average score assigned by humans through crowdsourcing reflecting an uncertainty in ground truth.

2 Data Analysis

The goal of the initial data exploration was to investigate the amount of non-figurativeness in the train and trial data. Our analysis revealed that 99% of the training data could be classified using a simple heuristic: a regular expression decision list, hereafter called Tweet Label System (TLS), to split the training data into different key-phrase subgroups. The system searches for the expression in a tweet and then assigns a label in a cascade fashion following the order in Table 2, which lists the 14 possible label types (plus NONE), their associated expressions along with the support for each category

in the training data. Table 1 shows that only a small fraction of the train and trial data could not be associated to a subgroup and it can be seen that the final test data was estimated to have a very different distribution with 25% of tweets presumably containing literal language use.

Dataset	Train	Trial	Test
Instances	7988	920	4000
% Non-figurative	1%	7%	25%

Table 1: Retrieved instances in each data set and estimated amount of non-figurativeness.

Since there are obvious subgroups in the data, our hypothesis is that this fact can be used to construct a more informed baseline. In fact (§ 4.1), simply predicting the mean per subgroup pushed the constant mean baseline performance considerably (from 0.73 to 0.81 Cosine, compared to random 0.59).

Figure 1 plots predicted scores (ridge model, §3.1) of three subgroups against the gold scores on the trial data. It can be seen that certain subgroups have similar behaviour, ‘sarcasm’ has a generally negative cloud and the model performs well in predicting these values, while other groups such as ‘SoToSpeak’ have more intra-group variance.

Label	Expression	Support	Label	Expression	Support
Sarcasm	#sarcas	2139	SoToSpeak	so to speak	135
Irony	#iron(y ic)	1444	Proverbial	proverbial	22
Not	#not	3601	JustKidding	#justkidding	-
Literally	literally	344	Not2	not	29
Virtually	virtually	8	about	about	8
YeahRight	#yeahright	47	Oh	oh	3
OhYouMust	Oh.*you	2	NONE	-	92
asXas	as.*as	83			

Table 2: Tweet Label Type and Expression.

The Effect of a General Sentiment System

The data for this task is very different from data that most lexicon-based or general sentiment-analysis models fare best on. In fact, running a general sentiment classifier (GSA) described in Elming et al. (2014) on the trial data showed that its predictions are actually slightly anti-correlated with the gold standard scores for the Tweets in this task (cosine similarity score of -0.08 and MSE of 18.62). We exploited these anti-correlated results as features for our stacking systems (cf. § 3.2). Figure 2 shows the




Figure 1: Label Plots for RR predictions.

distributions of the gold scores and GSA predictions for the trial data. It shows that the gold distribution is skewed with regards to the number of negative instances to positives, while the GSA predicts more positive sentiment.

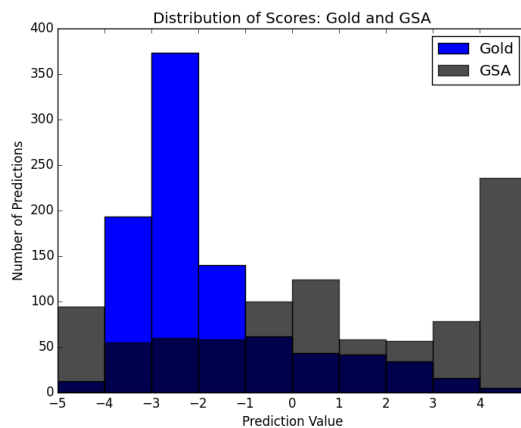


Figure 2: Distribution of Gold Scores and GSA Predictions for Trial Data.

3 System Description

We approach the task (Ghosh et al., 2015) as a regression task (cf. §4.4), combining several systems using stacking (§ 3.2), and relying on features without POS, lemma or explicit use of lexicons, cf. § 3.3.

3.1 Single Systems

Ridge Regression (RR) A standard supervised ridge regression model with default parameters.³

PCA_GMM Ridge Regression (GMM) A ridge regression model trained on the output of unsupervised induced features, i.e., a Gaussian Mixture Models (GMM) trained on PCA of word n-grams. PCA was used to reduce the dimensionality to 100, and GMM under the assumption that the data was sampled from different distributions of figurative language, k Gaussians were assumed (here $k = 12$).

Embeddings with Bayesian Ridge (EMBD) A Bayesian Ridge Regressor learner with default parameters trained on only word embeddings. A corpus was build from the training data and an in-house Tweet collection sampled with the expressions from the TLS. This resulted in a total of 3.7 million tweets and 67 million tokens. For details on how the word embeddings were built see §3.3.

3.2 Ensembles

We developed two stacking systems (Wolpert, 1992), *Stacking System 1* (ST1) and *Stacking System 2: Stacking with Backoff* (ST2). The systems used for these are shown in Table 3 and the Meta Learner used for both stacking systems is Linear Regression.

The systems used in ST1 and ST2 are not the only differences between the two. ST2 uses the TLS to identify the subgroup that each tweet belongs to. For any tweet with the NONE subgrouping, the system would back off to the predictions from the GSA. We built ST2 as a system that is not limited to sentiment analysis for a small subsection of language, the phenomenon of figurative language, but is applicable in situations covering many types of tweets including those in which literal language is used.

Single System / Stacking System	ST1	ST2
RR	X	X
GMM	X	
EMBD		X
GSA	X	X

Table 3: Systems in Ensemble Setups.

³<http://scikit-learn.org/>

3.3 Features

This section describe the features we used for the models in §3.1. Table 4 indicates the type of features used for the single models. Punctuation was kept as its own lexical item and we found removing stopwords and normalizing usernames to '@USER' increased performance and as such the preprocessing methods are the same across the models. Features were set on the trial data.

1. **Word N-Grams** Systems use different n-grams as features. In RR counts of 1 and 5 word grams, in GMM binary presence of 1,2, and 3 word grams.
2. **Uppercase Words** Counts of the numbers of word in a Tweet with all uppercase letters.
3. **Punctuation** Contiguous sequences of question, exclamation, and question and exclamation marks.
4. **TLS Label** The subgrouping label from TLS.
5. **Word Embeddings** Parameters for word embeddings:⁴ 100 dimensions, 5 minimum occurrences for a type to be included in the model, 5 word context window and 10-example negative sampling. Each tweet was represented by 100 features that represented the average of all the embeddings of the content words in the tweet.

Features/Systems	RR	GMM	EMBD
Word N-grams	X	X	
Uppercase	X		
Punctuations	X		
TLS Label	X		
Word Embeddings			X

Table 4: Features used in Single Models.

4 Results

4.1 Constant Baselines & Single Systems

We implemented the Mean, Mode, Median, Random and TSL (§2) baseline systems. TSL is the hardest baseline, and RR is the only system that beats it.

4.2 Results Stacking Systems

The performance of the stacking systems on the trial data can be seen below in Table 6. ST2 did not perform well on the trial data although a reason for this

⁴<https://code.google.com/p/word2vec/>

System	Cosine	MSE
TLS	0.81	2.34
Mean	0.73	3.13
Mode	0.73	3.13
Median	0.73	3.31
Random	0.59	5.17
RR	0.88	1.60
GMM	0.79	2.55
EMB	0.78	2.64

Table 5: Baseline and Single Systems On Trial Data.

is that only 7% of the trial data was found as not belonging to a known figurative type of tweet.

System	Cosine	MSE
ST1	0.86	1.88
ST2	0.79	2.57

Table 6: Stacking Model Results on Trial Data.

4.3 Final Results

Three models were submitted for final evaluation on the test data. The three models were RR, ST1, and ST2. For the final results we scaled back values outside the range [-5,5] to the nearest whole number in range. Tables 7 and 8 show the results for our systems on the final dataset and the performance of the overall winning system for the task (CLAC). Table 7 shows the overall cosine similarity and MSE for the systems on the test data and Table 8 shows the breakdown of the cosine similarity for the systems on the different parts of language. It is interesting to note that the performance of ST2 on the ‘Other’ type of language is identical as the performance for CLAC, this is also the best cosine similarity score ‘Other’ out of all submissions.

System	Test Cosine	Test MSE
RR	0.625	3.079
ST1	0.623	3.078
ST2	0.661	3.404
CLAC	<u>0.758</u>	<u>2.117</u>

Table 7: Submission System Test Results.⁵

System	Overall	Sarcasm	Irony	Metaphor	Other
RR	0.625	0.897	0.886	0.325	0.218
ST1	0.623	0.900	0.903	0.308	0.226
ST2	0.661	0.875	0.872	0.453	0.584
CLAC	<u>0.758</u>	0.892	<u>0.904</u>	<u>0.655</u>	<u>0.584</u>

Table 8: Cosine Test Results Breakdown.

4.4 The Case for Regression

Regression is less usual in NLP than classification. However for this data, it is desirable to use regression, because it incorporates the ordered relation between the labels, instead of treating them as orthogonal. It also keeps the decimal precision in the target variable when training, which is relevant when the target variable is the result of an average between several annotations. We ran classification experiments for this task but found that the best classification system’s⁶ performance (Cosine 0.82, MSE 2.51) is still far from the RR model (0.88,1.60).

5 Conclusions

We tested three systems for their abilities to analyse sentiment on figurative language from Twitter. Our experiments showed that a general SA system trained on literal Twitter language was highly anti-correlated with gold scores for figurative tweets. We found that for certain figurative types, sarcasm and irony, our system’s predictions for these phenomena fared well. Our system did not explicitly use a lexicon to define the sentiment of a tweet, but instead used machine learning and strictly corpus-based features (no POS or lemma) to place us 4th in the task. More effort may be needed to discriminate metaphorical from literal tweets to build a more robust system, although, even for humans the sentiment of tweets is hard to judge. This can be seen from the data where a number of tweets were repeated, but did not always share the same gold score.

⁵The numbers in bold indicate the best performance among our systems, underlined indicates the best performance between any of our systems and the winning system.

⁶Decision Tree with 7 classes and using the minimum score for instances in the classes in the training data to convert for class labels to scores.

References

- Enrique Amigó, Jorge Carrillo De Albornoz, Irina Chugur, Adolfo Corujo, Julio Gonzalo, Tamara Martín, Edgar Meij, Maarten De Rijke, and Damiano Spina. 2013. Overview of replab 2013: Evaluating online reputation monitoring systems. In *Information Access Evaluation. Multilinguality, Multimodality, and Visualization*, pages 333–352. Springer.
- Jacob Eisenstein. 2013. What to do about bad language on the internet. In *NAACL*.
- Jakob Elming, Barbara Plank, and Dirk Hovy. 2014. Robust cross-domain sentiment analysis for low-resource languages. In *Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*.
- A. Ghosh, G. Li, T. Veale, P. Rosso, E. Shutova, A. Reyes, and J. Barnden. 2015. Semeval-2015 task 11: Sentiment analysis of figurative language in twitter. In *Int. Workshop on Semantic Evaluation (SemEval-2015)*.
- Dirk Hovy, Barbara Plank, and Anders Søgaard. 2014. When POS datasets don't add up: Combatting sample bias. In *LREC*.
- Antonio Reyes and Paolo Rosso. 2012. Making objective decisions from subjective data: Detecting irony in customer reviews. *Decision Support Systems*, 53(4):754–760.
- Antonio Reyes, Paolo Rosso, and Tony Veale. 2013. A multidimensional approach for detecting irony in twitter. *Language Resources and Evaluation*, 47(1):239–268.
- Erik Tjong Kim Sang and Johan Bos. 2012. Predicting the 2011 dutch senate election results with twitter. In *Proceedings of the Workshop on Semantic Analysis in Social Media*.
- David H Wolpert. 1992. Stacked generalization. *Neural networks*, 5(2):241–259.