

Exploiting Class Labels to Boost Performance on Embedding-based Text Classification

Arkaitz Zubiaga

a.zubiaga@qmul.ac.uk

Queen Mary University of London

London, UK

ABSTRACT

Text classification is one of the most frequent tasks for processing textual data, facilitating among others research from large-scale datasets. Embeddings of different kinds have recently become the de facto standard as features used for text classification. These embeddings have the capacity to capture meanings of words inferred from occurrences in large external collections. While they are built out of external collections, they are unaware of the distributional characteristics of words in the classification dataset at hand, including most importantly the distribution of words across classes in training data. To make the most of these embeddings as features and to boost the performance of classifiers using them, we introduce a weighting scheme, Term Frequency-Category Ratio (TF-CR), which can weight high-frequency, category-exclusive words higher when computing word embeddings. Our experiments on eight datasets show the effectiveness of TF-CR, leading to improved performance scores over the well-known weighting schemes TF-IDF and KLD as well as over the absence of a weighting scheme in most cases.

1 INTRODUCTION

Word embeddings, or distributed word representations, have become one of the most common features for processing text. Word embeddings have been successfully used in numerous NLP and IR tasks such as sentiment analysis [3], machine translation [29], search [8] or recommender systems [18], as well as different domains such as biomedicine [4] or finance [5], outperforming traditional vector representation methods based on bags-of-words or n-grams.

In this work we focus on text classification [23], where word embeddings and derivatives are commonly used to represent the textual content of the instances to be classified [27]. While word embeddings are widely used features for text classification, they are generally used for vector representation of the textual content of the instances, independent of the importance each word can have within each category. We propose to incorporate information derived from category labels in the training data to improve vector representations. Here we propose Term Frequency-Category Ratio (TF-CR), a weighting scheme that exploits the category labels from training data to produce an improved vector representation using word embeddings, which is informed by category distributions in the training data.

There is a dearth of work in the literature trying to exploit the distribution of content across classes in the training data for improving word embedding representations. The few works that tackled the

problem have mainly focused on doing so for large-scale datasets, where it is possible to train separate word embedding models for each category thanks to the availability of abundant in-domain data. This is the case for problems such as sentiment analysis, where one can build such large annotated datasets by using distant supervision. In this paper, we aim to develop an improved word embedding representation for text classification where training data is not necessarily so abundant. To do that, we are the first to propose a novel weighting scheme, Term Frequency-Category Ratio (TF-CR), which can be applied on pre-trained, domain-agnostic word embedding models, only leveraging the training data available in the dataset at hand for dataset-specific weighting of the embeddings. The intuition behind TF-CR is to assign a higher weight to high-frequency, category-exclusive words as observed in the training data.

Our experiments on eight classification datasets show the consistent effectiveness of TF-CR, significantly improving performance of word embedding representations for text classification over the use of the well-known weighting schemes TF-IDF and KLD, as well as over unweighted word embeddings.

2 RELATED WORK

Early methods for learning distributed representations of words [2] through the so-called neural probabilistic language models have more recently gained momentum as embeddings [11, 20]. It sparked development of additional methods to reduce the dimensionality of traditional vector representation methods such as bags-of-words, by learning word embeddings. Two of the best-known methods to learn word embeddings include Word2Vec [17] and Glove [19], which enable dimensionality reduction as well as capturing semantic similarities across words. The key intuition is that, having a large corpus to train a model from, one can learn semantic characteristics of words by analysing their context, i.e. words surrounding other words. This leads to vectors of reduced dimensionality (word embeddings) to represent each word, which are normally between 100 and 500 dimensions. One of the widely adopted practices for sentence representation is then to get the sum or the average of the word embeddings in the sentence in question [28].

Use of word embeddings for text classification without specific weighting of words, however, ignores potentially useful information that can be extracted from class labels. While this problem has been tackled before, there is limited work exploring the utility of class labels to make the most of word embeddings for text classification. Previous work leveraging class labels to boost the performance of word embeddings on text classification has largely focused on sentiment analysis. The sentiment analysis task is suitable as it is possible to collect large, distantly supervised datasets [9] which are exploited to train sentiment-specific embeddings. Having large

annotated datasets, one can then train separate word embedding models for each class or learn models that incorporate class distributions in them. This has been achieved in different methods by combining multiple neural networks [14, 15, 24–26] or by using separate training processes [31] to train different word embedding models for each class in the dataset. This however requires availability of very large collections of labelled data to train separate models, which is possible for classification tasks exploiting distant supervision for data collection, as is the case with sentiment analysis. However, it is more limited for other text classification problems where gathering labelled data is expensive. In what follows, we propose a new weighting scheme to tackle this problem, TF-CR.

3 THE TF-CR WEIGHTING SCHEME

We propose a novel weighting scheme for word embedding representation in text classification tasks, which aims to determine the importance of each word for each particular category based on the distribution of the word across categories in the training data; this can provide additional information that word embeddings inherently disregard. This can be achieved by using well-known weighting schemes which are often used for text classification based on bags-of-words, such as TF-IDF [13, 22] and Kullback-Leibler Divergence (KLD) [16].

To suit the purposes of word embeddings in text classification, here we propose a new weighting scheme. The Term Frequency-Category Ratio (TF-CR) is a simple weighting scheme that combines the importance of a word within a category (Term Frequency, TF) and the distribution of the word across all categories (Category Ratio, CR). Both TF and CR are computed for each word w within each category c . TF is computed as the ratio of words in a category that are w , i.e. $TF_{wc} = \frac{|w_c|}{N_c}$, where $|w_c|$ is the number of occurrences of w in c , and N_c is the total number of words in c . CR is computed as the ratio of occurrences of w that occur within the category c , i.e. $CR_{wc} = \frac{|w_c|}{|w|}$, where $|w|$ denotes the number of occurrences of w across all categories.

The final TF-CR is the product of both metrics (Equation 1).

$$TF - CR = \frac{|w_c|}{N_c} * \frac{|w_c|}{|w|} = \frac{|w_c|^2}{N_c * |w|} \quad (1)$$

TF-CR ultimately gives a high weight to words that occur exclusively and with high frequency within a category. Low-frequency words exclusive to a category and high-frequency words that frequently occur across all categories will get lower scores.

3.1 Applying TF-CR on embeddings

In order to create a representation weighted using TF-CR, we first build category-specific word embedding representations of a text. This category-specific representation is created by summing up the embeddings of each of the words in a sentence, multiplied by their TF-CR score. This leads to k TF-CR-weighted embedding representations, where k is the number of categories in the dataset. We finally concatenate these k embedding representations to produce the final vector, which has a dimensionality of $k \times d$, where d is the number of dimensions of the word embedding model.

4 EXPERIMENTS

4.1 Datasets

We use eight different datasets:

- **RepLab polarity dataset [1]**: A dataset of 84,745 tweets mentioning companies, annotated for polarity as positive, negative or neutral.¹
- **ODPtweets [30]**: a large-scale dataset with nearly 25 million tweets, each categorised into one of the 17 categories of the Open Directory Project (ODP).
- **Restaurant reviews [12]**: a large dataset of 14,542,460 TripAdvisor restaurant reviews with their associated star rating ranging from 1 to 5.
- **SemEval sentiment tweets [21]**: we aggregate all annotated tweets from the SemEval Twitter sentiment analysis task from 2013 to 2017. The resulting dataset contains 61,767 tweets.
- **Distantly supervised sentiment tweets**: by using a large collection of tweets from January 2013 to September 2019 released on the Internet Archive², we produce a dataset of tweets annotated for sentiment analysis by using distant supervision following [9], leading to tweets annotated as positive or negative. The resulting dataset contains 33,203,834 tweets.³
- **Hate speech dataset [7]**: a dataset of 99,996 tweets, each categorised into one of {abusive, hateful, spam, normal}.
- **Newspace200 [6]**: a dataset of nearly 500K news articles, each categorised into one of 14 categories, including business, sports entertainment.http://groups.di.unipi.it/~gulli/AG_corpus_of_news_articles.html
- **20 Newsgroups**: a collection of nearly 20,000 newsgroup documents, pertaining to 20 different newsgroups, which are used as categories.⁴

For all datasets, we randomly sample 100,000 instances, except for those with fewer instances.

4.2 Word Embedding Models & Classifiers

We tested four word embedding models: (1) Google’s Word2Vec model (**gw2v**), (2) a Twitter Word2Vec model⁵ (**tw2v**) [10], (3) GloVe embeddings trained from Common Crawl (**cglove**) and (4) GloVe embeddings trained from Wikipedia (**wglove**).⁶

Different classifiers were tested. Due to limited space, we show here results obtained with a logistic regression classifier and *tw2v* embeddings, which consistently lead to optimal results. We report macro-F1 values as performance scores.

4.3 Weighting Schemes

We compare four different weighting schemes, all of which are applied following the methodology in §3.1:

- **No weighting (no wgt).**

¹<http://nlp.uned.es/replab2013/>

²<https://archive.org/details/twitterstream>

³<http://www.zubiaga.org/datasets/sentiment1319/>

⁴<http://qwone.com/~jason/20Newsgroups/>

⁵<https://fredericgodin.com/software/>

⁶<https://nlp.stanford.edu/projects/glove/>

- **TF-IDF**, which weights words with low document frequency higher. We compute TF-IDF scores for each word within each category, therefore calculating the importance of the word in each category.
- **KLD**, which determines the saliency of a word in a category with respect to the rest of the categories. Again KLD leads to a score for each word in each category.
- **TF-CR**, our weighting scheme defined in §3.

4.4 Varying Sizes of Training Sets

While we have up to 90,000 instances available as training data, we perform experiments with varying numbers of training instances. This allows us to assess the extent to which weighting schemes can help with varying sizes of training data, provided that calculations of weights using these schemes are done solely from the training data available in each case. We performed experiments for training sets ranging from 1,000 to 9,000. Training instances are randomly sampled in each training scenario, keeping the random sample consistent across different experiments with the same training size, and incrementally adding instances, i.e. a training set of 5,000 contains all of the training instances of that with 4,000 plus another 1,000. All performance scores reported are the result of averaging 10-fold cross-validation experiments.

5 RESULTS

Table 1 shows the results with varying numbers of training instances. We observe that TF-CR consistently outperforms the other weighting schemes, TF-IDF and KLD, regardless of the training size. The gap between TF-CR and the other weighting schemes generally becomes larger as the training data increases, showing that TF-CR exploits the class distributions more effectively. We can also observe that the unweighted approach outperforms TF-CR in six out of eight datasets when the training data is as small as 1,000 instances. However, TF-CR becomes more effective as the training data increases. TF-CR outperforms the unweighted method in five out of eight datasets with 10,000 training instances, and in seven out of eight datasets with 90,000 training instances. This reinforces the effectiveness of TF-CR for mid-sized training sets and above.

Figure 1 shows the tendency of all four methods as the training size varies from 1,000 to 90,000, with steps of 1,000. With the exception of the *hate speech* dataset, TF-CR outperforms all other methods for larger training sets. Moreover, TF-CR consistently outperforms all other methods for most training sizes in five datasets: *20newsgroups*, *newspace200*, *odptweets*, *restaurants* and *sentiment*.

6 DISCUSSION

We have introduced TF-CR,⁷ a first-of-its-kind weighting scheme that can leverage word distributions across categories in training data for text classification. The intuition behind TF-CR is to give higher weights to frequent words that exclusively or predominantly occur within a category. This leads to category-specific weights for each word, which allows an embedding representation that captures varying importances of words across categories. Experimenting on eight datasets, we show that (1) it improves over unweighted word embeddings in seven of the datasets with large training datasets,

⁷Code available at <https://github.com/azubiaga/tfcr>

	20ng	hs	ns200	odp	rl	rest	sem	sent
1K training instances								
no wgt	0.554	0.613	0.448	0.234	0.324	0.415	0.577	0.694
TF-IDF	0.653	0.534	0.436	0.184	0.265	0.403	0.465	0.634
KLD	0.668	0.543	0.440	0.179	0.264	0.400	0.452	0.620
TF-CR	0.783	0.566	0.456	0.196	0.268	0.426	0.474	0.665
2K training instances								
no wgt	0.591	0.620	0.473	0.264	0.333	0.443	0.592	0.698
TF-IDF	0.730	0.558	0.460	0.217	0.274	0.410	0.483	0.636
KLD	0.734	0.578	0.463	0.209	0.263	0.431	0.489	0.631
TF-CR	0.836	0.588	0.481	0.239	0.278	0.452	0.512	0.690
5K training instances								
no wgt	0.626	0.634	0.503	0.290	0.364	0.460	0.606	0.707
TF-IDF	0.645	0.545	0.476	0.266	0.308	0.408	0.503	0.633
KLD	0.646	0.590	0.491	0.262	0.296	0.459	0.537	0.642
TF-CR	0.811	0.600	0.516	0.296	0.332	0.479	0.562	0.708
10K training instances								
no wgt	0.596	0.641	0.519	0.301	0.370	0.475	0.612	0.713
TF-IDF	0.440	0.536	0.490	0.295	0.303	0.399	0.502	0.636
KLD	0.475	0.606	0.524	0.300	0.318	0.475	0.549	0.651
TF-CR	0.696	0.614	0.541	0.351	0.347	0.490	0.585	0.716
40K training instances								
no wgt	0.705	0.656	0.538	0.312	0.418	0.495	0.634	0.718
TF-IDF	0.893	0.548	0.493	0.335	0.367	0.390	0.528	0.648
KLD	0.860	0.632	0.570	0.349	0.368	0.505	0.577	0.660
TF-CR	0.930	0.635	0.575	0.423	0.427	0.513	0.632	0.736
90K training instances								
no wgt	0.705	0.661	0.544	0.325	0.422	0.503	0.635	0.721
TF-IDF	0.893	0.556	0.507	0.354	0.379	0.390	0.532	0.647
KLD	0.860	0.643	0.586	0.362	0.364	0.516	0.577	0.663
TF-CR	0.930	0.648	0.595	0.458	0.444	0.522	0.638	0.748

Table 1: Comparison of results using different weighting schemes for varying sizes of training data.

(2) it improves consistently for most training sizes in five of the datasets. TF-CR also consistently outperforms TF-IDF and KLD.

Our objective here has been to introduce and validate TF-CR. Additional tuning of classifier parameters, adding features, etc. for achieving state-of-the-art performance is beyond the scope of this work. We also aim to extend this work by further exploring the differences across datasets, to determine dataset characteristics that maximise the benefits of TF-CR.

ACKNOWLEDGMENTS

This research utilised Queen Mary’s Apocrita HPC facility, supported by QMUL Research-IT. <http://doi.org/10.5281/zenodo.438045>

REFERENCES

- [1] E. Amigó, J. Carrillo de Albornoz, I. Chugur, A. Corujo, J. Gonzalo, T. Martín, E. Meij, M. de Rijke, and D. Spina. 2013. Overview of RepLab 2013: Evaluating Online Reputation Monitoring Systems. In *Proceedings of CLEF*. 333–352.
- [2] Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. A neural probabilistic language model. *Journal of machine learning research* 3, Feb (2003), 1137–1155.
- [3] Danushka Bollegala, Tingting Mu, and John Yannis Goulermas. 2016. Cross-domain sentiment classification using sentiment sensitive embeddings. *IEEE Transactions on Knowledge and Data Engineering* 28, 2 (2016), 398–410.

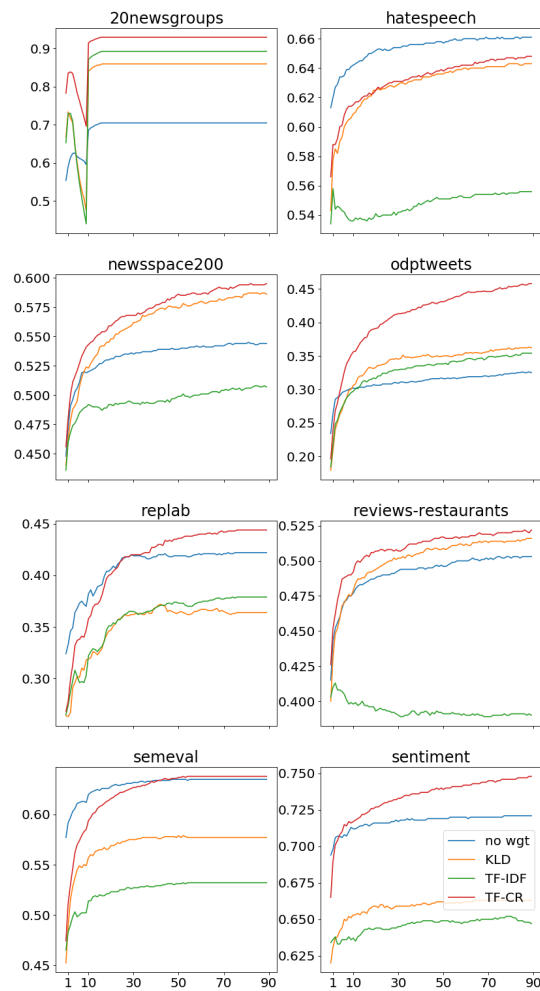


Figure 1: Macro-F1 performance scores for the eight datasets under study, with varying sizes of training data.

[4] Billy Chiu, Gamal Crichton, Anna Korhonen, and Sampo Pyysalo. 2016. How to train good word embeddings for biomedical NLP. In *Proceedings of the 15th Workshop on Biomedical Natural Language Processing*, 166–174.

[5] Keith Cortis, André Freitas, Tobias Daudert, Manuela Huerlimann, Manel Zarrouk, Siegfried Handschuh, and Brian Davis. 2017. Semeval-2017 task 5: Fine-grained sentiment analysis on financial microblogs and news. In *Proceedings of SemEval*, 519–535.

[6] Gianna M Del Corso, Antonio Gulli, and Francesco Romani. 2005. Ranking a stream of news. In *Proceedings of WWW*. ACM, 97–106.

[7] Antigoni-Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. Large Scale Crowdsourcing and Characterization of Twitter Abusive Behavior. In *Proceedings of ICWSM*. AAAI Press.

[8] Debasis Ganguly, Dwaipayan Roy, Mandar Mitra, and Gareth JF Jones. 2015. Word embedding based generalized language model for information retrieval. In *Proceedings of SIGIR*, 795–798.

[9] Alec Go, Richa Bhayani, and Lei Huang. 2009. Twitter sentiment classification using distant supervision. *CS224N Project Report, Stanford* 1, 12 (2009).

[10] Frédéric Godin, Baptist Vandersmissen, Wesley De Neve, and Rik Van de Walle. 2015. Multimedia Lab @ ACL WNUT NER Shared Task: Named Entity Recognition for Twitter Microposts using Distributed Word Representations. In *Proceedings of the Workshop on Noisy User-generated Text*, 146–153.

[11] Martin Grohe. 2020. word2vec, node2vec, graph2vec, X2vec: Towards a Theory of Vector Embeddings of Structured Data. *arXiv preprint arXiv:2003.12590* (2020).

[12] Aiqi Jiang and Arkaiz Zubiaga. 2019. Leveraging aspect phrase embeddings for cross-domain review rating prediction. *PeerJ Computer Science* 5 (2019), e225.

[13] Karen Spärck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation* 28 (1972), 11–21.

[14] Sicong Kuang and Brian D Davison. 2018. Class-specific word embedding through linear compositionality. In *2018 IEEE International Conference on Big Data and Smart Computing (BigComp)*. IEEE, 390–397.

[15] Sicong Kuang and Brian D Davison. 2019. Learning class-specific word embeddings. *The Journal of Supercomputing* (2019), 1–28.

[16] Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics* 22, 1 (1951), 79–86.

[17] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, 3111–3119.

[18] Cataldo Musto, Giovanni Semeraro, Marco de Gemmis, and Pasquale Lops. 2016. Learning word embeddings from wikipedia for content-based recommender systems. In *European Conference on Information Retrieval*. Springer, 729–734.

[19] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of EMNLP*, 1532–1543.

[20] Mohammad Taher Pilehvar and Jose Camacho-Collados. 2020. *Embeddings in Natural Language Processing: Theory and Advances in Vector Representation of Meaning*. Morgan & Claypool.

[21] Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. SemEval-2017 task 4: Sentiment analysis in Twitter. In *Proceedings of SemEval*, 502–518.

[22] Gerard Salton and Christopher Buckley. 1988. Term-weighting approaches in automatic text retrieval. *Information processing & management* 24, 5 (1988), 513–523.

[23] Fabrizio Sebastiani. 2002. Machine learning in automated text categorization. *ACM computing surveys (CSUR)* 34, 1 (2002), 1–47.

[24] Duyu Tang. 2015. Sentiment-specific representation learning for document-level sentiment analysis. In *Proceedings of WSDM*. ACM, 447–452.

[25] Duyu Tang, Furu Wei, Bing Qin, Nan Yang, Ting Liu, and Ming Zhou. 2016. Sentiment embeddings with applications to sentiment analysis. *IEEE Transactions on Knowledge and Data Engineering* 28, 2 (2016), 496–509.

[26] Duyu Tang, Furu Wei, Nan Yang, Ming Zhou, Ting Liu, and Bing Qin. 2014. Learning sentiment-specific word embedding for twitter sentiment classification. In *Proceedings of ACL*, Vol. 1, 1555–1565.

[27] Peng Wang, Bo Xu, Jiaming Xu, Guanhua Tian, Cheng-Lin Liu, and Hongwei Hao. 2016. Semantic expansion using word embedding clustering and convolutional neural network for improving short text classification. *Neurocomputing* 174 (2016), 806–814.

[28] Ruqing Zhang, Jiafeng Guo, Yanyan Lan, Jun Xu, and Xueqi Cheng. 2018. Aggregating Neural Word Embeddings for Document Representation. In *Proceedings of ECIR*. Springer, 303–315.

[29] Will Y Zou, Richard Socher, Daniel Cer, and Christopher D Manning. 2013. Bilingual word embeddings for phrase-based machine translation. In *Proceedings of EMNLP*, 1393–1398.

[30] Arkaiz Zubiaga and Heng Ji. 2013. Harnessing web page directories for large-scale classification of tweets. In *Proceedings of WWW*. ACM, 225–226.

[31] Arkaiz Zubiaga and Aiqi Jiang. 2020. Early Detection of Social Media Hoaxes at Scale. *ACM Transactions on the Web (TWEB)* 14, 4 (2020), 1–23.