CLASSIFICATION OF TRANSVERSAL DOMINIC SENSE USING SENSITIVE SENSITIVE EMBEDDING

1. N.KIRAN KUMAR 2. K.DEEPTHI
1.PG SCHOLAR, 2.ASSISTANTPROFESSOR
DEPARTMENT OF CSE, SCIENT INSTITUTE OF TECHNOLOGY HYDERABAD

ABSTRACT— Unsupervised Cross-domain Sentiment Classification is the task of adapting a sentiment classifier trained on a particular domain (source domain), to a different domain (target domain), without requiring any labeled data for the target domain. By adapting an existing sentiment classifier to previously unseen target domains, we can avoid the cost for manual data annotation for the target domain. We model this problem as embedding learning, and construct three objective functions that capture: (a) distributional properties of pivots (i.e., common features that appear in both source and target domains), (b) label constraints in the source domain documents, and (c) geometric properties in the unlabeled documents in both source and target domains. Unlike prior proposals that first learn a lower-dimensional embedding independent of the source domain sentiment labels, and next a sentiment classifier in this embedding, our joint optimisation method learns embeddings that are sensitive to sentiment classification. Experimental results on a benchmark dataset show that by jointly optimising the three objectives we can obtain better performances in comparison to optimizing each objective function separately, thereby demonstrating the importance of task-specific embedding learning for cross-domain sentiment classification. Among the individual objective functions, the best performance is obtained by (c). Moreover, the proposed method reports cross-domain sentiment classification accuracies that are statistically comparable to the current state-of-the-art embedding learning methods for cross-domain sentiment classification.

1 INTRODUCTION

THE ability to correctly identify the sentiment expressed in user-reviews about a particular product is an important task for several reasons. First, if there is a negative sentiment associated with a particular feature of a product, the manufacturer can take immediate actions to address the issue. Failing to detect a negative sentiment associated with a product might result in decreased sales. From the users' point-of-view, in online stores where one cannot physically touch and evaluate a product as in a real-world store, the user opinions are the only available subjective descriptors of the product. By automatically classifying the user-reviews according to the sentiment expressed in them, we can assist the potential buyers of a product to easily understand the overall opinion about that product. Considering the numerous applications of sentiment classification such as opinion mining [1], opinion summarisation [2], contextual advertising [3], and market analysis [4], it is not surprising that sentiment classification has received continuous attention. Sentiment classification can be considered as an instance of text classification where a given document must be classified into a pre-defined set of sentiment classes [5]. We use the term document to refer various types of user reviews. In binary sentiment classification, a document must be classified into two classes depending on whether it expresses a positive or a negative sentiment towards an entity. Alternatively, a document can be assigned a discrete sentiment score (eg. from one to five stars) that indicates the degree of the positivity (or negativity) of the sentiment. Once, a document has been identified as sentiment bearing, then further analysis can be performed, for example, to extract evidence for an argument.

2 RELATED WORK

Cross-domain sentiment classification methods can be classified as unsupervised versus supervised methods. In unsupervised cross-domain sentiment classification, the training data consist of (a) source domain labeled documents, (b) source domain unlabeled documents, and (c) target domain unlabeled documents. Supervised (or semi-supervised) cross-domain sentiment classification methods use a small set of labeled data for the target domain in addition to those three data sources. Unsupervised cross-domain sentiment classification can be considered as a much harder problem because of the lack of availability of labeled data for the target domain. Unsupervised domain adaptation methods assume that the output labels in the target domain are equally conditioned by the input, even though the input could be differently distributed in terms of marginal probability. Therefore, domain adaptation methods adjust for the differences in this conditional distributions between the two domains. Structural correspondence learning (SCL) [10] first selects a set of pivots, common features to both source and the target domains, using some criteria. One approach for selecting pivots is to select all features that occur more than a predefined number of times in both domains. Alternatively, a word association measure such as the mutual information (MI) could be used to measure the degree of association of a feature to a domain name, and select common features that have a high degree of association between both the source and the target domains [10]. The latter approach has shown to produce better results in cross-domain sentiment classification. Next, linear predictors are trained to predict the presence (or absence) of pivots in a document. Specifically, documents in which a particular pivot w occurs are considered as positive training instances for learning a predictor for w, whereas an equal number of documents in which w does not occur are selected as negative training instances. Unigram and bigram lexical-features are extracted from the selected training instances as features to train a binary logistic regression classifier with l2 regularisation. Finally, the weight vector learnt by the classifier is considered as the predictor for w. The predictors learnt for all pivots are arranged in a matrix on which singular value decomposition (SVD) is performed. The left singular vectors corresponding to the largest singular values are selected from the SVD result, and arranged as row vectors in a matrix. All source domain labeled training instances are multiplied by this matrix to predict the presence of pivots. Finally, a binary logistic regression model is trained using the predicted pivots and the original features. By first predicting the pivots, and then learning a classifier using those predicted pivots as additional features, SCL attempts to reduce the mismatch between features in the source and the target domains. Spectral feature alignment (SFA) [11] splits the feature space into two mutually exclusive groups: domain independent features (pivots), and domain specific features (all other features). Next, a bipartite graph is constructed between the two groups where the edge connecting a domain specific and a domain independent feature is weighted by the number of different documents in which the corresponding two features co-occur. Spectral clustering is performed on this bipartite graph to create a lower dimensional representation in which co-occurring domain specific and domain independent features are represented by the same set of lower dimensional features. Similarly to SCL, SFA trains a binary logistic regression model in this lower-dimensional space using the labeled documents from the source domain. Both SFA and SCL are similar to our proposed method in that first, a lower-dimensional feature representation is learnt, and second a binary sentiment classifie is trained on this embedded space. However, our proposed method is different from SCL and SFA in that, we consider not only the unlabeled data but also labeled data for the source domain when constructing the representation.

As we later see in Section 6, this enables us to learn customized representations that result in better performance on our final task of cross-domain sentiment classification. Bollegala et al. [14] created a Sentiment Sensitive Thesaurus (SST) that lists words that express similar sentiments in the source and target domains. For example, SST created from the two domains books and knives lists interesting as a related word for sharp. The thesaurus is automatically created using a sentiment sensitive asymmetric similarity measure that uses sentiment labels in the source domain documents. Analogous to the thesauri-based query expansion in information retrieval, SST is used to expand the source domain feature vectors by appending related features in the target domain. A binary logistic regression classifier is trained using the expanded feature vectors corresponding to the source domain labeled documents. Unlike, SCL or SFA, SST does not create lower-dimensional embeddings. Mapping Function The main strategy of mapping the words and documents to the space is to first compute the word embeddings, and then derive the document embeddings based on the word embeddings by considering the word occurrences. Linear projection is assumed to transform the original feature representation of words to their embedding presentation.Specifically, a d _ k projection matrix PA is used tomap words in domain A to a k-dimensional embedding space Rk, while a d _ h projection matrix PB is used to map words in domain B to the same embedding space. Given in total M þMA words in domain A including the M pivots appearing in both domains and MA non-pivot words only appearing in domain A, we let fezðAÞ i gMþMA i¼1 denote their corresponding word embeddings stored in an ðM þMAÞ _ k embedding matrix eZ A computed by the linear projection mapping given as Here, the pivots appear in both domains, thus possess two sets of feature representations UA and UB. Subsequently, they possess two sets of embedding representations after being mapped from the

two domains, which are UAPA and UBPB. Later on, we will show that according to rule 1
these two representations should be as similar as possible in order to reach the alignment between the two domains via pivots matching. After computing the word embeddings fezðAÞ I gMþMA i¼1 and i¼1 , it is straightforward to derive the embedding representations for documents. For example, it is possible to view the embedding vector of each document as a weighted sum of the embedding vectors of all the words that appear in the corresponding document. The occurrence frequencies (or scores) can be used as the summation weights to determine the contribution level of the words. Letting denote the document embeddings in domains A and B, respectively, we thus have Letting ZA and ZB denote the NA _ k and NB _ k document embedding matrices, the above formulations can be converted to their matrix presentation, given as where the notation D_1ðXÞ denotes a diagonal matrix compute from an input matrix X ¼ ½xij_ with its ith diagonal element obtained by ð P j xijÞ_1. Although a document is represented using the words that appear in that document, we are free to select any feature representation for the individual words. Specifically, the d features we use in the domain A (or the h features we use in the domain B) need not be words limited to the domain and can be, for example, bigrams of words or partof- speech tags of the words. The document feature space
and the word feature spaces are connected via Eqs. (3) and (4). This de-coupling of document and word representations allows us to incorporate semantically rich word representation such as the recently developed neural word embeddings. For the simplicity of the presentation, we limit the discussion in this paper to lexical features (unigrams and bigrams of words) and differ a study of rich semantic feature spaces to future work.

3 Model Construction

According computation of the word and document embeddings relies on the computation of the two projection matrices of PA and PB based on the input matrices of UA, UB, A, B, XA, XB and Y. In the following, we show how to derive PA and PB by solving an optimization problem constructed based on the three rules. DATASET We use the cross-domain sentiment classification dataset3 prepared by Blitzer et al. [10] in our experiments. This dataset consists of Amazon product reviews for four different product types: books, DVDs, electronics and kitchen appliances. Each review is assigned with a rating (0-5 stars), a reviewer name and location, a product name, a review title and date, and the review text. Reviews with rating > 3 are labeled as positive, whereas those with rating < 3 are labeled as negative. For each domain, there are 1;000 positive and 1;000 negative examples, the same balanced composition as the polarity dataset constructed by Pang et al.. The dataset also contains on average 17;547 unlabeled reviews for the four domains. Following previous work, we randomly select 800 positive and 800 negative labeled reviews from each domain as training instances (total number of training instances are 1;600 _ 4 ¼ 6;400),and the remainder is used for testing (total number of test instances are 400 _ 4 ¼ 1;600). Mutual information between a feature and the labeled reviews have been used in [10] for selecting pivots. However, in unsupervised domain adaptation we have labeled data only for the source domain. Therefore, there is no guarantee that we will obtain pivots that behave similarly in both the source as well as the target domains by this method. Moreover, source domain labeled data are only a fraction of all the data available for the adaptation task. Cooccurrence counts and probability estimates conducted using small datasets are likely to be sparse and unreliable. To overcome these issues, we use a pointwise mutual information (PMI)-based pivot selection method to select pivots that consider both source and the target domains. Specifically,

4 CONCLUSION

We considered three constraints that must be satisfied by an embedding that can be used to train a cross-domain sentiment classification method. We evaluated the performance of the individual constraints as well as their combinations using a benchmark dataset for cross-domain sentiment classification. Our experimental results show that some of the combinations of the proposed constraints obtain results that are statistically comparable to the current state-of-the-art methods for cross-domain sentiment classification. Unlike previously proposed embedding learning approaches for cross-domain sentiment classification, our proposed method uses the label information available for the source domain reviews, thereby learning embeddings that are sensitive to the final task of application, which is sentiment classification.

REFERENCES

[1] B. Pang and L. Lee, “Opinion mining and sentiment analysis,” Found. Trends Inf. Retrieval, vol. 2, nos. 1/2, pp. 1–135, 2008.

[2] Y. Lu, C. Zhai, and N. Sundaresan, “Rated aspect summarization of short comments,” in Proc. 18th Int. Conf. World Wide Web, 2009, pp. 131–140.

[3] T.-K. Fan and C.-H. Chang, “Sentiment-oriented contextual advertising,” Knowl. Inf. Syst., vol. 23, no. 3, pp. 321–344, 2010.

[4] M. Hu and B. Liu, “Mining and summarizing customer reviews,” in Proc. 10th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining, 2004, pp. 168–177.

[5] C. D. Manning and H. Sch€utze, Foundations of Statistical Natural Language Processing. Cambridge, MA, USA: MIT Press, 2002.

[6] H. Daum_e III, A. Kumar, and A. Saha, “Co-regularization based semi-supervised domain adaptation,” in Proc. Adv. Neural Inf. Process. Syst. 23, 2010. pp. 478–486.

[7] D. Lopez-Paz, J. M. Hernandez-Lobato, and B. Scholkopf, “Semisupervised domain adaptation with non-parametric copulas,” inProc. Adv. Neural Inf. Process. Syst. 25, 2012, pp. 674–682.

[8] H. Daum_e III, “Frustratingly easy domain adaptation,” in Proc. 45th Annu. Meeting Assoc. Comput. Linguistics, 2007, pp. 256–263.

[9] J. Blitzer, R. McDonald, and F. Pereira, “Domain adaptation with structural correspondence learning,” in Proc. Conf. Methods Natural
Language Process., 2006, pp. 120–128.

[10] J. Blitzer, M. Dredze, and F. Pereira, “Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification,” in Proc. 45th Annu. Meeting Assoc. Comput. Linguistics, 2007, pp. 440–447.

[11] S. J. Pan, X. Ni, J.-T. Sun, Q. Yang, and Z. Chen, “Cross-domain sentiment classification via spectral feature alignment,” in Proc. 19th Int. Conf. World Wide Web, 2010, pp. 751–760.

[12] D. Bollegala, D. Weir, and J. Carroll, “Cross-domain sentiment classification using a sentiment sensitive thesaurus,” IEEE Trans. Knowl. Data Eng., vol. 25, no. 8, pp. 1719–1731, Aug. 2013.

[13] X. Glorot, A. Bordes, and Y. Bengio, “Domain adaptation for large-scale sentiment classification: A deep learning approach,” in Proc. 28th Int. Conf. Mach. Learn., 2011, pp. 513–520.

[14] D. Bollegala, D. Weir, and J. Carroll, “Using multiple sources to construct a sentiment sensitive thesaurus for cross-

domain sentiment classification," in Proc. 49th Annu. Meet. Assoc. Comput. Linguistics: Human Language Technol., 2011, pp. 132–141.

[15] D. Bollegala, D. Weir, and J. Carroll, "Learning to predict distributions of words across domains," in Proc. Assoc. Comput. Linguistics, 2014, pp. 613–623.

[16] P. H. Calais Guerra, A. Veloso, W. Meira Jr., and V. Almeida, "From bias to opinion: A transfer-learning approach to real-time sentiment analysis," in Proc. 17th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining, 2011, pp. 150–158.

[17] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman, "Indexing by latent semantic analysis," J. Am. Soc. Inf. Sci., vol. 41, no. 6, pp. 391–407, 1990.

[18] I. T. Jolliffe, Principal Component Analysis. New York, NY, USA: Springer-Verlag, 1986.