



Deep learning for sentiment analysis: successful approaches and future challenges

Duyu Tang, Bing Qin* and Ting Liu

Sentiment analysis (also known as opinion mining) is an active research area in natural language processing. It aims at identifying, extracting and organizing sentiments from user generated texts in social networks, blogs or product reviews. A lot of studies in literature exploit machine learning approaches to solve sentiment analysis tasks from different perspectives in the past 15 years. Since the performance of a machine learner heavily depends on the choices of data representation, many studies devote to building powerful feature extractor with domain expert and careful engineering. Recently, deep learning approaches emerge as powerful computational models that discover intricate semantic representations of texts automatically from data without feature engineering. These approaches have improved the state-of-the-art in many sentiment analysis tasks including sentiment classification of sentences/documents, sentiment extraction and sentiment lexicon learning. In this paper, we provide an overview of the successful deep learning approaches for sentiment analysis tasks, lay out the remaining challenges and provide some suggestions to address these challenges.

© 2015 John Wiley & Sons, Ltd

How to cite this article:

WIREs Data Mining Knowl Discov 2015, 5:292–303. doi: 10.1002/widm.1171

INTRODUCTION

Sentiment analysis and opinion mining^{1–3} are the field that analyze people's opinions, sentiments, emotions from user generated texts. It is one of the most active research areas in natural language processing,^{4,5} and is also widely studied in data mining, web mining, and social media analytics as sentiments are key influencers of human behaviors. With the rapid growth of social media such as Twitter, Facebook, and online review sites such as IMDB, Amazon, Yelp, sentiment analysis draws growing attentions from both research and industry communities.

According to the definition by Liu et al.,² sentiment (or an opinion) is a quintuple, (e, a, s, h, t),

where e is the name of an entity, a is the aspect of e, s is the sentiment on aspect a of entity e, h is the opinion holder, and t is the time when the opinion is expressed by h. In this definition, the sentiment s can be a positive, negative, or neutral sentiment, or a numeric rating score expressing the strength/intensity of the sentiment (e.g., 1–5 stars) in review sites like Yelp and Amazon. The entity can be a product, service, topic organization, or event.^{6,7} Let us take a toy example to better explain the definition of 'sentiment.' Supposing an Amazon user called Tom posted a review 'The photos from my Samsung camera are not that great, but the battery life is great.' at June 4, 2015. In this example, there are two sentiment quintuples, namely (*Samsung*, *photo*, *negative*, *Tom*, *June 4, 2015*) and (*Samsung*, *battery life*, *positive*, *Tom*, *June 4, 2015*). Based on the definition of 'sentiment,' the objective of sentiment analysis aims at discovering all the sentiment quintuples in a document.

Sentiment analysis tasks are derived from the five components of the sentiment quintuple. For

*Correspondence to: qinb@ir.hit.edu.cn

Department of Computer Science and Technology, Harbin Institute of Technology, Harbin, China

Conflict of interest: The authors have declared no conflicts of interest for this article.

example, the task of document level sentiment classification⁸ targets at the third component (sentiment such as positive, negative, and neutral) while ignoring the other aspects. The task of fine-grained opinion extraction⁹ focuses on the first four components of the quintuple. In the past 15 years, machine learning driven methods almost dominate sentiment analysis tasks. As feature representation greatly affects the performance of a machine learner,¹⁰ a lot of studies in literature focus on effective features by hand with domain expertise and careful engineering.^{11,12} But this can be avoided by representation learning algorithms, which automatically discover discriminative and explanatory text representations from data.¹³ Deep learning is a kind of representation learning approach. It learns multiple levels of representation with nonlinear neural networks, each of which transforms the representation at one level into a representation at a higher and more abstract level.¹⁴ The learned representations can be naturally used as features and applied for detection or classification tasks.

In this study, we introduce successful deep learning algorithms for sentiment analysis. The notation of ‘deep learning’ in this article stands for learning continuous and real-valued text representation/feature automatically from data, mostly with neural network approaches. We first describe the methods to learn continuous word representation, also called word embedding, as words are the basic computational units of natural language. These word embeddings can be used as inputs to subsequent sentiment analysis tasks, so that we describe how word embeddings are investigated for different sentiment analysis tasks. In particular, we describe semantic compositional methods that compute representations of longer expressions (e.g., sentence or document) for sentence/document level sentiment classification task,^{15,16} followed by neural sequential models for fine-grained opinion extraction.¹⁷ As sentiment lexicon is an important resource for many sentiment analysis systems, we also describe neural methods to build large-scale sentiment lexicons.¹⁸ We finally conclude this study and provide some future directions.

WORD REPRESENTATION

Word representation aims at representing aspects of word meaning. For example, the representation of ‘cellphone’ may capture the facts that cellphones are electronic products, that they include battery and screen, that they can be used to chat with others, and so on.

A straight-forward way is to encode a word w_i as a one-hot vector. It has the same length as the size of the vocabulary, and only one dimension is 1, with all others being 0. However, the one-hot word representation only encodes the indices of words in a vocabulary, while failing to capture rich relational structure of the lexicon. One common approach to discover the similarities between words is to learn word clusters.^{19,20} Each word is associated with a discrete class, and words in the same class are similar in some respect. This leads to a one-hot representation over a smaller vocabulary size.

Instead of characterizing the similarity with a discrete variable based on clustering results which corresponds to a soft or hard partition of the set of words, many researchers target at learning a continuous and real-valued vector for each word, also known as word embedding. Existing embedding learning algorithms are typically based on the distributional hypothesis,²¹ which states that words in similar contexts have similar meanings. To this end, many matrix factorization methods can be viewed as modeling word representations. For example, Latent Semantic Indexing (LSI)²² can be regarded as learning a linear embedding with a reconstruction objective, which uses a matrix of ‘term-document’ co-occurrence statistics, e.g., each row stands for a word or term and each column corresponds to an individual document in the corpus. Hyperspace Analogue to Language²³ utilizes a matrix of ‘term-term’ co-occurrence statistics, where both rows and columns correspond to words and the entries stand for the number of times a given word occurs in the context of another word. Hellinger PCA²⁴ is also investigated to learn word embeddings over ‘term-term’ co-occurrence statistics. As standard matrix factorization methods do not incorporate task-specific information, it is not clear whether they are useful enough for a target goal. Supervised Semantic Indexing²⁵ tackles this problem and takes the supervised information of a targeted task (e.g., information retrieval) into consideration. They learn the embedding model from click-through data with a margin ranking loss.

With the revival of interest in deep learning and neural network,^{10,14,26} a surge of studies learn word embeddings with neural network. A pioneered work in this field is given by Bengio et al.,²⁷ which introduces a neural probabilistic language model that learns simultaneously a continuous representation for words and a probability function for word sequences based on these word representations. Given a word w_i and its preceding context words, the algorithm first maps each context word to its continuous vector with a shared lookup table. Afterward, context word

vectors are fed to a feed-forward neural network with softmax as output layer to predict the conditional probability of next word w_i . The parameters of neural network and lookup table are jointly estimated with back propagation. Following Bengio et al.'s²⁷ work, a lot of approaches are proposed to speed-up training processing or capturing richer semantic information. Bengio et al.²⁷ introduce a neural architecture by concatenating the vectors of context words and current word, and use importance sampling to effectively optimize the model with observed 'positive sample' and sampled 'negative samples.' Morin and Bengio²⁸ develop hierarchical softmax to decompose the conditional probability with a hierarchical binary tree. Mnih and Hinton²⁹ introduce a log-bilinear language model. Collobert and Weston³⁰ train word embeddings with a ranking-type hinge loss function³¹ by replacing the middle word within a window with a randomly selected one. Mikolov et al.^{32,33} introduce Continuous Bag-of-Words (CBOW) and Continuous Skip-gram, and release the popular word2vec (<https://code.google.com/p/word2vec/>) toolkit. The CBOW model predicts the current word based on the embeddings of its context words, and the Skip-gram model predicts surrounding words given the embedding of current word. Mnih and Kavukcuoglu³⁴ accelerate the word embedding learning procedure with Noise Contrastive Estimation.³⁵ There are also many algorithms developed for capturing richer semantic information, including global document information,³⁶ word morphemes,³⁷ dependency-based contexts,³⁸ word-word co-occurrence,³⁹ sense of ambiguous words,⁴⁰ semantic lexical information in WordNet,⁴¹ and hierarchical relations between words.⁴²

The aforementioned neural network algorithms typically only use the contexts of words to learn word embeddings. As a result, the words with similar contexts but opposite sentiment polarity like 'good'

and 'bad' are mapped into close vectors in the embedding space. This is meaningful for some tasks such as pos-tagging⁴³ as the two words have similar usages and grammatical roles, but it becomes a disaster for sentiment analysis as they have the opposite sentiment polarity. In order to learn word embeddings tailored for sentiment analysis tasks, some studies encode sentiment of texts in continuous word representation. Maas et al.⁴⁴ introduce a probabilistic topic model by inferring the polarity of a sentence based on the embedding of each word it contains. Labutov and Lipson⁴⁵ re-embed an existing word embedding with logistic regression by regarding sentiment supervision of sentences as a regularization item. Tang et al.⁴⁶ extend the C&W model³⁰ and develop three neural networks to learn sentiment-specific word embedding from Twitter. The tweets containing positive and negative emoticons are used as training data, regarding positive and negative emoticon signals as sentiment indicators. The learned word embeddings are applied for Twitter sentiment classification, and perform comparable performance with the state-of-the-art hand-crafted features¹¹ on SemEval 2013 dataset. They also build a system named 'Cooooo!!!' and participate in the Twitter sentiment classification evaluation in SemEval 2014.⁴⁷ The system yields the second place among 45 participants. Holding a similar idea, Tang et al.¹⁸ extend SkipGram method³³ and leverage sentiment of texts for word embedding learning. Brief illustrations of two neural models for learning sentiment-specific word embeddings are given in Figure 1.

SENTIMENT CLASSIFICATION

Sentiment classification is a fundamental and extensively studied area in sentiment analysis.^{1,2} It targets at determining the sentiment polarity (positive or




FIGURE 1 | Neural models for learning sentiment-specific word embeddings. The hybrid prediction model is proposed by Tang et al.⁴⁶ and the hybrid ranking model is introduced by Tang et al.¹⁸

negative) of a sentence (or a document) based on its textual content. In the literature, existing studies for sentiment classification are dominated by two mainstream directions: lexicon-based approach and corpus-based approach.

Lexicon-based approaches^{48,49} mostly use a dictionary of sentiment words with their associated sentiment polarity, and incorporate negation and intensification to compute the sentiment polarity for each sentence (or document). A representative lexicon-based method is given by Turney,⁴⁸ which consists of three steps. He first extract phrases if their postags conform to the predefined patterns. Afterwards the sentiment polarity of each extracted phrase is estimated through point-wised mutual information (PMI), which measures the degree of statistical dependence between two terms. In Turney's work, the PMI score is calculated by feeding queries to a search engine and collecting the number of hits. Finally, he averages the polarity of all phrases in a review as its sentiment polarity. Ding and Liu⁵⁰ apply negation words like 'not,' 'never,' 'cannot,' and contrary words like 'but' to enhance the performance of lexicon-based method. Taboada et al.⁴⁹ integrate intensifications and negation words with the sentiment lexicons annotated with their polarity and sentiment strength. Thelwall et al.⁵¹ develop SentiStrength using sentiment lexicons and linguistic rules to detect the sentiment strength of tweets. Reckman et al.⁵² develop an entirely rule-based system for Twitter sentiment classification. They use lexicalized hand-written rules, each of which is a pattern that matches words or sequences of words.

Corpus-based methods treat sentiment classification as a special case of text categorization problem.⁸ They mostly build the sentiment classifier from sentences with annotated sentiment polarity. The sentiment supervision can be manually annotated, or automatically collected by sentiment signals like emoticons in tweets⁵³ or human ratings in reviews.⁴⁴ Pang et al.⁸ pioneer to treat the sentiment classification of reviews as a special case of text categorization problem and first investigate machine learning methods. A brief illustration is given in Figure 2. They employ Naive Bayes, Maximum Entropy, and Support Vector Machines (SVM) with a diverse set of features. In their experiments, the best performance

is achieved by SVM with bag-of-words feature. Following Pang et al.'s work, many studies focus on designing or learning effective features to obtain a better classification performance. On movie and product reviews, Wang and Manning⁵⁴ present NBSVM, which trades-off between Naive Bayes and NB-feature enhanced SVM. Paltoglou and Thelwall⁵⁵ learn feature weights by investigating variants weighting functions from Information Retrieval, such as tf.idf and its BM25 variants. Nakagawa et al.⁵⁶ utilize dependency trees, polarity-shifting rules and conditional random fields (CRFs) with hidden variables to compute the document feature. On Twitter, Jiang et al.⁵⁷ use lexicon features and syntactic and POS tagging features. NRC-Canada¹¹ develop a state-of-the-art Twitter sentiment classifier in SemEval 2013 by using a variety of sentiment lexicons and hand-crafted features.

Feature engineering is important but labor-intensive. It is therefore desirable to discover explanatory factors from the data and make the learning algorithms less dependent on extensive feature engineering. With the rapid growing of deep learning (representation learning^{10,26}), many recent studies focus on learning the low-dimensional, dense, and real-valued vector as text features for sentiment analysis without any feature engineering. Existing deep learning methods for sentiment classification typically include two stages. In the first stage, they learn word embeddings from text corpus.^{33,46} In the second stage, word embeddings are applied to producing the representations of sentences/documents with semantic composition.⁵⁸ Embedding learning algorithms have been described in previous Section. Existing composition learning approaches are typically based on the principle of compositionality,⁵⁹ which states that the meaning of a longer expression (e.g., a sentence or a document) comes from the meanings of its constituents and the rules used to combine them. Specifically, Bessalov et al.⁶⁰ initialize the word embeddings by latent semantic analysis and further represent each document as the linear weighted of ngram vectors for sentiment classification. Glorot et al.⁶¹ use Stacked Denoising Autoencoder in an unsupervised fashion based on reconstruction, and apply it for domain adaptation. Autoencoder is a kind of neural network that is optimized by reconstructing the input itself.




FIGURE 2 | Supervised framework for sentiment classification of sentences/documents.

Denoising Autoencoding randomly masks the values of inputs and tries to reconstruct the noisy inputs. Socher et al.¹⁵ introduce a family of recursive deep neural models including Recursive Autoencoder (RAE),⁶² Matrix–Vector Recursive Neural Network (MV-RNN),⁶³ and Recursive Neural Tensor Network (RNTN)¹⁵ to learn the composition of variable-length phrases based on the representation of its children. Specifically, RAE first learns the structure of a sentence with greedy unsupervised reconstruction, and further conduct compositionality over the learned tree structure. In RAE, each word is encoded as a vector and the calculator is matrix multiplication plus nonlinear hyperbolic tangent function. In MV-RNN, each word is also associated with a matrix representation, and the tree structure is obtained from an external parser (like Stanford parser). In RNTN, they use neural tensor network as the compositional functions to better capture the interactions between elements. An example of RNN is given in Figure 3(b). RNN is one of the successful approaches in NLP community and extended with global feedbackward,⁶⁴ feature weight tuning,⁶⁵ deep recursive layer,⁶⁶ adaptive composition functions,⁶⁷ combined with Combinatory Categorical Grammar,⁶⁸ and used for opinion relation detection.⁶⁹ In particular, Hermann et al.⁶⁸ present Combinatory Categorical Autoencoders to learn the compositionality of sentence, which marries the Combinatory Categorical Grammar with Recursive Autoencoder. Dong et al.⁶⁷ develop adaptive RNN that uses more than one composition functions and adaptively select them depending on the input vectors. Li⁶⁵ extend RNN by tuning feature weight to control how much one specific unit contributes to the higher-level representation.

Another powerful neural network for semantic composition is convolutional neural network

(CNN).⁷⁰ A brief illustration is given in Figure 3(a). For example, Kalchbrenner et al.¹⁶ develop dynamic CNNs (DCNN) and introduce the dynamic k-max pooling for learning the sentence representation. Kim⁷¹ and Johnson and Zhang⁷² also investigate CNNs, and achieve several state-of-the-art performances on some benchmark datasets for sentiment classification. In Johnson and Zhang's⁷² model, the word order is effectively incorporated.

Sequential model like recurrent neural network and long short-term memory (LSTM) are also verified as strong approaches for semantic composition. Such method process a sentence in a sequential way from the beginning of a sentence to the end, one word at a time. Neural networks with shared parameters are used for calculation. The basic computational unit could be a simple matrix multiplication or a complex LSTM unit. Li et al.⁷³ compare the effectiveness of RNN and recurrent neural network on five NLP tasks including sentiment classification. Some recent studies^{74,75} investigate tree-structured LSTM for semantic composition.

Beyond sentence level compositionality, a few studies move eyes on document level semantic composition.^{76,77} Two recent studies exploit the semantic relationship between sentences for document level sentiment analysis. Tang et al.⁷⁸ compose sentence representations to document representation in a sequential way, without using external discourse parser. Bhatia et al.,⁷⁹ however, use RST discourse parser and integrate the parsed results with RNN for document level sentiment analysis.

OPINION EXTRACTION

Given a piece of text (e.g., a sentence or a document), fine-grained opinion extraction targets at finding the

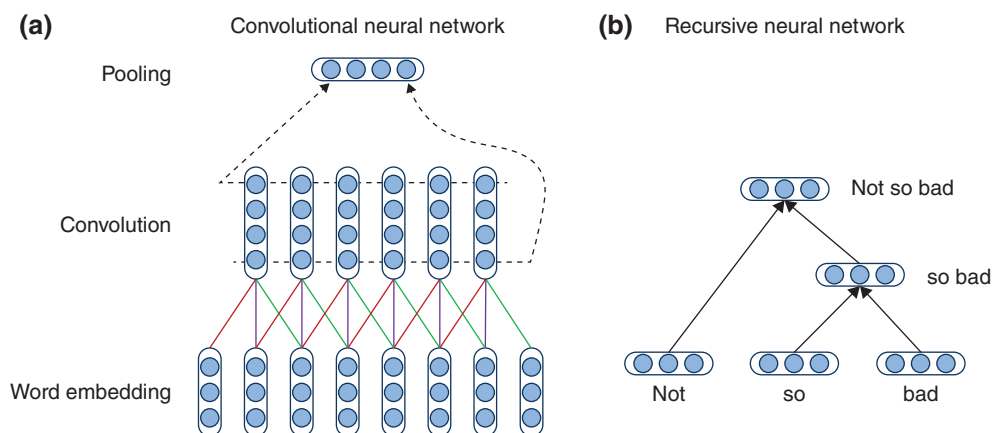


FIGURE 3 | A brief illustration about convolutional neural network and recursive neural network.

elements of a sentiment quintuple including the opinion holder, the entity/aspect, the sentiment of an affective expression, and so on.

This task is typically regarded as a sequential labeling problem.^{80,81} For example, the sentence ‘*the screen of apple is so amazing*’ should be tagged with ‘O B-Aspect O B-Entity O B-Sentiment I-Sentiment,’ where B indicates the beginning of an expression, I means the tokens inside an expression, and O stands for tokens outside any predefined class. To address this problem, a lot of studies use CRFs or semi-CRF with manually designed discrete features such as word features, phrase features, and syntactic features.^{12,80} As an alternative, Irsoy and Cardie¹⁷ use neural network without using hand-crafted features. An illustration of the approach is given in Figure 4. The method proceeds a word at a time from the beginning to the end of a sentence. For each index, the input is word embedding, which will be fed to one or more hidden layers to get more abstractive and discriminative representations of data. Each hidden unit is computed based on its history and the current input. The method can be naturally extended to a bi-directional way, where the ‘history’ of a hidden unit comes from its surrounding contexts without constraining to the preceding words. The last hidden layer of a position is regarded as the representation of the corresponding word, and is used for classifying the tag label of the word. This method outperforms strong CRF baselines on a benchmark opinion extraction dataset.¹⁷ Paulus et al.⁶⁴ investigate an analogous tree structured RNN for fine grained sentiment analysis task.

For fine-grained sentiment analysis, deep learning (neural network) approaches also achieve some promising results recently. For example, Vo and Zhang⁸² show that rich features including word embedding perform well on target-dependent sentiment classification. Zhang et al.⁸³ use word embeddings and integrate them into neural CRF for

open-domain sentiment analysis. Liu et al.⁸⁴ conduct empirical studies for fine-grained sequential labeling task, and show that LSTM recurrent neural network performs better than feature-based CRF on benchmark datasets.

BUILDING SENTIMENT LEXICON

A sentiment lexicon is a list of words and phrases such as ‘excellent,’ ‘awful,’ and ‘coooooll,’ each of which is assigned with a positive or negative score reflecting its sentiment polarity and strength. Sentiment lexicon is a fundamental component for sentiment analysis, which can be built manually, through thesaurus-based method and corpus-based method. As the manual method is time-consuming, it is mainly combined with automatic methods as the final check.² We describe thesaurus-based and corpus-based method.

Thesaurus-based method mainly utilizes semantic relationships (e.g., synonyms, antonyms, hypernyms, etc.) between tokens from an external thesaurus like WordNet. Under this direction, majority of existing studies regard word as basic unit,⁶ yet some researchers target at the synset⁸⁵ in WordNet. Kim and Hovy⁸⁶ use synonym and antonym relations from WordNet to build sentiment lexicon. The hypothesis is that the synonyms of a positive word should have a positive polarity, and vice versa for antonym relation. They manually label a small size of adjective and verb words as sentiment seeds, and then apply a bootstrapping method to expand the seed list. Esuli and Sebastiani⁸⁷ take the gloss information of a word in WordNet into consideration. They manually label some sentiment seeds, and use a semi-supervised method to classify the polarity of a word based on its gloss in WordNet. Rao et al.⁸⁸ use label propagation algorithm to detect the polarity of words in the graph induced from WordNet. Feng




FIGURE 4 | An illustration of recurrent neural network for sequential labeling.

et al.⁸⁹ propose learning connotation lexicon, which lists words with connotative polarity, i.e., words with positive connotation (e.g., 'award' and 'promotion') and words with negative connotation (e.g., 'cancer' and 'war'). Feng et al.⁹⁰ go a step further and focus on the words that are objective on the surface like 'intelligence,' 'human,' 'cheesecake,' and so forth. Besides detecting the sentiment of each word, many researchers focus on identifying the polarity of a WordNet synset (also referred to as sense). Baccianella et al.⁹¹ release the well-known SentiWordNet, in which each synset is associated with three numerical scores, describing how objective, positive, and negative the terms contained in the synset are. Each score in SentiWordNet is in range (0.0, 1.0) and the summation is 1.0. Esuli et al.⁹² use pagerank to rank senses of WordNet in terms of how strongly they are positive or negative. Su et al.⁹³ use a semi-supervised framework based on mincut to recognize the subjectivity of a word sense in WordNet.

As the thesaurus like WordNet cannot well cover the growing colloquial sentiment expressions on the web, many researches employ corpus-based methods that induce a sentiment lexicon from text corpora. Hatzivassiloglou et al.⁹⁴ pioneer this field by extracting the polarity of adjective words. The idea is that words conjoined with 'and' favor to have the same polarity, and words conjoined by 'but' favor to have opposite polarity labels. They start with a list of sentiment seeds, and then identify more subjective adjectives with pre-defined conjunction patterns. Qiu et al.⁹⁵ propose a semi-supervised method dubbed double propagation for opinion word expansion and target extraction based on dependency relations between sentiment words and aspect words. An enhanced double propagation approach is given by Liu et al.,⁹⁶ and the method shows strong performances on multiple datasets. Velikovich et al.⁹⁷ represent words and phrases with their syntactic contexts within a window size from the web documents, and use graph propagation for label inference. Chen et al.⁹⁸ utilize the Urban Dictionary and extract the target-dependent sentiment expressions from Twitter. Mohammad et al.¹¹ use pointwise mutual information between each phrase and hashtag/emoticon seed words, such as '#good,' and '#bad,' ':),' and ':(' Severyn and Moschitti⁹⁹ build sentiment lexicons using scores from an SVM model learned on a distant supervised corpora.

Tang et al.¹⁸ cast sentiment lexicon learning as a phrase-level classification task and use deep learning strategy. The method consists of two part: (1) an embedding learning algorithm to effectively learn the continuous representation of phrases, which are used

as features for word-level sentiment classification, (2) a seed expansion algorithm that expands a small list of sentiment seeds to collect training data for building the word-level classifier. The framework is illustrated in Figure 5. The generated sentiment lexicon is evaluated via applied as features for Twitter sentiment classification. The lexicon shows superior performances over traditional sentiment lexicon like MPQA¹⁰⁰ and large-scale sentiment lexicons like Sentiment140.¹¹

SOME SUGGESTIONS ON IMPLEMENTATION

In this part, we briefly talk about how to implement deep learning (neural network) approaches for the beginners in this area. Let us take the hybrid prediction model in Figure 1 as an example. There are two options when we implement the model. The first choice is to calculate the gradient of each parameter in terms of the loss function by hand, and use that to update the value of each parameter. However, it is not scalable because we need to calculate the gradients again even if we just slightly modify the neural architecture. In order to pay more attentions to developing powerful neural architecture rather than gradient calculating, we use another choice with layer-wise implementation. For example, we implement lookup table, linear layer, tanh layer, softmax layer, and so on, and we link them together to build up the final model. Once the basic neural layers are completed, we do not need to care about them anymore and we can test variations of neural architecture easily. The implementations by us can be found at <http://ir.hit.edu.cn/~dytang>. One could also use GPU supported toolkits including Theano (deeplearning.net/software/theano/), Torch (torch.ch), and Caffe (caffe.berkeleyvision.org/).

CONCLUSION

In this study, we introduce successful deep learning approaches for several sentiment analysis tasks involving word embedding learning, sentiment classification, opinion extraction, and sentiment lexicon learning.

We lay out the future challenges of deep learning for aforementioned sentiment analysis tasks in this part. As deep learning is an emerging topic over machine learning and artificial intelligence in both research and industry community and is currently evolving very quickly, so the challenges might soon be outdated.

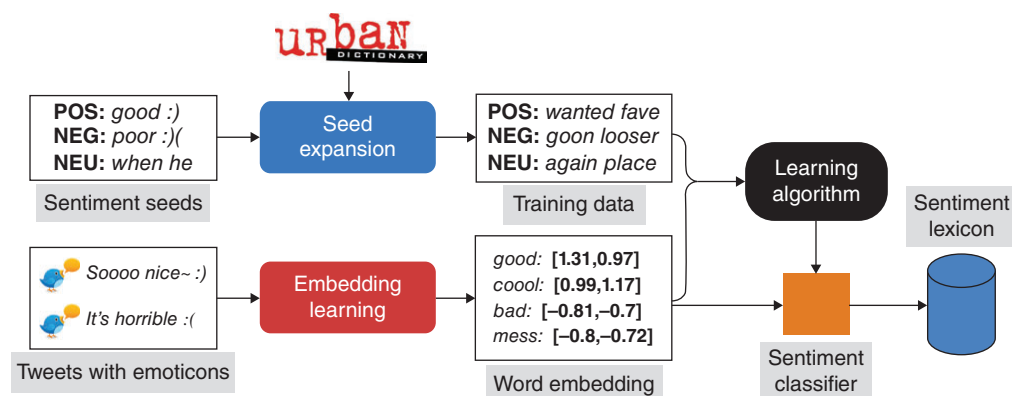


FIGURE 5 | A classification approach leveraging word embedding for building sentiment lexicon.

For learning better word representations for sentiment analysis tasks, Tang et al.^{18,46} incorporate sentence level sentiment signals as supervisions. In this line of research, different levels of sentiment signals may be also investigated such as external lexical level sentiment information. In addition, it has been show that the approached introduced by Tang et al.^{18,46} do not surpass context-based embeddings for document level classification task,^{101,102} which remains a potential challenge about how to use coarse- and fine-grained supervisions. Another interesting direction is how to interpret the meaning of a word embedding because the dense representations are uninterpretable. A recent study given by Faruqui et al.⁴¹ showed that transforming word embeddings into sparse vectors could yield promising results.

For semantic compositionality, most of existing studies focus on sentence level. Toward this goal, a

potential challenge is how to learn sentence structure^{103–105} together with the composed semantic representation as structure prediction¹⁰⁶ is a big challenge in both natural language processing and machine learning communities. Some studies also claim that the internal connections in CNN form an automatic structure of natural language. Another direction is how to interpret the effects of neural networks with visualization. Li et al.¹⁰⁷ provide several methods to visualize how neural models like LSTM are able to compose meanings in sentiment analysis task.

For building sentiment lexicon, the deep learning approach given by Tang et al.¹⁸ cannot infer the sentiment polarity of the phrases not covered by the existing vocabulary. How to generate new sentiment words/phrases from new corpus is a remained as a challenge. A self-studied life-long framework might be a practicable solution.

ACKNOWLEDGMENTS

We gratefully acknowledge the helpful discussions with Yaming Sun. We thank the editor and anonymous reviewers for their helpful comments and feedbacks. This work was supported by the National High Technology Development 863 Program of China (No. 2015AA015407) and National Natural Science Foundation of China (Nos 61133012 and 61273321). Duyu Tang is supported by Baidu Fellowship and IBM Ph.D. Fellowship.

FURTHER READING

Li J, Li R, Hovy E. Recursive deep models for discourse parsing. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, 2061–2069.

REFERENCES

1. Pang B, Lee L. Opinion mining and sentiment analysis. *Found Trends Inf Retr* 2008, 2:1–135.
2. Liu B. Sentiment analysis and opinion mining. *Synth Lect Human Lang Technol* 2012, 5:1–167.

3. Feldman R. Techniques and applications for sentiment analysis. *Commun ACM* 2013, 56:82–89.
4. Manning CD, Schütze H. *Foundations of Statistical Natural Language Processing*. Cambridge, MA: MIT Press; 1999.
5. Martin JH, Jurafsky D. *Speech and Language Processing*. 2nd ed. New Jersey: Prentice Hall; 2000.
6. Hu M, Liu B. Mining and summarizing customer reviews. In: *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Seattle, WA, August 2004, 168–177. ACM.
7. Deng L, Wiebe J. *MPQA 3.0: An Entity/Event-Level Sentiment Corpus*. NAACL; 2015.
8. Pang B, Lee L, Vaithyanathan S. Thumbs up?: sentiment classification using machine learning techniques. In: *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing*, vol 10, Morristown, NJ, USA, July 2002, 79–86. Association for Computational Linguistics.
9. Yang B, Cardie C. Joint inference for fine-grained opinion extraction. In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, vol 1, Sofia, Bulgaria, 2013, 1640–1649. Association for Computational Linguistics.
10. Bengio Y, Goodfellow IJ, Courville A. *Deep Learning*. Cambridge, MA: MIT Press (in preparation); 2015.
11. Mohammad SM, Kiritchenko S, Zhu X. NRC-Canada: building the state-of-the-art in sentiment analysis of tweets. In: *Second Joint Conference on Lexical and Computational Semantics (*SEM)*, vol 2, Atlanta, Georgia, USA, June 2013, 321–327.
12. Yang B, Cardie C. Extracting opinion expressions with semi-markov conditional random fields. In: *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, Jeju Island, Korea, July 2012, 1335–1345. Association for Computational Linguistics.
13. Bengio Y, Courville A, Vincent P. Representation learning: a review and new perspectives. *IEEE Trans Pattern Anal Mach Intell* 2013, 35:1798–1828.
14. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature* 2015, 521:436–444.
15. Socher R, Perelygin A, Wu JY, Chuang J, Manning CD, Ng AY, Potts C. Recursive deep models for semantic compositionality over a sentiment treebank. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Seattle, October 2013, 1631–1642.
16. Kalchbrenner N, Grefenstette E, Blunsom P. A convolutional neural network for modelling sentences. 2014, arXiv preprint arXiv:1404.2188
17. Irsoy O, Cardie C. Opinion mining with deep recurrent neural networks. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, 2014, 720–728.
18. Tang D, Wei F, Qin B, Zhou M, Liu T. Building large-scale twitter-specific sentiment lexicon: a representation learning approach. In: *Proceedings of COLING*, Dublin, Ireland, 2014, 172–182.
19. Brown PF, Desouza PV, Mercer RL, Pietra VJD, Lai JC. Class-based n-gram models of natural language. *Comput Linguist* 1992, 18:467–479.
20. Baker LD, McCallum AK. Distributional clustering of words for text classification. In: *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Melbourne, Australia, August 1998, 96–103. ACM.
21. Harris ZS. Distributional structure. *Word* 1954.
22. Deerwester SC, Dumais ST, Landauer TK, Furnas GW, Harshman RA. Indexing by latent semantic analysis. *J Am Soc Inf Sci* 1990, 41:391–407.
23. Lund K, Burgess C. Producing high-dimensional semantic spaces from lexical co-occurrence. *Behav Res Methods Instrum Comput* 1996, 28:203–208.
24. Lebre R, Legrand J, Collobert R. Is deep learning really necessary for word embeddings? EPFL Report No. 196986, Idiap Research Institute, 2013.
25. Bai B, Weston J, Grangier D, Collobert R, Sadamasa K, Qi Y, Chapelle O, Weinberger K. Learning to rank with (a lot of) word features. *Inf Retr* 2010, 13:291–314.
26. Hinton GE, Salakhutdinov RR. Reducing the dimensionality of data with neural networks. *Science* 2006, 313:504–507.
27. Bengio Y, Ducharme R, Vincent P, Janvin C. A neural probabilistic language model. *J Mach Learn Res* 2003, 3:1137–1155.
28. Morin F, Bengio Y. Hierarchical probabilistic neural network language model. In: *Proceedings of the International Workshop on Artificial Intelligence and Statistics*, January 2005, 246–252.
29. Mnih A, Hinton G. Three new graphical models for statistical language modelling. In: *Proceedings of the 24th International Conference on Machine Learning*, Corvallis, Oregon, USA, June 2007, 641–648. ACM.
30. Collobert R, Weston J. A unified architecture for natural language processing: deep neural networks with multitask learning. In: *Proceedings of the 25th International Conference on Machine Learning*, Helsinki, Fabianinkatu, July 2008, 160–167. ACM.
31. Smith NA, Eisner J. Contrastive estimation: training log-linear models on unlabeled data. In: *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, Ann Arbor, Michigan, June 2005, 354–362. Association for Computational Linguistics.
32. Mikolov T, Chen K, Corrado G, Dean J. Efficient estimation of word representations in vector space. 2013, arXiv preprint arXiv:1301.3781.

33. Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J. Distributed representations of words and phrases and their compositionality. In: *Conference on Advances in Neural Information Processing Systems*, 2013, 3111–3119.
34. Mnih A, Kavukcuoglu K. Learning word embeddings efficiently with noise-contrastive estimation. In: *Conference on Advances in Neural Information Processing Systems*, 2013, 2265–2273.
35. Gutmann MU, Hyvärinen A. Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *J Mach Learn Res* 2012, 13:307–361.
36. Huang EH, Socher R, Manning CD, Ng AY. Improving word representations via global context and multiple word prototypes. In: *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, vol 1, Jeju, Republic of Korea, July 2012, 873–882. Association for Computational Linguistics.
37. Qiu S, Cui Q, Bian J, Gao B, Liu TY. *Co-Learning of Word Representations and Morpheme Representations*. In: *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 141–150, Dublin, Ireland, August 23–29, 2014.
38. Levy O, Goldberg Y. Dependency-based word embeddings. In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, vol 2, Baltimore, Maryland, USA, 2014, 302–308.
39. Pennington J, Socher R, Manning CD. Glove: global vectors for word representation. In: *Proceedings of the Empirical Methods in Natural Language Processing (EMNLP 2014)*, vol 12, Doha, Qatar, 2014, 1532–1543.
40. Li J, Jurafsky D. Do multi-sense embeddings improve natural language understanding? 2015, arXiv preprint arXiv:1506.01070.
41. Faruqui M, Dodge J, Jauhar SK, Dyer C, Hovy E, Smith NA. Retrofitting word vectors to semantic lexicons. 2014, arXiv preprint arXiv:1411.4166.
42. Yogatama D, Faruqui M, Dyer C, Smith NA. Learning word representations with hierarchical sparse coding. 2014, arXiv preprint arXiv:1406.2035.
43. Ma J, Zhang Y, Zhu J. Tagging the web: building a robust web tagger with neural network. In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, vol 1, Baltimore, Maryland, 2014, 144–154.
44. Maas AL, Daly RE, Pham PT, Huang D, Ng AY, Potts C. Learning word vectors for sentiment analysis. In: *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, vol 1, June 2011, 142–150. Association for Computational Linguistics.
45. Labutov I, Lipson H. Re-embedding words. In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, Sofia, Bulgaria, August 2013, 489–493.
46. Tang D, Wei F, Yang N, Zhou M, Liu T, Qin B. Learning sentiment-specific word embedding for twitter sentiment classification. In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, vol 1, Baltimore, Maryland, USA, 2014, 1555–1565.
47. Tang D, Wei F, Qin B, Liu T, Zhou M. Coooolll: a deep learning system for twitter sentiment classification. In: *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, Dublin, Ireland, August 2014, 208–212.
48. Turney PD. Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In: *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, Philadelphia, PA, USA, July 2002, 417–424. Association for Computational Linguistics.
49. Taboada M, Brooke J, Tofloski M, Voll K, Stede M. Lexicon-based methods for sentiment analysis. *Comput Linguist* 2011, 37:267–307.
50. Ding X, Liu B, Yu PS. A holistic lexicon-based approach to opinion mining. In: *Proceedings of the 2008 International Conference on Web Search and Data Mining*, Palo Alto, California, USA, February 2008, 231–240. ACM.
51. Thelwall M, Buckley K, Paltoglou G. Sentiment strength detection for the social web. *J Am Soc Inf Sci Technol* 2012, 63:163–173.
52. Reckman H, Baird C, Crawford J, Crowell R, Micciulla L, Sethi S, Veress F. teragram: rule-based detection of sentiment phrases using SAS sentiment analysis. In: *Second Joint Conference on Lexical and Computational Semantics (*SEM)*, Atlanta, GA, 2013, 513–519.
53. Zhao J, Dong L, Wu J, Xu K. Moodlens: an emoticon-based sentiment analysis system for Chinese tweets. In: *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, August 2012, 1528–1531. ACM.
54. Wang S, Manning CD. Baselines and bigrams: simple, good sentiment and topic classification. In: *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, vol 2, Jeju, Republic of Korea, July 2012, 90–94. Association for Computational Linguistics.
55. Paltoglou G, Thelwall M. A study of information retrieval weighting schemes for sentiment analysis. In: *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, Uppsala, Sweden, July 2010, 1386–1395. Association for Computational Linguistics.
56. Nakagawa T, Inui K, Kurohashi S. Dependency tree-based sentiment classification using CRFs with hidden variables. In: *Human Language Technologies: The*

- 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics, June 2010, 786–794. Association for Computational Linguistics.
57. Jiang L, Yu M, Zhou M, Liu X, Zhao T. Target-dependent twitter sentiment classification. In: *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, vol 1, Portland, Oregon, USA, June 2011, 151–160. Association for Computational Linguistics.
 58. Mitchell J, Lapata M. Composition in distributional models of semantics. *Cognit Sci* 2010, 34:1388–1429.
 59. Frege G. On sense and nominatum. In: Feigl H, Sellars W, eds. *Readings in Philosophical Aristotelian Society*. New York: Appleton-Century-Crofts; 1949, 85–102.
 60. Beshpalov D, Bai B, Qi Y, Shokoufandeh A. Sentiment classification based on supervised latent n-gram analysis. In: *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*, Glasgow, United Kingdom, October 2011, 375–382. ACM.
 61. Glorot X, Bordes A, Bengio Y. Domain adaptation for large-scale sentiment classification: a deep learning approach. In: *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 2011, 513–520.
 62. Socher R, Pennington J, Huang EH, Ng AY, Manning CD. Semi-supervised recursive autoencoders for predicting sentiment distributions. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Edinburgh, Scotland, UK, July 2011, 151–161. Association for Computational Linguistics.
 63. Socher R, Huval B, Manning CD, Ng AY. Semantic compositionality through recursive matrix–vector spaces. In: *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, Jeju Island, Korea, July 2012, 1201–1211. Association for Computational Linguistics.
 64. Paulus R, Socher R, Manning CD. Global belief recursive neural networks. In: *Advances in Neural Information Processing Systems*, 2014, 2888–2896.
 65. Li J. Feature weight tuning for recursive neural networks. 2014, arXiv preprint arXiv:1412.3714.
 66. Irsoy O, Cardie C. Deep recursive neural networks for compositionality in language. In: *Advances in Neural Information Processing Systems*, 2014, 2096–2104.
 67. Dong L, Wei F, Tan C, Tang D, Zhou M, Xu K. Adaptive recursive neural network for target-dependent twitter sentiment classification. In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, 2014, 49–54.
 68. Hermann KM, Blunsom P. The role of syntax in vector space models of compositional semantics. In: *ACL (1)*, August 2013, 894–904.
 69. Xu L, Liu K, Lai S, Chen Y, Zhao J. Mining opinion words and opinion targets in a two-stage framework. In: *ACL (1)*, 2013, 1764–1773.
 70. Severyn A, Moschitti A. Twitter sentiment analysis with deep convolutional neural networks. In: *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, August 2015, 959–962. ACM.
 71. Kim Y. Convolutional neural networks for sentence classification. 2014, arXiv preprint arXiv:1408.5882.
 72. Johnson R., Zhang T. Effective use of word order for text categorization with convolutional neural networks. 2014, arXiv preprint arXiv:1412.1058.
 73. Li J, Jurafsky D, Hovy E. When are tree structures necessary for deep learning of representations? 2015, arXiv preprint arXiv:1503.00185.
 74. Tai KS, Socher R, Manning CD. Improved semantic representations from tree-structured long short-term memory networks. 2015, arXiv preprint arXiv:1503.00075.
 75. Zhu X, Sobhani P, Guo H. Long short-term memory over tree structures. 2015, arXiv preprint arXiv:1503.04881.
 76. Le QV, Mikolov T. Distributed representations of sentences and documents. 2014, arXiv preprint arXiv:1405.4053.
 77. Li J, Luong MT, Jurafsky D. A hierarchical neural autoencoder for paragraphs and documents. 2015, arXiv preprint arXiv:1506.01057.
 78. Tang D, Qin B, Liu T. Document modeling with gated recurrent neural network for sentiment classification. In: *EMNLP*, 2015, 1422–1432.
 79. Bhatia P, Ji Y, Eisenstein J. Better document-level sentiment analysis from RST discourse parsing. 2015, arXiv preprint arXiv:1509.01599.
 80. Choi Y, Cardie C, Riloff E, Patwardhan S. Identifying sources of opinions with conditional random fields and extraction patterns. In: *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, Vancouver, British Columbia, Canada, October 2005, 355–362. Association for Computational Linguistics.
 81. Breck E, Choi Y, Cardie C. Identifying expressions of opinion in context. In: *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-2007)*, vol 7, Hyderabad, India, January 2007, 2683–2688.
 82. Vo DT, Zhang Y. Target-dependent twitter sentiment classification with rich automatic features. In: *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)*, Buenos Aires, Argentina, 2015, 1347–1353.
 83. Zhang M, Zhang Y, Vo DT. Neural networks for open domain targeted sentiment. In: *Conference on*

- Empirical Methods in Natural Language Processing (EMNLP 2015)*, Lisbon, Portugal, 2015.
84. Liu P, Joty S, Meng H. Fine-grained opinion mining with recurrent neural networks and word embeddings. In: *Conference on Empirical Methods in Natural Language Processing (EMNLP 2015)*, Lisbon, Portugal, 2015.
85. Wiebe J, Mihalcea R. Word sense and subjectivity. In: *Proceedings of the 21st International Conference on Computational Linguistics and the 44th Annual Meeting of the Association for Computational Linguistics*, Sydney, Australia, July 2006, 1065–1072. Association for Computational Linguistics.
86. Kim SM, Hovy E. Determining the sentiment of opinions. In: *Proceedings of the 20th International Conference on Computational Linguistics*, August 2004, 1367. Association for Computational Linguistics.
87. Esuli A, Sebastiani F. Determining the semantic orientation of terms through gloss classification. In: *Proceedings of the 14th ACM International Conference on Information and Knowledge Management*, October 2005, 617–624. ACM.
88. Rao D, Ravichandran D. Semi-supervised polarity lexicon induction. In: *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics*, Athens, Greece, March 2009, 675–682. Association for Computational Linguistics.
89. Feng S, Bose R, Choi Y. Learning general connotation of words using graph-based algorithms. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Edinburgh, Scotland, UK, July 2011, 1092–1103. Association for Computational Linguistics.
90. Feng S, Kang JS, Kuznetsova P, Choi Y. Connotation Lexicon: a dash of sentiment beneath the surface meaning. In: *ACL (1)*, 2013, 1774–1784.
91. Baccianella S, Esuli A, Sebastiani F. SentiWordNet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining. In: *LREC*, vol 10, May 2010, 2200–2204.
92. Esuli A, Sebastiani F. Pageranking wordnet synsets: an application to opinion mining. In: *ACL*, vol 7, June 2007, 442–431.
93. Su F, Markert K. Subjectivity recognition on word senses via semi-supervised mincuts. In: *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, May 2009, 1–9. Association for Computational Linguistics.
94. Hatzivassiloglou V, McKeown KR. Predicting the semantic orientation of adjectives. In: *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics and Eighth Conference of the European Chapter of the Association for Computational Linguistics*, Madrid, Spain, July 1997, 174–181. Association for Computational Linguistics.
95. Qiu G, Liu B, Bu J, Chen C. Opinion word expansion and target extraction through double propagation. *Computat Linguist* 2011, 37:9–27.
96. Liu Q, Gao Z, Liu B, Zhang Y. Automated rule selection for aspect extraction in opinion mining. In: *International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
97. Velikovich L, Blair-Goldensohn S, Hannan K, McDonald R. The viability of web-derived polarity lexicons. In: *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, Los Angeles, California, USA, June 2010, 777–785. Association for Computational Linguistics.
98. Chen L, Wang W, Nagarajan M, Wang S, Sheth AP. Extracting diverse sentiment expressions with target-dependent polarity from Twitter. In: *ICWSM*, June 2012.
99. Severyn A, Moschitti A. On the automatic learning of sentiment lexicons. In: *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL HLT 2015)*, Denver, Colorado, USA, 2015.
100. Wilson T, Wiebe J, Hoffmann P. Recognizing contextual polarity in phrase-level sentiment analysis. In: *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, Vancouver, B.C., Canada, October 2005, 347–354. Association for Computational Linguistics.
101. Tang D, Qin B, Liu T. Learning semantic representations of users and products for document level sentiment classification. In: *Proceedings of the ACL*, 2015.
102. Tang D, Qin B, Liu T, Yang Y. User modeling with neural network for review rating prediction. In: *Proceedings of the IJCAI*, Buenos Aires, Argentina, 2015.
103. Tang D, Wei F, Qin B, Dong L, Liu T, Zhou M. A joint segmentation and classification framework for sentiment analysis. In: *Proceedings of EMNLP*, Doha, Qatar, 2014.
104. Tang D, Qin B, Wei F, Dong L, Liu T, Zhou M. A joint segmentation and classification framework for sentence level sentiment classification. *IEEE Trans Audio Speech Lang Process* 2015, 23:1750–1761.
105. Dong L, Wei F, Liu S, Zhou M, Xu K. A statistical parsing framework for sentiment classification. *Comput Linguist* 2015, 41:293–336.
106. Smith NA. Linguistic structure prediction. *Synth Lect Human Lang Technol* 2011, 4:1–274.
107. Li J, Chen X, Luong MT, Jurafsky D. Visualizing and understanding neural models in NLP. 2015, arXiv preprint arXiv:1506.01066.