ISSN: 2321-7782 (Online)
e-ISJN: A4372-3114
Impact Factor: 7.327

Volume 6, Issue 6, June 2018

International Journal of Advance Research in Computer Science and Management Studies

Research Article / Survey Paper / Case Study

Available online at: www.ijarcsms.com

Product Review Using Naive Bayes

Ashwini Mahendra Raut[1]
Student
Department of Computer Engg.
D. Y. Patil COE, Ambi-410 506 – India

Prof. Dhanashree Kulkarni[2]
Professor
Department of Computer Engg.
D. Y. Patil COE, Ambi-410 506 – India

Abstract: Product reviews are valuable for upcoming buyers in helping them make decisions. To this end, different opinion mining techniques have been proposed, where judging a review sentence's orientation (e.g. positive or negative) is one of their key challenges. Recently, deep learning has emerged as an effective means for solving sentiment classification problems. A neural network intrinsically learns a useful representation automatically without human efforts. However, the success of deep learning highly relies on the availability of large-scale training data. We propose a novel deep learning framework for product review sentiment classification which employs prevalently available ratings as weak supervision signals. The framework consists of two steps: (1) learning a high level representation (an embedding space) which captures the general sentiment distribution of sentences through rating information; (2) adding a classification layer on top of the embedding layer and use labeled sentences for supervised fine-tuning. We explore two kinds of low level network structure for modeling review sentences, namely, convolutional feature extractors and long short-term memory. To evaluate the proposed framework, we construct a dataset containing 1.1M weakly labeled review sentences and 11,754 labeled review sentences from Amazon. Experimental results show the efficacy of the proposed framework and its superiority over baselines.

Keywords: Deep learning, opinion mining, sentiment classification, weak-supervision.

I. INTRODUCTION

In the eras of social connectedness and social colonization, people are becoming increasingly enthusiastic about interacting, sharing, and collaborating through online collaborative media. In recent years, this collective intelligence has spread to many different areas, with particular focus on fields related to everyday life such as commerce, tourism, education, and health, causing the size of the Social Web to expand exponentially. The distillation of knowledge from such a large amount of unstructured information, however, is an extremely difficult task, as the contents of today's Web are perfectly suitable for human consumption, but remain hardly understandable to machines. Big social data analysis grows out of this need and combines multiple disciplines such as social network analysis, multimedia management, social media analytics, trend discovery, and opinion mining. For example, studying the evolution of a social network merely as a graph is very limited as it does not take into account the information flowing between network nodes.

Similarly, processing social interaction contents between network members without taking into account connections between them is limited by the fact that information flows cannot be properly weighted. Big social data analysis, instead, aims to study large-scale Web phenomena such as social networks from a holistic point of view, i.e., by concurrently taking into account all the socio-technical aspects involved in their dynamic evolution. Hence, big social data analysis is inherently interdisciplinary and spans areas such as machine learning, graph mining, information retrieval, knowledge-based systems, linguistics, common-sense reasoning, natural language processing, and big data computing.

Social emotion prediction is of value to market analysis and to political decision. With the free and convenient communication environment of internet, people show increasing enthusiasm of online communication. Meanwhile, the internet users prefer to pro-duce and convey online information through expressing personal opinions than just obtain online information.

There are a good number of companies, large and small, that include the analysis of social data as part of their missions. Big social data analysis can be exploited for the creation and automated upkeep of review and opinion aggregation websites, in which opinionated text and videos are continuously gathered from the Web and not restricted only to product reviews, but also to wider topics such as political issues and brand perception. Big social data analysis has also a great potential as a sub-component technology for other systems. They can enhance the capabilities of customer relationship management and recommendation systems allowing, for example, to find out which features customers are particularly happy about or to exclude the recommendations items that have received very negative feedbacks. Similarly, they can be exploited for affective tutoring and affective entertainment or for troll filtering and spam detection in online social communication

II. Objective

- To analysis of the emotion based on train data.
- We construct an opinion network to detect user-generated social emotion by the structures of opinion network

III. Literature Survey

E. Cambria, B. Schuller, Y. Xia, and C. Havasi, "New Avenues in Opinion Mining and Sentiment Analysis," IEEE Intell. Syst., vol. 28, no. 2, pp. 15–21, 2013.

Many companies use opinion mining and sentiment analysis as part of their research. For instance, companies use opinion mining to create and automatically maintain review and opinion-aggregation websites. Their systems continuously gather a wide array of information from the Web, such as product reviews, brand perception, and political issues. Other systems might also use opinion mining and sentiment analysis as subcomponent technology to improve customer relationship management and recommendation systems through positive and negative customer feedback. Similarly, opinion mining and sentiment analysis might detect and exclude "flames" (overly heated or antagonistic language) in social communication and enhance antispam systems.

Companies use sentiment analysis to develop marketing strategies by assessing and predicting public attitudes toward their brand. Research and development focuses on designing automatic tools that crawl online reviews and condense the information gathered. Numerous companies already provide tools that track public viewpoints on a large scale by offering graphical summarizations of trends and opinions in the blogosphere. Developing opinion-tracking systems is commercially important.

E. Cambria, "Affective Computing and Sentiment Analysis," IEEE Intell. Syst., vol. 31, no. 2, pp. 102–107, Mar. 2016.

Affective computing and sentiment analysis, hence, are key for the advancement of AI3 and all the research fi elds that stem from it. Moreover, they fi nd applications in various scenarios and companies, large and small, that include the analysis of emotions and sentiments as part of their mission. Sentiment-mining techniques can be exploited for the creation and automated upkeep of review and opinion aggregation websites, in which opinionated text and videos are continuously gathered from the Web and not restricted to just product reviews, but also to wider topics such as political issues and brand perception. Affective computing and sentiment analysis also have great potential as a subcomponent technology for other systems.

They can enhance the capabilities of customer relationship management and recommendation systems—for example, to reveal which features customers enjoy or to exclude from the recommendations items that received negative feedback. Similarly, they can be exploited for affective tutoring and affective entertainment or for troll fi altering and spam detection in online social communication. Business intelligence is also a main factor behind corporate interest in the fi elds of affective computing and sentiment analysis. Nowadays, companies invest an increasing amount of money in marketing strategies and are constantly interested in both collecting and predicting the attitudes of the general public toward their products and brands. The design of automatic tools capable of mining sentiments over the Web in real time and creating condensed versions of these represents one of the most active research and development areas. The development of such systems, moreover, is not only important for commercial purposes but also for government intelligence applications able to monitor increases in hostile communications or model cyber-issue diffusion.

E. Cambria, N. Howard, Y. Xia, and T.-S. Chua, "Computational Intelligence for Big Social Data Analysis [Guest Editorial]," IEEE Comput. Intell. Mag., vol. 11, no. 3, pp. 8–9, Aug. 2016.

In the eras of social connectedness and social colonization, people are becoming increasingly enthusiastic about interacting, sharing, and collaborating through online collaborative media. In recent years, this collective intelligence has spread to many different areas, with particular focus on fields related to everyday life such as commerce, tourism, education, and health, causing the size of the Social Web to expand exponentially. The distillation of knowledge from such a large amount of unstructured information, however, is an extremely difficult task, as the contents of today's Web are perfectly suitable for human consumption, but remain hardly understandable to machines.

Big social data analysis grows out of this need and combines multiple disciplines such as social network analysis, multimedia management, social media analytics, trend discovery, and opinion mining. For example, studying the evolution of a social network merely as a graph is very limited as it does not take into account the information flowing between network nodes. Similarly, processing social interaction contents between network members without taking into account connections between them is limited by the fact that information flows cannot be properly weighted.

C. Strapparava and A. Valitutti, "Wordnet-Affect: An Affective Extension of Wordnet," Proc. Fourth Int'l Conf. Language Resources and Evaluation (LREC '04), 2004.

The problem of social affective text mining, which aims to discover the connections between social emotions and affective terms based on user generated emotion labels. We propose a joint emotion-topic model b augmenting latent Dirichlet allocation with an additional layer for emotion modeling. It first generates a set of latent topics from emotions, followed by generating affective terms from each topic. Experimental results on an online news collection show that the proposed model can effectively identify meaningful latent topics for each emotion. Evaluation on emotion prediction further verifies the effectiveness of the proposed model.

Recent years have witnessed a rapid growth of online users and their increasing willingness to engage in social interactions. This inspires numerous social websites, e.g., sina.com.cn and people.com.cn, to provide a new service that allows users to share their emotions after browsing news articles. Figure 1 gives an example of eight social emotions from Sina, where each bar indicates the number of users choosing the corresponding emotion.

D.M. Blei, A.Y. Ng, and M.I. Jordan, "Latent Dirichlet Allocation," J. Machine Learning Research, vol. 3, pp. 993-1022, 2003.

RECENT years have witnessed a rapid growth of online users and their increasing willingness to engage in social interactions. This inspires numerous social websites, e.g., Sina.com.cn1 and People.com.cn,2 to provide a new service that allows users to share their emotions after browsing news articles. Figs. 1a and 1b give two examples of eight social emotions, which are collected from 378 Sina users and 32 People users, respectively. The user-generated social emotions provide a new aspect for document categorization, and they cannot only help online users select related documents based on emotional preferences, but also benefit a number of other applications such as contextual music recommendation. Several studies have been carried out to automatically predict the most probable emotions for documents. Strapparava and Mihalcea claimed that all words can potentially convey affective meaning. Every word, even those apparently neutral, can evoke pleasant or painful experiences because of their semantic relation with emotional concepts or categories. However, the way how text documents affect online users' social emotions is yet to be unveiled.

Y. Rao, "Contextual Sentiment Topic Model for Adaptive Social Emotion Classification," IEEE Intell. Syst., vol. 31, no. 1, pp. 41–47, Jan. 2016.

Social emotion classification is important for numerous applications, such as public opinion measurement, corporate reputation estimation, and customer preference analysis. However, topics that evoke a certain emotion in the general public are often context-sensitive, making it difficult train a universal classifier for all collections. We propose a multi-labeled sentiment topic model, namely, the contextual sentiment topic model (CSTM) for adaptive social emotion classification. The CSTM distinguishes context-independent topics from both a background theme that characterizes non-discriminative information, and a contextual theme that characterizes context-dependent information across different collections. Experimental results demonstrated the effectiveness of our model for the adaptive classification of social emotions.

Measuring the opinions of the general public about social events, company strategies, marketing campaigns, and product preferences has raised steady interest in scientific communities. Traditional methods required questioning a large number of people about their feelings using polls, which, although accurate, may be expensive or time-consuming. With the growth of the social web and the availability of reviews, statistical polling data, and other user-generated content, an increasingly large volume of significant information concerning user opinions is being stored online. The development of interactive services also booms the communication of emotions through news websites, blogs, microblogs/tweets, and so forth. For instance, many news websites now provide a service that allows users to convey their emotions after browsing news articles, with articles incorporating the emotional responses shared by its readers, in which they express by voting for a set of predefined emotion labels. The aggregation of such emotional responses from readers is known as social emotions.

S. Aral and D. Walker, "Identifying social influence in networks using randomized experiments," IEEE Intell. Syst., vol. 26, no. 5, pp. 91–96, 2011.

One important goal of applying statistical inference techniques to large networked datasets is to understand how behavioral contagions spread in human social networks. More precisely, understanding how people influence or are influenced by their peers can help us understand the ebb and fl ow of market trends, product adoption and diffusion, the spread of health behaviors such as smoking and exercise, the productivity of information workers, and whether particular individuals in a social network have a disproportionate amount of influence on the system.

However, if we are to truly understand how social interaction and peer infl uence shape behavioral dynamics in large networked populations, we must be able to separate correlation from causation. By now, there is abundant empirical evidence that human behaviors tend to cluster in network space and in time. Recent studies have shown behavioral clustering for trends in obesity, smoking, product adoption, happiness, economic development, and more. Still, several alternative explanations besides peer infl uence and social contagion could also explain these patterns. For example, people tend to make friends with those who are like themselves, a social process called homophily. As a result, preferences and behaviors cluster in networks because we tend to choose friends who like the same things and behave in the same ways that we do. Peers are also more likely to be exposed to the same external stimuli. We tend to make friends with people we work with or who live nearby. As a result, our exposure to changes in health benefit plans at work or new restaurants opening in our neighborhoods is correlated with that of our friends. Our correlated exposure to such external stimuli can in turn drive patterns of correlation in our preferences and behaviors over time.

N. Majumder and I. P. Nacional, "Deep Learning Based Document Modeling for Personality Detection from Text," IEEE Intell. Syst., vol. 32, no. 2, pp. 74–79, 2017.

Personality detection can also be exploited for word polarity disambiguation in sentiment lexicons,2 as the same concept can convey different polarity to different types of people. In mental health diagnosis, certain diagnoses correlate with certain personality traits. In forensics, knowing personality traits helps reduce the circle of suspects. In human resources management, personality traits affect one's suitability for certain jobs.

Texts often reflect various aspects of the author'spersonality. In this article, we present a method to extract personality traits from stream of-consciousness essays using a convolutional neural network (CNN). We trained five different networks; all with the same architecture, for the five personality traits (see the "Previous Work in Personality Detection" sidebar for more information). Each network was a binary classifier that predicted the corresponding trait to be positive or negative.

IV. PROBLEM DEFINITION

One fundamental problem in sentiment analysis is categorization of sentiment polarity. Given a piece of written text, the problem is to categorize the text into one specific sentiment polarity, positive or negative. Based on the scope of the text, there are three levels of sentiment polarity categorization, namely the document level, the sentence level, and the entity and aspect level. From a user's perspective, people can post their own content through various social media, such as forums, micro-blogs, or online social networking sites.

Sentiment analysis which is also known as opinion mining, studies people's sentiments towards certain entities. Internet is a resourceful place with respect to sentiment information.

Motivation

Noting the varied benefits of n-grams, we developed an algorithm that attempts to identify arbitrary-length substrings that provide "optimal" classification. We are faced with a tradeoff: as substrings become longer and generally more discriminatory, their frequency decreases, so there is less evidence for considering them relevant. Simply building a tree of substrings up to a cutoff length, treating each sufficiently-frequent substring as relevant, yields no better than 88.5 percent accuracy on the first test using both our baseline and Naive Bayes.

A more complicated approach, which compares each node on the tree to its children to see if its evidence-differentiation tradeoff is better than its child, sometimes outperforms n-grams. We experimented with several criteria for choosing not to pursue a subtree any further, including its information gain relative to the complete set, the difference between the scores that would be given to it and its parent, and its document frequency. We settled on a threshold for information gain relative to a node's parent. A second issue was how these features would be then assigned scores. Ways of matching testing data to the scored features are discussed later.

Goals

Main goal is to analysis the product review based on available dataset and generate result in terms of negative or positive. Also focusing on generate the graph of result in terms of pi-chart or bar-chart.

V. PROPOSED SYSTEM

Despite the promising performance of deep learning on sentiment classification, no previous work tried to leverage the prevalently available ratings for training deep models. In this work, we propose a novel deep learning framework for review sentence sentiment classification. The framework treats review ratings as weak labels to train deep neural networks. For example, with 5-stars scale we can deem ratings above/below 3-stars as positive/ negative weak labels respectively. The framework generally consists of two steps. In the first step, rather than predicting sentiment labels directly, we try to learn an embedding space (a high level layer in the neural network) which reflects the general sentiment distribution of sentences, from a large number of weakly labeled sentences. That is, we force sentences with the same weak labels to be near each other, while sentences with different weak labels are kept away from one another.

To reduce the impact of sentences with rating-inconsistent orientation (hereafter called wrong-labeled sentences), we propose to penalize the relative distances among sentences in the embedding space through a ranking loss. In the second step, a classification layer is added on top of the embedding layer, and we use labeled sentences to fine-tune the deep network. The framework is dubbed Weakly-supervised Deep Embedding.

Steps are as below:

1. We will be going to use Naive Bayes algorithm to sentiment analysis or Deep Embedding.
2. We proposed a novel deep learning framework named Weakly-supervised Deep Embedding for review sentence sentiment classification.
3. Construct a huge list of all occurring words per class.
4. Main aim is to analysis the product review based on available dataset and generate result in terms of negative or positive.
5. Calculate the relative occurrence of each word in this huge list.

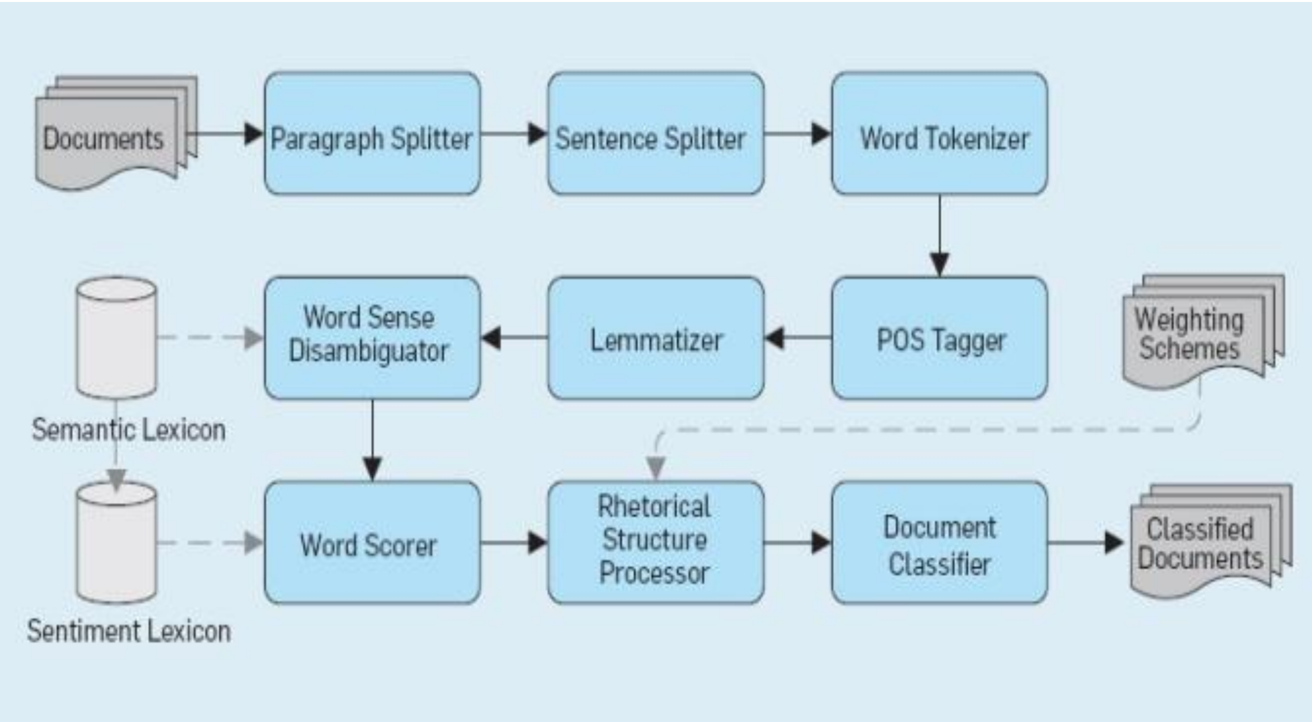

Fig 1: System Architecture

VI. Mathematical Model

Let S be the system that describes the tweet extraction, Preprocessing, Sentiment labeling, Sentiment Analysis-

S= {Tw, Pt, Sl}

Tw =Tweets extracted from Twitter.

Sl={Pv, Nv}

Pv= {P1, P2,…,Pn}= Positive Class

Nv={ N1,N2,…,Nn }= Negative Class

Where,

S= Sentimental analysis system.

Pt =Pre-processing of Tweets (Slang word translation, Non-English word removal, PoS tagging, URL and Stop word removal).

Sl=Sentiment Labeling using Sent Strength and Twitter Sentiment sentiment analysis tools (SVM to give more efficient and accurate results).

P1,p2..Pn positive tweets collection class

N1, N2...Nn Negative tweets collection class

The Naive Bayesian classifier is based on Bayes' theorem with independence assumptions between predictors. A Naive Bayesian model is easy to build, with no complicated iterative parameter estimation which makes it particularly useful for very large datasets. Despite its simplicity, the Naive Bayesian classifier often does surprisingly well and is widely used because it often outperforms more sophisticated classification methods.

Algorithm

The Nave Bayes classifier is a simple probabilistic model which relies on the assumption of feature independent in order to classify input data. Despite its simplicity, the algorithm is commonly used for text classification in many opinion mining applications. Much of it popularity is a result of its extremely simple implementation, low computational cost and it relatively high accuracy. The algorithm assumes that each feature is independent of the absence or presence of any other feature in the input data, because of this assumption it is known as nave. In reality words in a sentence are strongly related, their positions and presence in a sentence have a major impact on the overall meaning and sentiment in that sentence.

The Naive Bayesian classifier is based on Bayes theorem with independence assumptions between predictors. A Naive Bayesian model is easy to build, with no complicated iterative parameter estimation which makes it particularly useful for very large datasets. Despite its simplicity, the Naive Bayesian classifier often does surprisingly well and is widely used because it often outperforms more sophisticated classification methods. Social emotion prediction is of value to market analysis and to political decision.

Related Work

- **Part of Speech tagging**

One of the more powerful aspects of the NLTK module is the Part of Speech tagging that it can do for you. This means labelling words in a sentence as nouns, adjectives, verbs...etc. Even more impressive, it also labels by tense, and more.

- **Porter Stemmer**

The idea of stemming is a sort of normalizing method. Many variations of words carry the same meaning, other than when tense is involved.

- **Naïve Bayes for opinion mining**

The Naïve Bayes classifier is a simple probabilistic model which relies on the assumption of feature independent in order to classify input data. Despite its simplicity, the algorithm is commonly used for text classification in many opinion mining applications. Much of it popularity is a result of its extremely simple implementation, low computational cost and it relatively high accuracy.

The algorithm assumes that each feature is independent of the absence or presence of any other feature in the input data, because of this assumption it is known as ‘naïve’. In reality words in a sentence are strongly related, their positions and presence in a sentence have a major impact on the overall meaning and sentiment in that sentence. Despite this naïve assumption the classifier can produce high classification accuracy when used with quality training and in specific domains. A recent study addressed this assumption and presented strong evidence of how the algorithm could be so effective while relying on this assumption.

From the literature examined surrounding the Naïve Bayes it can be seen that despite its simplicity the algorithm has the ability to produce high classification accuracy on similar dataset to the one used in this project. It has many advantages over some of the more sophisticated algorithms, the major one being how simple it is to understand and implement, the low processing cost of the algorithm can be attributed to this simplicity.

VII. Result Analysis

Result is categorized into three different classes -Positive review, Negative review and neutral review. System will have generated p0-chart as well as bar chart for non-technical and technical users respectively.In this result first result is based on english language sentiment analysis and second result is based on hindhi language sentiment analysis.

a) Pi-chart

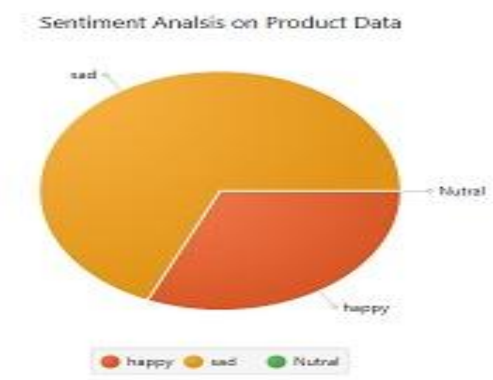

b)

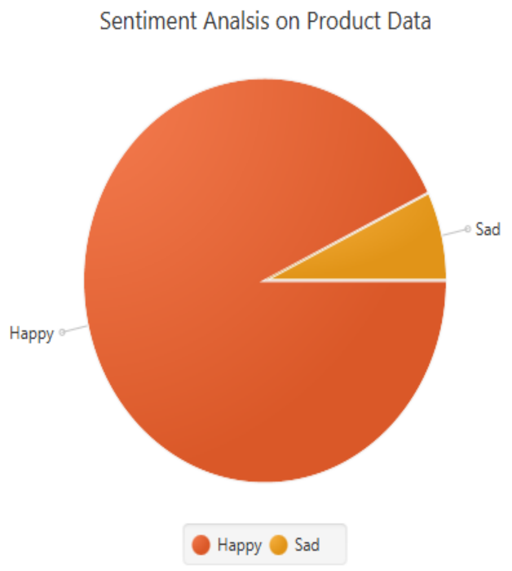

Acknowledgement

This is a great pleasure & immense satisfaction to express my deepest sense of gratitude & thanks to everyone who has directly or indirectly helped me in completing my Dissertation work successfully. Today on completion of this project dissertation, the persons I need to thank the most who have helped me throughout the making of this project dissertation and without whose help the project would not have seen the light of the day.

Primarily, I submit my gratitude and sincere thanks to my project guide and, Head, Department of Computer Engineering Prof. Dhanashree Kulkarni and ME Coordinator Prof. A. K.Bongale, Dr. D.Y. Patil College of Engineering, Ambi, Pune who guided & encouraged me in completing the Dissertation work in scheduled time. I would like to thank our Principal Prof. Dr.S.D.Shirbahadurkar, for allowing us to pursue my project in this institute.

No words are sufficient to express my gratitude to my family for their unwavering encouragement. I also thank all friends for being a constant source of my support.

References

1. C. M. Bishop. Pattern recognition and machine learning. springer, 2006. L. Chen, J. Martineau, D. Cheng, and A. Sheth. Clustering for simultaneous extraction of aspects and features from reviews. In NAACL-HLT, pages 789–799, 2016.
2. K. Dave, S. Lawrence, and D. M. Pennock. Mining the peanut gallery: Opinion extraction and semantic classification of product reviews. In WWW, pages 519–528, 2003.
3. S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman. Indexing by latent semantic analysis. Journal of the American society for information science, 41(6):391, 1990.
4. X. Ding, B. Liu, and P. S. Yu. A holistic lexicon-based approach to opinion mining. In WSDM, pages 231–240, 2008
5. J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. JMLR, 12:2121–2159, 2011.
6. R. Feldman. Techniques and applications for sentiment analysis. Communications of the ACM, 56(4):82–89, 2013.
7. X. Glorot, A. Bordes, and Y. Bengio. Domain adaptation for large-scale sentiment classification: A deep learning approach. In ICML, pages 513–520, 2011.
8. M. Hu and B. Liu. Mining and summarizing customer reviews. In SIGKDD, pages 168–177, 2004.
9. N. Kalchbrenner, E. Grefenstette, and P. Blunsom. A convolutional neuralnetwork for modelling sentences. In ACL, 2014.
10. Kim. Convolutional neural networks for sentence classification. In EMNLP, pages 1746–1751, 2014.
11. L. Qu, R. Gemulla, and G. Weikum. A weakly supervised model for sentence-level semantic orientation analysis with multiple experts. In EMNLP-CoNLL, pages 149–159, 2012.
12. Y. Shen, X. He, J. Gao, L. Deng, and G. Mesnil. A latent semantic model with convolutional-pooling structure for information retrieval. In CIKM, pages 101–110, 2014.
13. R. Socher, B. Huval, C. D. Manning, and A. Y. Ng. Semantic compositionality through recursive matrix-vector spaces. In EMNLP-CoNLL, pages 1201–1211, 2012.
14. Satoshi Morinaga, Kenji Yamanishi, Kenji Tateishi, and Toshikazu Fukushima. Mining product reputions on the web. In KDD 2002.
15. Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. Thumbs up? Sentiment classification using machine learning techniques. In Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 79–86.
16. Fernando C. N. Pereira, Naftali Tishby, and Lillian Lee. Distributional clustering of English words. In Meeting of the Association for Computational Linguistics, pages 183–190, 1993.
17. M.F. Porter. An algorithm for suffix stripping. In Program,volume 14, pages 130–137, 1980.
18. Deepak Ravichandran and Eduard Hovy. Learning surface text patterns for a question answering system. In ACL Conference, 2002.
19. Ellen Riloff. Automatically generating extraction patterns from untagged text. In Proceedings of AAAI/IAAI, Vol. 2, pages 1044–1049, 1996.
20. P. Subasic and A. Huettner. Affect analysis of text using fuzzy semantic typing. IEEE-FS, 9:483–496, Aug. 2001.
21. Loren Terveen, Will Hill, Brian Amento, David McDonald, and Josh Creter. PHOAKS: A system for sharing recommendations. Communications of the ACM, 40(3):59–62, 1997.
22. Richard M. Tong. An operational system for detecting and tracking opinions in on-line discussion. In SIGIR Workshop on Operational Text Classifiation, 2001.
23. P.D. Turney and M.L. Littman. Unsupervised learning of semantic orientation from a hundred-billion-word corpus. Technical Report ERB-1094, National Research Council Canada, Institute for Information Technology, 2002.
24. Janyce Wiebe. Learning subjective adjectives from corpora. In AAAI/IAAI, pages 735–740, 2000.
25. Janyce Wiebe, Rebecca Bruce, Matthew Bell, Melanie Martin, and Theresa Wilson. A corpus study of evaluative and speculative language. In Proceedings of the 2nd ACL SIGdial Workshop on Discourse and Dialogue, 2001.
26. Janyce Wiebe, Theresa Wilson, and Matthew Bell. Identifying collocations for recognizing opinions. In Proceedings of ACL/EACL 2001 Workshop on Collocation.

27. Mikio Yamamoto and Kenneth Church. Using suffix arrays to compute term frequency and document frequency for all substrings in a corpus. In Proceedings of the 6th Workshop on Very Large Corpora.
28. Manning, C. and Schütze, H. 1999. Foundations of Statistical Natural Language Processing, MIT Press. Cambridge, MA: May 1999.
29. Miller, G., Beckwith, R, Fellbaum, C., Gross, D., and Miller, K. 1990. Introduction to WordNet: An on-line lexical database. International Journal of Lexicography (special issue), 3(4):235-312.
30. Morinaga, S., Ya Yamanishi, K., Tateishi, K, and Fukushima, T. 2002. Mining Product Reputations on the Web. KDD'02.
31. NLProcessor – Text Analysis Toolkit. 2000. http://www.infogistics.com/textanalysis.html
32. Paice, C. D. 1990. Constructing Literature Abstracts by Computer: Techniques and Prospects. Information Processing and Management 26:171-186.
33. Pang, B., Lee, L., and Vaithyanathan, S., 2002. Thumbs up? Sentiment Classification Using Machine Learning Techniques. In Proc. of EMNLP 2002
34. Reimer, U. and Hahn, U. 1997. A Formal Model of Text Summarization based on Condensation Operators of a Terminological Logic. In Proceedings of ACL'97 Workshop on Intelligent, Scalable Text Summarization.
35. Sack, W., 1994. On the Computation of Point of View. AAAI'94, Student abstract.
36. Salton, G. Singhal, A. Buckley, C. and Mitra, M. 1996.
37. Automatic Text Decomposition using Text Segments and Text Themes. ACM Conference on Hypertext.
38. Sparck J. 1993a. Discourse Modeling for Automatic Text Summarizing. Technical Report 290, University of Cambridge Computer Laboratory.
39. Sparck J. 1993b. What might be in a summary? Information Retrieval 93: 9-26.
40. Tait, J. 1983. Automatic Summarizing of English Texts. Ph.D. Dissertation, University of Cambridge.
41. Tetreault, J. 1999. Analysis of Syntax-Based Pronoun Resolution Methods. ACL'99.