Annotated Bibliography

IMO: Sentiment Analysis Visualization on Mobile App Reviews

Table of Contents

Cluster 1 - Similar projects in different scopes

Authors： Bin Fu, Jialiu Lin, Lei Li, Christos Faloutsos, Jason Hong, Norman Sadeh

Title： Why People Hate Your App — Making Sense of User Feedback in a Mobile App Store

Venue: Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining

Year: 2013

Number of Citations: 250

Aims: User review is a crucial component of open mobile app markets such as the Google Play Store. The main aim of the paper is to create data analytics tools to automatically summarize millions of user reviews and make sense out of them beyond simple summaries such as histograms of user ratings. Besides the tools, they discussed in depth about how the techniques presented herein can be deployed to help a mobile app markets operator such as Google as well as individual app developers and end-users.

Conclusion: WisCom was a proposed and developed system that can analyze tens of millions of user ratings and comments in mobile app markets at three different levels of detail. This is an analytic tool that can provide insights into user reviews. The system was able to detect the inconsistencies between user comments and ratings, identify the major reasons why users dislike an app and learn how users' complaints changed over time.

Inspirations: Inspired by the methodology presented in the paper, we decided to conduct three levels of data exploration in our visualization design. Moreover, the scope of the project is limited to the comparison of games and other application in Android app on Google Play. We extend our analysis to the scope of both markets, discovering high-level knowledge and global trends in all sources.

Authors：Xiaodong Gu, Sunghun Kim
Title： “What parts of your apps are loved by users?”

Venue: 30th IEEE/ACM International Conference on Automated Software Engineering

Year: 2015
Number of Citations: 48

Aims: One of the most important questions software developers ask is “what parts of the software are used/loved by users.” User reviews provide an effective channel to address this question. However, most existing review summarization tools treat reviews as bags-of-words (i.e., mixed review categories) and are limited to extract software aspects and user preferences. They presented a novel review summarization framework, SURMiner. Instead of a bags-of-words assumption, it classifies reviews into five categories and extracts aspects in sentences which include evaluation of aspect using a pattern-based parser. Then, SURMiner visualizes the summaries using two interactive diagrams.

Conclusion: The final evaluation of 17 popular apps shows that SUR-Miner summarizes more accurate and clearer aspects than state-of-theart techniques significantly greater than that of ReviewSpotlight and Guzman's’ method. Feedback from developers shows that more developers agreed with the usefulness of the summaries from SUR-Miner.

Inspirations: This paper gives us an overview of the sentiment analysis on the app market in general. After exploring the resources, we decided to go with the old-way bags of words assumptions described in the paper and to present the results to the user in a more critical way. We are cumulative of the sophisticated words analysis and use a more intuitive way, which judges the reviews attitude based on the scores on it.

Cluster 2 - General Design Improvements

Authors： Katharina Reinecke, Tom Yeh, Luke Miratrix, Rahmatri Mardiko, Yuechen Zhao, Jenny Liu, Krzysztof Z. Gajos

Title： Predicting Users' First Impressions of Website Aesthetics With a Quantification of Perceived Visual Complexity and Colorfulness

Venue: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems

Year: 2013

Number of 182

Aims: This first impression is influential enough to later affect their opinions of a site's usability and trustworthiness. In this paper, a means is demonstrated to predict the initial impression of aesthetics based on perceptual models of a website's colorfulness and visual complexity. In an online study, the collected ratings of colorfulness, visual complexity, and visual appeal of a set of 450 websites from 548 volunteers.

Conclusion: Based on these data, developed computational models that accurately measure the perceived visual complexity and colorfulness of website screenshots. In combination with demographic variables such as a user's education level and age, these models explain approximately half of the variance in the ratings of aesthetic appeal given after viewing a website for 500ms only.

Inspirations: Based on the resulted presented by the paper, we decided to design the visualization in less than 5 major color schemes and eliminate the color legends when showing the different categories of the app. Which allow the user to focus on the vital sentiment analysis part.

Cluster 3 - Sentiment Analysis Methodology

Authors： Bo Pang, Lillian Lee

Title： Opinion mining and sentiment analysis

Venue: Foundations and Trends® in Information Retrieval

Year: 2008

Number of Citations: 8118

Aims: With the growing availability and popularity of opinion-rich resources such as online review sites and personal blogs, new opportunities and challenges arise as people now can, and do, actively use information technologies to seek out and understand the opinions of others. The focus is on methods that seek to address the new challenges raised by applications, as compared to those that are already present in the more traditional fact-based analysis.

Conclusion: This survey covers techniques and approaches that promise to directly enable opinion-oriented information seeking systems. Also, a discussion of available resources, benchmark datasets, and evaluation campaigns was provided in the paper.

Inspirations: We included material on the summarization of evaluative text presented in the paper. We use this as a guideline to conduct our own sentiment analysis dataset. Understanding how the data we selected is generated is helping us to understand the underlying weakness of the final deliverables.

Authors： Emitza Guzman, Walid Maalej

Title： How Do Users Like This Feature? A Fine-Grained Sentiment Analysis of App Reviews

Venue: IEEE 22nd International Requirements Engineering Conference (RE)

Year: 2014

Number of Citations: 269

Aims: For many apps, the amount of reviews is too large to be processed manually and their quality varies largely. The star ratings are given to the whole app and developers do not have a mean to analyze the feedback for the single features. This paper provides natural language processing techniques to identify fine-grained app features in the reviews and then extract the user sentiments about the identified features and give them a general score across all reviews. Finally, they use topic modeling techniques to group fine-grained features into more meaningful high-level features.

Conclusion: The proposed automated approach helps developers filter, aggregate, and analyze user reviews. The approach was evaluated with 7 apps from the Apple App Store and Google Play Store and compared its results with a manually, peer-conducted analysis of the reviews.

Inspirations: Based on the analysis in the paper, we extended our targeted user base on app developers and add a new task to systematically analyze user opinions about single features and filter irrelevant reviews.

Authors： Emitza Guzman, Walid Maalej

Title： How Do Users Like This Feature? A Fine-Grained Sentiment Analysis of App Reviews

Venue: IEEE 22nd International Requirements Engineering Conference (RE)

Year: 2014

Number of Citations: 269

Aims: For many apps, the amount of reviews is too large to be processed manually and their quality varies largely. The star ratings are given to the whole app and developers do not have a mean to analyze the feedback for the single features. This paper provides natural language processing techniques to identify fine-grained app features in the reviews and then extract the user sentiments about the identified features and give them a general score across all reviews. Finally, they use topic modeling techniques to group fine-grained features into more meaningful high-level features.

Conclusion: The proposed automated approach helps developers filter, aggregate, and analyze user reviews. They evaluated the approach with 7 apps from the Apple App Store and Google Play Store and compared its results with a manually, peer-conducted analysis of the reviews.

Inspirations: Based on the analysis in the paper, we extended our targeted user base on app developers and add a new task to systematically analyze user opinions about single features and filter irrelevant reviews.

Authors： Sebastiano Panichella, Andrea Di Sorboy, Emitza Guzmanz, Corrado A.Visaggioy, Gerardo Canforay, and Harald C. Gall

Title： How Can I Improve My App? Classifying User Reviews for Software Maintenance and Evolution

Venue : IEEE International Conference on Software Maintenance and Evolution (ICSME)

Year : 2015

Number of Citations : 149

Aims: Previous research showed that users feedback contains usage scenarios, bug reports, and feature requests, that can help app developers to accomplish software maintenance and evolution tasks. In this paper, they presented a taxonomy to classify app reviews into categories relevant to software maintenance and evolution, as well as an approach that merges three techniques: (1) Natural Language Processing, (2) Text Analysis and (3) Sentiment Analysis to automatically classify app reviews into the proposed categories.

Conclusion: The results showed that the combined use of these techniques allows achieving better results (precision of 75% and a recall of 74%) than results obtained using each technique individually (precision of 70% and a recall of 67%).

Inspirations: The paper showed the other use case of the app sentiment analysis, which is to help the developers maintenance on the features and set the priority on the tasks.

Authors： Mike Thelwall, Kevan Buckley, Georgios Paltoglou

Title： Sentiment Strength Detection for the Social Web

Venue: Journal of the American Society for Information Science and Technology banner

Year: 2011

Number of Citations: 729

Aims: Sentiment analysis is concerned with the automatic extraction of sentiment-related information from text. Most sentiment analysis algorithms are not ideally suited to this task because they exploit indirect indicators of sentiment that can reflect genre or topic instead. Hence, such algorithms used to process social web texts can identify spurious sentiment patterns caused by topics rather than affective phenomena. This article assesses an improved version of the algorithm SentiStrength for sentiment strength detection across the social web that primarily uses direct indications of sentiment.

Conclusion: The results from six diverse social web data sets (MySpace, Twitter, YouTube, Digg, Runners World, BBC Forums) indicate that SentiStrength 2 is successful in the sense of performing better than a baseline approach for all data sets in both supervised and unsupervised cases. Overall, the results suggest that SentiStrength is robust enough to be applied to a wide variety of different social web contexts.

Inspirations: This paper lay emphasizes on finding patterns not define a specific direction of one review. We decide to embrace this method to show a high-level generalization on the performance of the apps.

Authors： Changbo Wang, Zhao Xiao, Yuhua Liu, Yanru Xu, Aoying Zhou, and Kang Zhang

Title： SentiView: Sentiment Analysis and Visualization for Internet Popular Topics

Venue : IEEE TRANSACTIONS ON HUMAN-MACHINE SYSTEMS

Year: 2013

Number of Citations: *69*

Aims: There would be value to several domains in discovering and visualizing sentiments in online posts. This paper presents SentiView, an interactive visualization system that aims to analyze public sentiments for popular topics on the Internet. SentiView combines uncertainty modeling and model-driven adjustment.

Conclusion: Using a time-varying helix together with an attribute astrolabe to represent sentiments, The relationships of interest among different participants are presented in a relationship map and a new evolution model that is based on cellular automata, it is able to compare the time-varying features for sentiment-driven forums on both simulated and real data. Adaptable for different social networking platforms, such as Twitter, blog, and forum, the methods demonstrate the effectiveness of SentiView in analyzing and visualizing public sentiments on the Web.

Inspirations: By searching and correlating frequent words in text data, it mines and models the changes of the sentiment on public topics. It can visualize the changes of multiple attributes and relationships among demographics of interest and the sentiments of participants on popular topics.

Authors： Daniela Oelke , Ming Hao, Christian Rohrdantz, Daniel A. Keim, Umeshwar Dayal, Lars-Erik Haug, Halldór Janetzko

Title： Visual Opinion Analysis of Customer Feedback Data

Venue: IEEE Symposium on Visual Analytics Science and Technology

Year : 2009

Number of Citations : 110

Aims: This paper introduced several techniques to interactively analyze customer comments and ratings to determine the positive and negative opinions expressed by the customers. First, a new discrimination-based technique to automatically extract the terms that are the subject of the positive or negative opinion (such as price or customer service) and that are frequently commented on. Second, a Reverse-Distance-Weighting method to map the attributes to the related positive and negative opinions in the text. Third, the resulting high-dimensional feature vectors are visualized in a new summary representation that provides a quick overview.

Conclusion: Special thumbnails are used to provide insight into the composition of the clusters and their relationship. In addition, an interactive circular correlation map is provided to allow analysts to detect the relationships of the comments to other important attributes and the scores. These techniques were applied to customer comments from real-world online stores and product reviews from web sites to identify the strength and problems of different products and services and show the potential of techniques.

Inspirations: Given the way that all the sentiment data organized, we found a great pattern for the app store reviews. That is each category either contribute negatively or positively to the final scores of the app. We also use the same interactions in the paper to help the user tease out the insights.

Authors : Ming Hao, Christian Rohrdantz*, Halldór Janetzko*, Umeshwar Dayal
Daniel A. Keim*, Lars-Erik Haug, Mei-Chun Hsu

Title : Visual Sentiment Analysis on Twitter Data Streams

Venue : IEEE Conference on Visual Analytics Science and Technology (VAST)

Year : 2011

Number of Citations : 64

Aims: A large number of tweets include opinions about products and services. However, with Twitter being a relatively new phenomenon, these tweets are underutilized as a source for evaluating customer sentiment. To explore high-volume twitter data, three novel visual sentiment analysis techniques were introduced: (1) topic-based sentiment analysis that extracts, maps, and measures customer opinions; (2) stream analysis that identifies interesting tweets based on their density, negativity, and influence characteristics; and (3) pixel cell-based sentiment calendars and high density geo maps that visualize large volumes of data in a single view.

Conclusion: The techniques described above are applied to a variety of twitter data, (e.g., movies, amusement parks, and hotels) to show their distribution and patterns, and to identify influential opinions.

Inspirations: More work in the field not limited to app reviews give us more through comprehension on how to visualize the data in a more cohesive and effective way.

Authors： Michelle L. Gregory, Paul Whitney, Richard Carter, Elizabeth Hetzler, Alan Turner

Title: User-directed Sentiment Analysis: Visualizing the Affective Content of Documents

Venue: Proceedings of the Workshop on Sentiment and Subjectivity in Text

Year: 2006

Number of Citations: 81

Aims: Recent advances in text analysis have led to finer-grained semantic analysis, including automatic sentiment analysis— the task of measuring documents, or chunks of text, based on emotive categories, such as positive or negative. However, considerably less progress has been made on efficient ways of exploring these measurements. This paper discusses approaches for visualizing the affective content of documents and describes an interactive capability for exploring emotion in a large document collection.

Conclusion: The visualization tool includes a document viewer so that any selection of documents can be reviewed. In this paper, they have described the added functionality of exploring effect as one of the possible dimensions. As an exploratory system, it is difficult to define the appropriate evaluation metric. Because the goal of the system is not to discretely bin the documents into affect categories, traditional metrics such as precision are not applicable. However, to get a sense of the coverage of lexicon, comparisons on measurements to the hand annotations provided for the customer review dataset.

Inspirations: Besides the inspiration on the circular network diagram, this paper also shows a specific way to evaluate the expressiveness of the project.

Authors： Damian Borth, Tao Chen, Rong-Rong Ji, and Shih-Fu Chang

Title： SentiBank: large-scale ontology and classifiers for detecting sentiment and emotions in visual content

Venue: Proceedings of the 21st ACM international conference on Multimedia

Year: 2013

Number of Citations: 97

Aims: A picture is worth one thousand words, but what words should be used to describe the sentiment and emotions conveyed in the increasingly popular social multimedia? This paper includes a novel system which combines sound structures from psychology and the folksonomy extracted from social multimedia to develop a large visual sentiment ontology consisting of 1,200 concepts and associated classifiers called SentiBank. Each concept, defined as an Adjective Noun Pair (ANP), is made of an adjective strongly indicating emotions and a noun corresponding to objects or scenes that have a reasonable prospect of automatic detection.

Conclusion: Such large-scale visual classifiers offer a powerful midlevel semantic representation enabling high-level sentiment analysis of social multimedia. The demonstrated novel applications made possible by SentiBank including live sentiment prediction of social media and visualization of visual content in a rich intuitive semantic space.

Inspirations: This paper gives us the proofs on using visual marks. The size and the color should be conveyed different values of the attributes and be connected in the same way. The 24 emotions conveyed in the paper is interesting also.

Authors： Tao Chen, Damian Borth, Trevor Darrell, Shih-Fu Chang

Title： DeepSentiBank: Visual Sentiment Concept Classification with Deep Convolutional Neural Networks

Venue: Proceedings of the 21st ACM international conference on Multimedia

Year: 2014

Number of Citations: 142

Aims: This paper introduces a visual sentiment concept classification method based on deep convolutional neural networks (CNN). The visual sentiment concepts are adjective-noun pairs (ANPs) automatically discovered from the tags of web photos, and can be utilized as effective statistical cues for detecting emotions depicted in the images.

Conclusion: To deal with the biased training data which only contains images with strong sentiment and to prevent overfitting, they initialized the model with the model weights trained from ImageNet. Performance evaluation shows the newly trained deep CNN's model SentiBank 2.0 (or called DeepSentiBank) is significantly improved in both annotation accuracy and retrieval performance, compared to its predecessors which mainly use binary SVM classification models.

Inspirations: This is the following paper published by the team on SentiBank. More related works recently also show a more comprehensive way to present the information related to sentiment analysis.

Authors： Erik Cambria

Title： Affective Computing and Sentiment Analysis

Venue: 2016 IEEE Intelligent Systems

Year: 2016

Number of Citations: 314

Aims: Understanding emotions is an important aspect of personal development and growth, and as such, it is a key tile for the emulation of human intelligence. Besides being important for the advancement of AI, emotion processing is also important for the closely related task of polarity detection. The opportunity to automatically capture the general public's sentiments about social events, political movements, marketing campaigns, and product preferences has raised interest in both the scientific community, for the exciting open challenges, and the business world, for the remarkable fallouts in marketing and financial market prediction.

Conclusion: This has led to the emerging fields of affective computing and sentiment analysis, which leverage human-computer interaction, information retrieval, and multimodal signal processing for distilling people's sentiments from the ever-growing amount of online social data.

Inspirations: The paper shows that the better NLP performance by using the bag of narrative, which is the dataset segmentation we will use. This paper shore up our project.