REVIEWMINER: AN UNSUPERVISED METHOD OF ASPECT EXTRACTION AND ASPECT RATING FROM PRODUCT REVIEWS

A Thesis
Submitted to the Graduate Faculty
of the
North Dakota State University
of Agriculture and Applied Science

By

Anubrata Dutta

In Partial Fulfillment of the Requirements
for the Degree of
MASTER OF SCIENCE

Major Department:
Computer Science

May 2013

Fargo, North Dakota

North Dakota State University
Graduate School

Title

ReviewMiner: An Unsupervised Method of Aspect Extraction and Aspect Rating from Product Reviews

By

Anubrata Dutta

The Supervisory Committee certifies that this ***disquisition*** complies with North Dakota State University's regulations and meets the accepted standards for the degree of

MASTER OF SCIENCE

SUPERVISORY COMMITTEE:

Dr. Wei Jin
Advisor

Dr. Kendall Nygard

Dr. Simone Ludwig

Dr. Richard Rathge

Approved:

11/15/2013
Date

Dr. Kenneth Magel
Department Chair

ABSTRACT

One major piece of information available on the web is reviews about various products that are written by users. Some commercial websites provide additional information such as ratings about the products along with reviews. However, in the opinion mining research field, most existing methods have ignored this additional valuable information, thus influencing the accuracy of the mining results and the interpretation of various aspects related to the products.

In this thesis, we consider the reviews obtained from epinions.com related to cameras, and we propose *ReviewMiner*, an unsupervised method of automatically identifying useful aspects of a product and estimating the corresponding ratings for each aspect from the review texts. The method explores various linguistic patterns to extract potential aspects and context-dependent opinion phrases and employs a series of heuristic strategies and pruning techniques. Experimental results have demonstrated the effectiveness of the proposed techniques and shown their advantages over comparative baselines.

ACKNOWLEDGMENTS

I would like to express my sincere gratitude to everyone who has supported me during the completion of this research. I would like to thank Dr. Wei Jin for guiding me for all these semesters. Her support has been a significant motivating factor, and without her contribution, this thesis would not have been completed. I would also like to thank the Computer Science department for offering Information Retrieval as a course that helped me complete different chapters for my thesis The course also motivated my research interest in the field of opinion mining. I would like to thank Stephanie Sculthorp (Administrative Secretary) of the Computer Science department for timely reminders and encouraging me to complete my thesis on time. I would like to thank my family and friends for guidance, love, support and making it possible for me to study at North Dakota State University.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

CC......................Coordinating conjunction

CD......................Cardinal number

DT......................Determiner

EX......................Existential there

FW......................Foreign word

IN......................Preposition or subordinating conjunction

JJ......................Adjective

JJR......................Adjective, comparative

JJS......................Adjective, superlative

LS......................List item marker

MD......................Modal

NN......................Noun, singular or mass

NNS......................Noun, plural

NNP......................Proper noun, singular

NNPS......................Proper noun, plural

PDT......................Predeterminer

POS......................Possessive ending

PRP......................Personal pronoun

PRP$......................Possessive pronoun

RB......................Adverb

RBR......................Adverb, comparative

RBS......................Adverb, superlative

RP......................Particle

SYM......................Symbol

TO......................to

UH......................Interjection

VB......................Verb, base form

VBD......................Verb, past tense

VBG......................Verb, gerund or present participle

VBN......................Verb, past participle

VBP......................Verb, non-3rd person singular present

VBZ......................Verb, 3rd person singular present

WDT......................Wh-determiner

WP......................Wh-pronoun

WP$......................Possessive wh-pronoun

WRB......................Wh-adverb

CHAPTER 1. INTRODUCTION

A significant motivation for our information gathering behavior has always been to figure out the public's opinion about the product for which we are interested. The public's opinion has always played an important role in the decision-making process. Before the web, major sources to address these queries were our friends, relatives and consumer reports. With the post web these queries are satisfied with various blogs, e-commerce sites, review sites and discussion forums. There has been a sudden eruption of activity in the area of sentiment analysis that is attributable to the increasing number of Internet users and new products being launched into the market. The e-commerce and review websites emphasize customer reviews and feedback to improve the customer's shopping experience. These websites provide customers with a dedicated space to read and write reviews about the products the customers have purchased and used. With this ever increasing number of Internet users, more people are willing to share their perceptions about and experience about the product they have purchased and used. With this ever-increasing number of Internet users, more people are willing to share their perceptions about and experiences with a purchased product. The most popular products are garnering hundreds of reviews. A customer who is willing to purchase a popular product may go through these reviews before making a decision about whether to buy. Reading a few reviews might give a biased opinion about the product. Among these comments, there are long reviews which are time consuming to go through and make it much more difficult for the intended consumers to obtain useful information about the product. Furthermore, huge numbers of reviews are very difficult to maintain and, instead of being beneficial, become a burden for the websites maintaining them and the customers going through them. It is highly desirable to produce a summary of the reviews or to obtain an interpretable rating for the most talked-about aspects of the product.

In recent years, researchers have studied this emerging field called opinion mining or sentiment analysis [1]. The major challenge for this field is to detect the product features that have been mentioned in the reviewers' comments and to rate the comments based on the rating guideline [2].

Both tasks are very tough to automatically simulate from the free text reviews. This thesis proposes a method *ReviewMiner* for identifying the product features and rating the features from free text customer reviews. *ReviewMiner* gives the user a list of the important features about the product and people's opinion about the features. This feature list and rating will assist the consumers to make better choices while purchasing or using a product. A product feature is an attribute of the product which is mentioned in the review and given some opinion given by the reviewer. The opinion about the feature is not restricted to positive or negative judgment; instead, it rates the aspect on the scale of 1 to 5. The input data for the *ReviewMiner* are the free text reviews, a set of known aspects or known features and the rating guideline provided by epinions.com. The result is a set of new camera features, apart from the known aspects, that are listed in the reviews along with their computed ratings.

The feature emphasized view by *ReviewMiner* provides the users with multiple benefits:

- The results obtained from *ReviewMiner* provide new customer insights from unstructured content.
- New customers can compare different cameras based on the feature ratings.
- The output can be used as input for computer software to summarize the reviews.
- The features and their ratings can be utilized in recommendation systems to provide justifications for the suggestions.

- The features along with ratings companies keep on top of issues and respond to trends that impact business.

  Websites such as Epinions.com, Amazon.com, etc. provide additional information with the free text reviews and overall rating. This information includes

- A set of predefined aspects and their ratings. These predefined aspects, also called known aspects, are requested to rate by Epinions.com when users write the reviews.
- A rating guideline is provided by epinions.com which interprets the customer-satisfaction level from a range of 1 to 5 (e.g., excellent equals 5, good equals 4, average equals 3, poor equals 2, and terrible equals 1).

This supplementary information that is supplied by the reviewing websites is overlooked and not utilized by the prevailing methods. We created a dataset by crawling epinions.com and generated an XML-based dataset that contains the free text along with predefined aspects and their ratings. The new dataset is comprised of 1126 reviews and spans 130 products (cameras).

To resolve issues in the field of opinion mining and to give the customer a clear-cut picture about the product to be purchased, the contributions of this thesis are as follows:

- Extract aspects from free-text reviews for the cameras which are of utmost importance to the new customer
- Rate the aspects of the camera on a scale from 1 to 5 based on the rating guideline provided by epinions.com

  To extract the aspects, *ReviewMiner* goes through the opinion patterns of the known aspects from free-text reviews and determines the cutoff limit for the potential aspects. For aspect rating, *ReviewMiner* assesses the aspects based on the rating guideline. The known aspects are only used for mining opinion patterns, improving the accuracy of the aspect extraction, and are not included in the final list of potential aspects. Utilizing this additional information from epinions.com increases the preciseness of opinion mining.

- Following is the diagrammatic representation of the input for *ReviewMiner* and the output obtained after processing the data. The unstructured data and the known aspect, as shown in Figure 1.1, are captured in XML format from Epinions.com and supplied as an input to *ReviewMiner* which processes the data, applies the algorithms, and produces the output in text format as shown in Table 1.

Input

Unstructured review of Canon EOS 600D/Rebel T3i Digital Camera with 18-55 mm lens and known aspects extracted from Epinions.com.

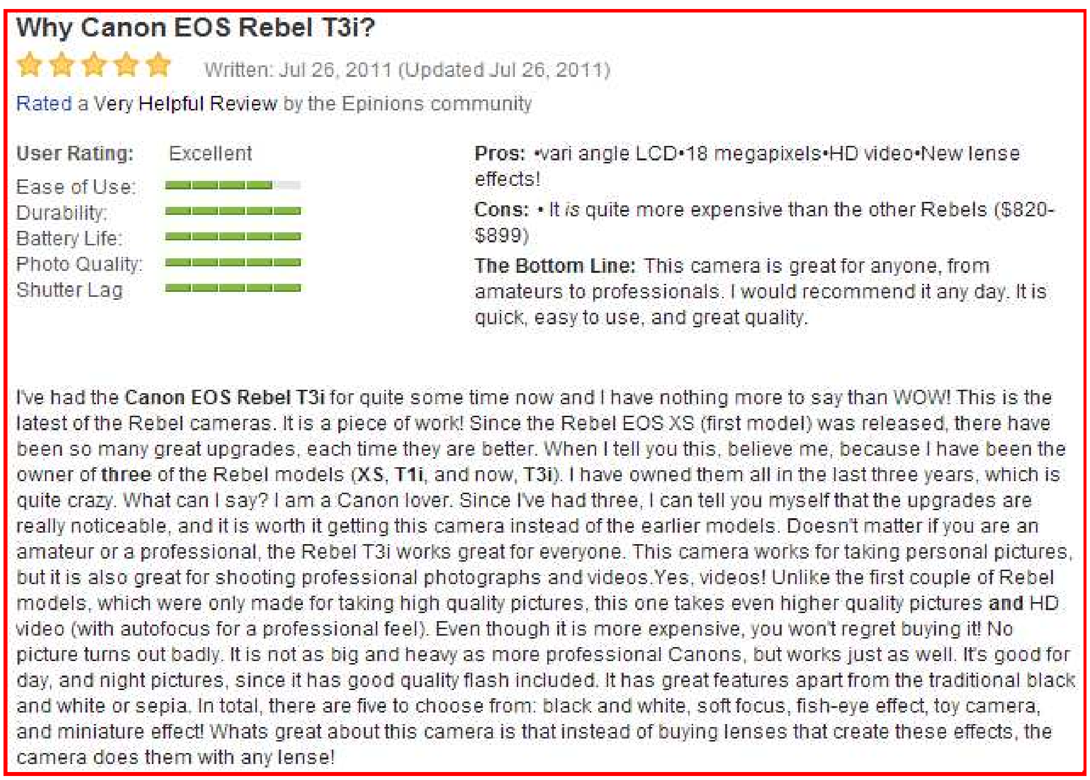

Why Canon EOS Rebel T3i?

Written: Jul 26, 2011 (Updated Jul 26, 2011)

Rated a Very Helpful Review by the Epinions community

User Rating: Excellent

Ease of Use:
Durability:
Battery Life:
Photo Quality:
Shutter Lag

Pros: •vari angle LCD•18 megapixels•HD video•New lense effects!

Cons: • It *is* quite more expensive than the other Rebels ($820-$899)

The Bottom Line: This camera is great for anyone, from amateurs to professionals. I would recommend it any day. It is quick, easy to use, and great quality.

I've had the **Canon EOS Rebel T3i** for quite some time now and I have nothing more to say than WOW! This is the latest of the Rebel cameras. It is a piece of work! Since the Rebel EOS XS (first model) was released, there have been so many great upgrades, each time they are better. When I tell you this, believe me, because I have been the owner of **three** of the Rebel models (**XS**, **T1i**, and now, **T3i**). I have owned them all in the last three years, which is quite crazy. What can I say? I am a Canon lover. Since I've had three, I can tell you myself that the upgrades are really noticeable, and it is worth it getting this camera instead of the earlier models. Doesn't matter if you are an amateur or a professional, the Rebel T3i works great for everyone. This camera works for taking personal pictures, but it is also great for shooting professional photographs and videos.Yes, videos! Unlike the first couple of Rebel models, which were only made for taking high quality pictures, this one takes even higher quality pictures **and** HD video (with autofocus for a professional feel). Even though it is more expensive, you won't regret buying it! No picture turns out badly. It is not as big and heavy as more professional Canons, but works just as well. It's good for day, and night pictures, since it has good quality flash included. It has great features apart from the traditional black and white or sepia. In total, there are five to choose from: black and white, soft focus, fish-eye effect, toy camera, and miniature effect! Whats great about this camera is that instead of buying lenses that create these effects, the camera does them with any lense!

Figure 1. Input to the *ReviewMiner* System

Output

Canon EOS 600D / Rebel T3i Digital Camera with 18-55mm lens

Table 1. Output from the *ReviewMiner* after Processing the XML Input

Aspect	Rating (scale of 5)
Frame lenses	4.33
Resolution	2.78
Pixel density	1.67
Video camera	4.17
Picture quality	4.0

CHAPTER 2. RELATED WORK

In recent years, researchers have studied opinion mining to detect the product features that have been mentioned by reviewers in their comments. The researchers have also tried to classify the polarity of the extracted feature, whether it is positive, negative, or neutral. The researchers already worked on the sentiment classification at the review level [16, 17]. In this thesis, sentiment classification is different from existing methods because the thesis concentrates on opinions expressed for each feature rather than providing an opinion for the entire product.

The sentence-level classification has been studied by researchers [18, 19] and differs from our approach because this thesis aims at a more granular level, i.e., identifying sentiments for a particular feature. A sentence can be comprised of multiple product features, so the opinions for each feature may be different; e.g., “the picture quality of the camera is good, but the battery life is short.” “Picture quality” and “battery life” are two different features. The “picture quality” has been given a positive feedback, and the reviewer’s opinion about “battery life” is negative. The sentence-level and product level classification can be found in [20].

In terms of the expressed opinions, the thesis also attempts to capture them at a more precise level. In previous studies, the opinions were considered to be positive, negative, or neutral [18, 22] and neglected the importance of extreme sentiments. For example, “the lens of the camera is very good,” and “the view finder of the camera is good”, both statements are positive opinions, but the sentiments “very good” and “good” are at different sentiment levels. The sentiment “very good” is equivalent to “great,” whereas “good” is at a level below “great.” The goal of this thesis is also to capture these subtle differences between opinion expressions.

Recent studies in the opinion-mining field have dealt only with one-word or two-word features of the product [2, 22], e.g., the “battery life of the camera is good,” where “battery life” altogether is the feature. In this thesis, more detailed analysis is done regarding the capture of multi-word features; e.g., “the rubber ring around the zoom helps in easy gripping,” and “flash coverage at close distances is excellent.” The previous works overlooked aspects such as “rubber ring around the zoom” or “flash coverage at close distances,” whereas in this thesis, word phrases ranging from 1 to 5 words are considered.

The reviewer’s opinions about each feature were measured as positive, negative, or neutral in the previous studies [18, 22]. In this thesis, a numerical value is assigned to all features so that a better evaluation can be done by the users or companies when they look at the final output generated by the *ReviewMiner*. For example, the *ReviewMiner* output “picture quality is 4.5.” signifies that within a range of 1-5, the picture quality has a score of 4.5 which interprets the “picture quality” of the camera as very close to “excellent”. Similarly, if a feature has a rating of 1.5, it is considered to be terrible.

CHAPTER 3. PROBLEM DEFINITION

In real life, opinions can be expressed about any kind of object. The term "object" can be a product, a person, an institution, a topic, etc. Here we consider opinions about a specific product, i.e., a camera. The product has specific features about which opinions are expressed by the customer reviews. For example, suppose a particular camera, Canon's EOS Rebel T3i, has specific features or aspects for the lens quality, video quality, and battery life about which opinions are expressed by the user reviews. The review text may contain opinions about these aspects, such as "the lens of the camera is good," or "the battery life is terrible." Specifically, the key here is the camera which has aspects about which the users express their opinions.

3.1. Assumptions

Let us suppose $C = \{C_1, C_2, C_3\ldots, C_n\}$ be the set of n cameras, such as Canon PowerShot A720, Canon PowerShot S20, Nikon D3000, etc. For every camera, C_i, we have a set of review, $R_i = \{ R_{i,1}, R_{i,2}, R_{i,3}\ldots, R_{i,k} \}$. Each review has a pair of known aspects and review texts. The known aspects also contain a rating, and an overall rating is assigned to the particular camera, C_i. For each camera, the dataset consists of free review text along with known aspects, their ratings, and an overall rating.

3.2. Terminologies

The following subsections explain the most commonly used terminologies for the thesis.

3.2.1. Aspect

An aspect is a feature or attribute of a particular product that denotes a distinctive characteristic of the product. The users have commented about the aspects in the review text.

Example 1

Sentence: “The lens of the camera is not so good.”

In the sentence, the word “lens” is the aspect, or attribute, of the product.

Example 2

Sentence: “The photo quality of the camera is impressive.”

In the sentence, the term “photo quality” denotes the aspect of the product which is an important attribute of the camera. Based on these qualities, the camera would be considered superior or not in comparison to others.

3.2.2. Known Aspect

Known aspects are those predefined aspects already present at epinions.com. These predefined aspects are requested to rate by Epinions.com when users write the reviews.

Example

According to epinions.com’s rating guidline for cameras, there are 5 known aspects: “battery life”, “durability”, “ease of use”, “photo quality” and “shutter lag”.

3.2.3. Sentiment

A sentiment is a view of or attitude toward a situation or event on the basis of which an opinion is derived. Sentiment is considered here as an adjective which quantifies an aspect of the camera [3].

Example 1

Sentence: “The photo quality of the camera is good.”

In the sentence, the adjective “good” quantifies the “photo quality” aspect of the camera.

Example 2

Sentence: “The shutter speed of the camera is poor.”

In the sentence, the adjective “poor” quantifies the “shutter speed” aspect of the camera.

3.2.4. Sentiment Orientation

A sentiment can be oriented in various degrees. In the basic, two-degree orientation, it can either be positive or negative. With a three-degree scale, it can be positive, negative, or neutral. This thesis follows the rating-guideline orientation suggested by epinions.com which is much more granular with a five-degree orientation level. The rating scale ranges from 1 to 5.

Example

The rating guideline supplied by epinions.com proposes that a rating of 5 is equivalent to “excellent”, a rating of 4 is equivalent to “good”, a rating of 3 is equivalent to “average”, a rating of 2 is equivalent to “poor”, and a rating of 1 is equivalent to “terrible”.

3.3. Problem Definition

For a given set of reviews R about multiple cameras, along with a set of predefined aspects P_i for each camera C_i and a rating guideline, we want to extract a set of aspects, $A_i = \{ A_1, A_2, A_3 \ldots, A_z \}$ for each camera C_i and estimate the corresponding ratings $X_{i,j} = \{ X_{i,j,1}, X_{i,j,2}, X_{i,j,3} \ldots, X_{i,j,m} \}$ for each A_i of C_i, based on the sentiments people expressed in the set of reviews $R_{i,j}$.

3.3.1. Problem

Both A and X are unknown. To resolve, we have to perform the following tasks:

Task 1

Extract aspects or features that have comments in the reviews.

Task 2

Collect the list of opinions that quantify the aspects.

Task 3

Compute a rating for the aspects by grouping the opinion word with synonymic words because different people use various linguistic terms to express their views.

3.4. Approach by *ReviewMiner* to Address the Issues

ReviewMiner is designed in a way which would address these issues and, finally, produce an output comprised of product-based aspects along with the ratings. The web crawler is an external system which searches epinions.com and extracts the unstructured product reviews along with the rating guidelines. The web crawler produces the output in XML format which is then input for Review Miner. Review Miner is comprised of various components, each of which performs a specific task.

Components of Review Miner

Following are the components of the Review Miner system which interact among each other. Each component performs its task in a sequential manner, sending and receiving a series of input and outputs.

POS Tagger and product-wise Tagged Reviews Generator

Aspect Finder and Nearest Sentiment Identifier

Pattern Mapping Filter and Noise Reduction

Aspect Rating according to the Rating Guideline

Figure 2 gives an overview for the different components of Review Miner. Their interactions, along with their inputs and output, are shown.

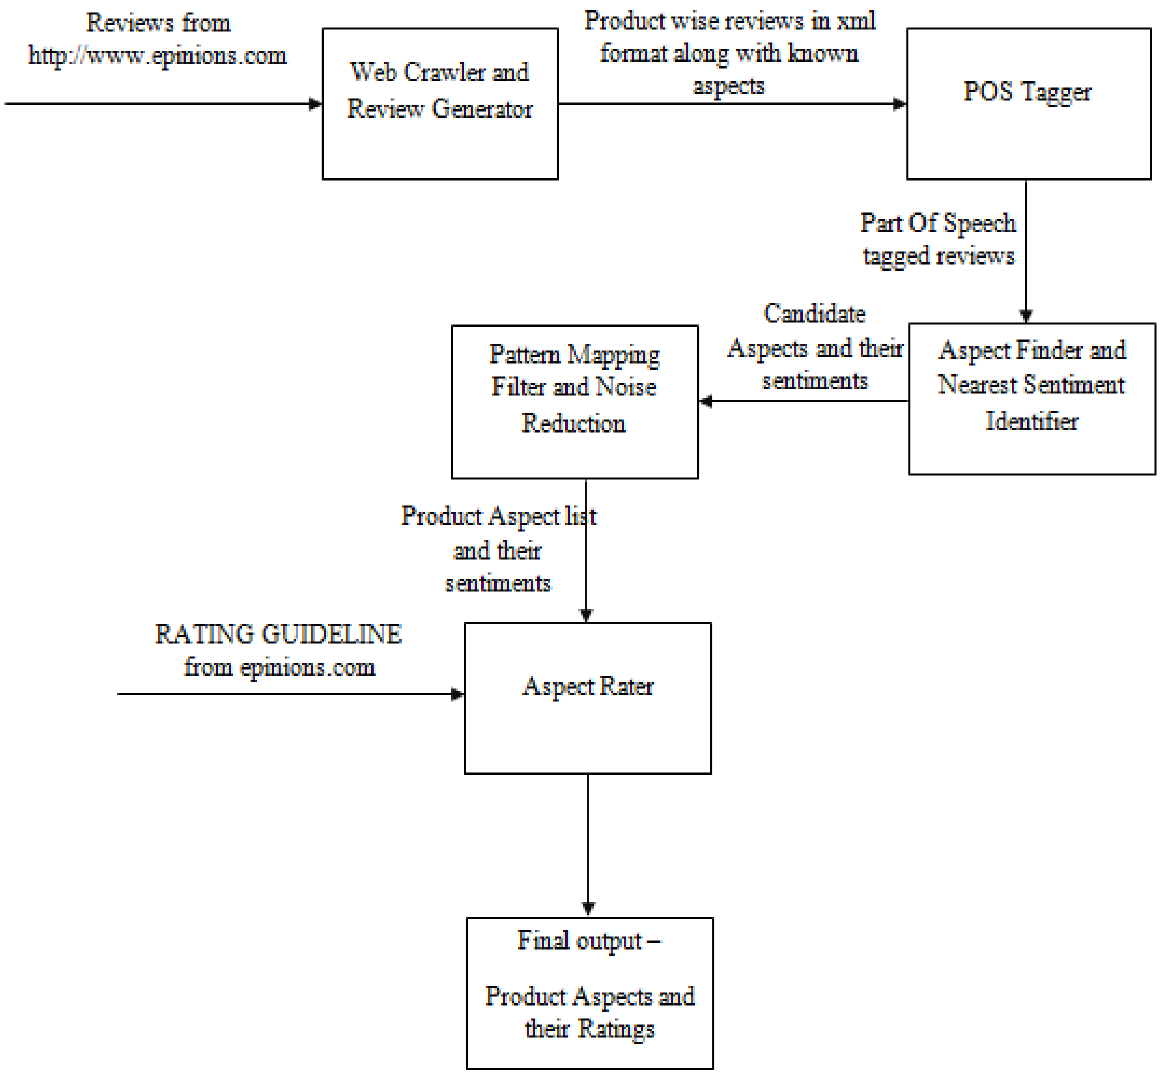

Figure 2. Components' Interaction for *ReviewMiner* with the Inputs and Outputs Shown

CHAPTER 4. ASPECT EXTRACTION AND SENTIMENT CAPTURING

Prior to extracting the aspects from the collection of reviews, the words' parts-of-speech (POS) are tagged. In corpus linguistics, part-of-speech tagging (POS tagging or POST), also called grammatical tagging or word-category disambiguation, is the process of marking a word in a text (corpus) as corresponding to a particular part of speech based on both its definition and its context, i.e., the relationship with adjacent and related words in a phrase, sentence, or paragraph [4].

For example, "This is a sample sentence" will be POS tagged as "This/DT is/VBZ a/DT sample/NN sentence/NN" [8], where DT is the determiner, VBZ is the verb (third-person singular present), and NN denotes noun [9, 10]. Review Miner uses the Stanford NLP Parser [2, 3, 7] for POS tagging of all the sentences in the review texts. An aspect can be a noun, an adjective, an adverb, or a verb. Research studies show that, in most cases, i.e., 60% to 70% of the cases, the aspects are noun [3]. Most likely, the nouns which are frequently mentioned in multiple reviews are supposed to be an aspect.

4.1. Determining type of Aspect

The most widely used method for finding the aspect in opinion mining is the frequent noun method [3]. The relevant content is mentioned in most of the reviews, whereas the irrelevant content is unlikely to be repeated in the reviews. The rare ones, which are not mentioned in most reviews, are unlikely to be a candidate for the camera's aspect.

4.1.1. Candidate Aspect

The candidate aspect is a feature of the camera which qualifies to be a candidate as a feature of the camera. It is a frequently used noun phrase and is found in the comments from most reviewers.

4.1.2. N-gram Model for Phrase Aspects

Review Miner uses the n-gram modeling approach to handle the phrase aspects which are most commonly found in the reviewers's comments [5]. Review Miner handles noun phrases with a maximum of 5 words. The noun phrases can be of 5 types. They are as follows:

<u>N type</u>

These single-word nouns can qualify as an aspect. They are in the form of single N (noun) type.

Example Sentence: The lens of the camera is good.

POS tagged Sentence: The_DT lens_NN of_IN the_DT camera_NN is_VBZ good_JJ ._.

Explanation: In the sentence, "lens" is a single-word noun which can be a candidate for an aspect.

<u>NN type</u>

These two-word aspects are consecutive nouns.

Example Sentence: The auto focus is wonderful for the new users.

POS tagged Sentence: The_DT auto_NN focus_NN is_VBZ wonderful_JJ for_IN the_DT new_JJ users_NNS ._.

Explanation: In the sentence, "auto focus" is an aspect comprised of consecutive nouns.

NX*N type

These four-word aspects are comprised of a noun at the start and the end. The second word is a preposition or an article, and the third word can be any part of speech.

Example Sentence: Indoor shots in poor lighting turn out well and it is fast enough to capture those spontaneous events that happen with kids and pets.

POS tagged Sentence: Indoor_JJ shots_NNS in_IN poor_JJ lighting_NN turn_VB out_RP well_RB and_CC it_PRP is_VBZ fast_RB enough_JJ to_TO capture_VB those_DT spontaneous_JJ events_NNS that_WDT happen_VBP with_IN kids_NNS and_CC pets_NNS ._.

Explanation: In the sentence, “shots in poor lighting” is an aspect where the first word (shots) and last word (lighting) are nouns. The second word (in) is a preposition. The third word (poor) can be any part of speech; in this case, it is an adjective.

NNX*N type

These five-word aspects are comprised of a noun in the first, second, and end positions. The third word is a preposition or an article, and the fourth word can be any part of speech.

Example Sentence: Nikon has better lenses and slightly better image processors for weird situations.

POS tagged Sentence: Nikon_NNP has_VBZ better_JJR lenses_NNS and_CC slightly_RB better_JJR image_NN processors_NNS for_IN weird_JJ situations_NNS ._.

Explanation: In the sentence, “image processors for weird situations” is an aspect where the first word (image), second word (processors), and last word (situations) are nouns. The third word (for) is a preposition. The fourth word (weird) can be any part of speech; in this case, it is an adjective.

4.2. Finding Candidate Aspects

4.2.1. Non-Alphabetic Word Handling

For a particular product, there is a set of reviews. Each review consists of the sentences which are all POS tagged. The *ReviewMiner* goes through the sentences and finds the nouns in the sentences. Then, it checks to see whether the noun only contains alphabetic letters.

Example

The word “EOS-350D” qualifies as a noun, but the word is unlikely to be an aspect. Therefore, they are filtered by this alphabet-only checker method.

4.2.2. Phrase Aspect Handling

First, the Phrase Aspect Handling method starts checking whether the first noun it encounters qualifies for the longest phrase, i.e., five words. Then the Phrase Aspect Handling method checks whether the adjacent word is also a noun. If the next word is a noun, then it checks whether the next word is a preposition or an article. If the next word is a preposition or an article, then the Phrase Aspect Handling method checks whether the last word is also a noun. If the last word is a noun, then there is a five-word phrase that is saved as a candidate aspect.

If the consequence of the first encountered noun does not qualify for the requirements of a five-word phrase but satisfies the consecutive nouns, then it is a two-word phrase. Otherwise, the Phrase Aspect Handling method moves on to check for a four-word phrase. With a four-word phrase, the Phrase Aspect Handling method first checks for a preposition or article in the second position and, finally, for noun in the fourth position. If all conditions are met, then it saves the phrase as a candidate aspect.

Lastly, if consequence of the first encountered noun does not satisfy the four-word phrase requirements, the Phrase Aspect Handling method checks whether there is a preposition or an article in the second position as well as a noun in the third position. If both conditions are met, it qualifies for a three-word phrase and is saved as a candidate aspect.

4.2.3. Reducing Derived Words to the Stem Word

Every word in the phrase aspect which qualifies for the candidate aspect is stemmed, or reduced, to the stem word. The Porter Stemmer algorithm [6] helps to narrow the number of aspects by reducing the derived words to their stems by removing the commoner morphological and inflexional endings from words in English.

4.2.4. Application of Porter Stemmer Algorithm and Phrase Aspect

The Porter Stemmer algorithm is applied prior to handling the phrase aspects, and each word of the phrase aspect is stemmed by the same method mentioned above and a stemmed phrase aspect comprising of the stemmed words as constitutes, is obtained. However, a generalized version of the phrase aspect is saved so that a meaningful interpretation can be obtained with the results instead of the stemmed aspect.

Algorithm getAspect()

1. for each review R_j in the cameras do
2.   for each word W_i in each review R_j do
3.   if W_1 is a noun then
4.   apply is-alphabet-only method for W_1
5. $W_1 = W_1$ is stemmed using Porter Stemmer
6.   if W_2 is a noun then
7.   if W_3 is a preposition or an article and W_5 is a noun then
8.   aspect phrase = from W_1 to W_5

else

aspect phrase = from W_1 to W_2

End if

if W_2 is preposition or an article and W_4 is a noun then

aspect phrase = from W_1 to W_4

End if

if W_2 is preposition or an article and W_3 is a noun then

aspect phrase = from W_1 to W_3

else

aspect phrase = W_1

End if

End if

End for

End for

4.3. Finding Sentiments for the Aspects

4.3.1. Nearest Adjective Capturing

The Review Miner, while creating the list of candidate aspects, searches for the nearest adjectives for all features in the reviews. The nearest adjective within an 11-word proximity of the fully qualified aspect phrase is considered as the sentiment of the aspect phrase. The adjective can be present before or after the aspect phrase to qualify for the sentiment.

4.3.2. Sentiment-Segment Handling

Review Miner considers sentence-segment handling while capturing the sentiments near an aspect phrase. Even if the sentiment is in the proximity, i.e., within 11 words boundary before or after the aspect, it is also checked to see if there is a semicolon or period between within that 11 words. If a semicolon or period is present, then the aspect sentiment is not qualified as an aspect-associated sentiment and is discarded.

Example Sentence:

Most features of the camera are good; the shutter speed has some issues.

POS Tagged Sentence:

Most/JJS features/NNS of/IN the/DT camera/NN are/VBP good/JJ ;/: the/DT shutter/NN speed/NN has/VBZ some/DT issues/NNS ./.

Although the aspect “shutter speed” is closer to the adjective “good,” the sentiment “good” does not quantify the aspect “shutter speed.” Instead, “good” quantifies the “other features” which is placed at a further distance than the aspect “shutter speed.” The aspect “shutter speed” is in a different sentence segment than “good,” with the segments separated by a semicolon. Thus, aspects in a sentence segment should not be quantified by adjectives from another sentence segment.

4.3.3. Handling Extreme Sentiments

Review Miner captures the degree of sentiment corresponding to an aspect. An aspect’s sentiments are the adjectives closest to the aspect. The adverb immediately before the adjective either quantifies the degree of the sentiment or reverses the sentiment.

Reversing the sentiment

This kind of adverb reverses the adjective's meaning. If the adjective is a positive sentiment, it becomes negative, and if it is a negative one, it becomes positive.

Example Sentence:

The auto focus is barely useful.

POS Tagged Sentence:

The_DT auto_NN focus_NN is_VBZ barely_RB useful_JJ ._.

Here, the aspect "auto focus" is quantified by the adjective "useful." "Useful" is positive sentiment. The adverb "barely" reverses the meaning of the adjective "useful," so, instead of being a positive sentiment, the "auto focus" aspect becomes negative because of the adverb's presence.

Quantifying the degree of the sentiment

This kind of adverb quantifies the degree of the sentiment. It either increases or decreases the degree for the sentiment of an aspect.

Example:

The zoom of the camera is very good.

Here, the aspect "zoom" is described by the adjective "good." The adjective "good" is a positive sentiment. The adverb "very" increases the degree of goodness to excellent according to the rating guideline. Therefore, "very good" is replaced by "excellent."

Special Cases for the adverbs "very" and "really":

Positive adverbs such as "very" and "really" are handled in a special way. If the adjective is preceded by a positive adverb (i.e., "very" or "really"), the degree of the adjective is increased one more level.

4.3.4. Opinion Pattern Mining

Review Miner utilizes the known aspects collected from epinions.com to mine the part-of-speech patterns that match the known aspects across reviews. Because these patterns are mined across the reviews, patterns can be applied to the potential candidate aspects to filter out false positives. To mine the pattern, Review Miner finds identical phrases in the reviews for each known aspect. Then, it searches for adjectives in the nearest proximity of that known aspect in the reviews as a corresponding sentiment. Whether the sentiment is before or after the known aspect, the parts of speech information from the aspect to the corresponding sentiment or from the sentiment to the associated aspect is captured. This captured POS sequence between the adjective and known aspect is saved as a pattern. The known aspect itself is generalized and replaced with "_ASP" in the pattern to identify the aspect section in the segment.

<u>Example Sentence for Post aspect sentiment</u>

The camera's battery life is superb.

POS tagged Sentence:

The_DT camera_NN 's_POS battery_NN life_NN is_VBZ superb_JJ ._.

POS tagged Pattern:

_ASP_VBZ_JJ

<u>Example Sentence for Pre aspect sentiment</u>

The camera has a good battery life.

POS tagged Sentence:

The_DT camera_NN has_VBZ a_DT good_JJ battery_NN life_NN ._.

POS tagged Pattern:

_JJ_ASP

In both cases, the known aspect (battery life) is tagged as _NN_NN and is replaced by _ASP so that it can be considered as a generalized pattern.

4.3.5. Filtering Non-Aspects: Based on Pattern Number

We further employ a cutoff value to filter out the non-aspects. Across all reviews, we obtained 913 occurrences of known aspects. Of these 913 occurrences, 728 appear in proximity to a sentiment. Mining these reviews for patterns of known aspects, the number of unique patterns is 304, so the average number of sentiments to which each pattern is mapped is 2.39(728/304). The average number of sentiments per pattern is taken as the “Pattern Average” factor which is rounded to the nearest integer and is taken into consideration for filtering the non-aspects. Patterns which have an appearance frequency less than this Pattern Average are ignored.

Example 1

The frequency of the pattern “_ASP_VBZ_JJ” in the reviews is 5. Because it is greater than 2, it is a valid pattern and is used to filter out the non-aspects.

Example 2

The pattern “_ASP_WDT_NNP_VBD_VBD_RBR_JJ”’s frequency in the reviews is only 1. Because it is less than 2, it is not a valid pattern and is discarded.

Now, the candidate aspects are searched for the nearest adjectives in the reviews. The POS segment between the aspect and corresponding adjective is saved in a similar manner as for the known aspects. For each candidate aspect, the number of matching POS patterns is calculated across the reviews. If the number of matching POS patterns is equal to or more than the predefined threshold, then it is kept; otherwise, it will be ignored.

Example

Suppose the threshold for matched patterns is 1.

The candidate aspect “picture quality” has matched with the following patterns:

1. _ASP_VBZ_JJ
2. _JJ_ASP
3. _ASP_VBZ_RB_JJ

Because the number of matching patterns is greater than 1, the “picture quality’ aspect is not filtered. This method reduces a major portion of the non-aspects from the list of candidate aspects. Still, there is noise in the candidate aspects because some non-aspects match the required pattern threshold. We leave the issue of setting more appropriate filtering technique for noise reduction in our future work.

4.3.6. Filtering Non-Aspects: Based on Sentiment Frequency

The frequent aspects are separated from the less-frequent ones based on the number of times they appear in the reviewers’ comments. In this strategy, a support value of 10% is applied to the number of sentiments belonging to each aspect in the reviews to filter out the non-aspects. The filtering of non-aspects is a potential area of improvement where the more complicated approach is left for future work. Applying the sentiment frequency as a measure to filter out potential non-aspects from the candidate aspects may still not eliminate the noise present in the reviews since some non-aspects may still remain in the list because they may have corresponding sentiment frequency values greater than the specified support value.

CHAPTER 5. ASPECT RATING

In contemporary research on opinion mining, most methods deal with a maximum of three levels of sentiment classification. Most research methods in this field classify sentiment to be either positive or negative. Some techniques classify sentiments as positive, negative, or neutral. The *ReviewMiner* handles sentiment classification up to five levels. The *ReviewMiner* also takes care of extreme goodness or badness and captures the rating with a 5-level orientation scale and rates the aspects of the camera on the basis of that scale.

The aspect rating is done for each camera independent of others. For each camera, C_i, *ReviewMiner* finds the nearest sentiment that quantifies the aspect, A_i, in the set of reviews. For each aspect, there is a list of sentiments associated with it. Pulling together the tally of all the sentiments finally gives an assessment of the aspect. For each sentiment, we will compute a rating and the rating guideline provided by Epinions.com is used for this purpose to assign a rating to the sentiment. To summarize, the aspect rating mechanism can be separated into the following two major steps.

5.1. Finding the Nearest Neighbor for Sentiments

5.1.1. Steps for Finding the Nearest Neighbor

<u>Step 1</u>

For each product aspect, the sentiment which is closest, either before or after, to the aspect within the same sentence segment is captured.

<u>Step 2</u>

The K nearest neighbor (KNN) algorithm is applied for each aspect sentiment using the Wordnet [12] hierarchical graph to find the nearest neighbors for the aspect sentiment.

Step 3

The rating guideline is used to compute the similarity between the aspect sentiment and the representative sentiments defined in the Epinions.com guideline.

Step 4

In this thesis, K is equal to 2, so two nearest neighbors of the aspect sentiment are obtained based on the rating guideline.

Step 5

Most sentiments fall in between two repetitive sentiments given in the rating guideline, i.e., having two nearest neighbors.

Example

For the “nice” aspect sentiment, the closest adjectives from the rating guideline are “good” and “average.”

Also, there are aspect sentiments which are above “excellent” or below “poor”. For example the aspect sentiment “worst” is placed below “poor” according to the rating guideline and having only one nearest neighbor.

5.2. Rating Mechanism

The nearest neighbor algorithm uses the rating guideline provided by Epinions.com. Each sentiment of the aspect is placed within the scale of rating guideline to get the nearest neighbors.

5.2.1. The Rating Guideline

The rating guideline is comprised of 5 level orientations as shown in Figure 3. The rating guideline has values ranging from excellent, good, average, poor, to terrible. Any sentiment is categorized within this range, and the two nearest neighbors of the sentiment are found by applying the 2-NN nearest-neighbor algorithm. For example, the word "nice" has two nearest neighbors, i.e., "good" and "average." Figure 3 shows the rating guideline with a range of values along with the placement of the "nice" sentiment within the range. The y-axis represents the numerical value of the rating range's values.

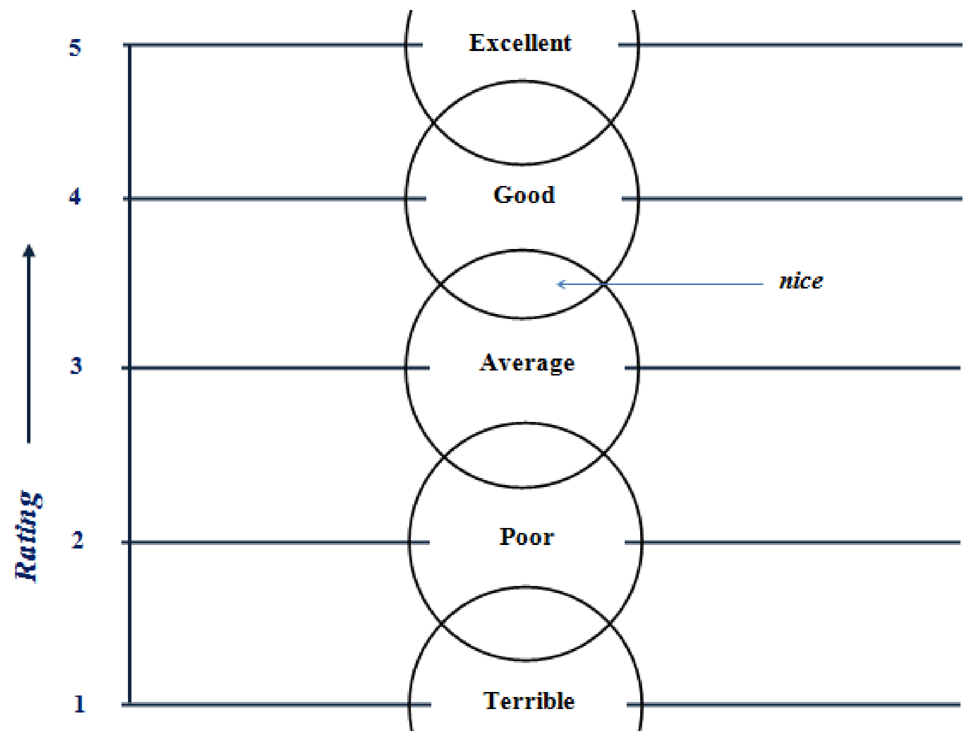

Figure 3. The Rating Guideline with the Placement of the "Nice" Keyword

5.2.2. Assumptions from the Rating Guideline (Epinions.com)

Epinions.com provides a rating guideline that assigns values to the 5-level orientation sentiments. A rating of 5 is assigned to “excellent,” with a rating of 4 as “good,” a rating of 3 as “average,” a rating of 2 as “poor,” and a rating of 1 as “terrible.”

For each aspect sentiment, the KNN algorithm finds the two nearest neighbors as the first comparable sentiment and the second comparable sentiment. The first sentiment is closest to the aspect sentiment. While rating the aspects based on the nearest neighbors, different weights are assigned to the found neighbors. The closest neighbor is given twice the weight compared to the second closest neighbor.

5.2.3. Weighted Rating

For example, the “picture quality” aspect has the sentiments “really good” and “very bad” identified from the reviews. The nearest neighbors of the “really good” sentiment, according to the rating guideline from the Wordnet database, are “good” and “excellent.” The first sentiment for “really good” is good, and the second sentiment is “excellent.” In this thesis, the first sentiment is given more priority than the second sentiment, i.e., twice the second’s assigned weight.

<u>Rating Example</u>

The sum of the rating for “really good” is 2 * good + 1 * excellent = 2 * 4 + 1 * 5 = 13.

The total weight for “really good” is 2 for good + 1 for excellent = 2 + 1 = 3.

Again, for the “very bad” sentiment, the nearest neighbors are “terrible” and “poor.” The first sentiment for “very bad” is “terrible,” and the second sentiment is “poor.” Similarly, The sum of the rating for "very bad" is 2 * terrible + 1 * poor = 2 * 1 + 1 * 2 = 4.

The total weight for "very bad" is 2 for terrible and 1 for poor = 2 + 1 = 3.

The sum of the ratings for the "picture quality" aspect is equal to 13 + 4 = 17. The sum of the weight is 3 + 3 = 6.

The final rating of the "picture quality" aspect is then 17/6 = 2.83 on a scale of 5. The rating 2.83 on a scale of 5 indicates that the camera's "picture quality" is in between average and poor.

In case of sentiments having only one nearest neighbor in the rating guideline, the sentiments like those placed above excellent or placed below poor are assigned full weights, i.e., 3 is assigned to the nearest sentiment in the rating guideline.

<u>Algorithm</u> computeRating()

Rating Guideline (Excellent, Good, Average, Poor, Terrible)

Sentiment 5-scale ratings: Excellent(5), Good(4), Average(3), Good(3), Poor(2), Terrible(1)

1. For each product P_i
2. For each aspect A_j of the product P_i
3. Breadth First Search (*BFS*) aspect A_j's sentiment S_k in Wordnet graph
4. FirstAdjective *FA* = the Nearest adjective for sentiment S_k in the Wordnet graph according to the Rating Guideline
5. SecondAdjective *SA* = Second closest adjective for sentiment S_k in the Wordnet graph according to the Rating Guideline
6. closestAdjectiveMap CAM_{ij} = { A_j , (*FA* , *SA*) }
7. End For

```
End For
For each product P_i
    For each CAM_ij for aspect A_j
        createProductRatingMap()
    End For
End For
```

Algorithm createProductRatingMap()

```
For each sentiment S_k in closetAdjectiveMap CAM_ij
    AdjectiveCount = 0;   Rating = 0;
    Matching with Sentiment guideline 5-scale rating values
        In  the case of two nearest neighbors
            Processing the first nearest neighbor FirstAdjective FA
                AdjectiveCount = AdjectiveCount + 2
                Rating = Rating + 2 * ScaleRating(FirstAdjective FA)
            Processing the second nearest neighbor  SecondAdjective SA
                AdjectiveCount = AdjectiveCount + 1
                Rating = Rating + 1 * ScaleRating(SecondAdjective FA)
        In  the case of single nearest neighbor
            AdjectiveCount = AdjectiveCount + 3
            Rating = Rating + 3 * ScaleRating(SingleAdjective)
    FinalRating = Rating / AdjectiveCount
End For
```

5.2.4. Sentiment Reversal Effect

If any adjective is preceded by the adverb “not,” the polarity, or orientation, of the sentiment is reversed.

Example Sentence 1

“The battery life of the camera is not good.”

Here, the “battery life” aspect is quantified by the adjective “good.” The adverb “not” just before “good” reverses the sentiment’s polarity. Here, “not good” is considered as “poor” according to rating guideline provided by Epinions.com.

Example Sentence 2

“The lens of the camera is barely useful.”

Here, the adjective “useful” is preceded by the adverb “barely” or “hardly.” Due to the presence of the adverb “barely,” the polarity of the sentiment is reversed.

Rating changes for the presence of “not”

These adjectives with “not” adverbs preceding them are substituted as follows:

1. not excellent = good
2. not good = poor
3. not average = good
4. not poor = average
5. not terrible = poor

5.2.5. Sentiment Changing in Effect

If any adjective is preceded by the adverb “very,” the sentiment’s degree is moved one level higher or lower in the same direction.

Example Sentence 1

"The shutter speed of the camera is very good."

Here, the "shutter speed" aspect is quantified by the adjective "good." The adverb "very," which just precedes the adjective "good," increases the orientation level of the "good" sentiment. Therefore, "very good" can be considered as the sentiment "excellent."

Rating changes due to a positive adverb effect

These adjectives with "very" adverbs preceding them are substituted as follows:

1. very good = excellent
2. very average = poor
3. very poor = terrible

Rating changes due to the presence of negative adverbs

These adjectives with "hardly" or "barely" adverbs preceding them are substituted as follows:

1. hardly good = average
2. barely average = poor

5.2.6. Combining Sentiment Reversal and Sentiment Changes

If any adjective is preceded by a combination of reversal and sentiment-changing adverbs, a change in the rating takes place. The following cases have been identified for the presence of both categories of adverbs.

Rating changes due to co-existence of the positive adverb and negation verb cases

1. not very good = average

2. not very poor = average

3. not very average = good

CHAPTER 6. EXPERIMENTAL RESULTS

The evaluation of our method is verified from both perspectives. The first perspective is the accuracy of the extracted aspects i.e. in comparison with manual extraction of aspects how well did the *ReviewMiner* do in extracting aspects. The second perspective deals with the rating of the product aspects i.e. in comparison with manual rating how well *ReviewMiner* rate the aspects of the product.

6.1. Dataset

Because data containing known aspects and a rating guideline in the required format were not readily available, a web crawler was designed and developed to extract camera reviews along with known aspects and a rating guideline. The dataset created by the web crawler consists of over 1.1K reviews for 130 cameras of various brands. The cameras which are used in the dataset are Cannon, DMC, Fujifilm, Nikon, Kodak, Olympus, Panasonic, and Pentax.

6.2. Baselines for Evaluating the Aspect Extraction

6.2.1. Adverb Preceding Adjectives

ReviewMiner captures the adverbs just preceding the aspect's sentiment. The *ReviewMiner* captures adverbs within the proximity of two words before the adjective because there can be sentiment phrases such as "not very good".

<u>Effect of Baseline</u>

The implementation of adverb handling method has a major impact on the aspect-rating calculation. Due to the presence of adverbs before the adjective, the effective sentiment might be reversed or changed (increased or decreased based on positive or negative adverbs). There were

130 cameras in total, altogether having 1,116 reviews. The *ReviewMiner* rated a total of 10,146 product aspects from 1,116 reviews.

When adverb handling method is applied across the reviews, it is seen that, of 7,571 rated aspects, there were 1,017 aspect-rating changes. The adverb handling method improves the accuracy of the rating as shown in Table 2 by 10.02% for all reviews.

6.2.2. Pattern-Mapping Noise Reduction

ReviewMiner employs the pattern-mapping technique to filter out the non-aspects from the candidate aspects. The pattern-mapping method uses the known aspects extracted from Epinions.com to filter out the non-aspects. The known aspects are matched for patterns across 1116 reviews.

It is observed that the count for the product aspects is reduced when the pattern-mapping method, compared to the naive one, is employed. When pattern mapping is not used, the aspect count is 13,735 across all 1,116 camera reviews. After implementing the pattern-mapping method, the number of aspects is reduced by 3,589, and the total number of aspects is reduced to 10,146. The result as shown in Table 2 improves by 26.13%, when filtering out the noises, over the naive method (without the pattern-mapping method applied).

6.2.3. Phrase Aspects

ReviewMiner handles phrase aspects of length 2, 3, 4, and 5 apart from the single-word aspects. The baseline is created by comparing the results obtained when the filter for capturing 2, 3, 4, and 5 words is turned off. While capturing phrases when the filter is turned off, only single-word nouns which have adjectives close to them are qualified as aspects.

While turning off the filter for capturing 2-, 3-, 4-, and 5-word aspects across all reviews, 14,662 one-word aspects are identified by *ReviewMiner*. When the method for capturing the phrase aspects is employed for capturing 2-, 3-, 4-, and 5-word aspects, the one-word aspect count is reduced to 8,192. Following are the counts for phrase aspects when the phrase-aspect method is employed.

1-word: 8,192

2-word: 1,744

3-word: 88

4-word: 118

5-word: 4

Therefore, the total number of aspects, including all the phase aspects (1-word + 2-word + 3-word + 4-word + 5-word), is 10,146 which is less than the 1-word aspects when the phrase-aspect filter is turned off. Thus, the 1-word aspects are reduced by 6,470, and the total number of aspects is reduced by 4,516. It is also noted that the 1-word aspects are based totally on the noun words. The 2-word aspect contains two nouns simultaneously, and other phrase aspects include multiple nouns apart from conjunction, preposition, or free part of speech. Thus, consecutive nouns which are near an adjective in a sentence qualify as a 2-word aspect rather than two 1-word aspects, reducing the total number of aspects. The aspect count is reduced to a large extent since the same sentiment patterns are grouped together. Thus an improvement by 30.8% is observed when the phrase-aspect handling mechanism is injected. Filtering out the 1-word aspects which are treated as phrase aspects with the introduction of the phrase-aspect method shows a reduction of 44.12%, or 6,470, in the 1-word aspect list as shown in Table 2.

6.2.4. Noise Reduction by Removing Non-Alphabetic Letters from the Aspect List

Number aspects that are captured on the basis of the “noun” part of speech contain numeric, or non-alphabetic, characters. *ReviewMiner* discards this noise by filtering out the aspects with non-alphabetic characters. When the non-alphabetic handler method is applied to *ReviewMiner*, the number of aspects is reduced from 10,246 to 10,146. Therefore, 100 aspects are filtered out as noise when the non-alphabetic handler method is employed by *ReviewMiner*. Thus, the *ReviewMiner* method as shown in Table 2 is improved by 1% when introducing this method.

6.2.5. Sentence-Segment Handling

ReviewMiner takes care of sentence segment handling to ensure that related pairs of aspects and sentiments are captured. If the aspect and the sentiment are present in different sentence segments separated by a comma, semicolon or a period then that aspect is very unlikely, qualified by the sentiment. *ReviewMiner* handles the comma, semicolon, and period separated sentence segments and ensures that, even if the aspect or sentiment is within the proximity and qualifies as an aspect-sentiment pair, the *ReviewMiner* discards the pair.

It is observed that, when the sentence-segment filtering is turned off, the number of aspects increases to 11,258. The number of actual aspects with the sentence-segment handling is 10,146. Therefore, 1,112 aspects are filtered out when the sentence-segment handling method is applied. Thus, the sentence-segment handling as shown in Table 2 leads to an improvement of 9.87% over the unhandled scenario.

Table 2. Aspect Extraction Evaluation Comparison: Review Miner and the Naive Method

Impact Method	Unfiltered Aspect Count	Filtered Aspect Count	Percent Improvement
Pattern-Mapping	13,735	10,146	26.13%
Phrase Aspects	14,662	10,146	30.8%
Non-Alphabetic	10,246	10,146	0.98%
Sentence-Segment	11,258	10,146	9.87%

6.2.6. Changes in the Aspect Rating Due to the Presence of Negative Adverbs

The presence of negative adverbs, such as "not," before the adjective which qualify an aspect reverses the direction of the aspect's sentiment. Thus, "not good" reverses aspect's opinion polarity and becomes "poor." The rating assigned to the aspect is reduced from 4 to 2, impacting the overall calculation of the aspect rating.

It is observed across all reviews that the number of rated aspects for 130 cameras is 10,146. There are 362 instances of rating changes for these aspects across the cameras. Hence, as shown in Table 3, an improvement of 3.56% is obtained with the introduction of the negative adjective-handling technique.

6.2.7. Changes in the Aspect Rating with the Presence of Positive Adverbs in the Sentiment

The presence of positive adverbs, such as "very," preceding the adjective which qualify an aspect changes the aspect's sentiment. The presence of positive adverbs either increases or decreases the aspect's rating value, in most cases pushing it to the extreme. Thus, the adverb "very" in the adverb adjective "very good" increments opinion value of the adjective "good" and it becomes “excellent,” so the rating assigned to the aspect increases from 4 to 5, impacting the overall calculation of the aspect rating.

The positive adverb-handling mechanism is applied across all 1,116 reviews, and the rating changes if the aspects are observed and counted. It is observed that there are 626 aspect-rating changes when the positive adverb-handling method is introduced in Review Miner. Thus, as shown in Table 3 an improvement of 6.17% is obtained with the method’s introduction.

6.2.8. Additional Weight for the Most Nearest Neighbor

A 2-NN algorithm is applied to find the nearest neighbor for each sentiment. In most cases, the two nearest neighbors for each sentiment are not placed at equal distances from the sentiment. Generally, one is closer than the other. The one which is closer has more impact, or is more similar to the sentiment, and is termed as the first adjective. The one which is placed a little further from the sentiment is called the second adjective. The second adjective has less impact, or is less similar to the sentiment, compared to first adjective. This emphasis on first and second adjectives is handled by assigning double the weight to first adjective compared to second one.

The impact due to the presence of the weight-handler method is observed in the rating aspects’ changes. The impact is studied across all 1,116 review and for all 130 products. A total of 10,146 aspects are rated for all the products. The introduction of the weight-handling methods causes 4,832 aspects to improve their ratings. As shown in Table 3, 47.62% improvement in the rating is observed across the reviews.

Table 3. Aspect Rating Evaluation Comparison between Naive and *ReviewMiner*

Impact Method	Aspect Count	Rating Changes	Percent Changes
Adverb Preceding adjectives	10,146	1,017	10.02%
Negate Word handling	10,146	362	3.56%
Positive Word Handling	10,146	626	6.17%
Weightage Calculation	10,146	4,832	47.62%

6.3. Overall Camera Evaluation and Comparison

Three digital cameras were randomly chosen (from 130 cameras) for manual tagging and aspect extraction. The reviews of the two digital cameras that were tagged manually by humans are considered as the “Golden Standard” for comparing the opinion-mining methods. The performance measures of precision, recall, and f score [11, 13] are used for aspect extraction. Table 4, Table 5 and Table 6 shows the comparison between the different methods and Review Miner for the measures of precision, recall, and f score.

<u>Naive Method</u>

The Naive method is the basic method without any algorithm applied for aspect extraction. This method involves single-word aspects.

<u>Base Method</u>

The Base method is the method with only phrase aspects applied for aspect extraction.

Non-Alphabetic Method

The Non Alphabetic Method is the method that filters the aspects consisting of non alphabetic letters.

Sentence-Segment Method

Sentence-Segment Method is the method that filters non-aspects by considering the aspect sentiment pair within the same sentence segment only.

Pattern-Mapping Method

Pattern-Mapping Mehtod is the method of filtering the non-aspects. The method filters the aspects which have a frequency less than the Pattern Number.

Non-Alphabetic and Sentence-Segment Method

Non-Alphabetic and Sentence-Segment Method is the method combining both non-alphabetic and sentence-segment methods. The combined method filters the aspects consisting of non alphabetic letters and aspect sentiment pair in different sentence segments.

Pattern Mapping and Non-Alphabetic Method

Pattern Mapping and Non-Alphabetic Method is the method combining both pattern mapping and non-alphabetic methods. The combined method filters the aspects which have a frequency less than the Pattern Number and consisting of non alphabetic letters.

Pattern Mapping and Sentence-Segment Method

Pattern Mapping and Sentence-Segment Method is the method combining both pattern mapping and sentence-segment methods. The combined method filters the aspects which have a frequency less than the Pattern Number and and aspect sentiment pair in different sentence segments.

Table 4. Comparison between Various Methods for Digital Camera 1

Methods	Digital Camera 1		
	Precision	Recall	F Score
Naive	28.0%	63.63%	38.88%
Base	51.06%	83.63%	63.4%
Non-Alphabetic	51.06%	72.72%	59.99%
Non-Alphabetic and Sentence-Segment	51.06%	76.36%	61.19%
Pattern Mapping	77.27%	82.22%	81.15%
Pattern Mapping and Non-Alphabetic	77.27%	82.22%	81.15%
Pattern Mapping and Sentence-Segment	70.83%	82.22%	77.45%
Sentence-Segment	79.66%	82.22%	82.45%
ReviewMiner	79.66%	85.45%	82.45%

Table 5. Comparison between Various Methods for Digital Camera 2

Methods	Digital Camera 2		
	Precision	Recall	F Score
Naive	17.85%	57.14%	27.20%
Base	39.39%	57.14%	46.63%
Non-Alphabetic	33.33%	60.0%	42.85%
Non-Alphabetic and Sentence-Segment	33.33%	60.0%	42.85%
Pattern Mapping	53.84%	60.0%	56.75%
Pattern Mapping and Non-Alphabetic	53.84%	60.0%	56.75%
Pattern Mapping and Sentence-Segment	53.84%	60.0%	56.75%
Sentence-Segment	31.48%	60.0%	41.29%
ReviewMiner	70.27%	74.28%	72.22%

Table 6. Comparison between Various Methods for Digital Camera 3

Methods	Digital Camera 3		
	Precision	Recall	F Score
Naive	18.30%	62.22%	28.28%
Base	41.81%	82.22%	55.43%
Non-Alphabetic	64.91%	82.22%	72.54%
Non-Alphabetic and Sentence-Segment	64.91%	82.22%	72.54%
Pattern Mapping	74.24%	82.22%	78.02%
Pattern Mapping and Non-Alphabetic	74.24%	82.22%	78.02%
Pattern Mapping and Sentence-Segment	71.42%	82.22%	76.44%
Sentence-Segment	65.74%	82.22%	73.06%
ReviewMiner	72.54%	82.22%	77.07%

CHAPTER 7. CONCLUSION

The thesis proposed a method for identifying and resolving the issues associated with plenty and long reviews for a particular camera. The long reviews, which are abundant in number, leave the customer perplexed and unable to extract useful information from them. The two major issues the thesis deals with are (1) aspect extraction and sentiment capturing, and (2) the aspect rating. These two issues, in turn, achieve the main goal of the customer who wants to buy a camera. The customer gets a clear picture about the camera's features, or aspects, along with their rating so that they can compare multiple cameras and select the best one.

Issue (1) is subdivided into two phases: (1.a) aspect extraction and (1.b) sentiment capturing. Review Miner takes a set of reviews for a camera, the rating guideline, and a ,set of known aspects as defined on epinions.com. For (1.a), Review Miner uses the n-gram model for candidate-phrase aspects along with the pattern-mapping technique which filters out the non-aspects. Non-alphabetic word handling is also used as a second level for filtering aspects. The Porter Stemmer algorithm is used to group similar aspects which differ by inflected endings. The aspect which does not have frequent sentiments near it is also filtered out and the nearest sentiments, along with the aspect, are saved. For (1.b), various methods have been applied to capture the exact sentiment mentioned in the reviews. Initially, the closest adjective for each aspect is captured by applying the extreme-sentiment handling method. Thus the adverbs preceding the adjective are included in the sentiment. Also, the sentiment segment-handling method is used for the sentiment capturing so that a more accurate aspect-sentiment pair can be obtained. For (2), the Wordnet library is used to quantify the sentiments based on the distance between the rating-guideline sentiments and the aspect sentiment. The rating guideline supplied by Epinions.com is taken as a standard for the computing aspect rating. The 2-NN algorithm is applied to find the nearest adjective from the rating guideline for each sentiment of the aspects. Then, the weight-handling method is applied to emphasize to the closest sentiment and enhance the computed rating's accuracy. In this work, detailed effort is taken to extract the phrase aspects which have been neglected by the existing methods.

Apart from extracting aspects from the reviews, the method also assigns a rating to the aspects. The previous methods only considered a two-level orientation of sentiment, i.e., either positive or negative. In comparison to previous methods, *ReviewMiner* shows an improvement for the opinion-orientation level by rating the aspects on the basis of a five-level orientation that takes extreme sentiments into consideration. These additional benefits make *ReviewMiner* more complete than the existing methods. Experimental evaluation demonstrates that the *ReviewMiner* technique performs much better than the contemporary methods in the field of sentiment analysis.

REFERENCES

[1] B. Liu, "Sentiment Analysis and Opinion Mining", Morgan & Claypool Publishers, CA, USA, 2012.

[2] M. Marcus, B. Santorini, and M. A. Marcinkiewicz, "Building a Large Annotated Corpus of English: The Penn Treebank", *Computational Linguistics*, vol. 19, number 2, pp. 313-330, 1993.

[3] M. Marcus, G. Kim, M. A. Marcinkiewicz, R. Macintyre, A. Bies, M. Ferguson, K. Katz, and B. Schasberger, "The Penn Treebank: Annotating Predicate Argument Structure", in *Advanced Research Projects Agency (ARPA) Human Language Technology Workshop*, pp. 114-119, 1994.

[4] D. M. Powers, "Evaluation: From Precision, Recall and F-Factor to ROC, Informedness, Markedness & Correlation", *Journal of Machine Learning Technologies 2 (1)*, pp. 37-63, 2011.

[5] J. Makhoul, F. Kubala, R. Schwartz, and R. Weischedel, "Performance Measures for Information Extraction", in *Proceedings of DARPA Broadcast News Workshop*, pp. 249-25, 1999.

[6] K. Crammer and Y. Singer, "Pranking with Ranking", in *Neural Information Processing Systems (NIPS)*, 2001.

[7] C. Fellbaum, "Wordnet: An Electronic Lexical Database", MIT Press, Cambridge, MA, USA, 1998.

[8] H. Guo, H. Zhu, Z. Guo, X. Zhang, and Z. Su, "Product feature categorization with multilevel latent semantic association", in *Conference in Information and Knowledge Management (CIKM)*, 2009.

[9] M. Hu and B. Liu, “Mining and summarizing customer reviews”, in *Knowledge Discovery from Data (KDD)*, 2004.

[10] B. Liu, “Web Data Mining: Exploring Hyperlinks, Contents, and Usage Data”, Springer, 2007.

[11] B. Liu, M. Hu, and J. Cheng, “Opinion observer: analyzing and comparing opinions on the web”, in *Proceedings of the 14th international conference on World Wide Web (WWW)*, 2005.

[12] M. F. Porter, “An algorithm for suffix stripping”, *Readings in Information Retrieval*, 1997.

[13] B. Snyder and R. Barzilay, “Multiple aspect ranking using the good grief algorithm”, in *Human Language Technology-North American Association for Computational Linguistics (HLT-NAACL)*, 2007.

[14] R. Srikant and R. Agrawal. “Mining sequential patterns: Generalizations and performance improvements”, in *International Conference on Extending Database Technology* (EDBT), 1996.

[15] P. D. Turney, “Thumbs up or thumbs down?: semantic orientation applied to unsupervised classification of reviews”, in *Association for Computational Linguistics (ACL)*, 2002.

[16] M. Gamon, A. Aue, S. Corston-Oliver, and E. K. Ringger, “Pulse: Mining customer opinions from free text”, in *Intelligent Data Analysis (IDA)*, 2005.

[17] E. Riloff and J. Wiebe, “Learning extraction patterns for subjective expressions”, in *Empirical Methods on Natural Language Processing (EMNLP)*, 2003.

[18] M. Hu and B. Liu, “Mining and summarizing customerreviews”, in *Knowledge Discovery from Data (KDD)*, 2004.

[19] A-M. Popescu and O. Etzioni, "Extracting Product Features and Opinions from Reviews", in *Empirical Methods on Natural Language Processing (EMNLP)*, 2005.

[20] P. Beineke, T. Hastie, C. Manning, and S. Vaithyanathan, "An Exploration of Sentiment Summarization", in *Proceedings of the AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications*, 2003.

[21] X. Ding and B. Liu and P. S. Yu, "A Holistic Lexicon-Based Approach to Opinion Mining", in *Web Search and Data Mining (WSDM)*, 2008.

[22] Samaneh Moghaddam and Martin Ester, "Opinion digger: an unsupervised opinion miner from unstructured product reviews", in *Conference on Information and Knowledge Management (CIKM)*, 2010.