

Linguistic Understanding of Complaints and Praises in User Reviews

Kavita Ganesan

Github Inc.
88 Colin P Kelly Jr St
San Francisco, CA 94107
kganes2@github.com

Guangyu Zhou

Department of Computer Science
University of Illinois at Urbana-Champaign
Urbana, IL 61801
gzhou6@illinois.edu

Abstract

Traditional sentiment analysis has been focused on predicting the polarity of texts as *positive* or *negative* at different granularity. This broad categorization does not account for informativeness of the underlying text. For many real-world applications such as social listening, brand monitoring and e-commerce platforms, the opinions that really matter are the informative opinions describing why something is good or bad. In this paper, we try to understand the properties of *complaints* and *praises* which is an informative subset of the negative and positive categories. Our analysis in the context of user reviews shows that complaints and praises have distinct properties that differentiate it from positive only or negative only sentences.

1 Introduction

Over the last two decades, sentiment analysis research has been focused on predicting *positive* and *negative* polarity ratings at different granularity - at the passage level, sentence level as well as aspect or phrasal level (Pontiki et al., 2015; Wang et al., 2015; Pontiki and Manandhar, 2014; Pang et al., 2002; Pang and Lee, 2008; Wilson et al., 2005; Wang et al., 2010; Snyder and Barzilay, 2007; Titov and McDonald, 2008; Lu et al., 2009; Wang, 2015; Diao et al., 2014). While such ratings provide a general sense of ‘*what people think*’, it does not take into account the informativeness of the comments that contribute towards those ratings as long as the comments have some form of desired subjectivity (e.g. should contain a noun and an adjective). Consider the following sentences about the *Xbox* :

1. *The Xbox is way too expensive!*
2. *I really hate the Xbox!*

In the traditional sense, both these sentences would be considered negative comments of equal

importance. However, sentence (1) is actually more informative than (2). This difference is critical in many application scenarios. For example, if a business analyst wants to analyze the complaints or pain points of a product, a comment such as ‘*I really hate the Xbox!*’ does not increase the understanding of the analyst. However, the comment ‘*The Xbox is way too expensive!*’ informs the analyst that one of the problems of the Xbox is that it is not affordable. Another example is in competitive intelligence. A company will learn more about a competitor’s product from a comment such as ‘*Wow! the Xbox is very user-friendly!*’ as opposed to ‘*The Xbox is just awesome.*’ as the former provides a concrete reason.

The goal of this paper, is to study the linguistic properties of *complaint* and *praise* sentences which we define as an informative subset of the more general negative and positive categories, providing reasons for a topic or aspect being positive or negative. We perform our study in the context of user reviews as user generated reviews tend to have informative subjective content inter-mingled with non-informative subjective content and factual or neutral utterances. We investigate several properties, including the length property, noun and adjective usage, past tense and negation usage and finally the usage of intensifiers and causal links. We contrast the properties of *complaint* and *praise* sentences with *negative only* or *positive only* sentences.

Our study shows several distinct properties of complaints and praises in contrast to positive only and negative only sentences. We believe that this study would set a foundation for improving existing sentiment classification and opinion summarization systems. The data set used for this study is publicly available at <http://kavita-ganesan.com/complaints-and-praises>¹.

¹You can also reference: <https://github.com/kganes2/complaints-and-praises>

negative only	complaint	positive only	praise
This is not a good company , stay away!	This company takes your payment but on the <i>day of the scheduled job, they don't appear.</i>	I really love that restaurant , its awesome.	This restaurant has <i>delicious tacos and the ambience is amazing!</i>
The phone's screen is very disappointing :(Unhappy with this phone , the <i>screen is not clear and the fonts are way too small.</i>	Nice phone , love it and totally recommend it!	I like the fact that the phone <i>fits right in your pocket</i>

Table 1: Examples of *negative only* and *positive only* sentences as well *complaint* and *praise* sentences. Bolded text are the topics and the italicized text answer ‘why’ the topic is positive or negative.

2 Related Work

While there are many systems attempting to predict finer granularity of sentiment ratings at the aspect or phrasal level (Pontiki et al., 2015; Pontiki and Manandhar, 2014; Wang et al., 2015; Wang et al., 2010; Snyder and Barzilay, 2007; Titov and McDonald, 2008; Lu et al., 2009; Wilson et al., 2005; Wang, 2015; Diao et al., 2014), these systems still do not have a clear understanding on what makes a sentence informative enough linguistically to be used for mining fine grained sentiments. The common assumption is that a subjective sentence should contain a *noun* and an *adjective*. In addition, these systems consider all negative and positive expressions as equal contributors. For example, the phrases ‘*screen is bad*’, ‘*the screen is way too small*’ and ‘*the screen is too big*’ would all equally affect ratings on the *screen* aspect. In reality, the first phrase is a general negative statement compared to the second and third phrases which are much more informative, providing reasons for the screen being bad. Having the option of analyzing only the informative subset would add significant value to sentiment analysis applications. For this purpose, there has to be a good understanding on how to distinguish between the different types of subjective comments.

In the work of (Kim and Hovy, 2006), the authors attempt to train a classifier to predict ‘pro’ and ‘con’ reasons in user reviews. However, there is a lack of definition on what ‘pros’ or ‘cons’ represent and how they can be linguistically identified. Our study thus bridges this gap by providing insights into key linguistic properties of *complaint* and *praise* sentences in contrast to plain *negative only* and *positive only* sentences.

3 Defining Complaints and Praises

In this section, we formally define the concept of a *complaint* and a *praise*. Given a sentence, S , we

refer to this sentence as a *positive* sentence if its connotation is positive and a *negative* sentence if its connotation is negative. We refer to a negative sentence as a *complaint* if it has a negative connotation with *supplemental information*, answering the question of why a topic or aspect is negative. We refer to a sentence as *negative only* if it is negative with no such supplemental information. Similarly, we refer to a positive sentence as a *praise*, if the sentence has a positive connotation with *supplemental information*, clearly indicating what makes the topic or aspect positive. A sentence is considered *positive only* if it is positive with no such supplemental information.

Our definition of supplemental information refers to any information in an opinionated sentence that answers the question of ‘why’ the like or dislike for a topic, improving the user’s understanding for that topic. For example, the sentence “*Xbox is just bad...I hate it.*” is considered *negative only* and not a *complaint* because it does not have information explaining why the user dislikes the Xbox. However, the sentence “*The Xbox is awfully expensive, I would not recommend it*” would qualify as a *complaint* as it answers ‘why’ the user has a negative opinion about the Xbox which in turn improves a user’s understanding about the Xbox (that it is expensive). Table 1, shows examples of negative only, positive only, complaint and praise sentences.

4 Dataset

To conduct our analysis, we collected 2500 reviews from various sources including TripAdvisor, Yelp, Walmart and Sephora. We then recruited 4 students to manually categorize sentences from the user reviews into 1 of 5 categories. The categories are: *NegativeOnly*, *Complaint*, *PositiveOnly*, *Praise*, and *Irrelevant*. The sentences were randomly assigned to the students. Students were asked to perform categorization based on the

Category	Avg. # of words	Avg. length
NegativeOnly	10.25	51.95
Complaint	15.75	80.88
Difference	+53.66%	+%55.69
PositiveOnly	10.33	53.13
Praise	15.54	82.45
Difference	+50.44%	+55.19%

Table 2: Average number words and average length per sentence in the dataset.

formal definition of a complaint and a praise sentence as described in Section 3. Non-opinion containing sentences and noise were placed in the *Irrelevant* category. We use the four main categories - *NegativeOnly*, *Complaint*, *PositiveOnly* and *Praise* for our study. For a fair comparison, we ensured that we only used 500 randomly picked sentences within each category.

5 Sentence Length Analysis

Our first analysis deals with understanding the general length of *complaints* and *praises* in contrast to *negative only* or *positive only* sentences. In Table 2, we report the average length and number of words in a sentence within each sentiment category. On average, a *praise* or *complaint* sentence is at least 50% longer than a *positive only* or *negative only* sentence. The average number of words in a *praise* sentence is 15.54 and a *positive only* sentence it is 10.33. The average number of words in a *complaint* sentence is 15.75 and a *negative only* sentence it is 10.25. Intuitively, this makes sense since *complaints* and *praises* require elaboration on why something is good or bad but in the case of *negative only* or *positive only* sentences, the statements can be fairly general.

6 Noun and Adjective Usage

Nouns and adjectives are essential parts of speech within subjective sentences as these together are key indicators of sentiment (Hu and Liu, 2004; Pang and Lee, 2008; Kim et al., 2011). For example, a negative only sentence such as ‘*the screen is bad*’ or a complaint such as ‘*the screen is not clear*’ both have nouns (‘screen’) and adjectives (‘bad’ and ‘clear’). Both the noun and adjectives play a role in indicating negative sentiment. To better understand if there is a difference in adjective and noun usage in a *complaint* or a *praise* versus a *negative only* or *positive only* sentence, we obtained the mean, mode and median of nouns and adjectives in our manually categorized dataset. We also noted the counts of nouns appearing near ad-

jectives within a 3 word window. We report the results in Table 3.

Noun analysis: Based on Table 3, we see that with the *NegativeOnly* and *PositiveOnly* categories, most sentences have 1 noun per sentence (see mode in Table 3). However, in the *Complaint* and *Praise* categories, most sentences use 3 nouns per sentence. This is because a *complaint* or a *praise* sentence describes ‘what’ was good or bad about a topic requiring more use of nouns. For example, if a praise was about the food at a restaurant, the *tacos*, *salsas* and *chips* could have been outstanding.

Adjective analysis: In terms of adjectives, the first point to observe is that most *praise* sentences use 2 or more adjectives while most *complaint* sentences use a single adjective. Upon further investigation, we realized that in a *praise* sentence, users tend to compliment more than one aspect of a topic within a single sentence. For example, consider the following praise sentence: ‘*This is a really lightweight machine and it is easy to assemble*’. The user is complimenting two aspects of the machine within a single sentence: (1) weight and (2) assembly. For this same reason, the adjective+nouns have the highest occurrence within the praise category as multiple positive sentiments are coupled within a single sentence.

This is different from *complaints*, where within complaints, users tend to elaborate why a single aspect of a topic is bad. For example, consider the following complaint: ‘*This machine was really hard to put together, the screws don’t fit so I sent it back*’. This sentence only describes a single aspect which is the fact that the product was hard to assemble, why that was the case and what the user did to address it. This appears to be the common nature of complaints. This is why most *complaints* use no more adjectives than *negative only* sentences.

7 Past Tense Analysis

One observation that we made while visually analyzing our manually constructed dataset is that complaints seemed to use more past tense than the other three categories. To validate our observation, we did a count of occurrence of past tense words in each category of our dataset. The average past tense counts per sentence and the top 6 past tense words used is reported in Table 4. From this, it is evident that complaints have the highest past tense usage compared to the other three cate-

Nouns	Mean	Mode	Median	Adjectives	Mean	Mode	Median	Noun + Adjective	Mean	Mode	Median
NegativeOnly	1.87	1.00	1.00	NegativeOnly	0.972	1.00	1.00	NegativeOnly	0.60	0.00	0.00
Complaint	3.36	3.00	3.00	Complaint	1.500	1.00	1.00	Complaint	1.00	0.00	1.00
Difference	+1.49	+2.00	+2.00	Difference	+0.528	0.00	0.00	Difference	+0.40	0.00	0.00
PositiveOnly	2.16	1.00	2.00	PositiveOnly	1.164	1.00	1.00	PositiveOnly	0.75	0.00	1.00
Praise	3.56	3.00	3.00	Praise	2.086	2.00	2.00	Praise	1.50	1.00	1.00
Difference	+1.40	+2.00	+1.00	Difference	+0.922	+1.00	+1.00	Difference	+0.75	+1.00	0.00

Table 3: Mean, mode and median of nouns and adjectives computed on a per sentence basis.

PositiveOnly	Praise	NegativeOnly	Complaint
0.55	0.68	0.63	1.26
Category	Top 6 past tense with counts		
PositiveOnly	was 56, had 23, been 14, were 13, loved 8, tried 8		
Praise	was 86, were 32, had 22, made 9, been 7, got 6		
NegativeOnly	was 61, had 24, did 19, got 11, were 11, been 9,		
Complaint	was 155, had 40, were 35, did 23, got 20, made 14		

Table 4: Average # of past tense words and top 6 past tense words in each category

gories. On average, every complaint sentence uses at least 1 past tense. As we investigated further, we found that this is related to our observation from the previous section where within a complaint, a user is often explaining away why something was bad and what their actions were in response to the situation, which is usually something in the past.

As we looked into the actual past tense words used, we noticed that there is no significant difference in the top past tense words used across categories which can be seen from Table 4. What remains evident is that the complaint category uses more of these words than any of the other categories.

8 Negation Analysis

When we want to say that something is not true or is not the case, we use negative words, phrases or clauses. Negation can happen in a number of ways, most commonly, when we use a negative word such as *no*, *not*, *never*, *none*, *nobody*, etc. In sentiment analysis tasks, negation words are typically associated with the negative category. However, since we are interested in a finer granularity of the negative class (i.e. negative only and complaint) as well as the positive class (positive only and praise), we try to get an understanding of how negations are used across the 4 categories. For this analysis, we used a list of common negation words (e.g. not, no, never) along with words that end

Category	Avg # Negations	%Sentences with Negations	Top 6 Negation Words
PositiveOnly	0.094	8.8%	not 16, no 9, don't 6, can't 5, never 4, didn't 3
Praise	0.208	18.6%	not 46, no 15, don't 13, didn't 8, wouldn't 6, can't 4
NegativeOnly	0.538	48.8%	not 131, don't 33, never 27, no 16, didn't 13, nothing 9
Complaint	0.428	37.8%	not 88, no 36, didn't 28, don't 17, never 9, wasn't 9

Table 5: Distribution of negations and top 6 negation words per category.

with “n’t” (e.g. doesn’t, hasn’t, haven’t) and did a count of these words to determine percentage of sentences containing negations as well as the average number of negations per sentence. We also noted the top negation words in each category and the results are reported in Table 5.

Based on Table 5, we see that while positive only sentences rarely use negations, the praise sentences use negations to a certain extent (~20% of sentences contain negations). Manual inspection revealed that negations were primarily used to describe a positive aspect of a topic. For example, consider the negation in the following sentence: *this lasts all night and feels really great on my skin not oily cakey or heavy*. The negation here is used to describe the fact that the product does not feel bad on the skin.

Another interesting finding is that the NegativeOnly category has the highest use of negations with almost 50% of the sentences containing at least one negation word. This number is even higher than the complaint category where the sentences are generally longer. Through visual inspection, we found that this happens because negative only sentences have limited description on why something is ‘not good’. Therefore, clear indication of *disapproval* is with the use of negations. The following sentences from our dataset

Category	% sentences containing intensifiers
PositiveOnly	14.80%
Praise	20.60%
NegativeOnly	13.40%
Complaint	16.80%

Table 6: Intensifier usage across categories.

are negative only sentences with clear indication of dislike with the use of negations:

“*i would not recommend dinner here at all*”

“*never going back*”

“*don't stay there*”

“*for the price it just wasn't worth it*”

9 Intensifier Usage

Intensifiers are words that strengthen the meaning of other expressions and show emphasis. Common intensifiers include *absolutely*, *completely*, *extremely*, *highly*, *rather*, *really*, and etc. We noticed that intensifiers are heavily used in user generated opinions to emphasize appreciation or in some cases dissatisfaction. For example, to express appreciation on some restaurant service one may say ‘*The service was extremely fast and the food was super delicious!*’. The intensifiers in this sentence clearly adds strength to the user’s opinion. To understand which category of sentences are more inclined towards using intensifiers, we computed the percentage of sentences that contain at least one intensifier. We used a list of 35 intensifiers, expanding on the list published in (Ganesan and Zhai, 2012).

Based on Table 6, we can see that intensifiers are mostly used in praise sentences with almost 20% of the sentences containing at least one intensifier. The use of intensifiers is less prominent in the other 3 categories. Based on manual analysis, we found two reasons for this. First, is the fact that praise sentences tend to couple multiple positive aspects into a single sentence as pointed out in Section 6. So there is more use of intensifiers with the adjectives. The second reason stems from the fact that users tend to over emphasize positive points and state negative points more in a matter of fact fashion. The words ‘*very*’ and ‘*really*’ are the top intensifier words used across all four categories.

10 Causal Transitions

As complaint and praise sentences contain explanation of ‘why’ a particular topic is good or bad,

Category	% sentences containing all causal transitions	% sentences containing ‘because’ and ‘since’
PositiveOnly	16.60%	2.4%
Praise	23.80%	2.4%
NegativeOnly	16.40%	2.2%
Complaint	28.20%	4.6%

Table 7: Causal transitions used across categories.

there will be natural occurrences of cause and effect, known as causal transitions. For example, the sentence ‘*Seems like this phone is poorly designed as the screen keeps blurring out and the buttons keep getting stuck*’ shows clear cause and reason. Causal transition words include *because*, *since*, *therefore*, *as*, *as a result*, *for this reason*, *consequently*, *thus* and etc. We used 20 such words covering *cause and reason*, *effect and result* as well as *consequence* to understand if such words have a stronger occurrence in one category than another.

Table 7 shows percentage of sentences containing causal transitions within each category. Notice that while approximately 17% of the *positive only* and *negative only* sentences use causal transitions, there is a much stronger relationship between causal transitions and the *complaints* category with 28% of the sentences carrying causal links. This tells us that complaints tend to have more explicit description on what caused something to be bad or reaction in response to something negative. For example, within user reviews, it would be more common to see an expression such as ‘*I returned the vacuum because it was broken*’ as opposed to ‘*I love the vacuum because it works really well*’. To further understand this behavior, we looked at occurrences of strong causal expressions (i.e. sentences with ‘because’ and ‘since’) to validate that there indeed is more explicit use of causation in the complaints category. Based on Table 7, we can see that there is clearly a higher use of strong causal expressions in the complaints category compared to the other three categories.

11 Conclusion and Future Work

In this paper, we sought to understand the linguistic properties of *complaints* and *praises* which we define as an informative subset of the more general negative and positive polarity with **reasons** or **explanation** on what makes a topic neg-

ative/positive. Our study in the context of user reviews has shown several interesting findings.

We first showed that, *complaint* and *praise* sentences are in general longer and use more nouns than adjectives compared to a *positive only* or a *negative only* sentence. Even though subjective sentences are assumed to contain nouns and adjectives we now have evidence that nouns appear more frequently than adjectives in more informative subjective sentences. We also showed that *praise* sentences tend to use more adjectives and intensifiers compared to *complaint* sentences. The higher use of adjectives can be attributed to the fact that people were more likely to compliment several aspects of a topic within a single sentence as opposed to complaints, where people explain away the reason for the dislike or disapproval. Intensifiers play a more significant role in praise sentences as users tend to over emphasize positive points. Our study also shows that there is a stronger link between causation and *complaints* compared to the other categories.

In the future, we would like to test the power of some of the prominent features in our study to understand the value of these features in developing a fine-grained sentiment classifier.

References

- [Diao et al.2014] Qiming Diao, Minghui Qiu, Chao-Yuan Wu, Alexander J Smola, Jing Jiang, and Chong Wang. 2014. Jointly modeling aspects, ratings and sentiments for movie recommendation (jmars). In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 193–202. ACM.
- [Ganesan and Zhai2012] Kavita Ganesan and ChengXiang Zhai. 2012. Opinion-based entity ranking. *Information retrieval*, 15(2):116–150.
- [Hu and Liu2004] Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177. ACM.
- [Kim and Hovy2006] Soo-Min Kim and Eduard Hovy. 2006. Automatic identification of pro and con reasons in online reviews. In *Proceedings of the COLING/ACL on Main Conference Poster Sessions, COLING-ACL '06*, pages 483–490, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [Kim et al.2011] Hyun Duk Kim, Kavita Ganesan, Parikshit Sondhi, and ChengXiang Zhai. 2011. Comprehensive review of opinion summarization (opinion mining survey).
- [Lu et al.2009] Yue Lu, ChengXiang Zhai, and Neel Sundaresan. 2009. Rated aspect summarization of short comments. In *Proceedings of the 18th international conference on World wide web*, pages 131–140. ACM.
- [Pang and Lee2008] Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. *Foundations and trends in information retrieval*, 2(1-2):1–135.
- [Pang et al.2002] Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up?: sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10*, pages 79–86. Association for Computational Linguistics.
- [Pontiki and Manandhar2014] Maria Pontiki and Suresh Manandhar. 2014. Semeval-2014 task 4: Aspect based sentiment analysis.
- [Pontiki et al.2015] Maria Pontiki, Dimitrios Galanis, Haris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. Semeval-2015 task 12: Aspect based sentiment analysis.
- [Snyder and Barzilay2007] Benjamin Snyder and Regina Barzilay. 2007. Multiple aspect ranking using the good grief algorithm.
- [Titov and McDonald2008] Ivan Titov and Ryan T McDonald. 2008. A joint model of text and aspect ratings for sentiment summarization. Citeseer.
- [Wang et al.2010] Hongning Wang, Yue Lu, and Chengxiang Zhai. 2010. Latent aspect rating analysis on review text data: a rating regression approach. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 783–792. ACM.
- [Wang et al.2015] Linlin Wang, Kang Liu, Zhu Cao, Jun Zhao, and Gerard de Melo. 2015. Sentiment-aspect extraction based on restricted boltzmann machines. In *Proceedings of ACL 2015*.
- [Wang2015] Hao Wang. 2015. *Sentiment-aligned Topic Models for Product Aspect Rating Prediction*. Ph.D. thesis, Applied Sciences: School of Computing Science.
- [Wilson et al.2005] Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the conference on human language technology and empirical methods in natural language processing*, pages 347–354. Association for Computational Linguistics.