ICED13
INTERNATIONAL CONFERENCE ON ENGINEERING DESIGN, ICED13
19-22 AUGUST 2013, SUNGKYUNKWAN UNIVERSITY, SEOUL, KOREA

CONSUMER PREFERENCE ESTIMATION FROM TWITTER CLASSIFICATION: VALIDATION AND UNCERTAINTY ANALYSIS

Thomas Michael STONE, Seung-Kyum CHOI
Georgia Institute of Technology, United States of America

ABSTRACT

In recent years, the membership and activity of Twitter, Facebook, blogs, and other user-generate content sites has experience significant growth. Users express their opinions regarding a wide range of topics, including consumer products and services. Thus, these sites have the potential to facilitate product design via the extraction of consumer opinion and sentiment regarding product features. A key challenge is how to appropriately extract consumer preferences from the messages. This challenge is addressed with respect to Twitter using a smartphone case study. Twitter messages regarding particular smartphone attributes are classified according to sentiment: positive, negative, or neutral. This sentiment information is then used to develop an estimate of consumer preference for particular smartphone attributes, such as battery life or screen size. Uncertainty analysis is conducted in order to assess the effects of sentiment classification accuracy. Validation techniques indicate that a revised framework would be useful for predicting consumer decisions and facilitating product design; however refinement in terms of comprehensiveness and accuracy or needed.

Keywords: decision making, new product development, design theory, sentiment classification, consumer behavior

Contact:
Thomas Michael Stone
Georgia Institute of Technology
George W. Woodruff School of Mechanical Engineering
Rincon
31326
United States of America
tstone7@gatech.edu

1 INTRODUCTION

Companies that develop consumer products must make significant efforts to understand the desires and opinions of their consumers. For decades, the theory of conjoint analysis and various other preference modeling techniques have been continually developed for this purpose (Green and Rao, 1971). The consideration of market structures and metrics is a prevalent topic in engineering design, especially since uncertainty in market metrics is relevant for robust product design (Besharati et al., 2006). Shiau and Michalek (2009) demonstrate how features of market systems are important for maximizing profit by incorporating consumer preference models in the product design process. In conjoint models, the consumers' utility for various levels of product attributes are estimated based on their ranking/scoring of hypothetical products in a conjoint survey (Akaike, 1973). The utility value, or part-worth, for each attribute level is then added together to determine the overall utility of a potential product. Through rigorous investigation and development, conjoint theory has proven itself invaluable in consumer product design. However, it is not currently capable of leveraging the rapidly growing online content. Moreover, it can be expensive in terms of both time and money as developers must find willing participants to complete surveys/questionnaires; whereas, the online content is freely available and continuously updated.

Netzer et al. (2009) note that while conjoint analysis may be considered a mature field, relatively new methods of preference measurement are open for exploration. Alternative sources of consumer opinion data include sites with product reviews/feedback such as Amazon.com and Epinions.com, blogs and news sites, micro-blogs such as Twitter.com, social networking sites such as Facebook.com, and other online forums. Twitter is one of the user-generated content mediums that remains largely unexplored with respect to rigorous econometrics or preference modeling. On the other hand, the information extracted from product reviews has already been shown to correlate with consumer preferences obtained via conjoint analysis (Decker and Trusov, 2010). Since thousands of product reviews are freely available for analysis, the venture could possibly provide a profitable alternative to the traditional conjoint techniques. Twitter also produces a large amount of user-generated content, and many of the analysis techniques for product reviews can be ported over to Twitter analysis.

1.1 Preference estimation from product reviews

Preference modeling based on product reviews has been a recent area of interest for design and market researchers. Decker and Trusov (2010) used over 20,000 product reviews to determine preference data for mobile phones. The attributes and brand names of the products were analyzed with respect to their effect on consumers' decisions. Archak et al. (2007) used a wide range of products in their study: 242 products from the 'Audio and Video' and 'Camera and Photo' categories at Amazon.com, analyzed over a 15-month timespan. By also incorporating the sales data from Amazon.com, hedonic regression was used to determine the utility that customers implicitly assign to each product attribute. In a similar work of analyzing pricing premiums among sellers on Amazon.com, descriptive phrases in the product feedback messages were each assigned a dollar value–thereby demonstrating how the consumers' feedback is related to the pricing power of sellers (Ghose et al., 2007).

1.2 Sales forecasting from reviews

Research using product reviews also seeks to directly correlate review contents with product sales/revenue. Li and Hitt (2008) investigated the variation in sales caused by review bias, which is when consumers see only the most recently posted reviews or some other subset of all reviews. Results showed that book sales were correlated with the consumers' review ratings, and the positively-biased reviews occurring immediately after the product release have a significant influence on product sales. Dellarocas et al. (2007) used traditional methods as well as online content to forecast the success of movies during opening week; incorporating the online content significantly increased the forecast accuracy. The influence of user-generated content on movie sales has been shown via various methodologies (Asur and Huberman, 2010; Pang and Lee, 2008; Liu, 2006). Similar studies exist in the context of music (Dhar and Chang, 2009) and TV ratings (Godes and Mayzlin, 2004).

1.3 Using Twitter messages in sentiment analysis

While the econometrics of reviews and product feedback have been explored considerably, the preference modeling benefits of Twitter messages has not. This may be due to the fact that Twitter is relatively new, increasing from zero to over 500 million users in less than a decade. Membership and user activity is increasing with the support of the current smartphone boom, resulting in as many as 400 million messages created in a single day. In the United States, over 50% of the population are Twitter users (Sysmos, 2010). Opinion and sentiment in Twitter messages is already being applied to political issues; the Twitter Political Index tracked the favorability of the US presidential candidates in 2012 (https://election.twitter.com/). Twitter messages were also used to produce a popular graphic depicting the changing mood of United States residents. Sentiment extracted from Twitter messages enabled the visualization of peoples' emotion in different regions of the US at various times during the day and week (Mislove et al., 2010). The fact that information concerning the opinions of Twitter users can be extracted from their messages has already been sufficiently demonstrated. Thus, leveraging Twitter information in the context of product design has the potential to be profitable for design companies.

One key challenge regarding the use of Twitter messages in preference modeling is the fact that the messages are not automatically associated with the author's sentiment. In the case of product reviews, the user is oftentimes required to designate comments as pros/cons and assign an overall rating to the product. This is helpful in trying to correlate users' comments regarding product features with their sentiment (e.g., positive/negative/neutral). This challenge requires some sort of classification algorithm, since classifying tens of thousands of messages manually is not practical. Sentiment classification has been successfully implemented in econometric studies based on product reviews where the pros and cons are not already labeled (Dave et al., 2003). Go et al. (2009) observed that many classification procedures achieve accuracies of around 80% when classifying Twitter messages as positive or negative in sentiment.

This research is an instance of a larger effort to explore and address various challenges associated with preference modeling via Twitter messages—ultimately developing a framework for leveraging the Twitter information in product design. The specific challenges addressed in this paper are the consideration of uncertainty in preference modeling and potential methods for validating preference models. A smartphone case study is used to demonstrate the uncertainty analysis and validation.

2 CONJOINT DESIGN OF QUERY TERMS AND DATA COLLECTION

For the smartphone case study, the attributes and attribute levels shown in Table 1 were selected—seeing as they are the commonly used attributes in smartphone conjoint studies. This is by no means a complete list of smartphone attributes or levels; attributes related to the smartphone brand, price, processor speed, etc., could have also been considered. However, the selection of the attributes in Table 1 was not arbitrary. The operating system is thought to be a significant driver in smartphone purchase decisions, as Android and iOS competitively compete for market share. This has also been a dynamic smartphone feature, seeing as multiple operating systems became obsolete in recent years. Developers are hopeful that the new Windows 8 OS can garner support among phone users. The screen size attribute has also experienced recent changes, seeing as the Note II has a 5.5 in. screen—38% larger than the popular iPhone 4. The iPhone had its screen size adjusted up to 4.3 in. with the iPhone 5 model, and many other smartphone designers followed suit—including a recently announced 6.1 in. smartphone (Rodriguez, 2013). Prompt identification of these shifting consumer preference trends has the potential to be a profitable endeavor for smartphone designers. Talk time, memory, and camera quality are also influential attributes that designers take into consideration, seeing as these attributes can severely affect the manufacturing cost, weight, and size of the final product.

Twitter messages were collected for each attribute level by using appropriate query terms. For example, a Twitter message that contains 'iPhone 4' and 'screen size' is considered to be a message related to the 4 in. screen size. Messages which contain query terms for more than one attribute level are omitted, since the subject of those messages is more ambiguous. Details regarding the search strategy are included in (Stone and Choi, 2013). The list of attributes and levels are manually created based on the designer's knowledge of the smartphone market. A more automated process is desired, but the current process involves iterative improvements to the attributes/levels based on the designer's assessment. As time passes, new attributes may need to be incorporated to represent new smartphone features. It is worth noting that the proposed framework is only applicable for a limited number of consumer products, namely, those which are actively discusses on Twitter. For example, a study regarding washing machines may not produce enough Twitter results to facilitate preference modeling. However, this may change as consumers increase their activity on social media.

Table 1. Attributes and levels considered in Twitter query

Attribute	Level									
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Screen Size (in.)	3.5	3.7	4	4.3	4.52	4.65	4.7	4.8	5	5.5
Talk Time (min)	200-300	301-500	501-600	601-700	701-800	801-1000	1001-1300	>1300		
Camera (megapixels)	3	5	8	16						
Memory (GB)	8	16	32	64						
Operating System	Android	iOS	Win							

Pre-processing Twitter messages

Over 7,000 Twitter messages were collected at from February 27, 2013 to March 22, 2013—not counting duplicate messages which were omitted. Since the classification algorithm considers each unique word as a feature of the message, significant feature reduction can be attained by performing simple pre-processing tasks. URLs, e-mail addresses, usernames, numbers, and prices are normalized; e.g., the username '@johndoe' is converted to 'username'. Also, only the 1,200 most frequent unigrams are used to identify the Twitter messages. Thus, each message is described by 1,200 features indicating the absence or presence of the selected unigrams.

3 SENTIMENT CLASSIFICATION

Algorithms used for sentiment classification typically need a large amount of training data—in this case, messages which have been labeled as positive/negative/neutral. Automatic labeling techniques are popular for easily obtaining a large amount of training data; however, these techniques typically only provide for two possible sentiment classes, positive and negative. Many of the Twitter messages are in fact neutral with respect to the subject of the message. Some sentiment analysis methodologies use a two-step system, where the message is first classified for subjectivity and then classified as positive or negative (if the message is first classified as subjective) (Barbosa and Feng, 2010). For this research, over 800 messages are labeled manually as positive/negative/neutral; then, a 3-class Support Vector Machine (SVM) algorithm is used to classify the remaining messages. Details on the classification method are included in (Stone and Choi, 2013). The results of the 3-class classification method are shown in Table 2.

Table 2. Classified Twitter messages for each attribute level

Level	Number of Messages (Pos,Neg,Neutral) Screen Size	Talk Time	Memory	OS
1	36, 48, 22	1, 3, 40	2, 1, 10	9, 9, 36
2	0, 0, 0	207, 640, 803	85,47,388	16, 22, 64
3	103, 92, 150	45, 83, 115	4,2,47	9, 4, 30
4	8, 6, 3	12, 21, 18	0,0,0	--
5	0, 0, 0	5, 7, 6	--	--
6	1, 0, 1	131, 118, 123	--	--
7	4, 2, 8	8, 8, 5	--	--
8	5, 5, 4	105, 19, 40	--	--
9	3, 1, 0	--	--	--
10	95, 43, 55	--	--	--

Many of the messages correspond to a select few of the attribute levels, such as the 4 in. screen size and the 16 GB memory. This is expected due to the popularity of those particular attribute levels in smartphone design; the 4 in. screen size is used in the iPhone 3GS, iPhone 4, and iPhone 4S. However, the low number of messages collected for other attribute levels limits the analysis possibilities for those particular attribute levels. Thus, attribute levels without at least 10 messages were not considered in the final analysis. Figure 1 shows the number of classified messages for each attribute, including all levels. The relatively high number of messages for talk time and the camera *may* indicate that these attributes are more important to consumers than memory, screen size, or operating system when considering a smartphone purchase. Relative attribute importance is a vital component of conjoint studies, and the classification results suggest that this feature may be possible to extract from Twitter messages.

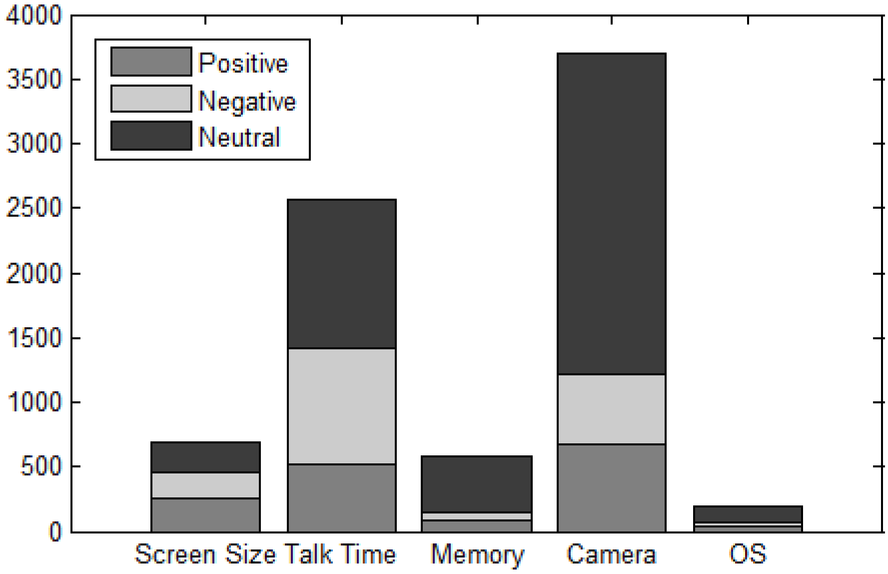

Figure 1. Classified message total for each attribute

4 PREFERENCE MODELING

The sentiment classification results suggest that some key features of consumer preference may be contained in the collection of Twitter messages, including relative importance of attributes, positive or negative sentiment regarding particular attribute levels, and also uncertainty related to the reliability of the classification algorithm and number of messages classified. The relative importance and sentiment are addressed in this section.

Let i=1,…,I indicate the attribute, and j=1,…,J_i indicate the attribute levels for attribute i. Pos_{ij}, Neg_{ij}, and $Neut_{ij}$ are the number of messages corresponding to attribute i and level j classified as positive, negative, and neutral, respectively. $Sent_{ij}$ in Eq (1) is an estimate of the consumer sentiment with respect to attribute i at level j.

$$Sent_{ij} = \frac{Pos_{ij}-Neg_{ij}}{Pos_{ij}+Neg_{ij}+Neut_{ij}} \sum_{j=1}^{J_i} \left[Pos_{ij}+Neg_{ij} + Neut_{ij}\right] \quad (1)$$

The first term (both the numerator and denominator) represents the overall polarity of the messages, while the second term (summation) accounts for the importance of each attribute. Before implementation in preference modeling, the $Sent_{ij}$ values are normalized such that they range from -1 to 1. The preference for a product is then estimated by summing all of the normalized sentiment values of its attributes.

Figures 2-6 show the sentiment values for each attribute and attribute level. Some individual attribute levels were removed from consideration because they had no more than 10 messages. In Figure 2, the talk time sentiment shows a strong upward trend as talk time increases, indicating that this attribute is important to consumers. The screen size sentiment in Figure 3 has a slight upward trend as the screen size gets larger. Consumer preference for screen size has been dynamic in recent years. The Note II, which has a 5.5 in. screen has been well-received by a section of the market, and other companies have followed suit—even though this decreases talk time, increases weight, and increases

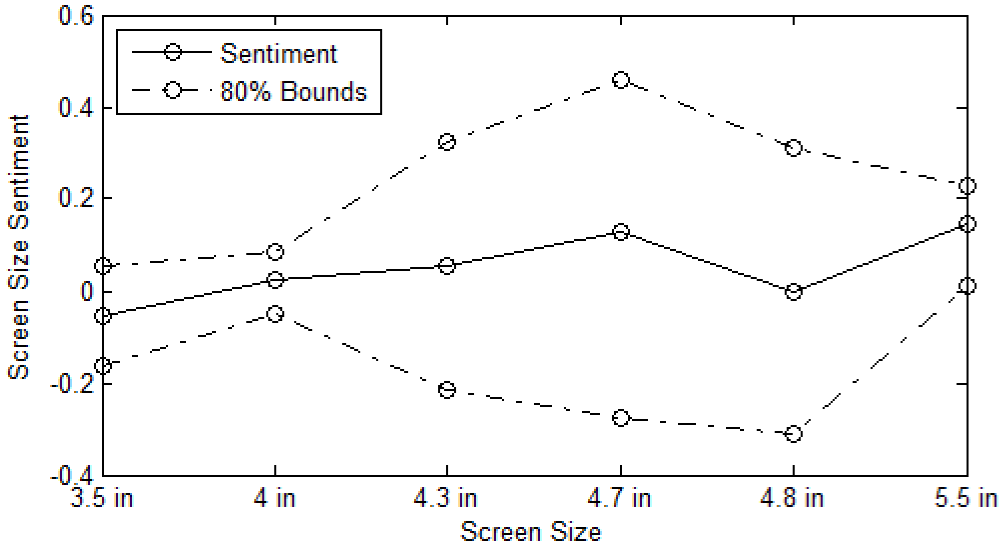

Figure 2. Sentiment values for screen size

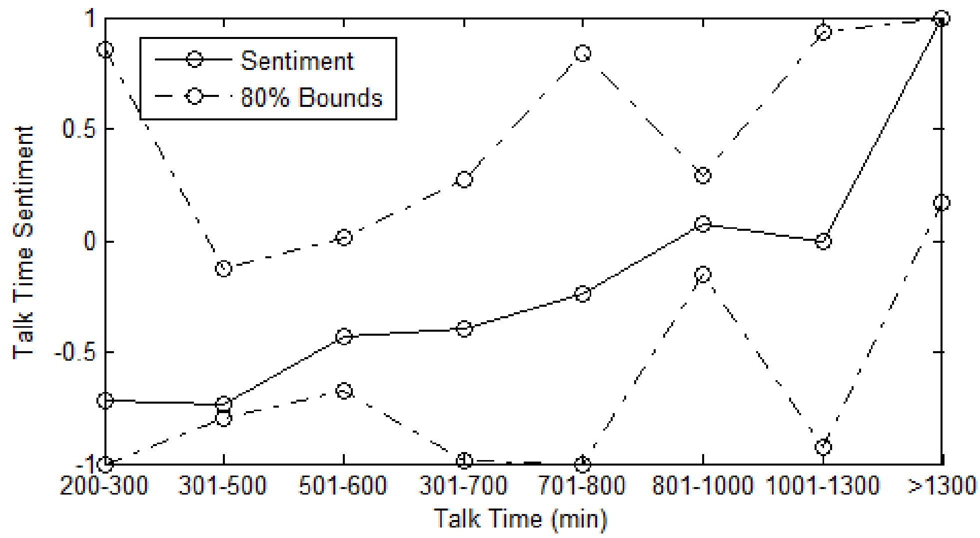

Figure 3. Sentiment values for talk time

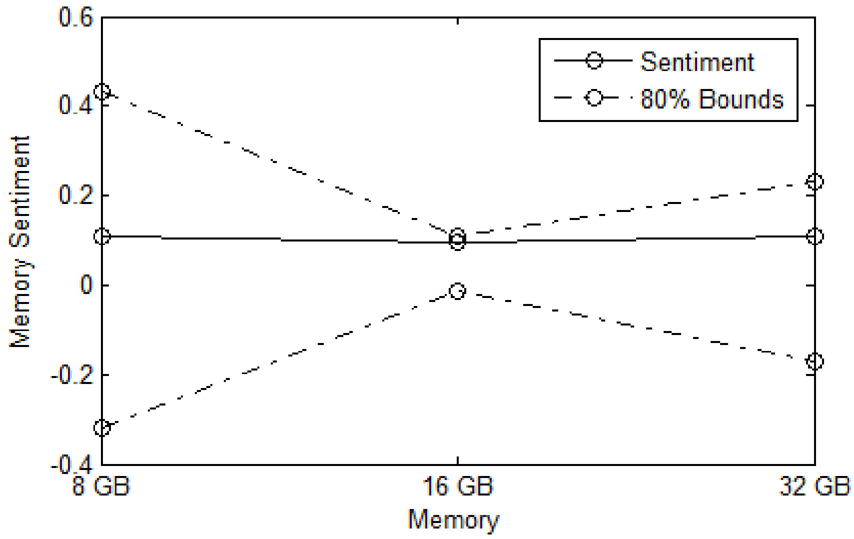

Figure 4. Sentiment values for memory

manufacturing costs. The screen size sentiment is something that designers need to attentively follow in order to optimize profits. The sentiment for memory and operating system (Figures 4 and 6) demonstrate relatively small changes in sentiment across levels. This indicates that consumers are not as concerned about these attributes when making purchases, relative to screen size. In Figure 5, the camera sentiment demonstrates a strong upward trend as megapixels increase.

For intermediate attribute levels that are not included in the figures (e.g., 5.0 in), linear interpolation is used to estimate the corresponding sentiment. The 80% bounds are determined according to the process in Section 5, and are helpful for assessing the relationship between the number of messages collected and the certainty of the sentiment values. As expected, the bounds 'widen' for levels with small numbers of collected messages, indicating limited confidence in those sentiment values.

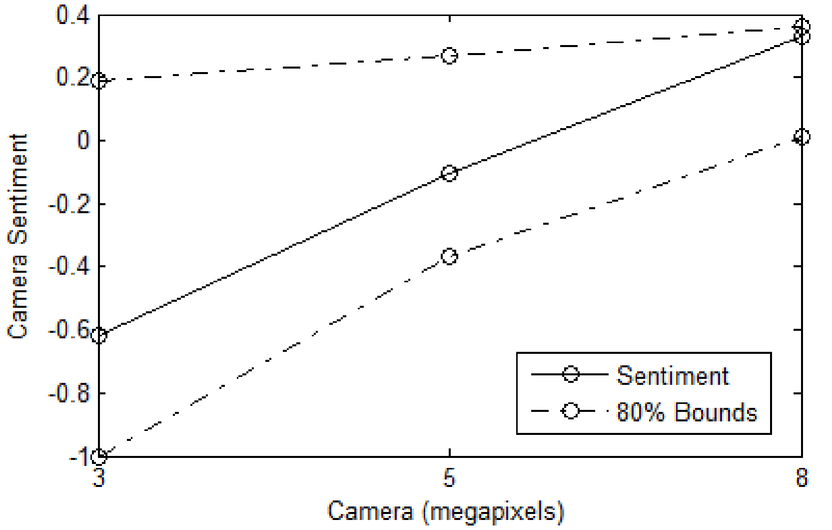

Figure 5. Sentiment values for camera

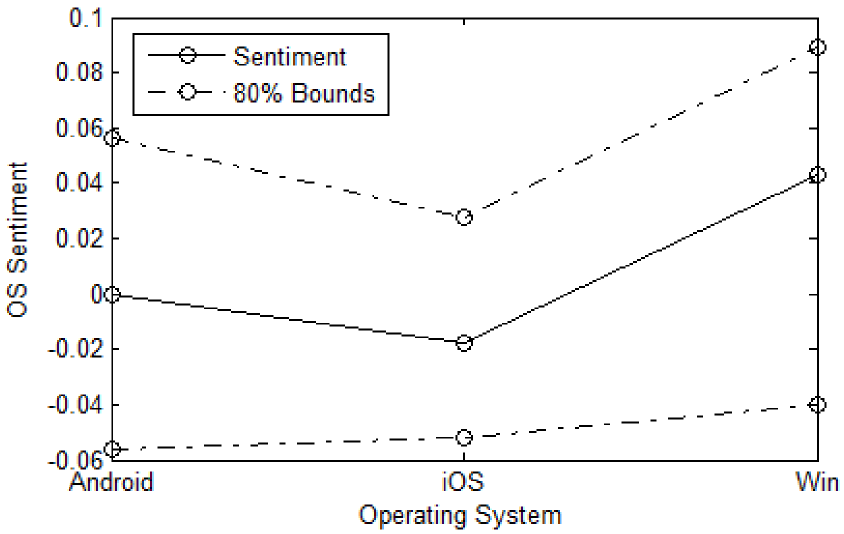

Figure 6. Sentiment values for OS

5 UNCERTAINTY ANALYSIS

The reliability of the sentiment classification algorithm is 66.7%, as determined by 5-fold cross-validation. In other words, the classifier correctly classifies the sentiment of a message with 0.667 probability, which is not abnormal for sentiment classification with three classes. This is a major factor in the uncertainty of the sentiment values, as well as the limited number of messages for some of the attribute levels.

The bounds on the sentiment values are determined as follows. Each of the approximately 7,000 message classifications is treated as a discrete random variable, with 0.667 probability of having the class that was originally assigned by the classification algorithm. Then, there is a 0.33 probability that the message actually belongs to another sentiment class. It is assumed that this 0.33 probability can be divided equally amongst the two remaining classes. For example, if a message is classified as positive, then there is a 0.135 probability that the message is actually negative, and a 0.135 probability that the message is actually neutral.

A Monte Carlo Simulation is performed with the discrete random variables according to Eq (1). 100,000 samples are simulated; thus, 100,000 values of $Sent_{ij}$ are calculated for each attribute i and attribute level j. The 80% bounds are found by calculating the 0.1 quantile and 0.9 quantile, which are treated as the lower and upper bounds, respectively. The 100,000 simulated $Sent_{ij}$ values are first sorted and then assigned to the (0.5/n), (1.5/n), …, ([n-0.5]/n) quantiles, where n corresponds to the n^{th} simulation result and the n^{th} simulated value of $Sent_{ij}$. Linear interpolation is used to determine the values of the 0.1 quantile and 0.9 quantile.

6 VALIDATION

Validation is perhaps the most critical feature of this research field moving forward. There is an abundance of online content and a virtually infinite number of ways to extract information and interpret data. Validation ensures that a preference model accurately predicts consumer decisions. If the framework produces a valid preference model, then the designer has access to a continuously updated design tool for as long as the online content continues to grow and the framework remains valid for that content. For the proposed preference model, a comparison to traditional conjoint studies could be made; however, not all of the attributes used in typical conjoint studies are present in this Twitter message analysis. Another method is to compare the model to actual sales numbers, which is accomplished by using the Amazon.com best-seller lists.

The best-seller list for ‘Cell Phones with Service Plans’ was used to approximate the sales rank for various smartphones (Amazon, 2013a). Of course, the best-seller list is not intended to be used as an econometric tool, so it does not provide perfect sales-rank accuracy. For instance, iPhone products did not appear in the top 25 phones on this particular list even though the product’s sales are high. An explanation is that consumers may simply choose to make iPhone purchases at other venders. However, the list is still a decent representation of US consumer purchases, and it is also updated continuously, so it captures current trends. In order to double-check the results, the best-seller list for ‘Unlocked Cell Phones’ was also used (Amazon, 2013b). The total sentiment of each smartphone on the list is calculated by simply adding all of the sentiment values corresponding to the products’ attributes.

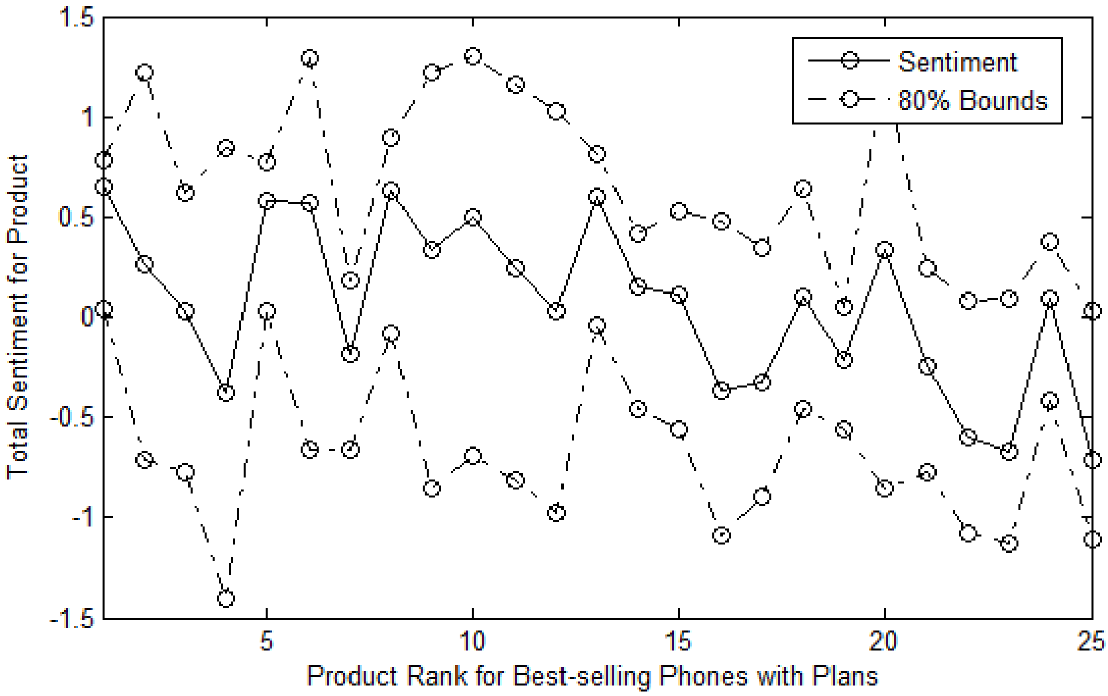

Figure 7. Sentiment for products in best-seller list for phones with plans (Amazon, 2013a)

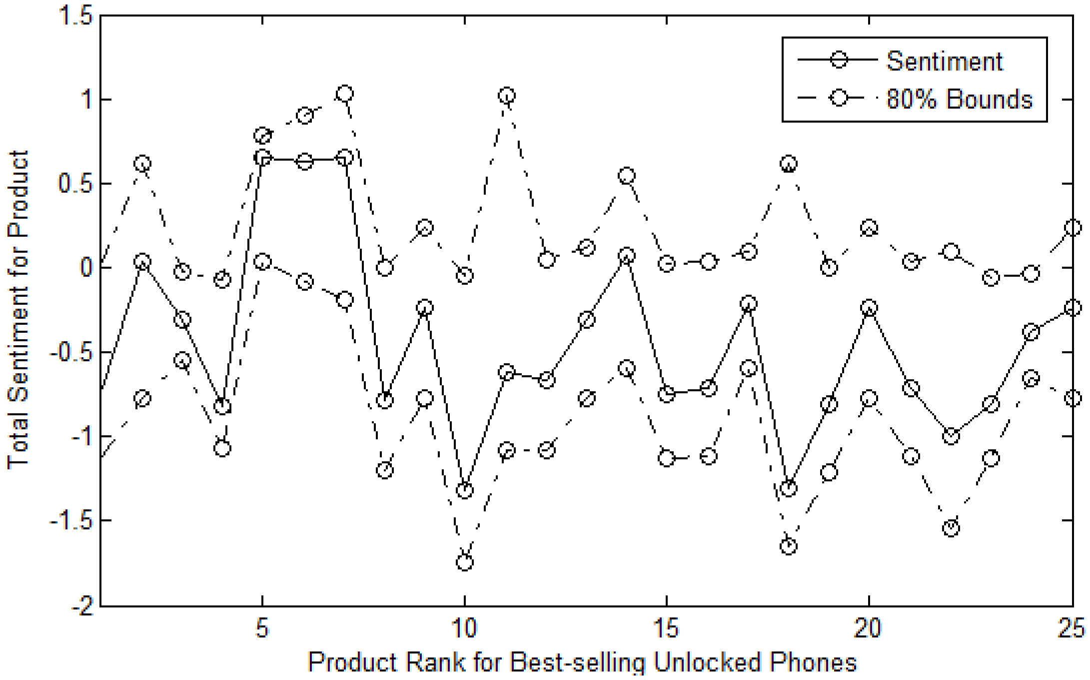

Figure 8. Sentiment for products in best-seller list for unlocked phones (Amazon, 2013b)

Figures 7 and 8 show the calculated sentiment of each product and its rank on the best-seller list. Non-smartphone products were omitted, along with similar products, such as the same phone in a different color. The sentiment for the top 25 smartphones with service plans (Figure 7) shows a strong downward trend as product rank increases (worsens). This is encouraging because it indicates that the sentiment values are useful for estimating the preferences and purchasing decisions of consumers. However, the sentiment for unlocked smartphones (Figure 8) only shows a slight downward trend. Perhaps, the unlocked smartphones are better-described by another set of attributes and levels. For example, many of the unlocked smartphones have much lower talk time, relatively, so more resolution may be needed for talk time less than 501 minutes. Official phone sales numbers would be most helpful in identifying the root causes of discrepancies. There may be problems with regards to using the Amazon.com data, instead of actual sales data.

7 CONCLUSIONS AND FUTURE WORK

The preference model derived from Twitter messages demonstrated some agreement with actual consumer behavior. Even though there is currently no rigorous, conclusive validation of the proposed preference model, the results are promising and call for the development of more comprehensive validation techniques and subsequent improvements to the framework. Namely, the message collection timespan needs to be varied, and the number of smartphone attributes should be increased. In this study, important attributes such as price, brand, and processor speed are not considered. Moreover, the reliability of the sentiment classification has the potential to improve with more training data and different classification algorithms. Many other validation techniques could be implemented, such as comparisons with traditional conjoint analyses. Also, including past Twitter messages from a Twitter database would allow for validation of the preference modeling framework over past product cycles (from which accurate sales data is available).

Overall, the proposed preference modeling method based on Twitter data would be useful for designers only after significant refinement of the framework. Also, the Twitter preference results could be combined with product review analysis—which is already being investigated by many researchers—in order to provide a more comprehensive analysis of the online Voice of the Consumer and consumer preference. The product review analysis will probably advance in reliability with regards to preference modeling well before the proposed Twitter analysis. The current preference results derived from Twitter should be viewed as a secondary source for decision-making, with traditional conjoint studies and market analysis taking precedence. However, the remarkable growth of online content calls for immediate investigation of preference extraction from user-generated, online content.

REFERENCES

Akaike, H. (1973) 'Information theory and an extension of the maximum likelihood principle', In *Second International Symposium on Information Theory,* Tsahkadsor, Armenian SSR, pp. 267-281.

Amazon (2013a), 'Best Sellers in Cell Phones With Service Plans,' [online], http://www.amazon.com/Best-Sellers-Cell-Phones-Accessories-Service-Plans/zgbs/wireless/2407747011/ref=zg_bs_tab_t_bs (accessed May 2013).

Amazon (2013b), 'Best Sellers in Unlocked Cell Phones,' [online], http://www.amazon.com/Best-Sellers-Cell-Phones-Accessories-Unlocked/zgbs/wireless/2407749011/ref=zg_bs_nav_cps_1_cps (accessed May 2013).

Archak, N., Ghose, A. and Ipeirotis, P. G. (2007) 'Show me the Money! Deriving the pricing power of product features by mining consumer reviews', Proceedings of the *13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,* San Jose, pp. 56−65.

Besharati, B., Luo, L., Azarm, S. and Kannan, P. K. (2006) 'Multi-Objective Single Product Robust Optimization: An Integrated Design and Marketing Approach', *ASME J. Mech. Des*., Vol. 128, No. 4, pp. 884–892.

comScore (2007) [online], http://www.comscore.com/Insights/Press_Releases/2007/11/Online_Consumer_Reviews_Impact_Offline_Purchasing_Behavior (accessed August 2008).

Dave, K., Lawrence, S. and Pennock, D. M. (2003) 'Mining the peanut gallery: Opinion extraction and semantic classification of product reviews', Proceedings of the *12th International World Wide Web Conference*, Budapest, pp. 519−528.

Decker, R. and M. Trusov (2010) 'Estimating aggregate consumer preferences from online product reviews', Vol. 27, No. 4, pp. 293–307.

Dellarocas, C., Zhang, X., and Awad, N. F. (2007) 'Exploring the value of online product reviews in forecasting sales: The case of motion pictures', *Journal of Interactive Marketing,* Vol. 21, No. 4, pp. 23−45.

Dhar, V. and Chang, E. A. (2009) 'Does chatter matter? The impact of user-generated content on music sales', *Journal of Interactive Marketing*, Vol. 23, No. 4, pp. 300−307.

Ghose, A., Ipeirotis, P. G. and Sundararajan, A. (2007) 'Opinion mining using econometrics: A case study on reputation systems', Proceedings of the *45th Annual Meeting of the Association of Computational Linguistics*, Prague, pp. 416−423.

Go, A., Bhayani, R. and Huang, L. (2009) 'Twitter sentiment classification using distant supervision', *CS224N Project Report*, Stanford Digital Library Technologies Project, pp. 1-12.

Godes, D. and Mayzlin, D., (2004) 'Using online conversations to study word-of-mouth communications', *Marketing Science*, Vol. 23, No. 4, pp. 545-560.

Green, P. E. and Rao, V. R. (1971) 'Conjoint measurement for quantifying judgmental data'. *Journal of Marketing Research*, Vol. 8, pp. 355–363.

Li, X. and Hitt, L. M. (2008) 'Self selection and information role of online product reviews', *Information Systems Research*, Vol. 19, No. 4, pp. 456−474.

Liu, Y. (2006). 'Word-of-Mouth for Movies: Its Dynamics and Impact on Box Office Receipts', *Journal of Marketing,* Vol. 70, pp. 74-89.

Mislove, A., Lehmann, S., Ahn, Y., Onnela, J. and Rosenquist, J. N. (2010) *Pulse of the Nation: U.S. Mood Throughout the Day inferred from Twitter* [online], http://www.ccs.neu.edu/home/amislove/twittermood/ (accessed January 2013).

Pang, B. and Lee, L. (2008) 'Opinion Mining and Sentiment Analysis', *Foundations and Trends in Information Retrieval*, Vol. 2, Nos. 1-2, pp. 1-111.

Rodriguez, S. (2013) *Los Angeles Times* [online], http://www.latimes.com/business/technology/la-fi-tn-ces-huawei-ascend-mate-d2-20130107,0,6986891.story (accessed January 2013).

Shiau, C. and Michalek, J. (2008) 'Should Designers Worry About Market Systems?', *ASME J. Mech. Des.*, Vol. 131, No. 1, pp. 1–9.

Stone, T. M. and Choi SK., (2013) 'Extracting Consumer Preference from User-generated Content Sources using Classification,' *ASME International Design Engineering Technical Conferences & Computers and Information in Engineering Conference,* Portland, OR, August 4-7, 2013.

Sysmos Inc. (2010) *Exploring the Use of Twitter Around the World* [online], http://www.sysomos.com/insidetwitter/geography/ (accessed March 2012)