Hit Songs' Sentiments Harness Public Mood & Predict Stock Market

Rachel Harsley¹, Bhavesh Gupta², Barbara Di Eugenio¹, and Huayi Li¹

¹University of Illinois at Chicago, Chicago, IL, USA

²Cerner Corporation, Kansas City, MO, USA

{rharsl2, bdieugen, hli47}@uic.edu

bhavesh.gupta@cerner.com

Abstract

This work explores the relationship between the sentiment of lyrics in Billboard Top 100 songs, stocks, and a consumer confidence index. We hypothesized that sentiment of Top 100 songs could be representative of public mood and correlate to stock market changes as well. We analyzed the sentiment for polarity and mood in terms of seven dimensions. We gathered data from 2008 to 2013 and found statistically significant correlations between lyrical sentiment polarity and DJIA closing values and between anxiety in lyrics and consumer confidence. We also found strong Granger-causal relationships involving anxiety, hope, anger, and both societal indicators. Finally, we introduced a vector autoregression model with time lag which is able to capture stock and consumer confidence indices $(R^2=.97)$, p < .001 and $R^2 = 0.72$, p < .01 respectively).

1 Introduction

Many would agree that the success of top songs is due to a complicated mix of marketing, popularity, and blockbuster theory in which companies invest big money into few products. Yet, we also hypothesized that the lyrical sentiment of top songs can be viewed as transient, but genuine snapshot of public mood. There is a plethora of research which unequivocally confirms both the influence of mood on music choice and the influence of music on mood and even buying behavior (Areni and Kim, 1993; Bruner, 1990; Chen et al., 2007; R McCraty, 1998)

While researchers have attempted to model public opinion indices and the stock market via sentiment analysis of news articles, microblogging and social media sites, no research has taken this correlation-seeking approach using popular song lyrics. We hypothesized that song lyrics from the Billboard Hot 100, a weekly listing of the top 100 songs, is representative of public mood. Moreover, we aimed to explore correlating, causal, and even predictive relationships between song lyrics, public opinion, and the stock market.

Thus, we aimed to explore the sentiment of top song lyrics in a manner similar to researchers who used Twitter to ascertain public mood and correlated this sentiment to public opinion polls (Bollen et al., 2011; O'Connor et al., 2010). While Twitter offers a fine-grained time-based approach to harnessing public expressions, other mediums, such as popular song lyrics, may offer the same insight while being less costly to obtain and less susceptible to the "burstiness" of Twitter. In other words, use of popular song lyrics would automatically filter out the noise of ephemeral, popular happenings which pervade Twitter.

The Hot 100 listing is calculated based on a single's selling performance, radio airplay audience impression, and online streaming activity (Trust, Gary, 2013). We explored if correlation would reflect in the Thomson Reuters/University of Michigan Consumer Confidence Index (ICC) and the Dow Jones Industrial Average (DJIA), a major U.S. stock market index.

For our work, we gathered the entire set of weekly Hot 100 songs between 2008 and 2013. We used OpinionFinder to analyze the positive and negative polarity of lyrics (Wilson et al., 2005). We then used a second tool, WordNet Affect, to perform sentiment analysis along ninedimensions (Strapparava and Valitutti, 2004). We assessed the strength of sentiment correlation to the DJIA and ICC. We then explored poignant Granger-causal relations and created a predictive model for both societal indicators. In this work, we discover, there are indeed correlating and causal relationships between the song sentiments and these societal indicators.

2 Related Work

Bollen et al. (2011) explored the notion that public mood can be correlated to and even predictive of economic indicators. They used sentiment analysis of large scale twitter feeds and compare it with the Dow Jones Industrial Average over time. High correlation results led them to create a neural network to predict the DJIA given their Twitter sentiment insights. They reached 87% accuracy in predicting the daily up and down changes of the DJIA. Similarly, O' Connor et al. (2010) connected measures of public opinion measured from polls with the results of sentiment analysis over text on twitter feeds. They analyzed several surveys on consumer confidence and political opinion between 2008 and 2009 and found correlation between sentiment word frequencies in twitter messages.

Acerbi et al. (2013) examined the usage of "mood" in the context of 20th century books written in English. They used WordNet Affect to perform sentiment analysis on the literature and found evidence for distinct historical periods of positive and negative moods in American Literature. Further, these periods often correlated to historical happenings.

Daas and Puts (2014) explored changes in the sentiment in Dutch public blogs and social media messages i.e. Twitter, Facebook and LinkedIn over a 3.5-year period. They performed sentiment analysis on the text and compared results with changes in Netherlands monthly consumer confidence. They discovered a high correlation (up to r=0.9) and that changes in social media sentiment precede the consumer confidence changes.

While there has been much interest in automatically determining the sentiment of songs from both acoustic and natural language processing communities, there has been far less success in performing the task. Xia et al (2008) proposed using a sentiment vector space model in conjunction with a support vector machine to overcome the hurdle of lyrical data sparseness.

Other research has focused on combining audio and lyrical data for ascertaining the mood of a given song (Hu and Downie, 2010; Zhong et al., 2012).

While past work has looked at correlations between Twitter, literature, and other social media in regards to stocks and public opinion, our work looks at the correlation between the sentiment of popular song lyrics and these societal measures. There is an abundance of research linking the effect of music on mood and social behaviors including buying decisions and even it's inverse; the role of mood in music choice(Areni and Kim, 1993; Bruner, 1990; North and Hargreaves, 1997; R McCraty, 1998; Sloboda, 2011). Due to the strong relationship between music and mood, we considered it to be a reasonable hypothesis that top music choice of the nation, via the Hot 100, could in some ways be representative of public mood.

3 Dataset

3.1 Lyrics

The first step in preparing our analysis was to collect the song listings themselves. Our dataset spans six years, 2008 through 2013. In order to perform this, we consulted the Ultimate Music Database (http://www.umdmusic.com/) which has the full listing of Billboard Music Charts songs available. We collected a total of 36,000 song listings (with some songs being repeated). For each listing, we queried and scraped LyricsWikia (http://lyrics.wikia.com/) for actual lyrics. The lyrical data along with full chart listings are available for public use by contacting the authors.

3.2 Societal Indicators

For the purpose of this research, we examined two societal indicators; Dow Jones Industrial Average (DJIA) and Thomson Reuters/University of Michigan Index of Consumer Confidence (ICC). The two were chosen due to their role in previously researched correlations as described in the Related Work section.

The DJIA shows how 30 large publicly owned companies based in the United States have traded during a standard trading session in the stock market. It is the second oldest U.S. market index and is influenced by many factors. The ICC aims to measure consumers' level of optimism towards their own financial situation, short term general economy and long term general economy, according to their share of spending and savings (Curtin, 2004).

4 Methods

In this section, we describe our approach to find correlations between the lyrical sentiment of top songs and societal indicators.

4.1 Sentiment Analysis

We chose to utilize two simplistic approaches to sentiment mining given the insight of past work showing the uniqueness of song sentiment classification. Positive results using the simplest of sentiment analysis techniques would be indicative of opportunities for fine-tuning the sentiment analysis method for enhanced results. We performed our first analysis of the lyrics using OpinionFinder. OpinionFinder performs sentiment analysis by labeling words as either positive, negative, or neutral. It generates a text file that tags the words in the document with respect to their contextual polarity. Using OpinionFinder results, we calculated the polarity for a song using the following ratio:

$$polarity = \frac{num_pos + num_neut + .1}{num_neg + num_neut + .1}$$
(1)

where *num_pos*, *num_neut*, and *num_*neg represent the number of words with positive, neutral, and negative sentiment valance. We used 0.1 to accommodate smoothing in the case of missing values. For example, we could not automatically retrieve lyrics for a song in year 2008 due to a mislabeling of artist in the lyrics data.

We performed a second sentiment analysis using multi-dimensional sentiment classification. We found this step to be necessary in order to capture holistic public mood, which is rich, multifaceted, and not limited to bipolarity. For this we use the text analysis tool WordNet Affect. WordNet Affect labels a given word with one of over 300 possible sentiment labels. Labels fall under distinct hierarchies. These hierarchies include: emotion, mood, trait, cognitive state, physical state, hedonic signal, emotion eliciting, situation, emotional response, behavior, attitude, and sensation. Our analysis exclusively considers labels branching under the emotion hierarchy.

We automatically labelled and aggregated the lyrical data on a weekly basis. Moreover, after retrieving labels, we narrowed our focus to the seven most frequently occurring given the sentiment analysis results. The seven labels are as follows: anxiety, anger, expectation, dislike, joy, negative fear, and sorrow. The ambiguous expectation sentiment can be seen as a measure of hope or fever in the lyrics. Further, *negative* *fear* is distinguished in contrast to *fear*, which may also signify reverence. Finally, we included positive emotion and negative emotion to compare to OpinionFinder's polarity results.

In order to effectively compare the range of sentiment counts and the societal indicators, each time series was normalized to their z-score using the overall mean and standard deviation. The normalization allows all time series to fluctuate around a zero mean and be expressed on the scale of a single standard deviation. The Z score of time series X is denotes Z_X is defined as:

$$Z_X = \frac{X_t - \bar{X}}{\sigma} \tag{2}$$

where \overline{X} and σ represent the mean and standard deviation of the X time series.

4.2 Correlation Analysis

We began by examining the Pearson Correlation Coefficient for the sentiment time series in comparison to the societal indicator data. Pearson's correlation coefficient (r) measures the strength of the association between two quantitative, continuous variables. The Pearson correlation analysis allowed us to quantitatively determine the relationship between the societal indicators and sentiments. Graphical plots of the time series for each lyrical sentiment method and each societal indicator afforded opportunity for visual correlation analysis and cross validation of our sentiment analysis findings to known socio-cultural events. Additionally,

We then measured the correlation between the trend obtained from our sentiment analysis and each societal indicator using multiple regression. The regression model is shown below:

$$Y_i = \alpha + \sum_{j=1}^{N} \beta_j X_j + \varepsilon_t \tag{3}$$

where i=1,2 and Y_i is the societal indicator trend and N=9, X_1 , X_2 ,... X_N represent the mood time series obtained from the multi-dimensional sentiment analysis. The regression model allows us to further quantify any relation between the lyrical sentiment and societal indicators.




Figure 1: Plot of time series data for the the DJIA and ICC societal indicators (top). Plot of the positive vs. negative polarity for song lyrics (bottom). Note, the z-scores of the societal indicators are used.

5 Results

5.1 Polarity-Based Sentiment Correlations

We quantified the correlation between the song polarity and the societal indicators using the Pearson Correlation Coefficient and t-tests (used in order to establish the significance of the correlation). As a baseline, we calculated the correlation of ICC, the standard of public mood, to the DJIA (r= .6563, p<.001). We wanted to measure how top songs' lyrical sentiments compared to this baseline as another measure of public mood. The results indicated that the lyrical polarity and DJIA have a significant negative correlation (shown in Table 1). However, the absolute value of the correlation coefficient is half that of the baseline, which indicates the association is not as strong.

The plot of lyrical polarity for each week from the year 2008 to 2013 is shown in Figure 1. We compared it to the time series of societal indicators z-scores also shown in Figure 1. There were little visual similarities between the plots. However, interestingly, we noted that the polarity plot is able to capture some of the trends of typical U.S. holidays. For example, the polarity exhibits local peaks of high positivity during the Christmas holiday time from 2010 through 2013. Similarly, the polarity reaches its lowest on Valentine's Day of 2011, suggesting a time of negative feelings. In order to gain an intuition that the sentiment analysis system was performing reasonably well and the dip was not due to inaccuracy of the system, we manually examined song changes between the two weeks. As expected, the Valentine's week Hot 100 added several

songs to the list which contained hints of negative sentiments in juxtaposition to themes of love. These included, *Loveeeeeee Song* by Rihanna and *Same Love* by Macklemore, and *As Long as You Love Me* by Justin Bieber.

	Correlation Coefficient	t	р
DJIA	-0.3679	-3.3108	0.0014*
ICC	-0.0212	-0.1777	0.8594

 Table 1: Pearson Correlation results for song polarity in comparison to DJIA and ICC.

5.2 Multi-dimensional Sentiment Correlations

In addition to polarity, we also plotted the trend series for the nine sentiments of interest as shown in Figure 2. Visual inspection of the plots suggested a high correlation between anxiety and ICC. The anxiety seems to be almost a shift of the ICC with a three-month lag. Moreover, visual analysis offered another interesting insight. The anger plot exhibits a large rise during the time of the November 2012 Presidential Election. Further, the peak of the anger occurs on the day following Election Day 2012 as shown in Figure 3. We examined the reason for this through manually comparing difference in top songs from the prior week. During election week, songs by Taylor Swift and Kendrick Lamar entered into the Hot 100 listing. They introduced a variety of sentiments including anger, for example, in Swift's song Stay Stay Stay she repeatedly uses the word *mad* to describe her emotional state in a dating relationship.




Figure 2: The trend series for the nine sentiments of interest. Z-score not shown.



Figure 3: The trend series for *anger* during weeks prior to and after the November 2012 Presidential Election.

As in the polarity analysis, we also examined the quantitative correlation between the nine sentiments and the societal indicators using the Pearson Correlation Coefficient and with t-tests in order to establish the significance of the correlation. The results of the analysis between the sentiments and ICC are shown in Table 2. The result of the comparison between sentiments and DJIA is shown in Table 3.

The analysis results indicated that anxiety is significantly correlated to the Michigan ICC (r=0.4761, p < 0.001) and the correlation is almost as strong as that between the ICC and the DJIA. Additionally, anger exhibits a correlation which is significant (p<.05) though not as strong and more difficult to ascertain upon visual inspection. Surprisingly, joy, which may be seen as the opposite of anger, did not significantly correlate with the Michigan ICS.

Sentiment	Correlation Coefficient	t	р
Anxiety	0.4761	4.5	0.000**
Anger	0.2974	2.6	0.011*
Ambiguous_ Expectation	-0.2056	-1.7	0.083
Dislike	-0.1043	-0.8	0.383
Joy	0.1494	1.2	0.210
Negative_ Emotion	0.1582	1.3	0.184
Negative_ Fear	0.1554	1.3	0.192
Positive_ Emotion	0.1967	1.6	0.097
Sorrow	-0.2233	-1.9	0.059

*p<.05 **p<.001

 Table 2: Pearson Correlation Results for WordNet Affect
 Sentiments and Michigan ICC

The correlation results for the DJIA and nine sentiments shows there is an especially strong correlation between the societal indicators and both ambiguous expectation and negative emotion (p<.001). As the DJIA values increases, the ambiguous expectation decreases. We suspected this trend may have been indicative of positive correlations with time lag. Further, both anxiety and joy exhibit strong correlation with the DJIA (p<.05). Though several correlations where significant and strong, they were not intuitive as the visual plots did not align with either of the societal indicators. Thus, we deemed it necessary to explore causal relationships and time lag as described in Section 5.3.

Sentiment	Correlation Coefficient	t	р
Anxiety	0.2504	2.1	0.033*
Anger	0.2248	1.9	0.057
Ambiguous _ Expectation	-0.4734	-4.4	0.000**
Dislike	-0.1312	-1.1	0.271
Joy	0.2582	2.2	0.028*
Negative_ Emotion	0.4328	4.0	0.000**
Negative_ Fear	0.0964	0.8	0.420
Positive_ Emotion	0.0024	0.0	0.983
Sorrow	-0.0038	0.0	0.974

*p<.05 **p<.001

 Table 3: Pearson Correlation Results for WordNet Affect

 Sentiments and DJIA

5.3 Modelling, Causality and Time-lag

In order to gain a greater understanding of the relationship between the multidimensional lyrical sentiments and the societal indicators, we performed multiple regression using the model described in Equation 3. The multiple regression models moderately captured stock and consumer confidence indices (R^2 =.61, p<.001 and R^2 =.52, p < .001). As expected, the sentiments with significant and strong Pearson correlations, anger, anxiety, ambiguous expectation, were all significant features for modelling each societal indicator. Additionally, sorrow was significant in modelling the ICC and DJIA while negative emotion was significant in modelling the DJIA. The results for ICC are shown in Table 4 while the results for DJIA are in Table 5.

We discovered several non-intuitive correlations, including the negative correlation between ambiguous expectation and the DJIA. However, this strange correlation could be indicative of positive correlation with a lag. Given this, we recognized that more fine-tuning of the model would better fit the societal indices. Further, due to the non-intuitive correlation between several sentiments and the societal indicators, we deemed it worthwhile to explore the role of time lag. With the addition of time based modelling, we also aimed to discover whether causal relationships existed among the sentiments and the societal indicators.

In order to accomplish this, we utilized the statistical concept of Granger causality in a similar manner as in Bollen's (2011) Twitterbased stock market predictions. The key intuition in Granger causality is as follows; if a variable X causes Y then changes in Y will be 1) be preceded by changes in X and 2) be better predicted by using information from time-lagged X and Y rather than information than solely Y. In effect, Granger causality then tests whether one time series has predictive information about the other by checking for statistically significant correlations between the time lagged X and the resulting Y.

	Coef- ficient	Stand- ard Error	t Stat	P-Value
Sorrow	-0.315	0.112	-2.8	0.006*
Pos-Emo	0.016	0.092	0.1	0.861
Neg-Fear	-0.156	0.135	-1.1	0.251
Neg-Emo	0.499	0.097	5.1	2.9E-06**
Joy	0.071	0.116	0.6	0.543
Dislike	0.070	0.102	0.6	0.493
Anxiety	0.114	0.119	0.9	0.339
Anger	0.515	0.141	3.6	0.0005**
Ambig-				
Expecta-	-0.623	0.104	-5.9	1.175
tion				
$P^2 = 0.5234$ n < 001 *n < 05 **n < 001				15 **n < 0.01

$R^{-}=0.5234 \text{ p}<.001$	*p<.05	**p<.001
Table 4: Multiple regression result	lts for WordN	let Affect
Sentiments vs l	ICC	

	Coeffi- cient	Stand- ard Error	t Stat	P-Value
Sorrow	-0.421	0.125	-3.3	0.001*
Pos-Emo	0.164	0.104	1.5	0.119
Neg-Fear	-0.107	0.151	-0.7	0.481
Neg-Emo	0.281	0.108	2.5	0.012
Joy	0.150	0.130	1.1	0.255
Dislike	0.270	0.115	2.3	0.022
Anxiety	0.519	0.133	3.8	0.000**
Anger	0.291	0.158	1.8	0.071
Ambig- Expecta- tion	-0.333	0.116	-2.8	0.006*
$R^2 = 0.6179 p$	o<.001	•	*p<.0	5 **p<.001

 Table 5: Multiple regression results for WordNet Affect Sentiments vs DJIA

We tested for Granger Causality given an X-month lag and found that there is significant Granger Causality between several of the sentiments and the DJIA along with the ICC. The Granger causality test rejects the null hypothesis that the ICC does not predict both anxiety and ambiguous expectation. In agreement with earlier findings, we confirmed that consumer confidence Granger-causes anxiety with a 5-month lag. By visualizing the time lag in more detail, the correlation between anxiety, ambiguous expectation and the ICC become quite clear. Moreover, we discovered that anger Granger-causes consumer confidence shifts.

The use of time lag and Granger causality brought clarity to the prior findings in which ambiguous expectation was negatively correlated to the DJIA and ICC. Both the DJIA and consumer confidence changes Granger-cause ambiguous expectation, or hope, in popular song lyrics. They require minimal time lag, especially for consumer confidence. In effect, consumers listen to more hopeful music when stocks and consumer confidence are high. The opposite can also be said. The full results for Granger causality are shown in Table 6 and Table 7.

We plotted several causal relationships with appropriate time lag in Figure 4. We note the visible alignment given the introduced time lag. For example, the time series of negative emotion and the DJIA frequently overlap or point in the same direction given that negative emotion Granger-causes the DJIA with a six month delay.

	S to I	I to S	P-Value	Months
Sorrow		Х	0.032	1
Pos-Emo		Х	0.039	16
Neg-Emo	Х		0.036	16
Joy	Х		0.074	17
Anxiety		Х	0.013	5
Anger	Х		0.048	4
Ambig- Expectation		Х	0.009	1

 Table 6: Sentiment (S) and Granger Causality ICC (I) along with amount of lag. S to I denotes that S Granger-causes I.

 While I to S denotes I Granger-causes S.

	S to D	D to S	P-Value	Months
Sorrow	Х		0.013	5
Pos-Emo	Х		0.049	1
Neg-Fear		Х	0.036	1
Neg-Emo	Х		0.035	6
Dislike	Х		0.041	17
Anger		Х	0.036	15
Ambig- Expectation		Х	0.000	3

Table 7: Sentiment (S) and Granger Causality DJIA (D) along with amount of lag. S to D denotes that S Granger-causes D. While D to S denotes D Granger-causes S.

As a final measure, we added further predictability to our analysis through the use of vector autoregression. Vector autoregression is an econometric technique that allows for multiple time series to be captured for linear interdependencies. Vector autoregression is a natural extension of univariate autoregression in which more than one variable is able to lag. We obtained tighter fit models for both the DJIA and ICC (R^2 =.97, p = 0.0 and R^2 =0.72, p<.01). We then computed the mean square error of the predictor time-series on data from 2014 using VAR. The model predicted the Z-score of either societal indicator and achieved greater success on the DJIA. The results are shown in Table 8.

	MSE
DJIA	0.109
ICC	0.748
	01110

Table 8: Mean squared error of VAR model prediction for	
2014 data of societal indicators.	

Sorrow, with a five-month lag, was the most significant of the time-lagged features for modelling the DJIA (p<.01). Unsurprisingly, sorrow also was shown to Granger cause the DJIA with a five-month lag. Additionally, ambiguous expectation with a three-month lag was significant (p<.05) in the DJIA model. The dual role of ambiguous expectation in both influencing and being influenced by the DJIA is unsurprising given the prior research in music and mood and it's inverse relationship. Similarly, ambiguous expectation (one-month lag), anxiety (one-month lag), and negative fear (two-months lag) were all significant features (p<.05) in modelling the ICC.

6 Conclusion

In this paper, we explored whether Billboard Hot 100 lyric sentiment is indicative of public mood. We measured this by comparison to correlations with the Michigan Consumer Confidence Index (ICC). Moreover, we investigated whether this type of public mood measure is causally related to both the ICC and the Dow Jones Industrial Average (DJIA). We analyzed lyrics based on positive and negative sentiment polarity. We also performed a multi-dimensional sentiment analysis that examines nine emotions. Visual analyses of trend plots showed notable correlation between song lyrics sentiment and societal indicators. Moreover, the Pearson Correlation analysis quantified the relationship between song polarity and DJIA to be statistically significant, though some correlations were non-intuitive.

We confirmed correlating relationships through use of Granger causality and the introduction of time lag. This revealed a Grangercausal relationship between ICC and between anger and ICC and anxiety. We also discovered




Figure 4: Plot of the X-month time lagged sentiments and DJIA or ICC

Granger-causal relationships between DJIA and hope. We presented both a multiple regression model and an improved VAR model that incorporate the multi-dimensional sentiment analysis components to DJIA and ICC. The model solidified our finding in regards to the role of anger in influencing the ICC and the role of sorrow in influencing the DJIA.

In future work, we plan to improve our prediction model by focusing on the features with known causal relationships. We also plan to explore other prediction models beyond VAR. We acknowledge that past work has shown that song lyrics use creative language, slang, and metaphors which typically are not accounted for by valance based sentiment analysis. We will integrate these factors in our modeling of the sentiment of the song. All lyrics data can be retrieved by contacting the authors.

References

Alberto Acerbi, Vasileios Lampos, Philip Garnett, and R. Alexander Bentley. 2013. The Expression of Emotions in 20th Century Books. *PLoS ONE*, 8(3):e59030, March.

Charles S. Areni and David Kim. 1993. The influence of background music on shopping behavior: classical versus top-forty music in a wine store. *Advances in consumer research*, 20(1):336–340.

Johan Bollen, Huina Mao, and Xiaojun Zeng. 2011. Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1):1–8, March. Gordon C. Bruner. 1990. Music, Mood, and Marketing. *Journal of Marketing*, 54(4):94–104.

Lei Chen, Shuhua Zhou, and Jennings Bryant. 2007. Temporal Changes in Mood Repair Through Music Consumption: Effects of Mood, Mood Salience, and Individual Differences. *Media Psychology*, 9(3):695– 713, May.

Richard Curtin. 2004. Consumer Sentiment Surveys. In *articolo presentato alla 27 a conferenza del CI-RET*, pages 15–18.

P. J. H. Daas and M. J. Puts. 2014. Social Media Sentiment and Consumer Confidence. Statistics Paper Series, No. 5. European Centeral Bank.

Xiao Hu and J. Stephen Downie. 2010. Improving Mood Classification in Music Digital Libraries by Combining Lyrics and Audio. In *Proceedings of the 10th Annual Joint Conference on Digital Libraries*, pages 159–168, New York, NY, USA. ACM.

Adrian C. North and David J. Hargreaves. 1997. Music and consumer behaviour.

Brendan O'Connor, Ramnath Balasubramanyan, Bryan R. Routledge, and Noah A. Smith. 2010. From tweets to polls: Linking text sentiment to public opinion time series. *ICWSM*, 11:122–129.

B. Barrios-Choplin R McCraty. 1998. The effects of different types of music on mood, tension, and mental

clarity. *Alternative therapies in health and medicine*, 4(1):75–84.

John A. Sloboda. 2011. Music in Everyday Life: The Role of Emotions. In Patrik N. Juslin and John Sloboda, editors, *Handbook of Music and Emotion: Theory, Research, Applications*. OUP Oxford.

Carlo Strapparava and Alessandro Valitutti. 2004. WordNet Affect: an Affective Extension of WordNet. In *LREC*, volume 4, pages 1083–1086.

Trust, Gary. 2013. Ask Billboard: How Does The Hot 100 Work? September.

Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. In *Proceedings of HLT/EMNLP 2005*, pages 34–35. Association for Computational Linguistics.

Yunqing Xia, Linlin Wang, and Kam-Fai Wong. 2008. Sentiment Vector Space Model for Lyric-Based Song Sentiment Classification. *International Journal of Computer Processing of Languages*, 21(04):309– 330, December.

Jiang Zhong, Yifeng Cheng, Siyuan Yang, and Luosheng Wen. 2012. Music sentiment classification integrating audio with lyrics. *Journal of Information and Computational Science*, 9:35–44.