Sound Opinions: Creating a Virtual Tool for Sound Art Installations through Sentiment Analysis of Critical Reviews

Anthony T. Marasco
School of Music and CCT
Louisiana State University
Baton Rouge, LA
amarasco@lsu.edu

ABSTRACT

The author presents Sound Opinions, a custom software tool that uses sentiment analysis to create sound art installations and music compositions. The software runs inside the NodeRed.js programming environment. It scrapes text from web pages, pre-processes it, performs sentiment analysis via a remote API, and parses the resulting data for use in external digital audio programs. The sentiment analysis itself is handled by IBM's Watson Tone Analyzer.

The author has used this tool to create an interactive multimedia installation, titled *Critique*. Sources of criticism of a chosen musical work are analyzed and the negative or positive statements about that composition work to warp and change it. This allows the audience to only hear the work through the lens of its critics, and not in the original form that its creator intended.

Author Keywords

NIME, Sentiment analysis, Node-Red.js, IBM Watson, virtual tool, web information retrieval, music, composition, installation, sound art

CCS Concepts

•Information systems → Sentiment analysis; •Applied computing → Performing arts; **Sound and music computing; •Human-centered computing** → *Natural language interfaces;*

1. INTRODUCTION

The initial concept for this virtual tool was born from the author's desire to use a critical analysis of an artwork to modify the way its audience experiences it. Before heading to the movies or attending a concert, our perception of the music we are about to hear or the film we are about to watch is often modified by the negative or positive reviews we read in blogs or entertainment media. In today's world of instant criticism due to unfettered access to forums devoted to posting our opinions, our experiences are often shaped by previous opinions other than our own[8]. Figure 1 illustrates how this concept would be implemented as an interactive sound art installation.

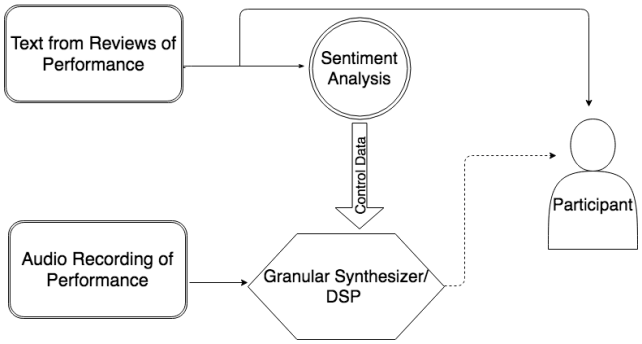

Figure 1: Diagram showing an artistic work flow where Sound Opinions can be utilized.

2. PRIOR WORK

Sentiment and emotion are intrinsically tied to the way human beings experience art and music, with countless research endeavors focusing on the impact media has on the emotional state of listeners and viewers[4, 6]. It would stand to reason then that emotions, in turn, could modify the way a listener perceives a work of art, contributing to their overall enjoyment of the work as a whole. With the use of sentiment analysis tools, the invective of critics could be used to change someone's critical and aural perception of a musical composition.

Sentiment analysis is a very active field of research. New tools for web crawling and machine learning have made this technology more accessible for artists[2]. Related prior work includes installations such as *AMYGDALA* (which uses the Synesketch library to analyze emotional content of text in Twitter posts as a means of generating audiovisual content)[1, 5], and research endeavors at the University Center of Belo Horizonte centered on using multimodal analysis methods in order to extract sentiment data from videos of newscasts[9].

3. IMPLEMENTATION AND USAGE

Sound Opinions [1] is built as a JavaScript applet within the Node-Red.js programming environment. The IBM Watson Tone Analyzer service (a cloud-based API that uses linguistic analysis algorithms to determine the emotional and social context of written text) handles the sentiment analysis and is used in the applet through the inclusion of IBM's Watson library modules for Node-Red.

Before the applet is run, users specify the URLs of the reviews they would like to analyze in the "Review Blog"

[1]Git repository located at `https://bit.ly/2v37knh`

NIME'18, June 3-6, 2018, Blacksburg, Virginia, USA.

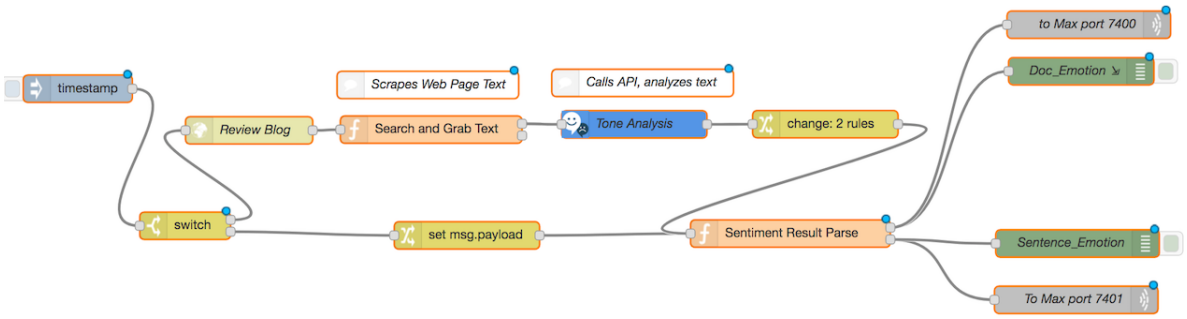

Figure 2: The core of the virtual tool, realized in the user-friendly Node-Red programming IDE. This core flow is duplicated for each source of text that needs to be analyzed.

modules of the Node-Red flow. Once the applet is deployed, text is grabbed from the designated web page, packaged as a string, and sent to the Tone Analyzer API. The API then returns the analyzed text as JSON containing each sentence and their corresponding emotional breakdowns and scores. The data is then passed into a custom "Sentiment Result Parse" module that separates the emotional scores for the entirety of the text from the scores for each individual sentence in the analyzed document. This function also derives numerical values by interpolating between the scores of each sentence's most prominent tone in the Emotional and Social language categories, flagging this newly generated value to be use as control data for digital signal processing parameters. The parsed results are then sent via UDP into any music performance coding environment that the user chooses. Figure 2 shows the core flow.

Users can also modify the JavaScript code inside of the the "Sentiment Result Parse" node in order to only send along particular sentences out of the entire document, or to sort or weigh particular emotional/social scores against others and only send along data that passes their chosen conditions. After the API is called once and the analysis scores are returned, the data is stored in a "Set" module and can be redeployed to the parsing module multiple times without needing to call the API again.

4. INTERACTIVE INSTALLATION

Using this tool, the author has created *Critique*, an interactive audiovisual installation[2]. Participants are stationed in front of a computer and are asked to choose an audio file to listen to from a list of musical compositions. Upon clicking on their choice, the Sound Opinions software scans a predetermined review of their chosen musical work and uses the various negative or positive statements in the text to warp and change the audio during playback. This causes the participants to hear the work through the lens of its critics, and not in the original form that its creator intended.

In order to sonically realize the effects of the review's emotional content on the audio, data from each sentence's sentiment analysis score is scaled and intuitively mapped to various digital signal processing and granular synthesis engine parameters during playback. Keywords in a sentence that mention musical elements (such as "rhythm" or "pitch") and the severity of a particular emotion tied to that word trigger preset modulations that emphasize those aspects of the critique (i.e. a mention of "rhythm" in a sentence with a higher value of negative emotions triggers erratic time stretching to be enacted upon the audio file, further emphasizing the critic's negative perception of this musical element). The corresponding text is displayed on screen one sentence at a time, along with each sentence's sentiment analysis results. This provides the participant with a clear presentation of how five primary emotions (Anger, Disgust, Joy, Sadness, and Fear) are balanced together throughout the review.

[2]Video examples can be found at `https://bit.ly/2JCp7Vx`

5. FUTURE WORK

Future developments will focus on converting the the tool's custom functions into a collection of stand-alone modules using the Node-Red SKD (allowing for easier distribution and use amongst the Node-Red user community), as well as the creation of Max/MSP and PureData abstraction versions. A more advanced set of parsing functions are also currently in development, which will allow for more nuanced correlation between the analyzed text and the audio parameters they are used to control.

6. REFERENCES

[1] Amygdala. `http://fuseworks.it/en/project/amygdala-en/`. Accessed: 2017-04-03.

[2] L. M. Gómez and M. N. Cáceres. Applying data mining for sentiment analysis in music. In *Trends in Cyber-Physical Multi-Agent Systems. The PAAMS Collection - 15th International Conference*, pages 312 – 325. PAAMS, 2017.

[3] IBM. IBM Watson Tone Analyzer documentation. 2017.

[4] A. Kawakami, K. Furukawa, and K. Okanoya. Music evokes vicarious emotions in listeners. *Frontiers in Psychology*, 5:431, 2014.

[5] U. Krcadinac, P. Pasquier, J. Jovanovic, and V. Devedzic. Synesketch: An open source library for sentence-based emotion recognition. In *IEEE Transactions on Affective Computing*, pages 198–205. IEEE, September 2013.

[6] L.-O. Lundqvist, F. Carlsson, P. Hilmersson, and P. N. Juslin. Emotional responses to music: experience, expression, and physiology. *Psychology of Music*, 37(1):61–90, 2009.

[7] Nodered.org. Node-red: Documentation guide. 2017.

[8] J. A. Pamies. "Seven Remarks and A Postscript on Music Criticism". 2016.

[9] M. H. R. Pereira, F. L. C. Padua, A. C. M. Pereira, F. Benevenuto, and D. H. Dalip. Fusing audio, textual, and visual features for sentiment analysis of news videos. In *Proceedings of the Tenth International AAAI Conference on Web and Social Media (ICWSM 2016)*, pages 659–662, May 2016.