Sonifying Twitter's emotions through music

Mariana Seiça[1], Rui (Buga) Lopes[2], Pedro Martins[1] and F. Amílcar Cardoso[1]

[1] CISUC, Informatics Engineering Department, University of Coimbra
marianac@student.dei.uc.pt, {pjmm, amilcar}@dei.uc.pt
[2] Independent Researcher
bugamail@gmail.com

Abstract. Sonification is a scientific field that seeks to settle sound as a mean to interpret data, whose development has been directly connected with the growth of technology in everyday life. This allowed the establishment of new on-line communication tools, such as Twitter, as a part of a daily routine that gave users the possibility to share their opinion and feelings about national and global events. This paper reports a work in progress with the purpose of uniting this social media phenomena with sonification, using Twitter data to extract the user's emotions and mapping them to musical variables. The purpose of this work is to explore the potential of music in translating data as personal and complex as human emotions, creating new ways of reading and understanding them.

Keywords: Musical Sonification, Emotion Detection, Twitter, Algorithmic Composition, Sound Design

1 Introduction

Sonification, defined by Kramer et al. [10] as "the use of nonspeech audio to convey information", has been establishing its place as a new field of communication, exploring new techniques to represent complex data through sound [7]. Since the birth of the International Community of Auditory Display (ICAD) in 1992, where the study of auditory displays was proposed as a scientific field, sonification tecnhiques have been developed significantly, with applications in areas such as Medicine or Seismology, and concepts from multiple areas, from Human Perception to Design and Engineering that form its interdisciplinary nature [7].

The development of the sonification field is directly connected to the significant growth that technology experienced in the last decade, with personal computers containing the hardware and software needed to manipulate sound [7], and auditory displays becoming a presence in everyday life. This technological growth and accessibility to the main population also allowed the establishment of new communication media as a daily routine: the social media. Facebook and Twitter are examples of these social tools that became the new mass media, not only for the common citizen, but also for companies, news industry and important figures. The study of social media data has gained new potential in several areas, such as marketing for extracting consumer's opinions, or social studies for understanding the user's moods and views about events. The field of sentiment analysis emerges from this potential, with the focus of studying computational analysis to extract opinion and sentiment from text, interacting with affective computing to explore the computer's ability in recognizing emotions [16].

This paper presents work in progress that handles with these three fields: sonification as the core, with data retrieved from social media, specifically Twitter, and analysed through sentiment analysis. The main goal is to explore new ways to read, through sound, data as personal and complex as human emotions. This study is primarily motivated by the potential of music in conveying emotions, and lies to this potential, exploring ways to transmit information through a melodic, harmonic and rhythmic composition. It involves two major challenges: the emotion extraction, implementing a system that properly analyses the tweets and classifies their emotions; and the musical mapping, choosing a set of parameters that can distinguish and embody the emotions. The main focus of this paper is on the musical mapping. However, we also briefly describe how the process of extracting emotional information from the tweets is implemented.

To explain the development made so far, this paper starts with an overview of Sonification works and projects that used social media data, which influenced and inspired this work. Section 3 discusses the model of the emotions chosen, and the sentiment analysis process implemented to classify tweets. Section 4 presents the mapping structure and musical implementation in course, showing the results obtained and concluding with a reflection on future development.

2 Related Work

There are many studies of sonifications developed in a vast number of areas, proven successful in practical and scientific terms [10].

The first two examples show the potential of sonification in scientific data. The work of Vicinanza, specifically his *Voyager 1 & 2 Spacecraft Duet* is a sonification of data gathered by the Voyager 1 & 2 NASA probes during 37 years of spatial exploration. It is composed by two melodies in different frequencies, with the measurements made at the same times, but billions of kilometers apart [20]. The second example is *The Climate Symphony* by Quinn [18], where he used data from the chemical composition of an ice block in Greenland to translate into music the climatic changes endured by the great continental ice sheets.

In the poetry field, Coelho, Martins & Cardoso [3] created a *A Musical Sonification of the Portuguese Epopee*, specifically of The Lusiads, by Luís de Camões. It is an interactive sonification where the user can explore the poem by choosing different levels of "zooming", listening to it as a whole, as a canto/subnarrative or a specific episode and therefore customising the experience.

Bulley and Jones developed *Living Symphonies*, a sound installation based on the fauna and flora of four ecosystems in the United Kingdom. The authors built a model that reflected the behavior, movement and daily patterns of every being in the wild, translating a network of interactions that formed the ecosystem [2].

#tweetscapes is the most similar project to this study, consisting in the sonification of German tweets in real-time. Developed by Hermann, Nehls, Eitel, Barri and Gammel, the goal was to create "a new sense of media awareness" [8]. Tweets were mapped according to the hashtags, replies and location, adding a visual geographic distribution that accompanied the sonification.

#Emotional Imaging Composer is an experience conducted in the Input Devices and Music Interaction Lab (IDMIL) at the McGill University, that aims to create a real-time audio expression of emotions, extracted from a vocal performance [21]. This interactive sonification is based on Russell's Arousal/Valence circumplex [19], positioning musical parameters over the two axis.

The last example is a website created by Harris and Kamvar, named *We Feel Fine* [6]. It is a visualization that collects human emotions from a vast number of blogs, searching for entries that contain the expression "I feel" or "I am feeling" and providing a social and demographic study.

3 Processing Data

The process implemented in our system comprises four steps (Fig.1) and uses three software tools to produce the sonification: Processing, to get and classify tweets, Max, to generate the musical composition, and Ableton Live, to play the composition using VST's plug-ins. The current section describes how the three first steps were implemented. The sound mapping step will be presented in detail in Section 4.

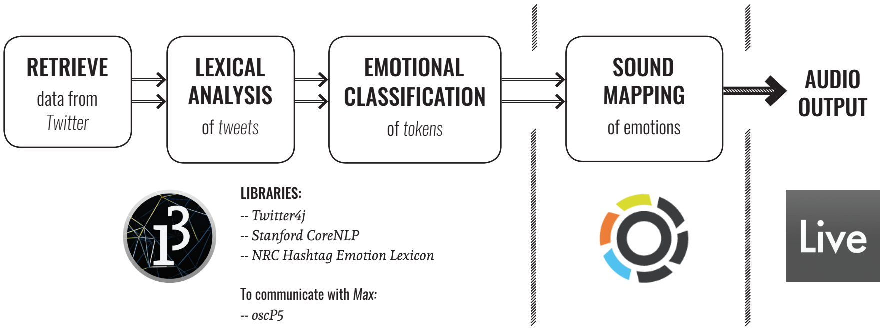

Fig. 1. Process flow diagram

3.1 Data Gathering

The emotional content of a tweet lies in two elements that will form the dataset: the hashtags, meta-data tags that establish the subject and mood of a tweet, and the main text, that elaborates the subject and expresses the user's opinion.

To retrieve the tweets, we are using the *Twitter4j* Java library. The data is filtered by language, receiving only tweets in English. To ensure data with some relevance, the tweets are also filtered by the number of followers of the tweet's author: only tweets whose author has more than 1000 followers are considered.

3.2 Emotion Lexicon

We implemented a system based on a lexicon of words, composed by associations of emotions to each word using a lexicon developed by Mohammad, the *NRC Hashtag-Emotion Lexicon* [15], with the aim of maintaining a simpler and more open approach It is based on a model of emotions created by Plutchik [17] (Fig.2), that is comprised of eight primary emotions: joy, trust, fear, surprise, sadness, disgust, anger and anticipation.

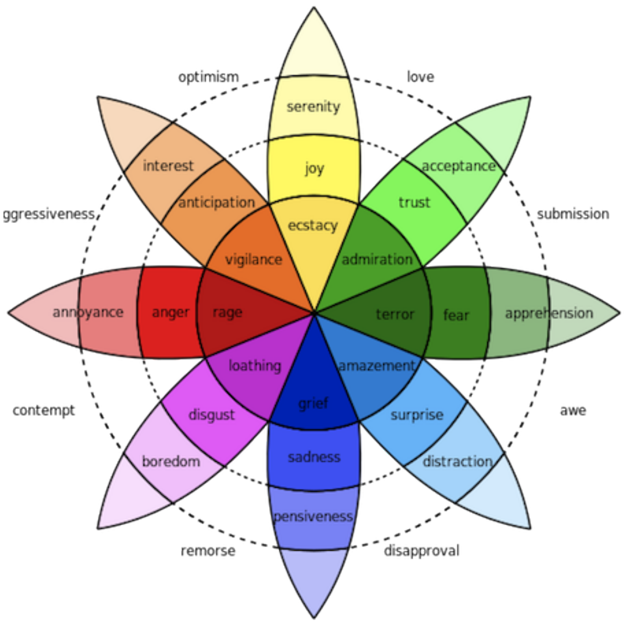

Fig. 2. Robert Plutchik's Model of Emotions

3.3 Lexical and Emotional Analysis

The next task was to implement a set of natural language processing (NLP) tools to parse the tweets, establishing the structure of sentences, word dependencies and "Part-Of-Speech Tagging", classifying each word with its root form, called lemma, and its grammatical category. Working in a Java environment, we have chosen the *Stanford CoreNLP* tool [13], developed by the Stanford NLP group.

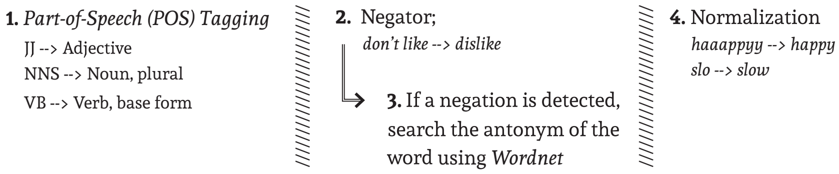

Fig. 3. Steps of the lexical analysis

The analysis of the receiving tweets is comprised of three steps (Fig.3):

1. Tweet's parsing and tagging: to define the main structure of the sentences and the words classes. Words classified as nouns, verbs, adverbs or adjectives are stored, due to these classes describing usually the emotion and intensity of a sentence.
2. Identify negations: find the antonyms of negated words using *Wordnet* [14].
3. Submit the remaining words to a combination of two normalisation lexicons [5] [12], converting the terms and abbreviations of the texting language to its correct writing form (example of the word "happy" in Fig.3).

The emotional classification is then applied to the resulting list of words. First, we search the existence of each word in the *NRC Hashtag-Emotion Lexicon.* If it is not found, a search for the lemma of the word is made. Without results, *Wordnet* is used to find a synonym of the word, repeating the first and second steps if a result is found. For the hashtags, the process is more simplified: as the *NRC Hashtag-Emotion Lexicon* contains a considerable amount of hashtags and its emotional association, we search the hashtags directly in the lexicon.

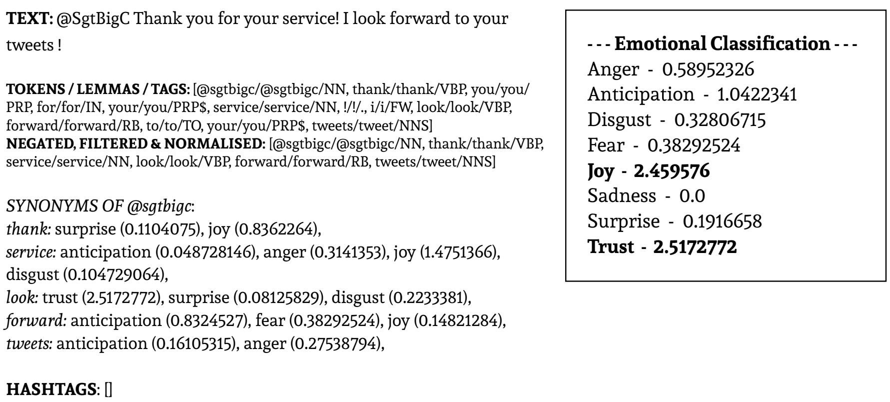

Fig. 4. Example of a tweet's emotional classification

3.4 Intermediate Results

All the steps were saved in a text file for evaluation, including the parsing (with each word, correspondent lemma and tag), the lexical analysis' resulting list, the emotions extracted from each word and the sum of every word's emotion, providing the tweet's classification. In the example shown in Fig.4, *Trust* and *Joy* have the highest and similar values. The majority of the words obtain a classification, allowing a more complete and differentiating categorisation of each tweet. This process takes an average time of 1.6 tweets per second, ensuring a consistent set of variables to sonify and establishing emotional tendencies.

4 Mapping Twitter's Emotions to Music

In 1936, Hevner [9] conducted a series of studies of the expressive elements in music, associating a list of adjectives to musical parameters, such as major/minor mode, dissonant/consonant harmonies and firm/flowing rhythms. The results determined a tendency in the associations, achieving an universal affective nature in musical forms. The main challenge of this project is to explore this expressive dimension to map emotions. Gabrielsson & Lindström [4] gathered over 100 studies made in the last century on this subject, that served as an initial foundation for this sonification. The majority focused on evaluating simple parameters, like tone quality, melody direction, loudness or tempo. Studies of more complex parameters, such as harmonic progressions or chords natures are very limited, exploring only differences in consonant/dissonant harmonies. The authors concluded that although each parameter can influence the emotional expression, it is almost never determined by one factor, but a combination of several.

In our project, we decided to organise the mapping into three main aspects of music: melody, rhythm and harmony, associating probabilities with each parameter for each emotion. At the start of the program, the main root note is defined, which provides the tonic for the harmonic progressions and the melody's scale.

For the melody, each note is raffled from the current scale, or from the current chord, or as a chromatism, following probabilities that change according to each emotion (See Fig. 5). Chromatic notes are dissonant notes that occur half-step above or below one of the chord's pitches. They travel outside a given scale and are generally used as transition notes to create tension before returning to consonance, releasing the tension. In our system, they are played on weak beats, lasting only half a beat, to keep a subtle dissonant and a tonal feeling. *Fear*, for example, has a higher chance of occuring a chromatic note (50%) than *Joy*, with only 5%.

Hevner also concluded that a slow tempo was associated with more solemn, sad or gentle sounds, whereas happy, exciting and vigorous sounds where likely translated by a fast tempo [4]. Translating this structure to rhythm figures, the duration of the melody notes were associated through probabilities for each emotion, using four rhythm figures: whole, half, quarter and eighth notes (Fig. 5). *Anticipation*, *Anger* and *Surprise* have higher changes of producing quarter

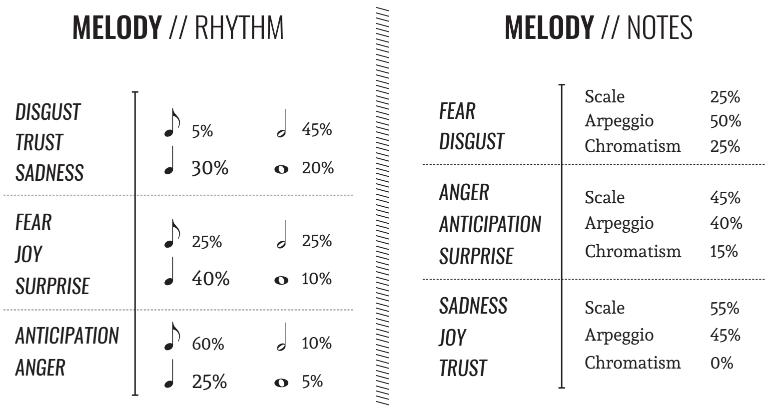

Fig. 5. Melody probabilities on notes and rhythm figures

notes, ensuring a more rapid and tense melody. In opposition, for *Trust* and *Disgust* there is a higher tendency to play longer notes.

MELODY // INTERVALS

	ANTICIPATION FEAR	ANGER	JOY TRUST SADNESS	SURPRISE	DISGUST
Tonic	0%	5%	5%	5%	20%
2nd	25%	30%	35%	10%	10%
3rd	25%	15%	30%	15%	10%
4th	10%	5%	5%	10%	10%
5th	25%	10%	10%	15%	10%
6h	5%	15%	5%	20%	15%
7th	5%	15%	5%	15%	15%
8th	5%	5%	5%	10%	10%

Fig. 6. Melody probabilities on intervals

The melodic interval between two consecutive notes is also raffled according to the emotion being conveyed (See Fig. 6). *Joy* and *Sadness* have a higher chance of producing a stepwise motion, using seconds and thirds to provide a more stable and comfortable flow. On the other hand, *Surprise* has a higher probability of sixths and sevenths in order to create sudden jumps that bring an unexpected feeling. An emotion like *Trust* has then a higher probability of producing a consonant sound, in opposition to *Disgust*, that produces a heavily dissonant sound.

Music is built over tension and release moments that define the harmonic sequence, which have an impact on the conveyed emotions. Our approach for distinguishing each emotion harmonically was based on Hevner's studies and Gabrielsson & Lindström's analysis, comparing major to minor modes and consonant to dissonant harmonies. Hevner concluded that "it is apparent that the use of the major or minor mode is of the most clear-cut significance in the expression of four different mood effects" [9], with the major mode strongly associated with happiness, gaiety and playfulness, and the minor with sadness, agitation and disgust". A simple/consonant harmony is defined as happy, graceful, serene and dreamy, connected to joy and tenderness, and a complex/dissonant harmony as vigorous and sad, connected to agitation, fear and unpleasantness. These two notions serve as the basis for exploring the progressions and correspondent chords. Due the complexity associated with harmony, we decided to define a set of 20 chords with different natures (Fig.7). These chords would serve as the structure for a series of progressions that translate each emotion (Fig.8), ensuring the coherence of the sequence and the affective nature associated. The chords may be played in three possible voicings, chosen randomly: in the root position, with an added tonic in the bass (an octave lower), with the tonic and the fifth in the bass, or with the tonic and the seventh (third if triad) in the bass.

HARMONY // CHORDS

TRIADS

Major	1 3 5
Minor	1 b3 5
Diminished	1 b3 b5
Augmented	1 3 #5
Suspended4	1 4 5
Suspended2	1 2 5

TETRACHORDS

Major 7	1 3 5 7	*Major 7#11 (Lydian chord)*	1 3 #4 7
Dominant (7)	1 3 5 b7	*Lydian Dominant*	1 3 #4 b7
Minor 7	1 b3 5 b7	*7 Suspended4 add 9*	1 4 b7 9
Half-Diminished	1 b3 b5 b7	*Major Six-Nine*	1 3 6 9
Diminished	1 b3 b5 bb7(6)	*Minor Major 7*	1 b3 5 7
Augmented Major 7	1 3 #5 7	*Minor 6*	1 b3 5 6
Augmented 7 (Altered)	1 3 #5 b7	*Minor Major 7#11*	1 b3 #4 7

Fig. 7. List of chords

For each emotion, a progression is chosen from the list, all with equal probability. Each progression is associated with a number of scales from which the melody is created. The scales are either built over the modes of the major scale, or the pentatonic scale, or harmonic and melodic scales.

The harmonic structure explores known progressions, such as the major *I (Major) - VI (Minor) - IV (Major) - V (Dominant)* or the minor *I (Minor) - I (Minor) - II (Half-Diminished) - III (Major)*, and combinations of different chord natures, building sequences that relate them with the tension associated with each degree. For example, the progression *I (Suspended4) - I (Six-Nine) - I (Suspended4) - I (Major)* combines different major chords with a suspended chord, always built on the first degree. The major chords maintain the stability, and the suspended chord adds tension and a more open sound, establishing the possible set of association with *Joy*.

HARMONY // PROGRESSIONS EXAMPLES

<table>
<tr><td>ANTICIPATION</td><td>|| III^{-7} VI^{7} | II^{-7} V^{7} ||
|| $I^{MAJ7\#11}$ | VII^{-7b5} $bVII^{7}$ | $bVI^{MAJ7\#11}$ ||</td><td>SADNESS</td><td>|| II^{HDIM7} | V^{AUG7} | I^{-7} ||
|| I^{-7} | I^{-7} | II^{HDIM7} | $III^{MAJ7\#11}$ ||</td></tr>
<tr><td>SURPRISE</td><td>|| I^{SUS2} | bV^{SUS2} | VII^{SUS2} ||
|| I^{-7} II^{-6} ||</td><td>FEAR</td><td>|| I^{DIM7} | I^{DIM7} $bII^{7\#11}$ ||
|| $I^{MAJ7\#11}$ $bIII^{MAJ7\#11}$ | $bV^{7\#11}$ $bV^{7\#11}$ bVI^{DIM7} $V^{7\#11}$ ||</td></tr>
<tr><td>JOY</td><td>|| $I^{MAJ7\#11}$ | $I^{MAJ6/9}$ | $I^{MAJ7\#11}$ | I^{7} ||
|| I^{MAJ} | VI^{-7} | IV^{MAJ7} | V^{7} ||</td><td>ANGER</td><td>|| V^{MAJ} | V^{7} | V^{AUG7} | V^{MAJ} ||
|| I^{-MAJ7} | bII^{AUG7} | VII^{DIM7} | IV^{DIM7} | bVI^{DIM7} ||</td></tr>
<tr><td>TRUST</td><td>|| V^{MAJ} | I^{MAJ} | V^{MAJ} | I^{MAJ} ||
|| I^{MAJ7} | II^{-} | III^{-7} | I^{MAJ} ||</td><td>DISGUST</td><td>|| I^{AUG7} ||
|| $I^{AUG\,MAJ7}$ | bVI^{DIM7} | $I^{AUG\,MAJ7}$ ||</td></tr>
</table>

Fig. 8. List of progressions for each emotion

The progression chosen is repeated a random number of times, with the system choosing a progression again randomly at the end of loop, ensuring a more dynamic and flowing composition. For each progression loop, Processing saves the values of the emotions extracted from each tweet and their sum. At the end of the loop, Max sends a message to Processing requesting the next emotion with the highest value, which represents the emotional tendency of the tweets during that time. The sonification is dynamically built in the communication of the two modules, with Ableton Live playing the composition through the framework Reaktor 6, from Native Instruments: the harmony with two VSTs, Kontour and Prism, and the melody with FM8. The tone quality is adapted to each emotion, associating certain sounds and technical aspects, such as reverb, attack, sustain and distortions, to each musical context provided by the harmony. The overall tone quality resembles the ambient music genre, evoking a more atmospheric and open sound.

Some of the resulting sounds produced so far are available at: `https://soundcloud.com/mariana-seica/sets/music-emotions-sonification`

5 Reflections and Future Development

The major challenge of this study, aimed at sonifying emotions from tweets, is to define and maintain a parametrization with reasonable complexity that allows the mapping from emotions to the musical output to be understandable and distinguishable in a compelling composition. The work developed so far proved this difficulty, and demands more experimentation. One of the impending tasks will be to add more progressions to each emotion, ensuring a more dynamic sonification. The tempo and loudness are also parameters that will be considered, exploring the possible affective results of their changing.

A significant amount of tests will be needed to understand key elements of this study: first, if the tweets are being correctly classified, i.e., if the emotions extracted from the tweets embody their emotional content; and second, if the composition reflects this emotional content, with a set of musical parameters that express it. The validation is imperative to interpret the effectiveness of this project, and how the users perceive the sonification.

References

1. Aldwell, E., Schachter, C.: Harmony & Voice Leading. Thomson-Schirmer (2003)
2. Bulley, J., Jones, D.: Living Symphonies (2014), `http://www.livingsymphonies.com/`
3. Coelho, A., Martins, P., Cardoso, A.: A Musical Sonification of the Portuguese Epopee. In: MUME 2016 - The Fourth International Workshop on Musical Metacreation (2016)
4. Gabrielsson, A., Lindström, E.: The Role of Structure in the Musical Expression of Emotions. In P.N. Juslin, J. Sloboda (Eds.): Handbook of Music and Emotion: Theory, Research, Applications, Oxford University Press, New York (2010)
5. Han, B., Cook, P., Baldwin, T.: Automatically Constructing a Normalisation Dictionary for Microblogs. In: Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, Jeju Island, Korea (2012)
6. Harris, J., Kamvar, S.: We Feel Fine (2006), `http://number27.org/wefeelfine`
7. Hermann, T., Hunt, A., Neuhoff, J.G.: The Sonification Handbook. Logos Verlag, Berlin, Germany (2011)
8. Hermann, T., Nehls, A.V., Eitel, F., Barri, T., Gammenl, M.: Tweetscapes - Real-time sonification of Twitter data stream for radio broadcasting. In: Proceedings of the 18th International Conference on Auditory Display. Atlanta, GA, USA (2012)
9. Hevner, K.: Experimental Studies of the Elements of Expression in Music. The American Journal of Psychology. 48(2), 246–268 (1936)
10. Kramer, G., Walker, B., Bonebright, T., Cook, P., Flowers, J.H., Miner, N., Neuhoff, J.: Sonification Report: Status of the Field and Research Agenda. Faculty Publications, Department of Psychology. Paper 444 (2010) `http://digitalcommons.unl.edu/psychfacpub/444`
11. Ligon, B.: Jazz Theory Resources: Tonal, Harmonic, Melodic and Rhythmic Organization of Jazz. Houston Publishing (2003)
12. Liu, F., Weng, F., Jiang, X.: A Broad-Coverage Normalization System for Social Media Language. In: Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics. pp. 1035–1044 (2012)
13. Manning, C.D., Surdeanu, M., Bauer, J., Finkel, J., Bethard, S.J., McClosky, D.: The Stanford CoreNLP Natural Language Processing Toolkit. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations, pp. 55–60 (2014)

14. Miller, G.M.: WordNet: A Lexical Database for English. Communications of the ACM. 38(11), 39–41 (1995)
15. Mohammad, S. M., Kiritchenjko, S.: Using Hashtags to Capture Fine Emotion Categories from Tweets. Computational Intelligence. 31(2), 301–326 (2013)
16. Pang, B., Lee, L.: Opinion Mining and Sentiment Analysis. Foundations and Trends in Information Retrieval, 2(1-2), pp. 1–135 (2008)
17. Plutchik, R.: The Nature of Emotions. American Scientist. 89, 344–350 (2001)
18. Quinn, M.: Research set to music: the Climate Symphony and other sonifications of ice core, radar, DNA, seismic and solar wind data. In: Proceedings of the International Conference on Auditory Display. Espoo, Finland (2001)
19. Russell, J.: A Circumplex Model of Affect. Journal of Personality and Social Psychology. 39(6), 1161–1178 (1980)
20. Vicinanza, D.: Voyager 1 & 2 Spacecraft Duet (2014), `http://www.geant.net/MediaCentreEvents/news/Pages/The-sound-of-space-discovery.aspx`
21. Winters, R.M., Hattwick, I., Wanderley, M. M.: Emotional Data in Music Performace: Two Environments for the Emotional Imaging Composer. In: Proceedings of the 3rd International Conference of Music & Emotion (ICM3), Jyväskylä, Finalnd (2013)