DATA MINING / LIVE SCORING - A LIVE ACOUSTIC ALGORITHMIC COMPOSITION BASED ON TWITTER

Alexandros Drymonitis

PhD Candidate Royal Birmingham Conservatoire
Birmingham City University
200 Jennens Road
Birmingham, UK
alexdrymonitis@gmail.com

Nicoleta Chatzopoulou

PhD Candidate Technical University of Crete (TUC)
School of Architecture
Campus, 73100
Chania, Greece
nicol.chatz@gmail.com

ABSTRACT

Data Mining / Live Scoring is a live algorithmic composition for a six member acoustic ensemble. The algorithm receives feed from the popular social platform Twitter, sonifying the tweets it receives by writing a score for each of the six musicians. The concept behind this performance is to create a musical composition which can serve as a soundtrack for Twitter, depending on what is written by the users. Based on the sentiment of the tweets and a musical library written especially for this project, the algorithm attempts to create a score which reflects the mood on Twitter in an abstract way. This project was commissioned by the Onassis Cultural Centre in Athens, and was presented there twice in April 2019. It was realized for the ARTéfacts Ensemble.

1. INTRODUCTION

Data Mining / Live Scoring is a project which combines data sonification with audience input. Data sonification occurs by using the Twitter social platform to retrieve tweets which are transformed into scores for an acoustic ensemble during the performance. Sentiment analysis of the tweets was the backbone of this project. It was chosen as an efficient method and way of demonstrating the close relationship between the meaning of the natural language of the tweets and the ambiance of the resulting music. By using Twitter, the audience could - and was encouraged to - participate by writing tweets during the performance. The resulting scores were rendered with Lilypond [1] and displayed on monitors, one for each performer, with an openFrameworks [2] program which included an integrated blinking cursor for the synchronization of the ensemble. Excerpts are available online at the following link https://vimeo.com/369534737

1.1. Data Sonification in Music

Sonifying data to create musical works is nothing new. Early works like *Music for Solo Performer* by Alvin Lucier [3] or *Pythoprakta* by Iannis Xenakis [4] have utilized data long before the digital age of big data. Aiming at an artistic output, the approaches

to realizing such works vary and depend heavily on subjective perceptions on how data relates to sound [5]. The same applies to the source of data used for such an artwork. Soil elements [6], georeference data [7], earthquake data [8], or even facial expressions [9] are a few of different data sources which have been used to create musical works. Various sonification musical works also differ in whether the data sonified are being collected in real-time during the performance, or beforehand, in which case it is common practice to scale the time of the data collection down to a shorter period suitable for a live performance for audience. While the latter approach provides more freedom in the mapping between the time frame of data collection and the duration of the music piece - Rob King collected data from gun license background checks [10] during a period of eleven years to create a piece of a few minutes [11] - the former creates a stronger bond between the two, where dense data create dense or long pieces, and sparse data create short pieces or notes/sounds with long durations.

1.2. Audience Input

Audience participation in music is another practice that has been celebrated by many artists. [12, 13, 14, 15] are examples that share a common approach, that of the smart phone as an interface for the audience to participate. Compared to *Data Mining / Live Scoring*, [13] stands out from these four as it is for an acoustic ensemble and audience participation. It has a different approach to the combination of ensemble and audience in that the audience creates sounds immediately by using a mobile application, plus there is a limit of three audience members. A similar project is Tin Men and the Telephone [16] which is an improvisational band which incorporates audience input from a mobile application during performance. The audience participation affects various aspects of the improvisation structure of the group as well as live visuals. Apart from the mobile phone, there are other interfaces used for audience participation. [17] suggests that the audience members wear an armband-based Musical Haptic Wearable with which they can interact with the performer. [18] uses brainwaves recorded by EEG recording devices, one worn by the performer and the other by one audience member.

1.2.1. Audience Engagement and Directness of Participation

There are many different ways with which one can incorporate audience participation, either direct or indirect. Depending on the level of commitment one wants the audience to make, one needs

to find a suitable way to let the audience be part of a performance. The reasons one chooses to actively involve the audience in a performance can also affect the way this involvement will happen. Direct commitment can be a delicate matter when it dictates the participants to be publicly visible or audible when they get involved, as such an approach should require the audiences consent, unless participation occurs only through the personal will of the audience members. Indirect commitment on the other hand can be more discrete, but can result in being not so clear as to how the audiences participation forms or influences the performance.

1.3. Twitter Input

Twitter is a tool used in a number of data sonification musical projects. The way the tweets are mapped to musical elements varies greatly from project to project. [19] uses Text-To-Speech to recite the downloaded tweets. [20] groups tweets based on similarity, placing them in a tree-like structure of nodes, where each node and sub-nodes are provided with a randomly generated melody. [21] is a web application where the users can send tweets including specific syntax which marks rhythmic elements. The syntax is of the form *a-a-a-abc-cc* and it has to be placed within square brackets in order for the algorithm to locate the rhythmic phrase within the whole tweet. [22] analyzes the sentiment of tweets about music from the top 50 current trending artists on Twitter (at the time of writing of the referenced paper). This work utilizes additive synthesis where the sentiment analysis determines the frequency and amount of harmonics. [23] focuses on geolocation where the distance from a focal point determines pitch changes in a granular synthesis program. [24] sets an impact value based on the number of followers of Twitter users, which determines timbre parameters, calculates panning based on users longitude, and reverberation based on the users distance from the center of Germany.

1.4. Choice of data source

Bearing in mind the various sonification sources mentioned above, as well as the impact the audience input can have and how this can be achieved effectively and at the same time discretely, *Data Mining / Live Scoring* utilized the Twitter streaming API in order to retrieve tweets during the performance. These tweets served as the sole input for the creation of the scores for the musicians of the acoustic ensemble. Twitter can provide a massive database with an immense volume of information, both from users who are part of the audience, and from users who are not physically present and are not aware of the performance or the fact that their tweets are being analyzed in this context. It also provides a way of direct involvement only through the personal will of each participant, but without exposing them, as the tweets are being projected without displaying any information about the user - name, location, etc.

2. TWEET ANALYSIS

In order to create music that reflects the sentiment of the tweets, text analysis of some sort had to be done. In the field of machine learning, classification of text according to sentiment is a practice commonly used in services where the opinions of customers are important [25], usually applied on product reviews. There are two main techniques applied in sentiment analysis. One is the semantic approach which uses lexicons with words and phrases paralleled with a sentiment valence. The other is based on

machine learning models like Support Vector Machines (SVM), Naive Bayes, and K-Nearest Neighbor (KNN) [26]. For *Data Mining / Live Scoring* we chose to use a lexicon with short phrases with opposing polarities often used in Twitter [27]. This lexicon consisting of 1,178 words or phrases was released in 2016 and has the following form:

$$\langle term \rangle \ \langle sentiment score \rangle \ \langle POS pattern \rangle \ \langle term freq \rangle$$

For example, the first line of this lexicon reads:

"seriously great" 1 R+A 17

The only information used in this project was the valence field. This line tells us that the phrase "seriously great" has a valence of 1 (which is 100% positive). All tweets used were compared to this table. Each word was tested against the whole list. When a match was found, its valence was accumulated to the valence of the whole tweet. This process was repeated until the whole tweet had been scanned. After this process was finished the overall valence was normalized to fit the -1 to 1 range, and then mapped to a scale between 1 and 5 which was used to index the five different music scales used in this project. Further exaplanation on these scales is provided in the Music Library section of this paper.

Another element of the tweets that were analyzed was the prosody of the text. This was done with the Prosodic module for Python [28]. This was occasionally used in order to define the rhythmic structure of the resulting scores. More information on how this was utilized is provided in the Music Library section of this paper.

3. CHOICE OF HASHTAGS

A few minutes before each performance I would log in on Twitter and check the currently trending hashtags. Each time we chose three of these hashtags together with the *#datamininglivescoring* hashtag we created for this project. The latter hashtag was used for input from the audience in case someone wanted to tweet something specifically for the performance. The other three trending hashtags were chosen in order to ensure flow of data, in case the audience would remain idle and not tweet. On top of that, using various hashtags with different thematics was our aim in order to focus on social behavior on Twitter, regardless of conversation subjects.

4. SYSTEM ARCHITECTURE

The architeture of the system for *Data Mining / Live Scoring* consisted of a computer running the main algorithm written in Python, and one Raspberry Pi computer with a monitor for each performer. We chose to use Raspberry Pis for the performers mainly due to low cost and ease of integration into the stage settings which consisted of a light installation with a thin metal sctructure. The Pis fit discretely behind their monitors and were not noticable by the audience. All computers were inter-connected via Ethernet through an Ethernet switch.

4.1. Order of Operations of Main Algorithm

The performance was separated in structure blocks. In the beginning of every performance, the algorithm would download tweets for three minutes or up to thirty tweets before it started creating the scores for the first block. While the first block was being played, the algorithm would download tweets for the next block, leaving enough time before the current block ended, in order to analyze the tweets and create the new scores. The first operation was to analyze the sentiment of these tweets with the technique mentioned in the Tweet Analysis section above, and send their sum to an SVM classifier together with the density of those tweets. The density was defined by the number of the downloaded tweets against the time it took to download them. This classifier was trained with sets of valence and density, each mapped to one of nine music form structures. These structures contained information on orchestration, melodic formation, dynamics, approach to rhythm construction, and tempo, and were composed intuitively by the second author so as to reflect different sentiment valences and density of tweets. The Music Library section of this paper provides a more detailed analysis on these structures. The SVM classifier predicted which of the nine form structures would be used for the upcoming music block.

The next operation was to create the rhythmic phrases for all instruments, and based on that the algorithm would then create the melodies. Afterwards the algorithm added various techniques, specific to each instrument, and dynamics, both according to the form structure which was predicted. Once all musical elements were chosen, the main algorithm created strings in the Lilypond syntax and sent them to the Raspberry Pis over Ethernet via the OSC protocol. Each Pi rendered the score as a .png file using Lilypond. The scores were being displayed on the monitors via an openFrameworks program which detected the beginning of each bar and included a blinking cursor on top of the bar currently being played. This was necessary to keep the performers synchronized and help them keep track of the score. Figure 1 is a flowchart of the entire process.
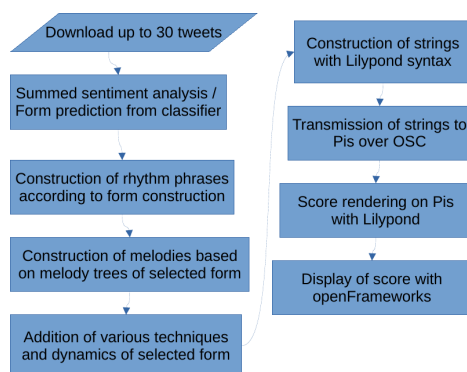

Figure 1: Flowchart of the process of *Data Mining / Live Scoring* from the initial phase of downloading tweets to the last phase of displaying the resulting scores

5. MUSIC LIBRARY

The sentiment of the tweets, mapped to a scale from 1 (negative) to 5 (positive) was the basis on which the musical scales were chosen for this algorithm, and the pitch material of the melodies which consequently formed were based on these scales. The rhythmical patterns used in those melodies were drawn either from the composed database of rhythms or from the prosody of the tweet. Nine different forms were composed, having certain criteria based on the tweets, and specifying the orchestration and articulations in percentages, the tempi and the dynamics. The criteria for the choices in scales and rhythmic structures, as well as the percussion groups were mainly aesthetic and intuitive. The aim was to create a structure for the input from the algorithm in order to reflect the sentiment of the tweets, from negative to positive.

5.1. Proto-elements and Patterns

In this context, the musical scales comprise the proto-elements. Five different scales, a wholetone, Iwato, Pelog and 2 mixed mode scales, were allocated to one sentiment each (1-5). The material of each scale is organized in the form of trees, a method which was consistently used both for rhythm and pitch. More specifically, every note of a scale serves as the top of a tree, which branches out twice, giving four options for pitch successions for each individual tree. Additionally, within each scale, there were 3 distinct roles in the melodic formation possibilities, entitled as : Solo, Mixed and Accompaniment. In the Solo option, the pitch movement within each tree is more agile, in the mixed option it is less agile and in the Accompaniment option, there is relative stillness in the pitch changes within each branch of a tree. Thus, for a scale of 7 tones, there were a total of 21 pitch trees for all three roles. Similarly, for the octatonic scale the total number of trees was 24. For the pelog pentatonic scale, one tone was doubled, but was branching in a different way.

As mentioned earlier, the rhythm is determined in two ways : the prosody of the tweets and trees. For the prosody, each strong syllable was mapped to one rhythmic symbol (i.e. one eight note). A table was then constructed with rhythmic patterns consisting of 2 - 9 rhythmic symbols, for five different meters (2/4, 3/4/ 4/4, 5/8, 6/8) and for the 3 different roles, solo, mixed and accompaniment. This was the predetermined material that the algorithm used for the strong syllables of the tweet. The concept behind this approach was to be able to recite the tweet along with the music, where each musical element, be it a note or a rest, would correspond to one strong syllable. There was no recitation of tweets during the performance, but our intention was to associate the tweets with the music in various ways.

The trees for the rhythm were constructed in an analogous way to those for the pitch, under the three different roles. In this case, the root of the tree and the first branch was in 4/4 metre, whereas the rest of the branches in 2/2. For each tree, there were four different options for the root and four for the first branching in order to have more variety in the melodic structures forming.

5.2. Orchestration

The instrumentation was clarinet in B flat, soprano or baritone saxophone, violin, viola and percussion. Six percussion groups were formed, which were different combinations of standard, world and junk instruments, played by two performers. A detailed study of the techniques and articulations of each instrument with respect

to dynamics was made in collaboration with the performers, and the most suitable findings for this work were organized in a table. The five columns of the table corresponded to dynamic gradations ranging from *fff-ff* to *pp-ppp*, and was hard-coded in the sub-blocks of the form, described in more detail below. A detailed, but concise description of each technique was included in this table and it appeared as such on the score, when it was chosen by the algorithm.

5.3. Form

A total of nine blocks of forms were created for this composition. Each form had a different character. The criteria that the algorithm used in order to choose a form were the sentiment and the density of tweets, ranging from 1 to 5. The characteristics included in each form-block were rhythm, tempo, orchestration combinations and dynamics. The first characteristic refers to whether the rhythm was based on the prosody or on tree structure and defined the percentages for those cases. It was possible to have one form-block using prosody and trees alternating or just one of the two. The second characteristic refers to the chosen tempo for the entire form block. The orchestration was the most analytical characteristic because it defined the techniques of the instruments used in percentages with respect to the dynamics. The orchestration did not remain fixed throughout most of the form-blocks, but changed also according to a given percentage. So, it was possible to have five or six different parts within one block of form.

6. ISSUES

In order to have a smooth flow between the music and the data that created it, the algorithm was downloading enough tweets to create one page of music at a time. The amount of tweets received while one page was being played, which were used to create the next page, was so big that we had to constrain the tweets used by the algorithm to a certain number, namely thirty tweets per page. This restriction was necessary in order to minimize the delay between the time a tweet was received and the time the music created by that tweet was heard (since the tweet was projected as text). If we were to use more than thirty tweets per page, the music created for each block would become very long, much longer than one page. As the performance progressed, this problem would get augmented every time since the tweets would be even more than those downloaded for the previous block. Since the audience was tweeting during performance, this issue became clear to some audience members who were tweeting but did not see their tweets projected.

Another issue was the choice of hashtags before the performance. In the first presentation we included the hashtag of a reality show from the Greek television because it started the same time as the performance. This was not a good choice because tweets with this hashtag overshadowed the rest of the hashtags chosen (including the *#dataminglivescoring* hashtag created by us for the performance). Combined with the problem mentioned above, the first thirty tweets chosen by the algorithm were almost entirely under this hashtag, and the audience could not interact with the music for a certain amount of time, until we removed this hashtag from the algorithm during the performance.

One more issue that arose was that sometimes the Raspberry Pis did not have enough time to create the next score. This occurred when the form the ensemble was playing at a certain point

was a short one (which would occur when there was not a lot of traffic on Twitter) and at the same time the tweets the algorithm was receiving to create the next score were much more. This difference in the amount of tweets resulted in different amounts of music generated, but also in different durations for the system to create a score. When the algorithm had to create a lengthy score while the previous score was short, some Raspberry Pis did not have enough time to render the next score before the main algorithm would initiate the start of this next score. This would cause a crash in the Pi that did not have time to render the next score.

The last problem was the average sentiment analysis of all the tweets for a given form block. Having chosen various hashtags (up to four for each performance) the tweets received were varying in sentiment. This caused the sentiment analysis to have a value around the center of the sentiment extreme values, which resulted in the SVM classifier predicting between the two central forms for long periods.

7. FUTURE WORK

One aspect of this project that we would like to develop further is the technique used to generate the scores. Currently the algorithm chooses elements from a music library written by the composer especially for the project. A future goal is to incorporate machine listening techniques based on neural networks. In this scenario the algorithm would learn from music composed for this project, not in the form of a library composed of musical elements, but in the form of finished works with varying ambiances. Through classification the algorithm would map the various music pieces to sentiment classes. It would then be able to compose its own music based on what it has learned from the composed music and the input it receives during performance. The problems mentioned in the previous section should also be tackled with.

8. CONCLUSIONS

We have presented the concept and process of *Data Mining / Live Scoring*, a live algorithmic composition for the ARTéfacts Ensemble, a six-member acoustic ensemble, based on Twitter. The aim of this project was to create a one hour-long musical piece composed on the spot and sight-read by the performers, embracing unpredictability from input from Twitter. Based on a musical library and nine different forms - or suites - laying out a series of decisions regarding the orchestration, dynamics and tempo programmed into the algorithm, the result was very coherent and abstract at the same time. In this respect, this project was successful since the resulting compositions of the dress rehearsal and the two performances were different, but at the same time shared common elements. Thus, these works gave the impression of having been composed by the same composer. We observed some patterns or tendencies which seemed to develop in parts of the composition. There were moments that there seemed to be a convergence towards certain sonorities or patterns, becoming almost repetitive. The explanation we offered was that this happened due to re-tweeting tweets that had already been analyzed and utilized in the composition.

By prompting the audience members to tweet during the performance they engaged more actively than in a traditional concert situation. By making abstract mappings of the tweets analysis to the resulting music, there was no clear indication as to how a tweet could affect the music other than the sentiment. This way no audience member was able to control the performance in its whole

or in part, thus the control was divided between the algorithm and randomness through the Twitter input.

9. REFERENCES

[1] http://lilypond.org/.

[2] https://openframeworks.cc/.

[3] V. Straebel and W. Thoben, "Alvin lucier's music for solo performer: Experimental music beyond sonification." *Organised Sound*, vol. 19, no. 1, pp. 17–29, April 2014.

[4] C. McKell, "Sonification of optically-ordered brownian motion," in *Proc. of the International Computer Music Conference (ICMC)*, Shanghai, China, 2017, p. 524.

[5] A. Schoon and F. Dombois, "Sonification in music," in *Proc. of the International Conference on Auditory Display (ICAD)*, Copenhagen, Denmark, 2009.

[6] J. Suchanek, "Soil choir v.1.3 soil moisture sonification installation," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, Birmingham, UK, 2020, pp. 617–618.

[7] S. Park, S., K. S., Lee, and W. S. Yeo, "Online map interface for creative and interactive music-making," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, Syndey, Australia, 2010, pp. 331–334.

[8] P. Lindborg, "Pacific bell tower, a sculptural sound installation for live sonification of earthquake data," in *Proc. of the International Computer Music Conference (ICMC)*, Shanghai, China, 2017, pp. 207–209.

[9] M. Funk, K. Kuwabara, and M. J. Lyons, "Sonification of facial actions for musical expression," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, Vancouver, BC, Canada, 2015, pp. 127–131.

[10] R. King, "Music of the people: Music from data as social commentary," in *Proc. of the International Conference on Auditory Display (ICAD)*, Newcastle-upon-Tyne, UK, 2019, p. 104.

[11] https://soundcloud.com/robking-2/gun-licenses-case-study-i.

[12] A. Xambo and G. Roma, "Performing audiences: Composition strategies for network music using mobile phones," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, Birmingham, UK, 2020, pp. 55–60.

[13] E. Egozy and E. Y. Lee, "*12*: Mobile phone-based audience participation in a chamber music performance," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, Blacksburg, Virginia, USA, 2018, pp. 7–12.

[14] M. Gimenes, P. E. Largeron, and E. R. Miranda, "Frontiers: Expanding musical imagination with audience participation," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, Brisbane, Australia, 2016, pp. 350–354.

[15] L. Zhang, Y. Wu, and M. Barthet, "A web application for audience participation in live music performance: The open symphony use case," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, Brisbane, Australia, 2016, pp. 170–175.

[16] https://tinmenandthetelephone.com/.

[17] L.Turchet and M. Barthet, "Demo of interactions between a performer playing a smart mandolin and audience members using musical haptic wearables," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, Blacksburg, Virginia, USA, 2018, pp. 82–83.

[18] J. Eaton, W. Jin, and E. Miranda, "The space between us: A live performance with musicalscore generated via affective correlates measured in eeg of one performer and an audience member," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, London, UK, 2014, pp. 593–596.

[19] A. Endo, T. Moriyama, and Y. Kuhara, "Tweet harp: Laser harp generating voice and text of real-time tweets in twitter," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, Ann Arbor, Michigan, USA, 2012.

[20] L. Dahl, J. Herrera, and C. Wilkerson, "Tweetdreams: Making music with the audience and the world using real-time twitter data," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, Oslo, Norway, 2011, pp. 272–275.

[21] B. Tome, D. D. Haddad, T. Mochover, and J. A. Paradiso, "Mmodm: Massively multiplayer online drum machine," in *Proc. of the International Conference on New Interfaces for Musical Expression (NIME)*, Vancouver, BC, Canada, 2015, pp. 285–288.

[22] K. Ash, "Affective states: Analysis and sonification of twitter music trends," in *Proc. of the International Conference on Auditory Display (ICAD)*, Atlanta, GA, USA, 2012, pp. 257–259.

[23] B. Boren, M. Musick, J. Grossman, and A. Roginska, "I hear NY4D: Hybrid acoustic and augmented auditory display for urban soundscapes," in *Proc. of the International Conference on Auditory Display (ICAD)*, New York, USA, 2014.

[24] T. Hermann, A. V. Nehls, F. Eitel, T. Barri, and M. Gammel, "Tweetscapes real-time sonification of twitter data streams for radio broadcasting," in *Proc. of the International Conference on Auditory Display (ICAD)*, Atlanta, GA, USA, 2012, pp. 113–120.

[25] B. Liu and L. Zhang, *A Survey of Opinion Mining and Sentiment Analysis*. Boston, USA: Springer, 2012.

[26] R. Poonguzhali, S. Vinothini, V. Waldiya, and K. Livisha, "Sentiment analysis on linkedin comments," *International Journal of Engineering Research & Technology IJERT (ICONNECT)*, vol. 6, no. 7, p. 415, April 2018.

[27] S. Kiritchenko and M. Saif M., "Happy accident: A sentiment composition lexicon for opposing polarities phrases," in *Proc. of the 10th edition of the Language Resources and Evaluation Conference (LREC)*, Portorož, Slovenia, 2016, pp. 1157–1164.

[28] https://pypi.org/project/prosodic/.