

# BIOMETRICALLY EVOLVED SITE-SPECIFIC MUSIC AS A RESPONSE TO LOCALISED ACOUSTIC CONDITIONS

Dr Paul Bavister\*

<sup>1</sup> Bartlett School of Architecture, University College London, UK

## ABSTRACT

This paper describes a series of experiments to engage intelligent systems and biometric sensing in a reciprocally creative relationship between computer composed music and physical space, accessing correlations between spatial volume, materiality, and performance.

The paper will review tests undertaken at UCL in 2022, in evolving site-specific music using virtual acoustics, evolutionary programming, and biometric sensing. The paper will describe and define the toolsets involved, neural networks to determine note generation and periodicity, evolutionary processes to determine sequencing and emotional response as a fitness function and summarises how these were applied. The analysis of the outputs is compared against the room's acoustic data for correlations and relationships. Metrics to seek correlations against tempo fluctuations are T30, T20, EDT and C80.

The output of the tests gives us clues as to what future music is likely to appeal emotionally in such spaces for differing listeners' demographics. As the spaces chosen were not typical acoustic musical venues, there are no preconceived ideas about what would or should not sound acceptable in each. If music can evolve to suit a space, then surely each space, however acoustically detrimental, can host something that can be viewed as aesthetically pleasing and site specific.

**Keywords:** *Evolutionary Computation, Biometric Sensing, Neural Networks, Virtual Acoustics*

---

\*Corresponding author: [p.bavister@ucl.ac.uk](mailto:p.bavister@ucl.ac.uk)

**Copyright:** ©2023 First author et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 Unported License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

## 1. INTRODUCTION

This work builds on existing work undertaken by the author looking at biometric evolutionary structures in sound and space. [1] However, the work documented in this paper expands this notion into a more structured process that involves more complex musical structures. Existing examples of evolutionary music have traditionally used interactive cues from a musician as to fitness function [2] and have not used biometric input to generate clues as to fitness. This study uses work undertaken in biometric sensing of galvanic skin response (GSR) to music and room shape [3] The propagation of sound in space and the bodies response to this phenomenon gives us a direct fitness function for evolutionary processes in musical development that directly responds to a rooms acoustic as a component of the generation of new music. This is directly analogous to the feedback loop used by musicians adapting performances to suit a space.

As observed in 'The influence of room acoustics on solo music performance: An empirical case study' there may be a correlation between spatial form, mainly room volume, and thus reverberation with a musician's performance. [4] The study shows that a musician's tempo decreases in reverberant spaces to allow notes to fade to a degree where articulation of melody can drive the form of a piece of music.

If musicians consciously and unconsciously adapt their performance to site, then it asks the question: would it be possible to simulate such conditions in a test where the unconscious emotional response of a listener and a non-musician can 'evolve' a piece of music to a site via evolutionary computation and virtual acoustics?

This work explores this idea and describes how participants listened to evolving music in four sites of differing acoustic qualities that are not constrained to

typical acoustic venues for the consumption of music. To listen to an evolutionary computational process generate a series of ‘aesthetically potent soundscapes’ that shared many qualities with composed music: timbre, pitch and pitch relationships, periodicity, and tempo. These qualities were played through a live convolution reverberation in Max/MSP in the UCL 12 speaker first-order ambisonic sound lab to a listener. Using a biometric sensing, the listener’s emotional change unconsciously sent a fitness value to the evolutionary system; this emotion change would be a response to the music reacting to the space. Over 20 successful generations, a best fit would be achieved between site and musical output.

This output was then be analysed in specialist data analysis software and compared against the room’s acoustic data for correlations and relationships, if any. Key metrics to seek correlations against would be T30, T20, EDT and C80, where there is already an inverse correlation between core acoustic metrics (e.g. C80 and T30 where, under general conditions, a high-clarity index cannot coexist with a high T30 measure).

The output of the tests possibly indicate what future music is likely to appeal emotionally in such spaces for differing listener demographics. As the spaces chosen were not typical acoustic musical venues, there are no preconceived ideas about what would or should not sound acceptable in each. If music can evolve to suit a space, then surely each space, however acoustically detrimental, can host something that will please a certain section of society no matter how small.

This was a project undertaken at the Bartlett School of Architecture in UCL, London UK, and is an architectural design exercise using ideas from acoustics and machine learning. The work described here is an exercise in research through design, a research exercise that borrows from the field of scientific enquiry; a ‘design’ exercise that acts as an exploration in new areas of interest for future development.

## 2. METHODS

### 2.1.1 Tests

**Sites:** Four sites were chosen for the experiment. Each was chosen to represent a differing set of acoustic qualities for music to adapt to. One commonality was that each space had an acoustic quality that was commensurate with a typical performance space and could engender a sense of feedback to a performer. The sites were all in the Barbican area of London and

surveyed by Arup Acoustics. The measurements strictly followed the standards set out in ISO 3382-1:2009 “Acoustics –measurements of room acoustic parameters – Part 1: performance spaces” and used the ‘integrated impulse response method’ by using a sine sweep source to measure the space. The sine sweep was generated by the software Dirac, which plays a ~20 second signal. This was recorded in the space by a Sound Field mic and captured by a 744T recorder. (See Fig 1 for results / data)

Metrics	T20(s)	T30(s)	EDT(s)	C80(db)	D50	LF
Barbican Cold Stores	2.69	2.64	2.62	-3.3	0.24	0.21
CharterHouse Cloisters	1.54	1.55	1.45	0.6	0.34	0.17
Fabric NightClub	0.99	0.97	0.95	4.6	0.61	0.31
Barbican Foyer	2.9	2.87	2.43	-1.3	0.29	0.16

**Figure 1.** Table showing acoustic metrics from the four sites. (Source ARUP Acoustics)

**Participants:** Over the pilot tests and full experiments, 53 healthy student volunteers were recruited from University College London by advertisement. The breakdown was 20 male and 33 female participants, and the mean age was 28. All test participants were informed of the procedures of the tests by written statement and invited to take part by advertisement. All participants gave their written informed consent and were paid for their participation.

**Ambisonic Sound Lab / Equipment:** All listening tests were undertaken in the UCL sound lab, a first-order ambisonic lab with 12 individually addressable Genelec speakers in a cube layout. The room was box-shaped, with dimensions of 3.49 m (width), 3.35 m (length), 3.16 m (height) A comfortable armchair on which the listener was seated was positioned centrally. The reverberation time (T30,500Hz–2kHz) measured in the test conditions with the interrupted noise method (6 microphone-source combinations, 2 source positions, 3 decays in each position) was 0.13s. The stimulus was presented by Cycling ‘74’s Max/MSP software running on a MacBook Pro 2020. Sound was delivered via USB to an RME Fireface USB interface, then spilt across 12 channels.

**GSR Dermal response to Stimuli:** Skin conductance was used as a measure of emotional response to be in line with previous tests undertaken in GSR & music, [5] and sound. [6] Sensors were placed on the medial phalanges of the non-dominant hand, just as in the other tests. The equipment used was Grove Seed’s GSR sensor running into an Arduino microcomputer, with the serial data running

into Max/MSP. Running the data into Max/MSP allowed for live analysis and to output triggers if the system detected a significant change in the data. The Max patch developed for dermal response analysis would send a trigger to the evolutionary programming code if the GSR stream had varied +/- 5 ms in the past 3 seconds. The system could be calibrated to deal with listeners with livelier skin responses, or flat responses. Max was only using the raw data from the Arduino which had been passed through the Grove Code and took out a little noise. The final stream was recorded in a text file and then imported into Matlab's Ledalab for final analysis. This would further triangulate the results and ensure that emotional responses had occurred when the system believed it had.

The listeners were seated in a comfortable chair with the test equipment on their fingers, positioned as discreetly as possible. The non-dominant hand was resting on an adjacent table for comfort, and the participants were asked to find a comfortable position and avoid as much unnecessary movement as possible. During the tests, the observers sat outside of the sound lab to monitor the stimulus presentation and record the physiological data. The experiments lasted for 30 to 40 minutes in total.

**Evolutionary Programming:** The major optimisation process undertaken in both tests was evolutionary programming (EP). EP has benefits and drawbacks and must be tailored to any given situation if a degree of success is to be achieved.

The most significant issue that arises in the tests is that of the search space. Having a larger number of variables in the initial genotypes gives more opportunities for diversity in the results, and a greater number of usable solutions. But this increase in variables requires more time to go through the options available and test each one according to a success criterion. A typical computational optimisation search, such as a genetic algorithm, will undertake thousands of generations before it reaches any degree of cover in a large search space. This is because all the variables are known and the system can simply crunch through the options until an optimum is reached, either user-defined or time-based. In an interactive situation, this is impossible; a user will reach saturation after a limited number of options, and further testing will be meaningless. Thus, the search space must be limited to a reasonable limit where a meaningful result can be achieved from limited variables.

In both tests, the search space was limited as much as possible to reduce the search space and limit user fatigue.

In a 'search' where epochs are limited to tens rather than thousands, it is also advantageous to generate a meaningful result as quickly as possible. All the existent works in

interactive evolutionary systems in music use very simple systems of fixed genotypes to generate simple melodies, [7] [8] simple systems with simple inputs. The architectural and acoustic searches did not use user interaction and could benefit from expanded search fields of thousands of generations. To resolve this problem, it was decided to tighten up the mutation coefficient every epoch, so that the tests would start off with a wide search field, and for every emotional response recorded, the mutation would get smaller, so that the range of options offered to the participant would reduce. Over 20 generations, this would result in an exponentially reduced search space, and allows trends to be seen over the results of the listeners. Common to both tests was a limit of 20 responses which would reduce the mutation from 50% variability to 1%.

**Post Test Questionnaire:** After each test, the listener was questioned on their thinking during the test to detect positive and negative valence. This was important as the GSR output can detect an 'emotional response' but cannot detect if the reaction was positive or negative; both emotions would be read as the same. Thus, a sound that elicited a negative valence in a listener would be read as a false positive and the test would have offered the listener more options based on something they felt was incorrect. The post-test questions removed this risk, and tests with a negative valence were removed from the results. The post-test questionnaire also allowed the listener to say if they had lost concentration during the tests, and if they were worried about some outside factor that may have influenced the readings. Again, if any listener commented that they were not fully engaged, the results would have been removed from the system. However, in practice this never happened, and all listeners were engaged.

### 2.1.2 Sound Generation

**Note Generation & Neural Networks:** To train the system, into generating a meaningful output a pool of notes was defined and assembled in a self-organising map (SOM). This uses Smith and Garnett's ML Star series of objects for Max/MSP. [9] The SOM is an unsupervised clustering and classifying tool, mapping high-dimensional input data on to a two-dimensional output space. This is a type of artificial neural network, or ANN. It was developed to visually show interrelationships between input data sources so that it was quick and intuitive to see the organisation of the source data. When an input is presented to the SOM - in this case, midi data of a selected piece of music - a search is performed to locate similarities in the data. By repeatedly training the map, distribution patterns can emerge and be clearly visible.

Use of the map is undertaken in three steps. Firstly, the midi data must be read to encode the notes into the neural network. The encoding process takes a stream of notes and turns them into a vector that tells the neural network how recently each note has been seen. This training is repeated several times to ‘enrich the’ training. Secondly, the neural net finds similarities in the note vectors and arranges them across the map in a two-dimensional lattice. Lastly, a note generator will use spatially-encoded data as a probability distributor to generate new notes. Whilst this process is running, selecting different areas of the map will change the pitch set.

The above system is implemented in the SOM.MIDI demonstration patch that comes with the ML Star set of objects. It requires a manual input to change the pitch set. In this test, in this case the SOM was adapted so that the pitch selection process is undertaken by an evolutionary system and not a mouse click.

The input source for the map is a MIDI file of an established piece of music traditionally associated with a set space.



**Figure 2.** A self-organising map (SOM) with 32x 32 nodes showing note vector distribution and similarities of De Bussey’s Clair De Lune.

**Pitch Data Sets & Site:** Because of its variation and its ability to straddle classical and contemporary tastes, the

data set used across all the tests was Debussy’s ‘Clair de Lune’. The work has a wide sonic range, from notes C0 to D4, which gives the data set a wide range of options to generate an evolution. The pitch relationships between the notes are proven to induce emotional response in that the work is a fundamental part of the Western canon, crossing the late Romantic period and modern movements in music. The voices used in these tests were based on electronic musical instruments (VSTs) in Apple’s Logic Pro software. There were 10 voices used with a wide range of tonal capacity, from pads creating drones and tones to percussive hits and short sharp notes. The intention was that any form of ‘music’ could be generated via this method.

To limit computational memory, five active voices were being generated in Max, each selected by one of the chromosomes in the evolutionary process (see following section on evolutionary sequencing below.) This made it impossible to have all ten voices being used in Logic, so each piece of evolved music was limited to 5 components or sections. This limitation actually increased the variability of the output and allowed a greater range of sound to emerge.

**Periodicity:** Rhythm and periodicity functions were managed by Karlheinz Essl’s Markov rhythm object, which generates a table of differing rhythmic sequences and uses a Markov matrix to ensure that there is no repetition of each sequence. The object requires an entry delay (or ED) to be defined by the user on a scale of minima and maxima. This is undertaken in milliseconds, so an ED scale of 500 to 5000 would mean that rhythmic events could take place every ½ a second or up to 5 seconds, and generating non-repeating spaces between notes. This timing can generate musical and non-musical passages, so in the case of typical classical music - and ‘Clair de Lune’ - the scale of entry delays was kept between 1s and 4s.

To ensure that a degree of emergent polyphony could be generated, five independent rhythmic systems were in place, each governed by a single evolutionary process. Each system would share the note pool generated by the SOM but send differing rhythmic pulses to the ten independent voices set up in Logic.

### 2.1.3 Evolutionary sequencing

The sounds generated by the system had to be a musical and acoustic ‘best fit’ the chosen space as fully as possible, and this could only be undertaken by a search process that allowed a sequence of trial and error to finally find the appropriate notes, amplitudes and durations. The aim of this process was to match the aspirations of a composer and meet the emotional

expectations of an audience. This is analogous to a performer seeking to modulate a performance to fit the natural acoustics of a room. However, as this was not in any way considered to be a live event and durations were not fixed, time could be taken to find a musical mean that would meet an end goal. The search was replicating a series of performances that in reality would have been undertaken over years, not hours.

An evolutionary algorithm was chosen to undertake the sequencing of the notes and rhythmic structures. By using the pitch set from the SOM, voice selection and the entry delays from the Markov rhythm object, a series of elements were in place that could be modulated by a heuristic process to develop a musical structure that could be a 'fit' for an existing space and a listener.

**Assignment of Musical Properties:** To enable an analysis of musical and acoustic performance, properties related to pitch set selection and other musical items had to be defined in the genetic code of the algorithm. Core components of the 12-genotype chromosome were aligned with modulating note selection and periodicity functions.

For simplicity, each integer in the population was a variable of 1 to 20. Undertaking the exercise in hexadecimal integers, using 0s and 1s (like Sato and others) would have resulted in extremely long chromosome sequences and increased the search time considerably. Using floating point integers (like Spaeth et al.) would also have been acceptable, running from 0.1. But that limits the initial population to only 10 integers, in turn limiting the scaling function that follows.

**Algorithm Output & Selection:** The algorithm outputted 4 genotype chains, and this was fed into the system for evaluation. Each genotype had 5000ms of time before moving on to the next, so that each 'generation' of the algorithm took 20s. If any of the genotypes met the fitness function, then that was selected by the algorithm to seed the next generation.

The algorithm was a single-parent type of evolutionary process, so that one genotype could seed subsequent generations. A single-parent algorithm removes the need for extensive evaluation and choosing of suitable parents, and is essential in interactive algorithms where a user may not have time to undertake significant evaluation - although in self-contained systems such as this, a more traditional pairing of successful candidates may be useful. [10]

As the timing of each generation was a duration of 20s, use of a genetic algorithm and recombining successful phenotypes would have been extremely time-consuming. A typical generation sequence of 10,000 generations would have resulted in a recording of 12 hours and very large file sizes, where the first 8 hours would be largely random combinations before anything of value could be gained. For this proof of concept, it was decided to remove the recombination aspect of the algorithm and use a simple evolutionary search process to speed up the process and allow easier evaluation of the final recordings. This process was faster than a traditional genetic algorithm, but at the cost of fully evaluating the search space. However, as the system was faster, the system could be run more times, covering the search space just as efficiently. Repetition of the system could show important trends in the search direction, and this was of value.

**Fitness Function:** An evolutionary process such as this required a goal, or optima, that the algorithm was seeking to achieve. This is difficult with music, as 'good' music is a subjective quality, requiring human ears and cognition, as well as experience and knowledge to determine what was worth keeping, and what was worth discarding. As this test was automated with no conscious human input to determine fitness, another optima was required. In the context of composers like Haydn writing for rooms [11] it was decided to develop an optima that would satisfy contextual acoustic conditions rather than more subjective musical conditions. Analysing unconscious emotional response to the musical output in the context of the space would provide a metric that would give an indication of the performance in the space, and how well it was adapting to site-specific conditions.

Dermal response was used as a fitness function, allowing a level of interaction from a listener to modulate the output of the tests. This allowed users to interact with the tests on an unconscious level and give an indication of engagement with the test procedure.

**Mutation & Termination Condition:** To reduce user fatigue, the test was designed to last for 3 minutes so that the participants' attention was at an optimum. During this time, any number of emotional responses could be recorded. As in the Evolutionary Space test, mutation was exponentially decreased so that after 20 successful generations, no further mutation would take place and the test would be at an optimum. Note that due to the limitations of the test a full and final optimum will never be fully reached. It is only after analysis of multiple

results and trend patterns will a direction of travel be determined.

**Outputs:** As the voices in all the tests were being generated in Apple's Logic Pro software, the sounds were routed back to the main patch in Max/MSP for analysis and recording. The outputs were recorded in stereo AIFF-C format at a bit rate of 48kHz. These recordings were then analysed in Adobe's Audition software, Sonic Visualiser and Partiels, an analysis tool developed by IRCAM. Core metrics under review were the mean tempo of any generated piece.

### 3. RESULTS

#### 3.1.1 Questionnaire Output

Before any results were meaningfully analysed, the full questionnaire data for each participant was fully reviewed, as any participant who had displayed a considerable number of emotional responses but had commented that they did not find the experience enjoyable, and expressed a negative valence in the questionnaire, would have to be removed from the results.

The mutation coefficient allowed 20 generations before there would be no further change. As the mutation was inversely exponential, there was much variation in the lower integers, before finally reaching a stable conclusion around the 18th generation. This meant that any participant who scored a figure of <18 emotional responses could also be ruled out as having not reached a definitive point in the process.

What was observed in the emotional responses alone is that the Fabric nightclub was hugely popular, with 231 emotional responses; the other spaces did not even reach half that figure. It should be noted that the Barbican foyer also proved a highly popular space even though it was only introduced for four tests. As the responses of the four tests closely followed that of Fabric, it could be assumed that this relationship would follow had the space been used for all tests, resulting in a final emotional response similar to, or exceeding that, of Fabric.

#### 3.1.2 Analysis

All the sound files generated in the tests were reviewed in IRCAM's Partiels software, created and developed by Pierre Guillot at IRCAM. This is a very powerful suite of analysis tools that allows a user to analyse sonic information and export it as a text file for further analysis and representation in software such as

Microsoft's Excel spreadsheet software, or Hierarchical Regression Analysis, as undertaken here.

Analysis of the results showed that the larger volume rooms had the slowest outlying tempos, and the drier, less reverberant-sounding rooms had the fastest tempos; however, the mean tempos showed only a slight variation of +/- 6 BPM across all tests. The difference in outliers is +/- 28 BPM, which is musically significant.

The output of the BPM analysis was put into IBM's SPSS data analysis software, where a Pearson correlation coefficient was sought between the mean BPM for each test and the acoustic data for each room. This is a measure of the linear association of two variables, say tempo and C80. Values for each correlation vary between -.1 and +1. Positive values of correlation coefficient reflect a tendency for one variable to increase or decrease together with another variable. A negative value of correlation coefficient indicates a tendency that the increase in the value of one variable is associated with the decrease in another, and vice versa. Values close to zero reflect little or no correlation.

Key aspects worthy of note were a negative Pearson Correlation against tempo of -.965 for EDT, -.956 for T20, -.996 for T30, reflecting an inverse relationship. In instances where these metrics were lower, in drier environments the tempo increased, and when these metrics were higher, the tempo decreased. There were positive correlations between C80 at 0.999\* (the star indicates a significant Pearson Correlation) and C50.

#### 3.1.3 Discussion

All the resultant output of the tests was broadly similar in aesthetic feel; there were no immediately obvious differences between each output on which to make an objective observational assumption. In-depth analysis was required to reveal any inconsistencies and correlations between the recorded output and the host spaces.

Statistically, the more reverberant spaces generated the lowest BPMs and the drier spaces generated higher BPMs. At 26BPM, this variance is significant, and a tempo difference of such proportions is clearly audible.

EDT(s)	Pearson Correlation	-.965
	Sig. (2-tailed)	.169
	N	3
T20(s)	Pearson Correlation	-.956
	Sig. (2-tailed)	.189
	N	3
T30(s)	Pearson Correlation	-.966
	Sig. (2-tailed)	.167
	N	3
LF	Pearson Correlation	.829
	Sig. (2-tailed)	.378
	N	3
LFC	Pearson Correlation	.599
	Sig. (2-tailed)	.591
	N	3
C80(dB)	Pearson Correlation	.999*
	Sig. (2-tailed)	.035
	N	3

**Figure 3.** Pearson Correlation between tempo and acoustic measures. Significant Correlations are marked with an asterix.

The result of all analysis demonstrates that there was a clear observable correlation between room acoustic, metrics and tempo. Conditions reviewed were EDT, T20, T30 and C80. The tests replicated the observations of Kalkandjiev and Weinzierl in seeing a relationship between T30 and tempo. However, it should be noted that Kalkandjiev and Weinzierl were strictly analysing music played by an active musician, a musician who was engaged with the music to a high degree to ensure that the performance was as rehearsed and previously performed. The participants in these tests were inert and had no prior experience of the repertoire and convention. They were happy to let the music flow naturally, and the number of emotional responses shows this. There was clearly a correlation between emotional response and content, as well as correlations between room size and content. This is highly significant and shows a level of success in the tests.

The process outlined in this document was limited to a study of tempo correlations between listener and space, but there are many more metrics to extract from the

source audio, and future work should be undertaken to analyse other data. This will need to be adapted to a series of integers that will allow for appropriate statistical analysis in SPSS.

#### 4. CONCLUSION


This test was developed to explore the possibility of generating a bespoke and site-specific piece of ‘music’ to a space. Building on Kalkandjiev and Weinzierl’s observations that musicians vary tempo to relate to the reverberance of a space, it would theoretically be possible to undertake a test where tempo is unconsciously modulated through an evolutionary process by non-musicians. This was achieved and correlations can be observed between tempo and room size emerging from the tests. The evolutionary process has in effect replicated the intuitive response of a musician to a site, and the emerging music that arose from this process is linked to the listener and the site via an evolutionary process.

#### 5. ACKNOWLEDGMENTS

I would like to acknowledge the support given to me in this work by Professors Penelope Haralambidou, Stephen Gage & Raf Orłowski of the Bartlett School of Architecture, Tin Oberman & Soundscape Indices for the kind loan of the UCL Sound lab and Priya Tadinada for assistance in processing the outputs from the tests.

#### 6. REFERENCES

- 
- [1] Bavister, P. (2018) Artificial intelligence and the generation of emotional response to sound and space. *Proceedings of the Institute of Acoustics*
- [2] Biles, J. (2002), ‘GenJam in Transition: from Genetic Jammer to Generative Jammer’, *International Conference on Generative Art, Milan, Italy*
- [3] Lokki, T. & Pätynen, J. (2011), ‘Lateral reflections are favorable in concert halls due to binaural loudness’, *The Journal of the Acoustical Society of America*, 130(5), EL345–EL351.  
<https://doi.org/10.1121/1.3647866>



---

[4] Kalkandjiev, Z. S. & Weinzierl, S. (2013), 'The influence of room acoustics on solo music performance: An empirical case study', *Acta Acustica United with Acustica*, 99(3), 433–441.  
<https://doi.org/10.3813/AAA.918624>

[5] Pätynen, J. & Lokki, T. (2016), 'Concert halls with strong and lateral sound increase the emotional impact of orchestra music', *The Journal of the Acoustical Society of America*, 139(3), 1214–1224  
<https://doi.org/10.1121/1.4944038>

[6] Benedek, M. & Kaernbach, C. (2010), 'A continuous measure of phasic electrodermal activity', *Journal of Neuroscience Methods*, 190(1), 80–91.  
<https://doi.org/10.1016/j.jneumeth.2010.04.028>

[7] Nomura, K. & Fukumoto, M. (2017), 'Asynchronous distributed interactive genetic algorithm for creating music melody reflecting multiple users' feelings', *Proceedings - 18th IEEE/ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing, SNPD 2017*, 645–650.  
<https://doi.org/10.1109/SNPD.2017.8022791>

[8] Tokui, Nao & Iba, Nao & Iba, Hitoshi. (2000). Music Composition with Interactive Evolutionary Computation. *International Conference on Generative Art*.

[9] Smith, B. D. & Garnett, G. E. (2012), '*Unsupervised Play: Machine Learning Toolkit for Max*', retrieved from <http://cycling74.com>

[10] Rata, I., Shvartsburg, A. A., Horoi, M., Frauenheim, T., Siu, K. W. M. & Jackson, K. A. (2000), 'Single-parent evolution algorithm and the optimization of Si clusters', *Physical Review Letters*, 85(3), 546–549.  
<https://doi.org/10.1103/PhysRevLett.85.546>

[11] Meyer J. (1978), 'Raumakustik und Orchesterklang in den Konzertsälen Joseph Haydns', *Acustica* 41, no.3 pp 145-162