<https://helda.helsinki.fi>

Analysis of Noisy Biosignals for Musical Performance

Paalasmaa, Joonas

Springer-Verlag

2012-10-25

Paalasmaa , J , Murphy , D J & Holmqvist , O 2012 , Analysis of Noisy Biosignals for Musical Performance . in Advances in Intelligent Data Analysis XI : 11th International Symposium, IDA 2012, Helsinki, Finland, October 25-27, 2012. Proceedings . vol. 7619 , Lecture Notes in Computer Science , vol. 7619 , Springer-Verlag , pp. 241 252 , International Intelligent Data Analysis (IDA) , Helsinki , Finland , 25/10/2012 . https://doi.org/10.1007/978-3-642-34156-4_23

<http://hdl.handle.net/10138/37400>

https://doi.org/10.1007/978-3-642-34156-4_23

submittedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Analysis of Noisy Biosignals for Musical Performance

Joonas Paalasmaa¹, David J. Murphy², and Ove Holmqvist

¹ Beddit.com Ltd, Espoo, Finland;

Department of Computer Science, University of Helsinki, Finland

`joonas.paalasmaa@beddit.com`

² Meld Ltd, Helsinki, Finland

`david.murphy@meld.fi`

`ove@ove.fi`

Abstract. Biosignal sensors are now small, affordable, and wireless. We desire to include such sensors (e.g. heart rate, respiration, acceleration) in a live musical performance, which sets requirements on the reliability and variability of the data. Unfortunately the raw signals from such devices are unable to meet these requirements. We contribute our solutions for overcoming the shortcomings of these sensors in two parts. The first is an online data processing and analysis system, including on-line generative models that describe the signals but add consistency. The second is the end-to-end system for capturing wireless signal data for the analysis system and integrating the resulting output into a popular digital audio workstation in a very flexible manner conducive to live performance. We also explore the role of “analysis supervisor”—a member of the performing act who ensures that the results of biosignal analysis fall within the desired ranges to contribute to the music effectively.

Keywords: ECG, respiration, accelerometer, music, performance

1 Introduction

This paper presents a novel way to integrate biosignal measurement into live musical performance. We present a complete system from signal acquisition, signal analysis, data routing, to integrating the signals into instruments and effects in a digital audio workstation (DAW).

The motivation for our work comes from the desire to utilize inexpensive and convenient biosignal sensors in musical performance. This raised two technical challenges. The sensors themselves provide noisy and unreliable measurement that cannot be directly utilized in musical use. For example, directly using the heartbeat markers from a noisy electrocardiography (ECG) signal to trigger notes is not possible, because missing and spurious heartbeats will disturb the rhythm. The second challenge is that biological data has high-variability in both timing and magnitude, and incorporating biosignals into a composition can entail massaging the data. We needed an online solution to both of these challenges to facilitate live performance.

The contribution of this paper is twofold. First, we present biosignal analysis methods that turn noisy measurements of irregular biosignals into consistent and uninterrupted outputs, which are suitable for musical use. Second, we present a practical implementation for integrating the analysis into live musical performance.

We have tried to build the analysis system so that it fulfills the following requirements. First, it should be usable by a musician: it should have a sufficient range of output to be expressive, it should be robust in the face of noisy biosignals, and its working mechanism must be learnable, with a clear mental model of operation.

Second, it has to be possible to compose for it, such that a performer can communicate the composers intent faithfully. There is a tendency to use physiological signals in electronic music as audio generators, that is, the signals themselves are directly translated into sounds as synthesized audio or triggered samples. It is very difficult to purposefully convey intent within that framework. We see a distinction between the direct sonification of biosignals, and their integration into a musical performance. If a conflict arises between expressing the biosignals with high fidelity, or achieving the intent of the composer with high fidelity, our bias is towards the latter.

Finally, biosignal measurement should contribute towards the musical performance, not dictate it: we have elected to make the output of signal analysis malleable, and envisage one member of the group in the role of the “analysis supervisor”, carefully monitoring and if necessary shaping the physiological data to suit the desired expression of the piece. For example, signals may be of higher magnitude than anticipated due to the excitement of performance, and require scaling down to fall within the required range for modulating a parameter such as a synthesizer cut-off frequency.

A video showcasing our approach with a solo performance is available online³. In the video, it is shown how a synthesizer can be controlled with biosignal measurements.

2 Background

Generating audio and music from physiological signals has been the subject of research since the invention of the encephalophone in the 1940s at the University of Edinburgh [1]. It generated audio from measured electroencephalography (EEG) signals. Various artists have since used brainwaves in music. Krzysztof Penderecki’s *Polymorphia* (1963) used pitch notation derived from EEG data from patients listening to a recording of his *Threnody for the Victims of Hiroshima* (1961).

³ <http://vimeo.com/42032074> — The performer triggers the synthesizer with heartbeats and modulates the filter envelope amount with respiration. Tilting the body back and forth adjusts the low frequency oscillation rate, while bending sideways controls the arpeggiator distance.

Alvin Lucier’s *Music for Solo Performer* (1965) is considered to be the first music performance to use real-time biosignals, where EEG triggered percussive instruments through sympathetic resonance [2]. Richard Teitelbaum’s *In Tune* and *Organ Music* (1967/1968) used EEG, ECG and respiration measurement [3]. The research in biomusic by Manfred L. Eaton [4] inspired Erkki Kurenniemi to build instruments based on galvanic skin response (GSR) and EEG in the early 1970s [5]. In 1973, Pierre Henry performed his EEG-based piece *Cortical Art III*. David Rosenboom used neurofeedback and algorithmic composition in *Brainwave Music* (1976) and *On being invisible* (1977), which compared performers’ brain activity with previously stored patterns in real time [6].

The first commercial biomusic production environment, *BioMuse*, was introduced in 1992 [7,8].

Various contemporary projects have used EEG and ECG in music generation. *The Multimodal Brain Orchestra* measures EEG event related potentials, with an “emotional conductor” controlling the “affective expression” of the performance [9]. Brain computer interfaces (BCI) enable disabled people to perform music with EEG [10]. In *DECONcert*, EEG and ECG of 48 participants were monitored and methods such as signal averaging and detection of collective alpha synchronization were used in order to facilitate collaborative biofeedback in regeneration of the music [11]. A traditional chamber music performance has been augmented by EEG and ECG signals [12].

With the interactive cinema environment *Biosuite*, the audience’s emotional responses are measured using ECG and GSR in order to influence cinematic events [13]. In another project, ECG measurement from healthy and diseased hearts are mapped into musical notes, in order to highlight their differences⁴. In [14], respiration and cardiac events are translated into synthesized sound. The timbral brightness of the sound responds to respiratory sinus arrhythmia.

Apart from the academic research, there are performing musicians who utilize biosignals to varying degrees. *Lucky Dragons* use a GSR-touch interface to connect with their audience. *The Heart Chamber Orchestra* measures the orchestra’s ECG and presents the data to them as notation, which they subsequently perform. The dance and performance company *Manifold Motion* uses real-time ECG biofeedback in some of their shows.

3 Biosignal Analysis System

3.1 Signal Acquisition

Our system captures biosignals with two Bluetooth chest sensor belt models from Zephyr Technology Corporation⁵. Zephyr BioHarness provides a range of data including heartbeat events, a respiration effort signal and a 3-axis accelerometer signal. Zephyr HxM provides only heartbeat events and a rough estimate of activity level, but is more affordable. The signals from the sensors are transmitted

⁴ <http://polymer.bu.edu/music/>

⁵ <http://www.zephyr-technology.com/>

over Bluetooth to a computer where they are processed and forwarded to the DAW. We have implemented drivers for the Bluetooth protocols of the devices in Python (source code is available for download⁶). The implementation turned out to be surprisingly complex, because issues such as varying latency and connection breaks need to be taken into account.

The Zephyr devices transmit the signal data in packets. For example, the acceleration signal is sampled at 50 Hz and transmitted over Bluetooth in 20-sample packets. That creates a 0.4-second measurement latency for acceleration data. The Bluetooth connection can occasionally drop out completely, which has required us to implement reconnection logic, so that gaps in measurement during a performance would be as short as possible.

The latency problem could be avoided with a device that has a smaller packet size. However, the latency is in practice tolerable so we have not tried to find alternative devices. Replacing Bluetooth with a wired connection would improve the reliability of the connection, but would also impair the usability of the system in musical performance use too much.

3.2 Dealing with Noise and Missing Data

The way the physiological parameters are used in a performance motivates a signal analysis approach where noise and missing values are filled in with generated information.

For example, heartbeat events are primarily used as sample (e.g. drum) triggers. If the measurement of a heartbeat is missed, a triggering of the drum will be missing, and that disturbs the performance. Therefore, missing and corrupted measurements need to be filled in in a way that is musically sensible. That process is mostly automatic, but may at times require human intervention.

The mediation of signal analysis is performed by an “analysis supervisor”, who is one of the musicians in the performance and knows a musically optimal way to massage the data, having practiced the performance with the other musicians, and being keenly aware of the moment-by-moment intent.

The flow of physiological measurement information during a performance is summarized in Figure 1. It can be described in three steps:

1. Biosignals are measured (ECG, respiration effort and acceleration) from one or more people.
2. The biosignals are analyzed so that noise and missing data is taken into account. If needed, the analysis is mediated by the analysis supervisor. This produces consistent results for the physiological parameters (e.g. heart rate, respiration signal and activity level).
3. A musician uses the physiological parameters in real-time as part of the performance.

⁶ <https://github.com/jpaalasm/zephyr-bt>

Fig. 1. The flow of information during a performance.

3.3 Heartbeat Interval Analysis

A healthy human heart beats at intervals that vary relatively little. For example, at rest, the interval between successive heartbeats might vary, say, between 0.8 and 1.2 seconds, corresponding to around 60 beats per minute (BPM). When heartbeats are detected from a noisy ECG signal, there can be incorrectly detected heartbeats and several-second long gaps in heartbeat detection. When a chest heart rate belt is used, measurement errors are frequent since signal quality is relatively low.

The aim is to output a consistent sequence of heartbeat events from ECG measurement, so that the sequence of events does not have incorrect heartbeats or measurement gaps. The analysis is mostly automatic, and the analysis supervisor just sets upper and lower bounds for allowable heartbeat intervals. That information is needed in cases where signal quality is bad.

The heartbeat interval analysis procedure is visualized in Figure 2. First, heartbeat intervals are analyzed by an automatic ECG heartbeat detection algorithm, which tries to find heartbeat intervals from the signal. The result of that analysis may contain some erroneous values, but the detected heartbeat intervals should mostly be correct.

Then, the intervals are manually limited to a range by the analysis supervisor. Two erroneous short heartbeat intervals are shown in the second row of Figure 2 at time 13 seconds. They are discarded by the manual limitation, because they are outside the specified range. If the previous step in the analysis works correctly, no manual limitation is needed.

Last, a consistent heartbeat sequence is generated to fill the gaps in the heartbeat interval data. A consistent heartbeat sequence is created by generating synthetic heartbeat positions to the gaps in the heartbeat sequence from the previous step. The heartbeat positions are generated from an Inverse gaussian distribution with a manually specified variability parameter λ [15]. The end result is in Figure 2d.

With a good-quality measurement, the heartbeat events go through the three steps of analysis unchanged, as no heartbeat intervals need to be discarded and no gaps need to be filled.

Fig. 2. The analysis of heartbeat intervals. In (a), an ECG signal with detected heartbeat positions is shown. The corresponding heartbeat intervals are in (b). The analysis supervisor limits heartbeat intervals to range 0.55...1.1 seconds, to discard erroneous heartbeat intervals, and the result of that is in (c). The resulting consistent heartbeat event sequence is in (d).

3.4 Respiration Signal Analysis

The BioHarness sensor measures a chest respiration effort signal. Variation in signal quality is large. The signal sometimes tracks respiration in a sinusoidal manner, but, at other times, has no clear structure. We have not found out why the signal quality varies so much.

The analysis of the respiration signal consists of defining the output as a linear combination of the measured respiration effort signal and a generated synthetic sinusoidal. The amplitude, frequency and phase of the synthetic signal are estimated from the measured signal, so that it is possible to transition as smoothly as possible between the measurement and the synthetic signal.

The task of the analysis supervisor is to continuously adjust a weight parameter that defines the linear combination. The manually chosen weight is a number between 0 and 1, and values in the output signal are simply generated as

$$output = measured \times weight + synthetic \times (1 - weight).$$

This means that the measurement is output when $weight = 1$ and the sinusoidal when $weight = 0$. Values between 0 and 1 output a mixture of the signals.

Manual weighting between the real measurement and a synthetic signal enables smooth swithing between them during the performance, when signal quality changes. That is visualized in Figure 3.

3.5 Accelerometer Signal Analysis

The three-axis accelerometer signals from BioHarness include information about the posture and activity level of the measured person. Posture is measured as two

Fig. 3. A measured respiration signal is shown in the topmost plot, and a corresponding synthetic sinusoidal in the second plot. The output of the analysis is a weighted sum of the signals. Smooth switching between the signals happens at time 12...14 seconds.

parameters: tilt (back and forth) and sway (sideways), by simply calculating the angles of the accelerometer’s axes in space. The activity level is calculated with a moving average of the high-pass filtered and squared accelerometer signal. The high-pass filtering and squaring is done for each axis separately, and the resulting three values are summed to give the activity level. The length of the moving average is controlled by the analysis supervisor, and can vary between 0.1 seconds to 10 seconds. The activity analysis is visualized in Figure 4.

3.6 Networking

The biosignal analysis and music workstation software may run on different computers, so there needs to be a reliable way to transmit the physiological parameters between them. Our goal was to create a system where all parameters are available to all music workstations in the performance, regardless of the number of devices involved in capturing the sensor data.

The solution we have settled on is broadcasting the data as Open Sound Control (OSC) messages [16] over User Datagram Protocol. In the composition and synthesis environment, performers “tune in” to a physiological parameter of interest based on OSC message tags. Using OSC requires us to create a wired or wireless network for our participating computers. Basic multi-computer performance functionality such as synchronizing MIDI clocks depends on being networked, so this does not introduce a new requirement to the performing environment.

Fig. 4. Activity analysis example. The upper plot contains the three-axis acceleration signals, with orthogonal measurement directions X, Y and Z. An activity value signal calculated with a three-second moving average is shown in the lower plot.

4 Performance

The performance is the ultimate goal of our work. As digital technology makes recorded music ubiquitous, the importance of the performance as a unique event becomes elevated. We hope that linking the music to the musician’s physiological signals will result in a performance that feels intimate, unique and a little magical.

A challenge with performances utilizing biosignals is the evolving and unpredictable nature of the signals, and the burden that monitoring for bad signals places on a performer. Viewed over an entire performance, the rates and magnitudes of a given signal can change drastically, or the signal may be lost entirely. Our solution to the problem of managing the data flow in the service of the performance is the role of the analysis supervisor, who ensures that other musicians do not have to worry about problems related to the signals.

We feel it is important to make the link between the physiological signals and the resulting sounds as direct as possible, so an audience is not mystified by how the sound is being produced, which has led us to favour straightforward translations from signal to sound. Mappings that we consider intuitive include:

- triggering drums with heartbeat events
- matching the performance tempo to heart rate
- modulating synthesizer parameters to reflect respiration phase
- modulating synthesizer parameters such as pitch-bend using pose
- modulating the frequency spectrum of the performance using activity level

We have created a functional music production environment that allows us to make performances based around the mappings we have defined as being intuitive to us. The system uses Cycling 74’s MaxForLive⁷ plug-ins for the popular

⁷ <http://cycling74.com/products/maxforlive/>

Ableton Live digital audio workstation. Ableton Live organizes music performances into tracks, with one track per instrument. Each track can have a rack of modulating effects. Using MaxForLive, we can create instruments that can be embedded in Live's effect racks, and we can stack as many effects per rack as we like. Thanks to our "tunable" broadcast network architecture, any performer's signals can be used by any instrument.

The most basic instrument we have simply assigns the BPM of the track to the BPM of a measured heart. In this manner it is very trivial for us to control the tempo of the performance using one of the performers' heart rates. By carefully controlling their exertion, a performer is able to raise or lower the tempo within reasonable bounds.

The Modulator instrument accepts data from our signal sources and performs some basic mapping: a higher and lower bound for the output can be set, and a squashing function is provided, which allows a signal to be mapped to all or part of the MIDI range for a selected parameter, and furthermore be amplified or diminished depending on the desire of the performer. The output can be dynamically assigned to any suitable parameter in the rack. This is very flexible—we can use it to assign a setting value to any exposed user interface element on any instrument or effect in any track, and we can have multiple instances per track. In this way we have been able to create instruments where performer body position can be used to alter the cut-off frequency on a modelled analog synthesizer, and passages where instrument volume is effected by performer activity level. See Figure 5.

Fig. 5. An example phrase composed for keyboard and biosensor - the horizontal lines indicate two held notes, and the rising graph indicates the composed value of the parameter "activity level". The parameter is linked to instrument volume. In the performance the performer must try to increase their activity for the duration of this phrase, so the instrument swells in. The analysis supervisor must scale the output of "activity level" so it measures from -50% to +50%

The Beat Train instrument is designed to work with heartbeat events and generates an assignable MIDI event when a heartbeat event is received. If the heartbeat events are not aligned with the "beat grid" of the track, the input will not be "in time" (in a musical sense). To compensate for this, the Beat Train has a mode for automatically quantizing the incoming events to the nearest

measure in the track. The value given for measure effectively sets the width of the quantization grid: assigning a quantization measure of one measure when the piece is being performed in Common or 4/4 time (4 quarter beats to a bar) means events will be delayed to occur “on the beat”. Progressively finer quantization grids allow more freedom. We find using a grid of 8th or 16th notes pleasing, as at that resolution the natural heart rate variability introduces some randomness and the heartbeat events still feel biological rather than robotic, yet they stay aligned with the track.

The Selector instrument allows us to choose between pre-recorded phrases of music depending on a given input. We have designed this instrument for alternating between looped backing tracks based on the most easily controlled posture parameters. In one instance we have made two variations of a phrase, one with an ascending intonation at the end, the other that descends and resolves. By assigning phrase selection to whether a performer is leaning left or right, a performer can freely improvise on another instrument using other physiological signals, and control whether the phrase repeats, or resolves based on body position at the end of the phrase, in a manner that is readily observable by fellow performers.

The biggest single issue in performing with the system is latency. Due to the Bluetooth protocol, there is an unavoidable minimum latency of about a second. Normally, latencies in musical inputs are measured in tenths of seconds, so at first glance a latency of this magnitude would seem to be a major difficulty. However, we have found that by performing in a roundelay manner, where the same pattern or a similar pattern is repeated at intervals, it is possible (and natural) for musicians to get “in the groove” using the system. In effect the musician plays the same phrase, making the bodily movements necessary for the physiological instrument to produce the desired intent, with the signals being effectively delayed one whole phrase. Posture and activity data are therefore used as “ambient” biosignals, as opposed to being an effort towards gestural control.

Finally, we have found that it is good to structure the performance to introduce the physiological elements one by one, to educate an audience, typically leading from heartbeat-triggered samples to activity to respiration.

5 Conclusion

Our work lies at the intersection of biosignal analysis, music performance, and music theory. We find this a very exciting area to explore. There is potential for signal analysis techniques to aid us in our quest to compose and perform music informed by live physiological data.

In summary, we have created a working system comprising multiple body-sensors, data-analysis components, and software that integrates the signals into Ableton Live for live performance. Our system functions wirelessly, and is built in as decoupled a manner as possible. The data analysis is driven entirely by the motivations of live performances, which has led to novel approaches to the problem of turning raw sensor data into data streams that are robust, consistent

and suitable for musical performance, while balancing a faithfulness to the composition with fidelity to the source signals. Working with the system in practice has led us to consider a role for one of the performers, the analysis supervisor, in shaping and editing the data in the service of the performance.

Future development may involve the addition of other wireless sensors, such as galvanic skin response sensors (for emotional responses) and gyroscopes (for more accurate posture analysis), but availability and costs have been limiting factors. Including some existing technologies, such as Kinect or video based systems, would allow for immediate gestural expressions to complement existing biosignals.

Our initial concept sketches envisaged monitoring audience members, for use as another input or as a feedback mechanism. The focus to date has been on making the system work for the performers, so we have neglected this aspect of the system. It is interesting to consider what parameters could be extracted from an audience using the same recording equipment as the performers.

References

1. Henry, T.K.: Invention locates hurt brain cells. *New York Times* (March 2, 1943) 21
2. Lucier, A.: Statement on: Music for solo performer. In Rosenboom, D., ed.: *Biofeedback and the arts, results of early experiments*. Aesthetic Research Centre of Canada Publications (1976)
3. Teitelbaum, R.: In tune: Some early experiments in biofeedback music. In Rosenboom, D., ed.: *Biofeedback and the arts, results of early experiments*. Aesthetic Research Centre of Canada Publications (1976)
4. Eaton, M.L.: *Bio-music: biological feedback experimental music systems*. Something Else Press (1971)
5. Ojanen, M., Suominen, J., Kallio, T., Lassfolk, K.: Design principles and user interfaces of Erkki Kurenniemi's electronic musical instruments of the 1960's and 1970's. In: *NIME'07: Proceedings of the International Conference on New Interfaces for Musical Expression*. (2007) 88–93
6. Rosenboom, D.: Extended musical interface with the human nervous system. Number 1 in *Leonardo Monograph Series*. International Society for the Arts, Sciences and Technology (1990)
7. Knapp, R.B., Lusted, H.S.: A bioelectric controller for computer music applications. *Computer music journal* **14**(1) (1990) 42–47
8. Knapp, R., Jaimovich, J., Coghlan, N.: Measurement of motion and emotion during musical performance. In: *ACII 2009: 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*. (2009) 1–5
9. Le Groux, S., Manzolli, J., Verschure, P.F.M.J.: Disembodied and collaborative musical interaction in the multimodal brain orchestra. In: *NIME'10: Proceedings of the International Conference on New Interfaces for Musical Expression*. (2010)
10. Miranda, E.R., Brouse, A.: Interfacing the brain directly with musical systems: On developing systems for making music with brain signals. *Leonardo* **38**(4) (2005) 331–336
11. Mann, S., Fung, J., Garten, A.: DECONcert: bathing in the light, sound, and waters of the musical brainbaths. In: *ICMC 2007: International Computer Music Conference*. (2007)

12. Barrass, S.: Sonifications for concert and live performance. *AI & Society* **27**(2) (2012) 281–283
13. Pérez, M., Knapp, R.: BioTools: a biosignal toolbox for composers and performers. In Kronland-Martinet, R., Ystad, S., Jensen, K., eds.: *Computer Music Modeling and Retrieval. Sense of Sounds*. Volume 4969 of *Lecture Notes in Computer Science*. Springer (2008) 441–452
14. Neumark, N., Khut, G.: Cardiomorphologies: An inner journey through art. *IEEE Multimedia* **14**(4) (2007) 5–7
15. Barbieri, R., Matten, E.C., Alabi, A.A., Brown, E.N.: A point-process model of human heartbeat intervals: new definitions of heart rate and heart rate variability. *American Journal of Physiology - Heart and Circulatory Physiology* **288**(1) (2005) H424–H435
16. Freed, A., Schmeder, A.: Features and future of Open Sound Control version 1.1 for NIME. In: *NIME'09: Proceedings of the International Conference on New Interfaces for Musical Expression*. (2009)