Emotion in Motion: A Reimagined Framework for Biomusical/Emotional Interaction

Brennon Bortz
Institute for Creativity, Arts, and Technology
Virginia Polytechnic Institute and State University
190 Alumni Boulevard
Blacksburg, VA 24060
brennon@vt.edu

Javier Jaimovich
Departamento de Música y Sonología
Universidad de Chile
Compañía 1264, Santiago, Chile
javier.jaimovich@uchile.cl

R. Benjamin Knapp
Institute for Creativity, Arts, and Technology
Virginia Polytechnic Institute and State University
190 Alumni Boulevard
Blacksburg, VA 24060
benknapp@vt.edu

ABSTRACT

Over the past four years Emotion in Motion, a long running experiment, has amassed the world's largest database of human physiology associated with emotion in response to the presentation of various selections of musical works. What began as a doctoral research study of participants in Dublin, Ireland, and New York City has grown to include over ten thousand emotional responses to musical experiences from participants across the world, from new installations in Norway, Singapore, the Philippines, and Taiwan.

The most recent iteration of Emotion in Motion is currently underway in Taipei City, Taiwan. Preparation for this installation gave the authors an opportunity to reimagine the architecture of Emotion in Motion, allowing for a wider range of potential applications than were originally possible with the initial development of the tools that drive the experiment. Now more than an experiment, Emotion in Motion is a framework for developing myriad emotional/musical/biomusical interactions with co-located or remote participants. This paper describes the development of this flexible, open-source framework and includes discussion of its various components: hardware agnostic sensor inputs, refined physiological signal processing tools, a public database of data collected during various instantiations of applications built on the framework, and the web application frontend and backend. We also discuss our ongoing work with this tool, and provide the reader with other potential applications that they might realize in using Emotion in Motion.

Author Keywords

Experiment design, interaction design, biosignals, affective computing, music and emotion

ACM Classification

H.5.5 [Information Interfaces and Presentation] Sound and Music Computing J.5 [Computer Applications] Arts and Humanities — Music G.3 [Mathematics of Computing] Probability and Statistics — Experimental design H.5.1 [Information Interfaces and Presentation] Multimedia Information Systems — Audio input/output H.2.8 [Information Systems] Database Management — Database Applications

1. INTRODUCTION

Emotion in Motion is a study designed to improve understanding of people's emotional responses to musical experiences. Emotion in Motion explores these links between emotion and music through the lens of psychophysiology. Electrodermal activity and pulse are recorded from every participant during the presentation of a number of musical excerpts, and each is followed by a number of self-report questionnaires designed to explore constructs such as emotional response to the excerpt, and familiarity with and liking of the excerpt. The study was initially launched in the summer of 2010, and has continued to run in various locations around the world. The original musical excerpts used in the study were selected in order to elicit a range of positive and negative emotions and intensities of emotion, as well as to be representative of a range of musical styles and genres. To date, thousands of subjects have participated in the study. As the study continues to run in additional locations, we expect to have more than twenty thousand additional participants over the next two years.

This promise of such a large number of additional participants, as well as range of locations from which we have been able to collect data (the existing dataset currently holds data from runs in Ireland, the United States, Norway, the Philippines, Singapore, and Taiwan; and, data from runs in Chile and runs in new locations in the United States are forthcoming) provide opportunities not only to explore the role one's culture plays in his or her emotional response to musical experiences, but also to construct more focused studies of specific musical/emotional phenomena that were revealed in initial analysis of the original study data.

However, the original implementations of the tools the study uses did not lend themselves well to rapid design and deployment of smaller 'substudies' aimed at exploring more focused areas of inquiry with smaller groups of participants. This led us to reconsider the overall implementation of the experimental apparatus. While this process did result in such a tool, it is no longer merely a tool for designing studies of music and emotion. In fact, what has emerged is an entire framework for the rapid development of a range of musical interactions capable of capturing data from a range of sensors (physiological or otherwise), recording and acting on performer and audience members' active input through a web interface, presenting media excerpts to audience members, providing real-time feedback on performers' and audience members' emotional responses, controlling and communicating with any Open Sound Control-capable (OSC) tool, and more.

In this paper, we describe this new software framework for musical/physiological/emotional interaction. We begin by discussing how our own research has motivated this framework, and then elaborate on the various components of the framework itself: hardware agnostic sensor inputs, refined physiological signal processing tools, a public database of data collected during various instantiations of applications built on the framework, and the web application frontend and backend. We also discuss our ongoing work with this tool, and provide the reader with examples of other potential applications that they might realize in using Emotion in Motion. This new framework is available to the public and continuously being improved.[1]

2. ON MUSIC, PHYSIOLOGY AND EMOTION EXPERIMENTS

Emotion plays a fundamental role in music. Composers create pieces to convey or induce a particular mood in their audiences, the music scored to accompany a film can help the audience to interpret a character's emotion, or the music selected by a worker can boost attention and productivity. Overall, there is a widespread belief, both in the scientific and artistic communities, that music has the ability to evoke emotional responses in listeners [11]. However, the actual mechanisms of how music modulates human emotion are still not well understood. Disciplines as varied as musicology, psychology, computer science, neuroscience, music therapy, music information retrieval, and human computer interaction have studied these mechanisms, influencing applications and performances relevant to NIME by creating novel forms of interactive art and musical interfaces that incorporate measurement of emotion via physiological indicators for performers and audiences.

Even though the research community has come to no consensus on even the definition of emotion [3], there is agreement that physiological adjustments are one of the factors that constitute an emotion [1], and there is evidence of multiple bodily responses to musical stimuli that correlate with self-reported emotional states [6]. Unfortunately, experiments presented in the literature are often conducted in a laboratory setting, with a relatively small number of participants (usually college students or the researchers themselves), and utilize dissimilar methodologies that can explain the lack of consistency between results of different studies [2], [10].

In response to this, Emotion in Motion was designed with the objective of measuring changes in physiological responses of participants from the general public outside a laboratory setting. The aim was also to collect a large number of cases, partly inspired by Juslin and Sloboda's suggestion of creating a large sample that can serve as a pool of participants; measuring different variables in order to test multiple hypotheses [9].

While the public experiment/installation format may not be a 'typical' music listening environment, it is an informal, open, and of a non-mediated nature, which at the very least provides an interesting counterpoint to lab-based studies. This design then, has the potential of producing a more natural set of responses to music stimuli.

3. MUSIC AND EMOTION APPLICATION FRAMEWORK

In the first eight iterations of Emotion in Motion, the design of the experiment was loosely structured around the following stages:

- Collection of participant demographic and musical background information
- Presentation of musical excerpts, each followed by a series of questions designed to gauge the participant's emotional response to the excerpt, as well as liking, and familiarity, among other questions (These questions varied between iterations, and were excluded altogether in one.)
- Collection of responses to questions about the overall experience ("Which song did you enjoy the most?", "With which song were you most engaged?", etc.)

Originally, the experiment was designed as a Max/MSP 5 patch that coordinated the presentation of musical excerpts, questionnaire presentation and response processing, signal acquisition, and data persistence. The system was packaged with a companion Max/MSP patch that allowed for the reordering of segments in the experiment design, as well as insertion of new segments. However, extending the system was difficult and error-prone.

Responsibilities in the new framework are more clearly defined and broken across individual discrete components. There are three major components to the framework: signal capture and processing, the database, and the web application.

3.1 Signal Capture and Processing

3.1.1 Hardware-Agnostic Sensor Input

When working with physiological indicators of emotion, the literature shows a range of signals that have been used both for experimental and performance applications. These can be as varied as electrocardiography, respiratory features, electroencephalography, electrodermal activity (EDA), muscle tension, functional magnetic resonance imaging and pupillary response. Moreover, each feature can be obtained using different technologies and methodologies, which do not necessarily share similar protocols and data structures.

In order to address this uncertainty, Emotion in Motion can receive sensor data structured as an OSC message, where each channel of data is time-stamped with the media playback time. This allows for sample accurate synchronization between the media stream and the sensor data, which is essential in the analysis of physiological signals that cannot be synchronized manually afterwards [8]. Additionally, the user can label the different channels of data used in the experiment, and these labels are embedded into the sensor data file. This design allows for heterogeneous signal sources to be incorporated into the framework and recorded synchronously alongside the experiment stimulus.

3.1.2 Signal Processing Tools

One of the key challenges of working with physiological signals for real-time applications is being able to continuously assess the status of the signal (e.g. check for artifacts, disconnections, etc.) and extract significant features that can inform and drive the desired process/application.

Drawing from data available in the extensive Emotion in Motion database, two robust real-time signal processing tools for heart rate (HR) and EDA have been developed. These tools were calibrated using the physiological signals of over 4,000 participants, learning from aspects such as user interaction with the sensors, variability of signals, and expected ranges for different demographics.

Each tool consists of a pre-processing stage, where the signal is formatted and analyzed for artifacts that can be

[1] `https://github.com/brennon/eim`

caused by motion or connection issues. This analysis results in the production of a quality index, which represents the percentage of the signal that corresponds to artifacts. The signals are then processed to extract significant features, such as phasic and tonic components from EDA, and HR from electrocardiography or pulse oximetry signals.

These tools are presented in detail in [7] and are available online[2].

3.2 Database

A long-standing goal of this project is to develop a massive database of listener/audience physiological responses to music. Such a database could be used in the development of research contributions to fields as diverse as interactive performances and installations, affective computing, music information retrieval, and psychology, to name only a few. Ideally, this database should be able to capture and represent information coming from a range of sonic listening and interaction experiences, including further designed experiments, live performances, and individual listening experiences in one's home. The representations of data collected from situations like these may differ widely. For instance, in a designed experiment setting, researchers may gather responses to a number of questions, real-time physiology data, and references to media that was played. In a live performance scenario, on the other hand, it may be useful to persist live recordings and timing information of the performance and physiology from the audience members or performers. Representing these disparate types of data in a centralized location with traditional relational databases, for instance, presents a challenge.

3.2.1 *Traditional Approach*

Much like their counterparts in a spreadsheet application, in a traditional relational database there are three levels of granularity with which both the designers and users of the database are concerned: tables and the rows and columns that combine to form these tables. Beginning at the least granular level, the tables in a well-designed database contain information about a collection of a specific type of data entity. In a traditional application, one table may hold information about a store's customers while another holds information about that store's inventory, and yet another the store's orders. In an interactive installation, perhaps, one table may hold information about lights used in the installation, another about sounds, and another about physical locations in the installation. At the next level of granularity, a row in a table represents a specific instance of one of the entities that the table represents. Continuing the traditional example given above, a row in the customer table may hold information for one specific customer. More appropriate to our purposes, a row may hold the demographic details for the subject of an experiment, the settings for a light in an interactive installation, or the patch settings for a single instrument in a performance. The columns of a table, on the other hand, stand at the finest level of granularity in a relational database. In a customer table, one column may represent the street address of the customers while another contains postal codes. On the other hand, the columns in a table of lights may contain red, green, or blue values for lighting settings, or the responses of experiment subjects to different questions.

This separation of concerns can provide the users of traditional databases with a high level of flexibility in designing a database schema that meets their particular needs. For the purposes of Emotion in Motion and other similar applications, there are several issues with this approach.

Differing Structures Between Rows.

One obvious mapping of raw qualitative experimental data to a database schema in designs such as those used in Emotion in Motion is to have a table of trials. Each trial in the trials table may contain, among other things, the responses to answers asked after each instance of musical playback. Across different versions of Emotion in Motion, however, the number, type, and content of these questions have all varied. A table in a traditional relational database requires that all rows in a table contain values for all columns in that table. It may be the case that null values are allowed for certain columns, but in a well-designed, normalized database, such instances will be rare, if they are indeed present [4]. In a database the purpose of which is to amass data from experiments of differing designs and various interactive scenarios, it is difficult, if not impossible, to design table structures that allow for both unified and normalized data representation.

As another example, it is meaningful in this database to store physiological signals together from test subjects, audience members, and performers alike. The conceptual grouping of all such signals together might, for example, give one the ability to state that all signals in this collection of signals represent a a group of people's physiological response to some musical experience. Grouping them in separate tables by experiment, performance, or some other criterion necessitates the collection and analysis of data from multiple sources (tables) when exploring broader questions of how people respond to different musics, in general. Furthermore, it is often the case that during analysis of physiological data, meaningful features are extracted from physiological signals, and it is helpful to store these features alongside the raw signals. The rigid table structure of traditional relational databases provides no clean way to do so.

Differing Structures Within Rows.

Continuing to consider the representation of physiological data in a database, the means by which one single signal is represented in a traditional database present a separate challenge. Given the description of tables, rows, and columns, how does one represent a discrete time-sampled physiological signal (or a sensor, audio, or any other discrete time-sampled signal) in a traditional database? Three approaches to this problem are routinely employed. First, each column may represent a particular sample in the signal. This calls for as many columns as are necessary to represent the entirety of the of signal that the table represents. While a large number of columns in a table is not a problem in and of itself, how does one store signals of differing lengths? Here, the problem of null values for columns at the end of all signals but the longest presents itself again. The second common solution is to store signals (which are, after all, nothing more than lists of numbers), as textual data (for instance, `1, 1, 2, 3, 5, ...`) in a single column. This approach provides for the possibility of representing, say, various signal types for a single subject in one row. Here, efficiency quickly becomes the problem. When attempting to analyze, in real-time or otherwise, signals from any respectably sized group of signals becomes a process of agglomerating all such textual representations of signals together, converting them back to numeric representations (with the added risk of loss of precision), and only then performing the analysis—a very inefficient process, indeed. Efficiency problems, especially in large datasets, plague the third approach, as well. Here, the data for each signal are

[2] `http://www.musicsensorsemotion.com`

stored in a file on disk, and entries in columns in a signal table simply point to those files. In a large dataset, the inefficiency here is clear, as well. Furthermore, new approaches for in-database analysis of large datasets are continually becoming available, but most often these techniques require the data to live 'in' the database, not in separate files referenced from within the database [5]. Taking any of the these approaches may be acceptable in datasets of a moderate size, but in massive (and ever-growing) datasets such as those in the Emotion in Motion database, these approaches are far from feasible.

3.2.2 Our Approach

Partially in an effort to address these problems, document-oriented NoSQL databases have been developed (for example, Cassandra, CouchDB, MongoDB, and Redis). Instead of storing individual data entities in the rows of tables, a document-oriented database stores them as documents. Documents are often grouped into logical collections, and documents and arrays can be embedded inside of other documents. There are no predefined schema for documents in a collection (though enforcement of structural "rules" can be performed at the application level, if required). This architecture allows for the straightforward creation of hierarchical relationships within and between documents, as well as the addition and removal of fields to documents with no implications for other documents within the same collection. Flexible schema and the ability to embed documents and arrays within documents exactly solve the problems discussed in Section 3.2.1.

Our approach uses the popular MongoDB database. In MongoDB, documents are stored in Binary JSON[3] (BSON) format, a binary-encoded serialization of JSON-like documents. Like JSON[4], BSON supports the embedding of documents and arrays within other documents and arrays. BSON also contains extensions that allow representation of datat types that are not part of the JSON spec. For example, BSON has a Date type and a BinData type.

In the Emotion in Motion database, trial documents represent audience responses and signal documents represent audience physiological signals, in general. For the data that have come from the Emotion in Motion experiments, a typical trial document looks something like this:

```
{
    "_id" : ObjectId("5410eee108ad6ee3090e2509"),
    "metadata" : {
        "terminal" : 0,
        "location" : "bergen",
        "session_number" : 1489
    },
    "answers" : {
        "music_styles" : [ "dance", "hip_hop"],
        "most_engaged" : 3,
        "most_enjoyed" : 1,
        "emotion_indices" : null,
        "sex" : "female",
        "dob" : ISODate("2002-01-01..."),
        "musical_background" : true,
        "ratings" : {
            "activity" : [ 3, 1, 5 ],
            "engagement" : [ 3, 1, 5 ],
            "familiarity" : [ 1, 4, 5 ],
            "chills" : [ 2, 2, 5 ],
            "positivity" : [ 5, 2, 5],
            "power" : [ 5, 2, 5 ],
            "like_dislike" : [ 5, 3, 5 ]
        },
        "nationality" : "norwegian",
        "visual_impairments" : false,
        "musical_expertise" : 4,
        "hearing_impairments" : false,
        "age" : 11
    },
    "timestamps" : {
        "start" : ISODate("2012-02-23..."),
        "test" : ISODate("2012-02-23..."),
        "end" : ISODate("2012-02-23..."),
        "media" : [ ISODate("2012-02-23...", ... ]
    },
    "media" : [ ObjectId("537ea1..."), ... ],
    "date" : ISODate("2012-02-23..."),
    "valid" : true,
    "random" : 0.7618690760428368,
    "experiment" : ObjectId("5432bc..."),
    "signals" : [ ObjectId("5410ee..."), ... ]
}
```

[3]http://bsonspec.org

[4]http://www.json.org

In this abbreviated document, ellipses indicate an abbreviated array of values or a truncated single value.

Using a document-oriented database does not preclude one from establishing relationships between documents in different collections. For example, the first signal in this trial corresponds to this typical signal document in the signals collection:

```
{
    "_id" : ObjectId("5410ee..."),
    "label" : "T015",
    "signals" : {
        "eda_status" : [ 1, 1, 1, ... ],
        "timestamps" : [ 0, 10, 20, ... ],
        "hr" : [ 103.448, 103.448, 103.448, ... ],
        "eda_filtered" : [ 428.7, 428.7, 428.8, ... ],
        "hr_status" : [ 0, 0, 1, ... ],
        "eda_raw" : [ 427, 428, 429, ... ],
        "pox_raw" : [ 0, 0, 1, ... ]
    },
    "metadata" : {
        "location" : "bergen",
        "terminal" : 0,
        "session_number" : 1489
    }
}
```

In this signals document, the raw signals (`eda_raw` and `pox_ raw`) are present, as well as derived signals that have been extracted and stored alongside the signals from which they were derived (`hr`, `eda_filtered`, etc.)

Everything in an interaction driven by the framework is stored as an object in the database. These may be signal and trial documents, stored media, or documents that describe the structure of a trial itself. These structure documents are used by the web component of the framework in order to dynamically construct interactions on-demand as specified by the structure documents (see Section 3.3.1). On the other hand, these structure documents may represent something entirely different: the structure of a piece or the performance parameters for an instrument, for example.

3.3 Web Application

The web component of Emotion in Motion dynamically structures interactions, coordinates with Max in the presentation of media and acquisition of sensor signals, persists real-time feedback from the audience, and provides a graphical display those interacting with the application. This component is written entirely in Javascript, using AngularJS for the frontend, and Node.js for the backend.

3.3.1 AngularJS Frontend

The frontend of Emotion in Motion is a single-page web application written in AngularJS.[5] AngularJS is a Javascript model/view/controller (MVC) framework that supports web development by allowing one to extend the syntax of HTML; providing real-time, two-way data binding between the models and views; and, handling complex dependency injection scenarios. In practice, this allows for any portion of an interaction built on Emotion in Motion to be written as an AngularJS module and dropped into the framework. Components for such tasks as media playback, physiological data collection, questionnaires, and real-time continuous response inputs are included in the framework.

When a component is selected (or a new one developed) for use in a new interaction, once a reference to it is added to the framework's 'routing' file, it can be simply be referenced and used in new interaction designs. These designs are specified in documents that are stored in the database. For the Emotion in Motion experiments, an excerpt of the design document looks something like the following:

```
{
  "structure" : [
    {
      "name" : "sound-test"
    },
    {
      "name" : "eda-instructions"
    },
    ... ,
    {
      "name" : "media-playback",
      "mediaType" : "random"
    },
    {
      "name" : "questionnaire",
      "data" : {
        "title" : "Media Questions",
        "introductoryText" : "This questionnaire ...",
        "structure" : [
          {
            "questionType" : "likert",
            "questionId" : "engaged",
            "questionLikertMinimumDescription" : "...",
            "questionLikertMaximumDescription" : "...",
            "questionLikertLeftImageSrc" : "a.png",
            "questionLikertRightImageSrc" : "b.png",
            "questionLabel" : "How involved...?",
            "questionStoragePath" : "answers.eng",
            "questionIsAssociatedToMedia" : true,
            "questionOptions" : {
              "choices" : [
                {
                  "label" : "1"
                },
                {
                  "label" : "2"
                },
                ...
              ]
            }
          },
          ...
        ]
      }
    },
    ...
  ]
}
```

[5] https://angularjs.org/

This excerpt specifies several views that the user will see during the interaction. Of particular interest are the media-playback and questionnaire views. The `media-playback` component is designed such that pre-specified or randomly selected media can be played back for the participant. The `questionnaire` component allows for the flexible creation of questionnaires of a number of different formats. Here (though only one is shown in the excerpt), a range of rating scales are presented for the participant to provide their self-report ratings for the media that was just played.

This structural description method allows the designer to implement a very broad range of designs with simple design documents. The questionnaire can be used for something like the following:

```
{
    "name" : "questionnaire",
    "data" : {
        "title" : "Filter Parameters",
        "structure" : [
            {
                "questionType" : "slider-scale",
                "questionId" : "center-frequency",
                "questionMinimumValue" : 0,
                "questionMaximumValue" : 20000,
            },
            {
                "questionType" : "slider-scale",
                "questionId" : "gain",
                "questionMinimumValue" : -96,
                "questionMaximumValue" : 6,
            },
            ...
        ]
    }
}
```

In this way, the framework can be used for much more than static experimentation. Here, for instance, it can become a means of connecting multiple interfaces to a live performance patch. Different audience members or performers might control different patch settings through dynamically generated web interfaces available on their phones or tablets, while at the same time their physiology is recorded across an array of sensors—the possibilities that this infrastructure enables really are quite far reaching.

3.3.2 Node.js Backend

Node.js is a platform built on Chrome's JavaScript runtime for easily building fast, scalable network applications.[6] Node.js uses an event-driven, non-blocking I/O model that makes it lightweight and efficient, perfect for data-intensive real-time applications that run across distributed devices.

The primary roles of the Emotion in Motion Node.js backend are the coordination of OSC communication with Max for sensor control, recording of physiology, and media presentation; and communication with the MongoDB database. With the help of Socket.IO, any custom OSC message can be passed through transparently from the AngularJS frontend, through the Node.js backend, and on to a running in-

[6] http://nodejs.org/

stance of Max.[7] In the same way, any custom OSC message can be sent from Max and arrive at the AngularJS frontend as a neatly-packaged JSON object. Similarly, data of nearly any structure can be passed to and retrieved from the MongoDB database. With this structure, for example, designing a Max patch that respond to inputs from a dynamically-generated web interface becomes a trivial process.

4. POTENTIAL APPLICATIONS AND FUTURE WORK

The framework presented here enables a wide variety of potential applications. The following are three NIME-related applications that could be realized with this framework:

- *Audience feedback and interaction* The difficult work of developing a means for audience interaction with a live performance has already been done in the development of Emotion in Motion. Designing numerous views and widgets for the audience to manipulate and mapping them to OSC messages that one might handle in Max, Pd, or SuperCollider, is a trivial process.
- *Distributed emotional/musical interaction* This is a simple extension of the above application. What was once a complex development process reduces to the description of interfaces in the JSON structure described in 3.3.1 and modification of a 'patch' to respond to the appropriate OSC messages.
- *Further experimentation* While the above two application areas already enable a broad range of new NIME applications, we intend to continue to use Emotion in Motion for a number of additional studies around music and emotion, to contribute our results to openly available datasets, and to encourage our peers to do so, as well. This work will motivate the development of new enhancements to and components for the framework. While we have already built an enormous (and enormously valuable) dataset for our and others' research into music and psychophysiology, we intend to use this framework to expand our work and contributions to the wider academic community.

5. CONCLUSION

This paper describes the ongoing evolution of Emotion in Motion: once a tool for experimentation in music and emotion that has now grown to become a flexible framework both for continued experimentation in the same, as well as the development of a wide range of musical, emotional, and physiological interactions. This framework has been split into a number of logical components. Signal acquisition and processing continue to be handled by Max, though new and refined tools are available and used for these purposes. The database of collected data has been rebuilt from top to bottom and is now available for research purposes to the academic community at large. The web frontend and backend have been specifically engineered to allow for easy design and deployment of new and varied interactions built on the framework. We sincerely hope that this framework enables and spurs on such interactions built by our peers. It is freely available and under continuous improvement.

6. REFERENCES

[1] J. T. Cacioppo, D. J. Klein, G. G. Berntson, and E. Hatfield. The Psychophysiology of Emotion. In M. Lewis and J. M. Haviland, editors, *The Handbook of Emotion*, pages 119–142. Guilford Press, New York, New York, USA, 1993.

[2] J. T. Cacioppo, L. G. Tassinary, and G. G. Bernston, editors. *Handbook of Psychophysiology*. Cambridge University Press, Cambridge, England, UK, 3rd edition, 2007.

[3] N. J. Christenfeld and G. Mandler. Emotion. In I. B. Weiner and D. K. Freedheim, editors, *Handbook of Psychology*, volume 1. John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 2nd edition, 2012.

[4] E. F. Codd. A Relational Model of Data for Large Shared Data Banks. In M. Stonebraker and J. M. Hellerstein, editors, *Readings in Database Systems*, pages 5–15. Morgan Kaufmann Publishers Inc., San Francisco, California, USA, 3rd edition, 1998.

[5] J. Dean and S. Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. *Communications of the ACM*, 51(1):107–113, Jan. 2008.

[6] D. A. Hodges. Psychophysiological Measures. In P. N. Juslin and J. A. Sloboda, editors, *Handbook of Music and Emotion: Theory, Research, Applications*. Oxford University Press, New York, 2010.

[7] J. Jaimovich. Creating Biosignal Algorithms for Musical Applications from an Extensive Physiological Database. In *Proceedings of the 2015 Conference on New Interfaces for Musical Expression (NIME 2015)*, Baton Rouge, LA, 2010.

[8] J. Jaimovich and R. B. Knapp. Synchronization of Multimodal Recordings for Musical Performance Research. In *Proceedings of the 2010 Conference on New Interfaces for Musical Expression (NIME 2010)*, pages 372–374, Sydney, Australia, 2010.

[9] P. N. Juslin. Music and Emotion: Seven Questions, Seven Answers. In I. Deliege and J. Davidson, editors, *Music and the Mind: Essays in Honour of John Sloboda*. Oxford University Press, Oxford, England, UK, 2011.

[10] S. D. Kreibig. Autonomic Nervous System Activity in Emotion: A Review. *Biological Psychology*, 84(3):394–421, July 2010.

[11] L. B. Meyer. *Emotion and Meaning in Music*. University of Chicago Press, London, England, UK, 2008.

[7] http://socket.io/