



HAL
open science

Toward Augmented Familiarity of the Audience with Digital Musical Instruments

Olivier Capra, Florent Berthaut, Laurent Grisoni

► **To cite this version:**

Olivier Capra, Florent Berthaut, Laurent Grisoni. Toward Augmented Familiarity of the Audience with Digital Musical Instruments. International Symposium on Computer Music Multidisciplinary Research, Sep 2017, Matosinhos, Portugal. hal-01577953

HAL Id: hal-01577953

<https://hal.science/hal-01577953>

Submitted on 28 Aug 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Toward Augmented Familiarity of the Audience with Digital Musical Instruments

Olivier Capra¹, Florent Berthaut¹, and Laurent Grisoni¹

Univ. Lille, INRIA, CRIStAL, France

olivier.capra@etu.univ-lille1.fr

florent.berthaut@univ-lille1.fr laurent.grisoni@univ-lille1.fr

Abstract. The diversity and complexity of Digital Musical Instruments often lead to a reduced appreciation of live performances by the audience. This can be linked to the lack of familiarity they have with the instruments. We propose to increase this familiarity thanks to a trans-disciplinary approach in which signals from both the musician and the audience are extracted, familiarity analyzed, and augmentations dynamically added to the instruments. We introduce a new decomposition of familiarity and the concept of *correspondences* between musical gestures and results. This paper is both a review of research that paves the way for the realization of a pipeline for augmented familiarity, and a call for future research on the identified challenges that remain before it can be implemented.

Keywords: digital musical instrument, performance, familiarity, causality, augmentations, correspondences

1 Introduction

Whether as a musician or as a spectator, the experience of live music has very particular characteristics. It is often immersive, both intimate and collective, implies different modalities, both low and high level cognitive engagements. Moreover, musical experience is often a corporeal, aesthetic as well as an emotional commitment and is therefore difficult to define. Despite their great diversity and the parallel treatments they require, we assimilate all these aspects in a natural way, without any apparent effort. We naturally build this intense feeling that we all have experienced by attending a concert. Beyond the unfathomable subjective part of this musical experience (i.e. tastes and colors), we can identify objective characteristics that influence how we perceive live music, and among these, the way musicians interact with their instrument and how we perceive and integrate these interactions.

On the one hand, in the production of music with acoustic instruments, gestures and sounds are intrinsically linked by the laws of physics or at least by intuitive connections that do not need any prior explanation. Beside the music they produce, every interaction is visible. Thereby the expressive intentions, as well as the intensity of emotions, are particularly vivid and underline the role of

a multimodal integration [25]. Along this line, the degree of perceived control, influenced by the additional information usually extracted from the musician’s interactions, participates to the emergence of the liveness of the performance. On the other hand, the experience of live music produced by digital musical instruments (DMIs) may suffer from a deterioration of this intuitive link between the behavior of the musician, the inner mechanisms of the instrument and the sounds actually produced. When what we hear is not directly linked to what we see or what we may infer from a given gestural behavior, when music is not the direct and causal consequence of the specific gestures of the musician, then we may lose the multimodal integration that contributes to the immersive experience of live music.

These considerations have led to recent research on measuring the audience experience [2, 14], often used as a way to evaluate the instruments themselves [27]. The less familiar we are with an instrument, the less we are able to perceive the fine relationships musicians build with their instrument. This lack of coherence in multimodal information might downgrade the attributed agency, that is the perception of how much the musician is controlling their performance. Thus the knowledge we have about a DMI or its obvious behavior are crucial features to understand and integrate the interactions and thus fully experience live music. How do we ensure familiarity with DMIs ?

1.1 Improving familiarity with Digital Musical Instruments

Several attempts have been made at solving this issue of familiarity. In this section, we propose a quick review of these solutions and their limitations.

As mentioned before, familiarity is the knowledge we have about the behavior and possibilities of an instrument. Quite naturally, a first way to increase the familiarity of people with a particular instrument is to promote its use and dissemination by **building a repertoire** of compositions around it. Once a majority of spectators have seen the instrument played by many musicians in different contexts, or even practiced it themselves, they are aware of its potential for musical expression and of its behavior. Thereby, the familiarity issue does not exist anymore. This kind of “natural familiarity” is definitely effective but requires a large amount of time and energy to be achieved. Moreover, this method is not compatible with the very idea behind DMIs and the exponential creativity they embed. Whether a musician wants to evolve its instrument or let other musicians modify it, as soon as the instrument changes, the whole process of natural familiarity needs to start all over again.

When the problem of familiarity lies partially in the understanding of the operations of an instrument, demonstrating the behavior of the instrument can be a valuable solution. Building the familiarity with **a pedagogical method** is a simple way to make an audience understand what is going on on stage. Before, after and even in breaks during the performance, the musicians can explain how their instrument works. Prior hands-on demos, where the audience can actually play the instrument, are also a good way to increase the familiarity. However, both these methods trigger some reservations as the technical

understanding of an instrument is not necessarily linked to a better appreciation of the performance. For example, Bin et al. [7] demonstrate that explanations before the performance do improve the understanding of the instruments and its mechanisms but do not increase either the appreciation or interest. Besides these results, the audience may forget important details during the performance and the demos may not be possible when dealing with a large audience.

Another strategy to increase the experience of spectators watching a musician playing an instrument they have never seen before lies in the design of the instrument itself. The idea is to expose a clear link between the gestures and the sound modifications operated by these gestures, what is called the *mapping* of an instrument. The fluency of perception of these mappings is often called “transparency”, and is defined by Fels et al. [9] as follows : *For the lay audience, this understanding is derived from cultural knowledge, including percepts of physical causality relationships.* **Taking care of the transparency** from the very first steps of the design of a musical device can lead to “easier to perceive” instruments. Following that lead, an interesting way to increase the transparency is the use of metaphors [9]. In that case, the common background of inexperienced audience, their general knowledge, are used as a mold to grow new knowledge about an instrument mapping. For instance, the timbre of a sound can be modified as the shape of an associated graphical representation that is getting sculpted. The instruments designed according to these specifications tend to maintain the familiarity we have with acoustic instruments and physical laws. Thus, despite the gain in transparency, this method may narrow the design possibilities usually available for DMIs.

In more recent research, Berthaut et al. [4, 5] proposed to increase the familiarity by recreating a link of causality through **visual augmentations of the instruments**. Based on the attribution of causality by spectators to a music performance, the visual augmentations provide insights on the musician’s interactions with its instrument. Animated 3D objects overlapping the device reveal the relationship between the gestures and the musical result to the audience. These augmentations are based on Wegner’s criteria of apparent mental causation [26] (See 2.1). By exposing the details of the interactions, the setup contributes to re-link the gestures of the artist to their intentions and expressiveness. Besides, the augmentations enable the distinction between automated part and actual live music production. Such a discrimination between live and pre-recorded music appears to be crucial in the experience of live music precisely because it is supposed to be a live performance. However, these augmentations do not take into account the audience reaction or their expertise. Indeed, the familiarity with an instrument is a personal characteristic that cannot be generalized. A concert in front of specialists or naive people does not imply the same requirement of explanations or augmentations.

These different methods offer a rather effective way to increase the understanding of the behavior of DMIs. However they don’t fully cover a major aspect of the music experience which is the real time multimodal integration of visual and auditory cues. Furthermore, while the visual augmentations gave good re-

sults but they do not take into account the audience expertise that may appeal for specific levels of details in the augmentations. Moreover, in some cases, non-visual augmentations may be more relevant. Thereby, we propose new insights to improve the familiarity in a more adaptable and reliable way.

Our contribution is two-fold. First, we propose a conceptual pipeline to improve the familiarity of the audience with DMIs in real-time by analyzing spectators' reactions, musicians' actions, musical outputs and by augmenting the instrument. To do so, we introduce the concept of *correspondences*. Second, for each module of this pipeline, we briefly review relevant knowledge in computer and cognitive sciences. We also identify interdisciplinary challenges that need to be solved to reach a functional software implementation of the pipeline.

2 Toward augmented familiarity

In this section we propose a novel approach for improving familiarity, formalized by a software pipeline. While it remains conceptual, it was designed so that it could be implemented by addressing the challenges described in Section 2.2.

Here is a scenario that we envision with our pipeline :

Patricia attends an electronic music concert. At the entrance, she is given a small device equipped with physiological sensors (a choice of either a bracelet or a special glass that she holds). During the concert, she has trouble understanding what is happening, in particular what the musician's action on the sound is. The device senses a change in a set of physiological signals, that corresponds to a loss in familiarity, and sends the data to a server. Patricia may also directly indicate her loss in familiarity with a graphical slider on an app on her smartphone. Simultaneously, this server has been analyzing the musician's gestures, the flow of data inside the instrument and the musical output. When it receives Patricia's familiarity signals, the server, with settings defined by the musician, selects the adequate augmentations to be displayed. They aim at compensating for the familiarity disruption caused by the musical interaction context. Consequently, visual augmentations are displayed around the musician either for Patricia alone or for the group of people around her using a mixed-reality display. They provide information that improve her degree of familiarity, allowing her for example to perceive the link between the musician's gesture and the resulting sound, and to enjoy the performance to a larger extent.

2.1 General approach

To improve the liveness of music performances with DMIs, we suggest to increase the familiarity thanks to a trans-disciplinary approach in which both human and technological signals are analyzed. To achieve this goal, we first clarify the notion of the familiarity. We then propose the concept of *correspondences*, which describes musical interactions and the way they are perceived. Finally we present a potential pipeline that would extract and process signals from the musician and from the audience to compose augmentations.

Familiarity dimensions As a high-level cognitive ability, familiarity is not a unitary notion. It can be decomposed into components that differ by their nature, their inner rules as well as the type of information they target. Besides, familiarity relates to expertise or attention and thus impacts on brain and motor system activities. We will describe in Section 2.2 how previous research show that these activities can be detected in more detail. Familiarity also relates to parts of the musical interaction that is perceived. Spectators can be more or less familiar with : musical gestures, relations between gestures and sound, instruments capabilities or musical genres. All these factors contribute to the general familiarity one has with an instrument. Taking this into account, our approach is to decompose familiarity into dimensions that we can both evaluate and improve. Based on results in HCI, NIME, and cognitive sciences that have explored various aspects of familiarity, we propose to decompose familiarity in five components : *Causality*, *Instrumentality*, *Instrument expertise*, *Musical Culture* and *Musical Genre Expertise*. These in turn contain dimensions that can be evaluated independently. *Causality* relates to the apparent mental causation, e.g. the judgment of causality of one’s action, defined by Wegner et al [26], where each of the three following dimensions are required to establish a judgment of causality of one’s action: *Priority* - the thought should precede the action at a proper interval; *Consistency* - the thought should be compatible with the action; *Exclusivity* - The thought should be the only apparent cause of action. As proposed in [4], we transpose this model from the perception of one’s own actions to the perception of others’ actions. *Instrumentality* relates to common knowledge that allows one to predict the range of sound possibilities of a musical instrument from its appearance. It is composed of two dimensions : the *composition*, i.e. shape and material, and *behavior*, i.e. mechanisms and degree of autonomy from the musician’s actions. *Instrument expertise* relates to the exposure the spectator has had to the instrument, from a first-time observer to an expert player. *Musical culture* corresponds to a basic knowledge of musical theory, that can for example be used to represent pitch as vertical position of a graphical element, i.e. mimicking a staff. Finally, *Musical Genre Expertise* is composed of dimensions that correspond to the familiarity with the specifics of a genre, such as of *structure* and *constraints*. The effect of this category is obvious in the study by Bin et al. [7], where the same instrument played in two genres has a different impact on familiarity. Notice that each component has a specific weight on the overall process. Furthermore, this model is consistent with the partial results that previous strategies have achieved. Culture and expertise for the “building a repertoire” and the “prior explanations” strategies, instrumentality and culture for “the design for transparency” and mental apparent causation for the augmentation strategy. This decomposition might evolve according to findings from the implementation phase. Some dimensions could prove harder to evaluate or others may emerge.

Correspondences In order to handle the heterogeneous data (physiological, behavioral, musical, visual, mechanical) associated with the dimensions of fa-

miliarity, we introduce the notion of *correspondence*: a conceptual object that stands as a digital multidimensional representation of a musical interaction. This object associates several *properties* : musical output (audio), gesture data (time series of 3D coordinates), video of the movement, control data (time series of sensor values) and the *source* of the interaction. This last property specifies the author of the interaction. Indeed, interactions are not only produced by one musician but can be triggered and mastered by another one in a collaborative performance or, more usually, by autonomous prepared processes, e.g. automations or playlists.

Each property is composed of three elements : a raw extracted element, a semantic element, and a classification element. The raw format element is a pointer to a collection of relative files. For instance, the raw element of the visual property of a clap correspondence is a pointer to a collection of short videos showing a clap from different angles and velocities. The semantic element is composed of annotated descriptors, including the taxonomies developed in the different analyses of gesture and sound. The classification element is dedicated to the classification of the raw part through machine learning methods. As an interaction may be more complex and decomposed in a sequence of more basic interactions, a correspondence can also be linked to other, simpler correspondences. In addition, correspondences hold a *score* for each of the familiarity dimensions described earlier (causality, instrumentality, ...). For example, two correspondences with the same musical result property, e.g. the fade-in of an audio loop, can have different gesture and control properties depending on the mapping chosen for the DMI. While a continuous gesture on a fader would have a high score (meaning, the most natural way to fade) for the *consistency* dimension, a discrete tap on a pad would have a low one, since the effect would not be consistent with the cause for a spectator, e.g. discrete input and continuous output.

Pipeline Our envisioned pipeline (Fig 1) is composed of five modules that handle the extraction of the data from the musician and their instrument (*EXT_M*), the extraction of physiological and subjective data from the audience (*EXT_A*), the processing of the data (*IA*) and the selection of fitting augmentations (*AUG*) based on a database of correspondences (*DB_C*). The pipeline is used at three different moments : before, during and after a performance.

Before the performance, correspondences, coming from a shared online database, or recorded specifically for the instrument, are saved in the *DB_C* module. ① During the performance, *IA* receives musical interaction data extracted by *EXT_M*. This data consists of both dynamic values such as gestural parameters, audio features and control values which will be used directly by *IA*, and of physical (position of gestures and sensors), logical (tracks, effects, synthesizers, ...) and structural information about the musical interaction, which will be used by *AUG*. ② *IA* builds and manages a set of ongoing correspondences from these signals. ③ Simultaneously, *IA* receives the familiarity evaluation (with the identification of the sensed individual or group) from *EXT_A*. If the familiarity evaluation is low, ④ *IA* finds in *DB_C* correspondences similar to the




Fig. 1. The proposed pipeline with its main modules and data flows

ongoing one, and select the familiarity dimensions that need to be compensated depending on their scores in these. *IA* then sends *AUG* the data required for the augmentation : live correspondences, associated signals from *EXT_M* with dynamic and structural data and identification of the source of the familiarity evaluation (in order to display the augmentations only to the correct person or group). ⑤ *AUG* creates (or selects if it already exists) the augmentation that matches the received structural data, for example a visual augmented-reality link between the physical position of a sensor and a virtual representation of an audio track. ⑥ Augmentations are then connected to *EXT_M* and listen to the signals required to update the augmentation, e.g. control values for the sensor, loudness of the track. After the performance, the familiarity extracted from the audience can be reused to refine the scores for the familiarity dimensions in each detected correspondence of *DB_C*.

2.2 Modules

Database of correspondences (*DB_C*) *DB_C* manages all pre-existing correspondences. These can be generic, or specific to an instrument or performance. It receives queries from *IA* to select correspondences matching the ones detected during the performance. A very promising approach in dealing with heterogeneous data is the use of databases. Even if the efficiency of machine learning and analysis tools is still evolving, a numerous amount of initiatives, especially in

analysis of emotion (DEAP[12], RECOLA[19]) but also in music-related actions [10], contribute to shared databases that compile multimodal and synchronized experimental data. The main goal of these databases is to predict complex and abstract states, e.g. the emotional state of an individual, thanks to the analysis of their physiological and behavioral signals such as face expression, electrodermal activity or heart rate variation. Most of the existing databases are composed of 15 to 30 entries referencing data of diversified nature. To allow the gathering of a more relevant amount of cross data, we first need to facilitate the indexing thanks to the more and more intuitive and effective front-end technologies of the web. We propose to develop a web interface that could provide the specific tools to aggregate the data required to constitute the properties of a correspondence. A typical correspondence would require a short video footage of a gesture, the motion capture of this gesture, sensor value and audio output of the instrument. Tags could also be manually added for each of these properties. The interface could be accessible from an open web platform where artists as well as researchers could populate the database, to constitute their own correspondences and therefore optimize their pipeline with more personal choices of gestures, mappings or sound processes. The same online platform can then be used for crowdsourced online evaluation sessions in which correspondences are exposed to participants with different levels of expertise. Their task is to indicate their understanding, using a questionnaire along the familiarity dimensions, and tag the correspondences.

Extraction of the musician’s interactions (*EXT_M*) As explained in Section 2.1, *EXT_M* extracts data from the instrument and musician’s gesture which is then sent to *IA* in order to detect ongoing correspondences. At the instrument level, *EXT_M* extracts sensors states, mapping values as well as musical result of the musician interactions. At the musician level, *EXT_M* extracts the control gestures, the body movements and physiological signals. In addition to these signals, structural information needs to be provided for further use in the augmentations, such as the position of the physical sensors of the DMI and the position of the musician’s hands, the list of tracks, effects or other sound processes with their names, or the mappings between sensors and sound parameters. While some of these obviously need to be defined by the musician manually, or sensed by devices external to the instrument, others can be extracted through a trans-disciplinary approach. Regarding the instrument input, research on gesture recognition, especially concerning hand gestures as demonstrated by Rautaray & Agrawal [18], can be used to identify the performed musical gestures. Regarding the instrument output, research in music information retrieval (MIR) provides tools for segmenting music from the audio signal only [17] using spectral, tonal, rhythmical descriptors and methods. We identify two main challenges for the implementation of this module. The first is the access to the data from the DMIs. In fact, while extraction from the audio signal provides many features that can be used to detect correspondences, it might not be enough for precise analysis. In order to access pre-mappings and post-mappings data, to differentiate be-

tween manual and automated changes and to analyze the output of individual tracks or other sound processes, one can not rely solely on the DMIs inputs and outputs, i.e. additional software components will need to be integrated. In most DMIs, plug-ins can be added at various stages of the instrument. However, the API might not provide enough information on the instrument to a single plugin. For example, one plugin per track might be needed to access and send the audio output features for each separately. The integration of *EXT-M* will be simpler if DMIs are built using patching environments such as PureData or Max/MSP, where the musician has more control over the architecture of the instruments. The second challenge is to combine detailed but costly and slow analysis of features for the detection of correspondences with maximum accuracy, and fast but less accurate analysis of features for the update of augmentations in *AUG*.

Extraction of the audience familiarity (*EXT-A*) The role of this module is to extract the audience subjective and objective information in order to inform *IA*. It is now commonly accepted that the live music experience engages multiple complementary processes of low level perception, en-action and embodiment (see Leman & Maes for a review [15]), processing of hierarchical and sequential information [16, 13], as well as strong affective and social aspects. This perspective of complementarity is also included in recent studies relative to music produced with DMIs [21]. The evaluation of such a rich experience triggers multiple methodological difficulties. Subjective assessment is therefore a common method that can be conducted through post-performance questionnaires or emotion rating during the concerts [22]. Post-performance questionnaires are a good source of information but may be less precise than live subjective reports. In addition, there are initiatives that do not require the participants to directly communicate their introspective evaluation. These more objective measurements require very specific equipment as eye trackers [3] and, to date, suffer from a lack of out-of-the-lab physiological measurements. Familiarity is one of the key aspects of the experience of live music and, as its other dimensions, relies on multiple underlying mechanisms. Therefore, as the mentioned studies above, and to initiate ecological (i.e. out of the lab, in “real life”) measurements of physiological signals, we propose a dual methodology to extract the familiarity of the audience. First, we base the subjective assessment on a familiarity application for mobile phone. Its main purpose is to supply *IA* with data about the ongoing familiarity of the audience, from a graphical familiarity slider that spectators activate. We believe this continuous survey could be an answer to the reservations we exposed about questionnaires after the performance. The second part of the extraction relies on physiological signals. Neuroscience studies show that the expertise, a key component of familiarity, influences the perception of action [8]. Those findings, applied to music expertise, may lead to a better understanding of its role in the live music experience. By measuring the peripheral signals, we aim at discovering potential patterns that could correlate with the subjective data we extract. Widely used in the emotion studies, and rather suitable to extract in natural condition, four signals are particularly interesting in our musical

context : Heart rate variability [20], electrodermal activity [23] and oculometry (eye tracking + pupillometry). The complementary analysis of extracted features of these signals already gave interesting result in the emotion classification by machine learning algorithms [11] and need to be further extended in music experience studies. Moreover, these signals can already be acquired by wearable devices and the quantified self movement [24] will surely provide more accurate and affordable devices in the near future.

Integration and Analysis (IA) *IA* is the central hub that connects to all the other modules. Its role is to : 1) compute the live correspondences with the data extracted from the musician and their instrument; 2) match the computed live correspondence with a correspondence from the database to calculate the familiarity dimensions that need to be augmented; 3) supply the *AUG* module with the information needed for the relevant augmentation selection; To fulfill these tasks, the module can rely on the classification element of each property of correspondences. Alongside the raw and the semantic parts, the classification element is a machine learning model dedicated to the classification of raw data. This model is pre-computed with the raw files registered in the correspondence. Its goal is to discriminate new stimuli and detect those who match with the recorded one. Considering the heterogeneity of the modalities, each property might require a specific machine learning model and specific extracted features. For example, MIR descriptors for a sound element and a deep convolutional network for picture classification. The main idea is to use the set of models as a global digital representation of the correspondence that can either be projected, depending of the context, on a single and more easy to handle property or be represented as a multidimensional vector that allows similarity comparison of whole correspondences (the matching process). Considering the variety of data and processes it has to handle, this module needs to be regularly updated with recent findings in signal processing, machine learning or movement models. Without a strongly modular structure of the available tools that the pipeline has at its disposal, the framework may not be able to evolve and thereby join a long list of deprecated initiatives.

Augmentation (AUG) This module manages both a database of available augmentations and a set of active ones. When a correspondence needs to be augmented, *AUG* receives the data required to create a new or select an existing augmentation, such as the physical position of the gesture and sensors and logical components of the instruments, the familiarity dimensions that need to be compensated, the destination of the augmentation (individual or group) and the data from *EXT_M* that the augmentation should listen to. The augmentations are selected from a database of augmentations designed to compensate the various dimensions of familiarity.

A number of research have shown the opportunities opened by augmentations to provide information on DMIs to the audience. In previous work, we proposed an augmented-reality (AR) approach where the visualizations are per-

ceptually consistent with the physical instrument, e.g. visual links attached to the physical sensors. We designed a display where multiple spectators can reveal the augmentations [6] and all perceive them consistently. Finally, we proved the effect of 3D visual augmentations on the causality aspect of familiarity [4], i.e. the degree of control perceived by the audience. On the haptic side Armitage [1] has experimented with using haptic feedback to provide information on musician’s interaction during a live-coding performance. We believe that the first challenge is the creation of a framework that allows one to design augmentations according to the specific dimensions of familiarity that they compensate. Rules will need to be defined so that one can adapt an augmentation to the artistic specificity of a particular performance. A second challenge is the design of augmentations that provide just enough content to fill the multimodal gap without distracting the audience from the musical performance because of a too heavy cognitive load.

3 Conclusion

In this paper, we presented a novel approach for augmenting familiarity of the audience with Digital Musical Instruments and reviewed associated research results and challenges. Among the perspectives, our first future work will be the construction of the database of correspondences. This will also help to further inform and refine the decomposition of familiarity and the structure of correspondences. We hope that this paper will trigger exciting new research in both NIMEs, computer science and cognitive science.

References

1. J. Armitage. Revealing timelines: Live coding and its gestures. In *Proceedings of ICLC*, 2016.
2. J. Barbosa, F. Calegario, V. Teichrieb, G. Ramalho, and P. McGlynn. Considering audience’s view towards an evaluation methodology for digital musical instruments. In *Proceedings of NIME*, 2012.
3. M. A. Baytas, T. Göksun, and O. Özcan. The perception of live-sequenced electronic music via hearing and sight. In *Proceedings of NIME*, 2016.
4. F. Berthaut, D. Coyle, J. Moore, and H. Limerick. Liveness Through the Lens of Agency and Causality. In *Proceedings of NIME*, 2015.
5. F. Berthaut, M. T. Marshall, S. Subramanian, and M. Hachet. Rouages: Revealing the Mechanisms of Digital Musical Instruments to the Audience. In *Proceedings of NIME*, 2013.
6. F. Berthaut, D. Martinez Plasencia, M. Hachet, and S. Subramanian. Reflets: Combining and Revealing Spaces for Musical Performances. In *Proceedings of NIME*, 2015.
7. S. A. Bin, N. Bryan-Kinns, and A. P. McPherson. Skip the pre-concert demo: How technical familiarity and musical style affect audience response. In *Proceedings of NIME*, 2016.

8. B. Calvo-Merino, D. E. Glaser, J. Grèzes, R. E. Passingham, and P. Haggard. Action observation and acquired motor skills: an fmri study with expert dancers. *Cerebral cortex*, 15(8):1243–1249, 2005.
9. S. Fels, A. Gadd, and A. Mulder. Mapping transparency through metaphor: towards more expressive musical instruments. *Organised Sound*, 7(2):109–126, 2002.
10. R. I. Godøy, A. R. Jensenius, A. Voldsund, K. H. Glette, M. E. Høvin, K. Nymoen, S. A. v. D. Skogstad, and J. Tørresen. Classifying music-related actions. 2012.
11. J. Kim and E. André. Emotion recognition based on physiological changes in music listening. *IEEE transactions on pattern analysis and machine intelligence*, 30(12):2067–2083, 2008.
12. S. Koelstra, C. Muhl, M. Soleymani, J.-S. Lee, A. Yazdani, T. Ebrahimi, T. Pun, A. Nijholt, and I. Patras. Deap: A database for emotion analysis; using physiological signals. *IEEE Transactions on Affective Computing*, 3(1):18–31, 2012.
13. E. Kohler, C. Keysers, M. A. Umiltà, L. Fogassi, V. Gallese, and G. Rizzolatti. Hearing sounds, understanding actions: action representation in mirror neurons. *Science (New York, N.Y.)*, 297(5582):846–8, 2002.
14. C.-H. Lai and T. Bovermann. Audience experience in sound performance. In *Proceedings of NIME*, 2013.
15. M. Leman and P.-J. Maes. The role of embodiment in the perception of music. *Empirical Musicology Review*, 9(3-4):236–246, 2014.
16. I. Molnar-Szakacs and K. Overy. Music and mirror neurons: from motion to 'e'motion. *Social cognitive and affective neuroscience*, 1(3):235–241, 2006.
17. J. Paulus, M. Müller, and A. Klapuri. State of the art report: Audio-based music structure analysis. In *Proceedings of ISMIR*, 2010.
18. S. S. Rautaray and A. Agrawal. Vision based hand gesture recognition for human computer interaction: a survey. *Artificial Intelligence Review*, 43(1):1–54, 2015.
19. F. Ringeval, A. Sonderegger, J. Sauer, and D. Lalanne. Introducing the recola multimodal corpus of remote collaborative and affective interactions. In *Automatic Face and Gesture Recognition (FG), 2013 10th IEEE International Conference and Workshops on*, pages 1–8. IEEE, 2013.
20. D. Sammler, M. Grigutsch, T. Fritz, and S. Koelsch. Music and emotion: electrophysiological correlates of the processing of pleasant and unpleasant music. *Psychophysiology*, 44(2):293–304, 2007.
21. J. C. Schacher and P. Neff. Skill development and stabilisation of expertise for electronic music performance. In *Proceedings of CMMR*. Springer, 2015.
22. E. Schubert, S. Ferguson, N. Farrar, D. Taylor, and G. E. McPherson. The six emotion-face clock as a tool for continuously rating discrete emotional responses to music. In *Proceedings of CMMR*, 2012.
23. H. Sequeira, P. Hot, L. Silvert, and S. Delplanque. Electrical autonomic correlates of emotion. *International journal of psychophysiology*, 71(1):50–56, 2009.
24. M. Swan. The quantified self: Fundamental disruption in big data science and biological discovery. *Big Data*, 1(2):85–99, 2013.
25. B. W. Vines, C. L. Krumhansl, M. M. Wanderley, I. M. Dalca, and D. J. Levitin. Music to my eyes: Cross-modal interactions in the perception of emotions in musical performance. *Cognition*, 118(2):157–170, 2011.
26. D. M. Wegner and T. Wheatley. Apparent mental causation: Sources of the experience of will. *American psychologist*, 54(7):480, 1999.
27. J. C. Wu, M. Huberth, Y. H. Yeh, and M. Wright. Evaluating the audience's perception of real-time gestural control and mapping mechanisms in electroacoustic vocal performance. In *Proceedings NIME*, 2016.