

2019-01-01

i-Berlioz: Towards Interactive Computer-Aided Orchestration with Temporal Control

Miranda, Eduardo

<http://hdl.handle.net/10026.1/12738>

International Journal of Music Science, Technology and Art
Accademia Musicale Studio Musica

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

This is the authors' own edited version of the accepted version manuscript. It is a prepublication version and some minor errors and inconsistencies might be present. The full published version of this work appeared after minor amendments and revisions in liaison with the editorial and publication team. This version is made available here in accordance with the publisher's policies.

i-Berlioz: Towards Interactive Computer-Aided Orchestration with Temporal Control

E.R. Miranda^{1,*}, A. Antoine¹, J-M. Celerier² and M. Desainte-Catherine²

¹ Interdisciplinary Centre for Computer Music Research (ICCMR), Plymouth University, UK

² Laboratoire Bordelais de Recherche en Informatique (LaBRI), University of Bordeaux, France

Abstract

This paper introduces i-Berlioz, a proof-of-concept interactive Computer-Aided Orchestration (CAO) system that suggests combinations of musical instruments to produce timbres specified by the user by means of verbal descriptors. The system relies on machine learning for timbre classification and generative orchestration, and tools to write and execute interactive scenarios. The paper details the two main parts of i-Berlioz: its generative and exploratory engines, respectively, accompanied with examples. The authors discuss how i-Berlioz can aid the composition of musical form based on timbre.

Keywords: Computer-Aided Orchestration, Interaction, Generative Orchestration, Machine Learning.

Received on 27 / 06 / 2018, accepted on 08 / 11 / 2018, published on DD MM YYYY

Copyright © YYYY Author *et al.*, licensed to IJMSTA. This is an open access article distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/3.0/>), which permits unlimited use, distribution and reproduction in any medium so long as the original work is properly cited.

doi: 10.4108/_____

1. Introduction

Computers have been programmed to make music as early as the beginning of the 1950's when the CSIR Mk1 computer was programmed in Australia to play back popular musical melodies [1]. The piece *Illiac Suite* for String Quartet, composed in late 1950's by Lejaren Hiller, in collaboration with mathematician Leonard Isaacson, at the University of Illinois, USA, is often cited as a pioneering piece of algorithmic computer music; that is, a piece involving materials composed by a computer. Its fourth movement, for instance, was generated with a probabilistic Markov chain [2]. A few years later, Hiller collaborated with Robert Baker to develop a piece of software for composition named MUSIC Simulator Interpreter for COmpositional Procedures, or MUSICOMP [3]. MUSICOMP probably is the first system ever developed for computer-aided composition: "... *it is a facilitator program. It presents no specific compositional logic itself, but it is capable of being used with nearly any logic supplied by the user.*" [4, p.1].

The burgeoning field of computer-aided composition has advanced considerably since MUSICOMP [5]–[10]. The computer has become ubiquitous in many aspects of musical composition. Nowadays musicians have access to a variety of software tools for composition, from user-friendly programming languages [11]–[13] and AI-based generators of musical ideas [14]–[16], to systems for generating and managing musical events interactively in real-time [17], [18].

Whereas the great majority of computer-aided composition systems currently available provide valuable tools for processing music in terms of pitches, rhythms, tempo and loudness, there is a generalized lack of tools for processing orchestration: that is, computer-aided processing of multi-instrumental properties and creation of unique timbres using acoustic instruments. Historically, this deficiency most probably is a legacy of the highly popular MIDI (Musical Instrument Digital Interface) communication protocol [19]. Originally developed in the 1980s to connect digital synthesizers (e.g., to control various synthesizers with one single keyboard controller), MIDI was quickly adopted by the computer music research community. Rather than representing musical sounds directly, MIDI encodes musical notes in terms of their pitches, durations, loudness, and labels

*Corresponding author. Email: eduardo.miranda@plymouth.ac.uk

indicating which instruments of the MIDI sound-producing device should play them. Although the musical possibilities of MIDI processing are vast, MIDI does not encode sounds per se, which renders it unsuitable for processing timbre.

We are interested in developing technology for computer-aided orchestration, or CAO. Despite the existence of a significant and diverse body of research into timbre and its defining acoustic features [20]–[25], there has been relatively less research into orchestral timbre emerging from the combination of various musical instruments playing simultaneously [26], [27]. Our research is aimed at furthering our understanding of orchestral timbre and building systems for CAO informed by such understanding. This paper focuses on the latter: it introduces *i-Berlioz*, a proof-of-concept CAO system that generates orchestrations from verbal timbre descriptors.

Before we proceed, we would like to clarify what is meant by the term ‘orchestration’ in the context of this research. Orchestration here is a combination of musical notes produced simultaneously by different instruments; playing techniques are also considered (e.g., pizzicato or bowed, for string instruments). It is necessary to bear in mind, however, that the art of musical orchestration is much more sophisticated than merely combining notes. It also involves extended playing techniques, dynamics, harmonic context, instrumental idiomatity, and so on. Also note that those resulting combinations of musical notes are referred to as ‘chords’ or ‘clusters’.

Currently, *i-Berlioz* is capable of processing five timbre descriptors: breathiness, brightness, dullness, roughness and warmth. Given a timbre descriptor, the system generates chords of musical notes with the instruments that would produce the required timbre, plus indications of how the notes should be played (e.g., usage of a specific bowing style for a string instrument). Moreover, *i-Berlioz* is able to generate sequences of such clusters for transitions between one timbral quality to another; for example, from very bright to less bright, or from dull to warm. The user can listen to the orchestrations and see their respective music notation. The resulting notation can be saved into a file, which can be edited by most music notation software.

One important characteristic of *i-Berlioz* is its ability to support interactive design of musical form based on timbre. The system is able to hold multiple solutions for specific timbral targets on a time line. And it includes a conditional branching mechanism, which enables composers to specify different branching options and explore them interactively. The system supports the specification of orchestration strategies in a hierarchical fashion and the execution of various solutions can be inspected on the fly, while maintaining a global consistency.

i-Berlioz combines the work on analysis and classification of timbre within orchestral audio and machine learning developed at Plymouth University’s ICCMR [28]–[30] and the research into methods and tools to write and execute interactive scenarios developed at University of Bordeaux’s LaBRI, conducted under the OSSIA (Open Scenario System for Interactive Applications) project [17], [31].

By way of related work, we cite the system *Orchids*, developed at IRCAM, Paris [32]. *Orchids* is aimed at producing orchestrations to imitate a given target sound; e.g., given the sound of a thunder the system would produce suggestions for imitating the timbre of a thunder with the orchestra. The system maintains a database of instrumental sounds. It extracts spectral information from the target sound and from each instrument of its database. This information is feed into a combinatorial search algorithm, which searches for groups of instruments whose combined spectrum matches that of the target sound [33]. *Orchids* outputs dozens of solutions that satisfy the matching criteria, but these solutions can still sound very different from each other. This requires the user to listen to dozens of solutions in order to select one. This process can be very tedious, ineffective and time-consuming, in particular when the user has a particular sound colour, or a perceptual sound quality, in mind. We believe this pitfall can be addressed by designing a constraint-based filtering mechanism to narrow the solutions to a more specific sound quality. *i-Berlioz*’ use of verbal sound description is one approach to achieve this.

The remainder of this paper is structured as follows. The next section presents an overview of the system’s architecture and briefly describes how orchestrations are produced from a given sound descriptor. In order to ascertain that the system outputs an orchestration that satisfies a required characteristic it needs a method for categorizing trials, which will be detailed next. Then, we present the system’s ability to generate sequences and transitions between timbral qualities.

2. System’s Overview

The system comprises two main modules referred to as generative and exploratory engines, respectively. The functioning of the generative module is depicted in the flowchart shown in Figure 1.

The system holds a comprehensive Instrument Database, which contains audio samples of single notes played by orchestral instruments. There are various versions for each note, with different playing techniques and dynamics. Given a Timbre Description, the system generates candidate solutions by assembling combination of notes from the database. With the Timbre Description the user can also provide additional information to constraint the combinatorial search space, such as a list of required instruments and number of notes. An example of a description could be: {bright, string, 4}, which means, a bright timbre using up to 4 string instruments.

The Generate Candidate Solution module generates an audio file, which is rendered using the samples contained in the database. In order to estimate the perceptual quality of the candidate solution, the Timbre Estimation Function module analyses the spectrum of candidate audio file and the results of the analysis are relayed to the Timbre Classification Model, which establishes whether the candidate audio has the required timbre characteristics or not. This process is repeated until a candidate solution matches the requirement.

The exploratory engine is depicted in Figure 2. It is based on an extension plug-in to the *i-score* software that integrates the generative engine with the interaction scoring capabilities of *i-score*, which is detailed in section 4. The general usage process is as follows: first, the user creates a score in which generative processes can be positioned in time. Such a score can contain branches which are specified graphically; these branches allows behaviours such as: “after the execution of a bright-to-warm transition, perform a warm-to-rough transition in one case and a warm-to-dull transition in another”. Then, during playback, cases can be selected according to external controls in real-time. Such controls can come from OSC or MIDI commands.

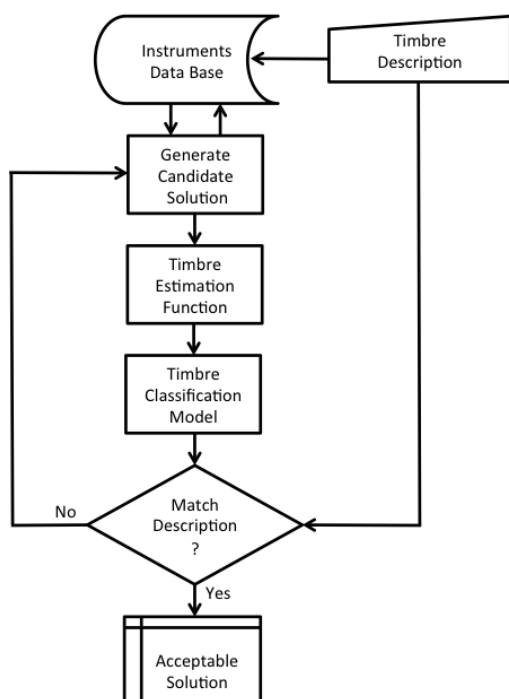


Figure 1. Flowchart of the generative engine.

The chords produced by the generative engine are passed through the audio output ports of the process for direct listening. Besides, each i-Berlioz process provides the list of chords it played during the play-back in a graphical view, so that they can be leveraged by the composer.

3. Timbre Descriptors and Classification Model

The system supports five descriptors of timbre, namely: breathiness, brightness, dullness, roughness and warmth. These descriptors can be retrieved from audio sources by performing various analyses methods, which are detailed below:

- **Breathiness:** in order to measure the level of breathiness of an audio file, we calculate the amplitude of the fundamental frequency against the noise content, and also the spectral slope [34], [35]. The larger the ratio between the amplitude of the fundamental and the noise content, the breathier the sound.
- **Brightness:** the acoustic correlates for the attribute brightness are the spectral centroid and the fundamental frequency [36], [37]. The higher the spectral centroid and the fundamental frequency, the brighter the sound.
- **Dullness:** as with brightness, in order to measure the dullness of a sound we need to calculate its spectral centroid. However, in this case, the lower the value of the spectral centroid the duller is the respective sound [38].
- **Roughness:** is measured by calculating the distance between adjacent partials in critical bandwidths and the energy above the 6th harmonic. A rough sound is characterized by a short distance between critical bandwidths; theoretically, the shorter the distance the rougher the sound [39]–[41].
- **Warmth:** the warmth of a sound is measured by calculating its spectral centroid and retrieving the energy of its first three harmonics. A low spectral centroid and a high energy in the first three harmonics suggest that the sound is warm [42].

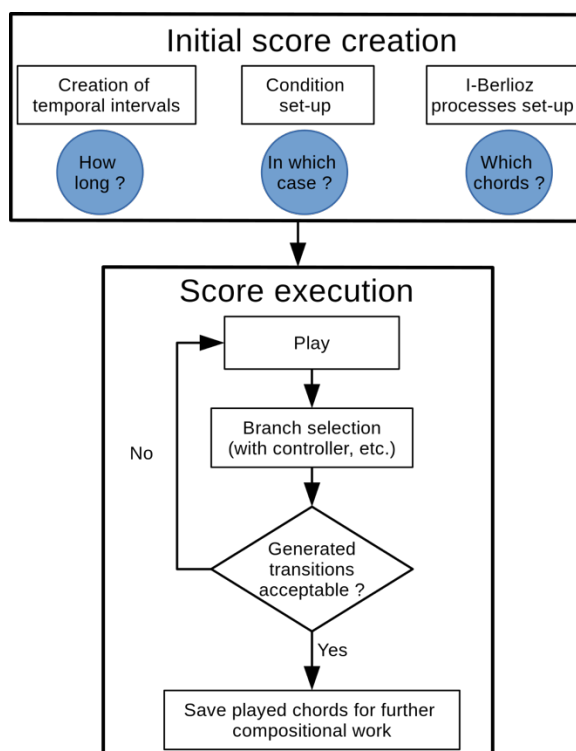


Figure 2. Flowchart of the exploratory engine.

In order to establish whether a solution candidate possesses the required characteristics or not we developed a method to automatically classify audio samples according to timbre description as described above. The difficulty here is that there is no agreed metrics for classifying orchestral audio

samples in terms of timbre properties. Therefore, we developed a bespoke comparative scale for each descriptor as follows: 250 audio recordings of various different well-known orchestral pieces have been analysed; e.g., Beethoven’s *5th Symphony*, Vivaldi’s *Four Seasons*, Debussy’s *Suite Bergamasque*, and Saint-Saëns’ *Carnival of the Animals*, to cite but four. Each of these pieces was sliced five times into audio samples lasting for 1, 2, 3, 4, and 5 seconds, respectively. Here, the analysis of longer audio samples would not provide accurate values as some acoustic features are time related. Therefore, it is essential to split them into short audio samples for analysis purposes. The different lengths have been chosen arbitrarily; these will have to be revisited as this research progresses. The data gathering resulted in performing timbre estimations onto 236,632 audio files, thus, compiling a dataset composed of 236,632 values for each descriptor. The analysis of this large number of samples enabled us to establish a scale for each attribute, and, thus, be able to normalise the data among the five timbre descriptors. Furthermore, the scale for each timbre attribute is continually calibrated as new audio files are analysed.

The development of a comparative scale for the different descriptors enabled us to input timbre values into a machine-learning algorithm, using a Support Vector Machine (SVM) supervised learning model [43]. SVM methods have been successfully applied in various applications, such as face detection [44] and speaker recognition [45] to name but two. SVM models try to find the separation between the different categories with a gap that is as wide as possible to create a delimited space for each category. Then, when a new value is presented, SVM models estimate the space in which the value sits, thus predicting the category it belongs to.

Supervised learning algorithms are dependent of a labelled training dataset. Therefore, the initial step was to create a set of examples that will then be used to train the SVM classification model. Here, examples consisted of calculated timbral values of short audio files as input data, and their dominant perceptual quality, represented by verbal attributes, as the desired output category. The training samples have been selected from the large training dataset created for the comparative scale mentioned previously. In this case, the 250 orchestral audio files that have scored the highest values for each verbal attribute have been selected and manually labelled by the authors with their corresponding attribute. For instance, the 250 samples with the highest values for the attribute brightness have been chosen and labelled ‘brightness’. In total, the corpus training contained 1,250 samples labelled accordingly.

The SVM algorithm has been implemented using the `svm.SVC` function, which is a SVM method for classification task, taken from the Scikit-Learn v0.18 library [46], with parameters kernel type = RBF (for Radial Basis Function), RBF kernel coefficient $\gamma = 0.2$, and penalty parameter $C=1.0$. Figure 3 shows a normalized confusion matrix created to estimate the performance of the classification model generated by the SVM algorithm. Here, the training samples consisted of 90% of the training corpus (1,125 samples), and 10% of the training corpus (125 samples) were selected as testing samples. Using this training dataset, the `svm.SVC`

function produced a success rate of 0.976, which means that the classification model predicted the correct verbal attribute 97.6% of the time. Users have the ability to calibrate the classification models by listening and labelling a selection of audio samples, which are then processed by the SVM algorithm. This method, inspired by the reinforcement learning techniques, allows users to input their own perception levels into the learning process, thus improving the accuracy of the suggested classifications.

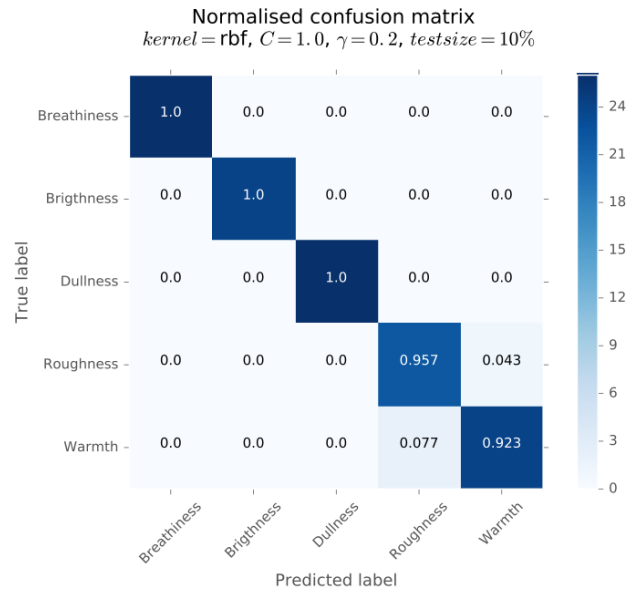


Figure 3. Normalised confusion matrix of the SVM classification model created from rescaled training dataset, with test size = 10% (1,125 training samples, 125 testing samples).

3.1. Solution Candidate Examples

This section illustrates the definition and generation of a solution candidate. As mentioned previously, i-Berlioz utilizes timbre descriptors as output decision parameter in the generative engine. Other parameters can be manipulated by the user to define their musical ideas and guide the generative processes. As additional constraints, the user can specify the group of instruments and the types of interval allowed between the sets of notes of a candidate solution; for instance, a type of chord such as major, minor, whole tones, semi-tone clusters, and so on. These parameters could also be randomly assigned should the user wish so.

For ease of description and discussion, the following examples are constrained to string instruments only: contrabass, violoncello, viola and violin. For the first example, the system was instructed to generate a solution with a bright timbre produced by three different string instruments. The input to the system looks like this: {bright, string, 3}. Figure 4 shows the spectrogram of a candidate solution generated by i-Berlioz. It produced a chord with the following 3 notes: {A#3, bass, note-lasting}, {D, viola, pizzicato-secco} and {F6, violin, note-lasting}. The first

parameter inside brackets is the name of the note, the second is the name of the instrument to produce the note and the third is the playing technique.

In the spectrogram of the audio file generated for this solution candidate we can note that this audio file contains considerable energy (intensity) at its high frequencies, a feature of the ‘brightness’ quality. Figure 5 presents a

spectrogram of a second example, generated from the following request: {warm, string, 3}. In this case, the system produced the following solution candidate: {D#2, bass, staccato}, {B2, cello, pizzicato-secco} and {F#4, violin, pizzicato-secco}. The spectrogram of this solution (Figure 5) shows that energy is more prominent at the lower frequencies of the spectrum; i.e., below 2,500 Hz.




Figure 4. Spectrogram of the audio file generated for a candidate solution matching the descriptor ‘brightness’.




Figure 5. Spectrogram of the audio file generated for a candidate solution matching the descriptor ‘warmth’.

4. Sequencing and Transitions

At the basis of the scheme for making transitions introduced above is a sophisticated engine for writing and executing temporal constraints, based on a formalism for the authoring of interactive scores developed during the course of the OSSIA research project at LaBRI [17]. This formalism manifests itself mainly in the *i-score* software, which is both a visual editor and an execution engine for such interactive scores.

The system is based on the following building blocks: temporal interval, temporal condition, instantaneous condition, and process. Temporal intervals are organized in directed acyclic graphs, with instantaneous conditions being the vertices of such graphs, and the direction of time being the direction of the graph. Temporal conditions allow to trigger groups of instantaneous conditions, which in turn stops and starts previous and following intervals. During the execution of an interval, processes can take place: these can be automations, sound files and effects, scripts, and so on.

The elements of the OSSIA model is shown in Figure 6. The execution takes place as follows: the interval A runs for a fixed duration. When it ends, an instantaneous condition is evaluated: if it is false, the branch with B will not run. Otherwise, after some time and the playback of an audio file, the execution of B ends up in a flexible area centred on a temporal condition. If there is an interaction, B stops and D starts. Else, D starts when the maximum bound of B is reached. Like after A, a condition chooses whether G will execute. If G executes, an automation and another computation are both executed for a fixed duration. Meanwhile, C started executing at the same time than B. C is waiting for an interaction, without any time-out. When the interaction happens, the two conditions following C are evaluated; the truth value of each will imply the execution of E and F. Finally, when H executes, a hierarchical sub-score starts.

The compositional process proposed in this paper only requires the use of instantaneous conditions in order to choose a specific sequence amongst the possible ones, and of the *i-Berlioz* process. This process leverages the generation engine to create and listen to chords according to a specific transition. The following parameters can be specified by the composer:

- The start and end target descriptors
- The duration of each orchestration cluster, or chord
- The number of instruments
- The set of instruments available for generating an orchestration cluster

When the process runs, at regular intervals fixed by the requested chord length, the generation engine receives a chord query. When a correct chord has been generated, its sound file is reported to the *i-Berlioz* process which plays it.

In order to perform a cross-fade between the start and end descriptors, the following algorithm is applied:

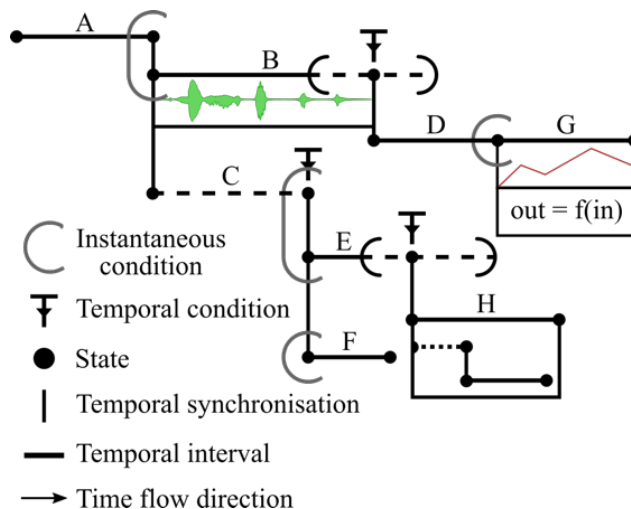


Figure 6. Elements of the OSSIA model.

- On the first tick, generate ten chords and use their attribute value, for instance ‘brightness’ to find a minimal and maximal expected value for the rest of the generation.
- Then, at every tick, generate ten other chords and find the one with the attribute that matches the expected progression the closest over all chords generated, so that the first generated chord has a maximal attribute, and the chord generated at the middle of the process has a minimal attribute.
- For the second half of the process, perform the same method but going from the minimum to the maximum instead.

4.1 An Example of Making Transitions with *i-Berlioz*

This section proposes to illustrate the processes for creating sequences of instrument combination transitions with *i-Berlioz* with a detailed example. First, considering a single run of the process shown in Figure 7, which consists of generating a sequence of chords from ‘brightness’ to ‘dullness’. The corresponding timbral values for each chord composing the sequence are displayed in Table 1. Here, the sequence comprised of ten chords divided in two parts. First, *i-Berlioz* generated five brass instrument chords matching the timbral descriptor ‘brightness’, which were then concatenated in descending order. Then, five chords matching the attribute ‘dullness’ were generated and concatenated in ascending order. In other words, the system was required to generate a

sequence going from ‘very bright’ to a ‘little bright’, then from a ‘little dull’ to ‘very dull’.

Next, consider the scenario example presented in Figure 8. Here, i-Berlioz was asked to generate three sequences using

the number of instruments defined by the curve drawn in the *i-score* automation box displayed at the top—in this example, using two to four instruments.

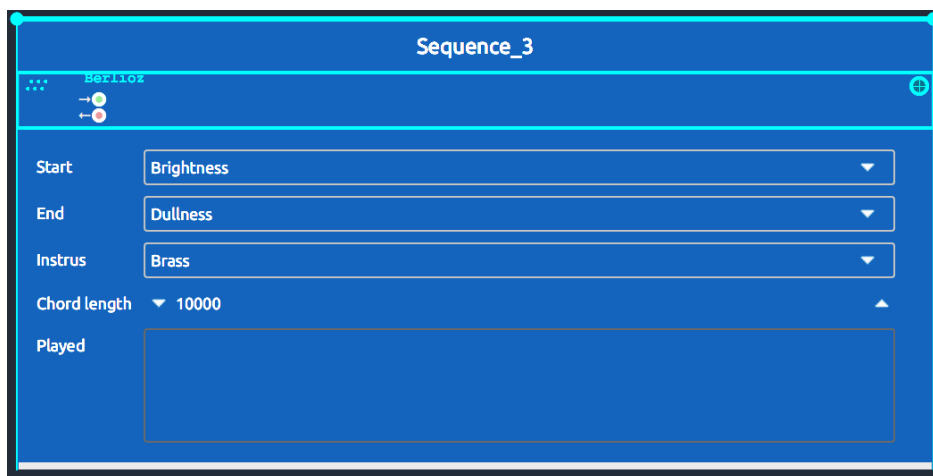


Figure 7. The i-Berlioz process embedded in an *i-Score* object.

Table 1. Calculated timbral values for the sequence shown in Figure 7.

Descriptor	Timbral Values				
Brightness	0.8129	0.7805	0.6685	0.5872	0.497
Dullness	0.6247	0.7170	0.8001	0.8159	0.875

The first sequence consisted of generating instrument combinations from dullness to breathiness with brass instruments. The second sequence had an instantaneous condition (shown in the OSSIA model in Figure 6). Here, the selection of the i-Berlioz box to generate the sequence 2 depended on a condition. In this example, we used a MIDI to specify the MIDI values of a fader. If the fader was up, the box named *Sequence_2a* was selected. Otherwise, if the fader was down, *Sequence_2b* was selected. Here, the fulfilment of the condition was manually defined by moving a fader on a MIDI controller. However, conditions can also be the result of different parameters, such as temporal or number of instruments to combine for example. Then, for the third sequence, the i-Berlioz box *Sequence_3* had a temporal condition which meant to start the generation of this sequence after 15 seconds. This last sequence was composed of brass instrument combinations from brightness to dullness.

5 Concluding Remarks

We are interested in developing technology for computer-aided orchestration, or CAO. Our research is aimed at furthering our understanding of orchestral timbre and building systems for CAO informed by such understanding. This paper introduced i-Berlioz, a proof-of-

concept interactive CAO system that generates orchestrations from verbal timbre descriptors.

Currently, i-Berlioz is capable of processing only five timbre descriptors: breathiness, brightness, dullness, roughness and warmth. Obviously, this is only a starting point. Our goal is to provide flexibility for user-defined vocabularies. However, this is a challenging task because timbre is a difficult concept to formalize, even more so orchestral timbre, as there has been relatively less research into orchestral timbre emerging from the combination of various musical instruments playing simultaneously.

Still more work need to be developed with respect to the Exploratory Engine (Figure 2). Here we introduced how we are exploring the potential of the OSSIA time-based formalism to aid the design of sequences and transitions of timbral events in time, which we believe would ultimately aid the composition of musical form based on timbre. Interactive scores, as presented in [47], allows a composer to write musical scores in a hierarchical fashion and introduce interactivity by setting interaction points. This would enable different executions of the same score to be performed, while maintaining a global consistency by the use of constraints on either the values of the controlled parameters, or the time when they must occur.

Acknowledgements

This research was developed thanks to a scholarship from the Arts and Humanities Research Council (AHRC) in the UK, through the 3D3 Centre for Doctoral Training at the University of Plymouth. The authors also thank the financial support for Initiative d’excellence de l’Université de Bordeaux (IdEx Bordeaux) in France.



Figure 8. A scenario example operation of i-Berlioz.

References

- [1] DOORNBUSCH, P. (2004) Computer Sound Synthesis in 1951: The Music of CSIRAC. *Computer Music Journal* **28**(1): 10-25.
- [2] HILLER, L. and ISAACSON, L. M. (1959) *Experimental Music: Composition with an Electronic Computer*, (New York, USA: McGraw-Hill).
- [3] HILLER, L. and BAKER, R. (1963) MUSICOMP: MUsic Simulator-Interpreter for COMpositional Procedures for the IBM 7090. *Technical Report 9*, University of Illinois at Urbana-Champaign, Experimental Music Studios.
- [4] BAKER, R. (1963) MUSICOMP: MUsic Simulator-Interpreter for COMpositional Procedures for the IBM 7090. *Technical Report 9*, (Urbana, USA: University of Illinois Experimental Music Studio).
- [5] AGOSITINI, A. and GHISI, D. (2013) Real-Time Computer-Aided Composition with Bach, *Contemporary Music Review* **32**(1): 41-48.
- [6] ASSAYAG, G., RUEDA, C., LAURSON, M., AGON, C. and DELERUE, O. (1999) Computer-Assisted Composition at IRCAM: From PatchWork to OpenMusic, *Computer Music Journal* **23**(3): 59-72.
- [7] BRESSON, J., AGON, C. and ASSAYAG, G. (2016) *The OM Composer's Book 3*, (Paris, France: Éditions Delatour).
- [8] FERNÁNDEZ, J. D. and VICO, F. (2013) AI Methods in Algorithmic Composition: A Comprehensive Survey, *Journal of Artificial Intelligence Research* **48**: 513-582.

- [9] MIRANDA, E. R. (2011) *Composing Music with Computers*, (Oxford, UK: Focal Press).
- [10] ROADS, C. (1996) *The Computer Music Tutorial*, (Cambridge, USA: The MIT Press).
- [11] BRESSON, J., AGON, C. and ASSAYAG, G. (2011) OpenMusic: Visual Programming Environment For Music Composition, Analysis And Research. In *Proceedings of the 19th ACM International Conference on Multimedia*, Scottsdale, Arizona, USA, 743-746.
- [12] FARNELL, A. (2010) *Designing Sound*, (Cambridge, USA: The MIT Press).
- [13] WINKLER, C. (2001) *Composing Interactive Music: Techniques and Ideas Using Max*, (Cambridge, USA: The MIT Press).
- [14] BILES, J. A. (2007) Improvizing with Genetic Algorithms: GenJam. In MIRANDA, E. R. and BILES, J. A. [ed.], *Evolutionary Computer Music* (London, UK: Springer).
- [15] GIMENES, M. and MIRANDA, E. R. (2011) An Ontomemetic Approach to Musical Intelligence. In MIRANDA, E. R. [ed.], *A-Life for Music: Music and Computer Models of Living Systems* (Middleton, USA: A-R Editions).
- [16] PACHET, F. (2002) The Continuator: Musical Interaction With Style. In *Proceedings of International Computer Music Conference (ICMC2002)*, Gotheborg, Sweden.
- [17] CELERIER, J-M., BALTAZAR, P., BOSSUT, C., VUAILLE, N. and COUTURIER, J-M. (2015) OSSIA: Towards A Unified Interface For Scoring Time And Interaction. In *Proceedings of TENOR 2015 Conference*, Paris, France.
- [18] TORO-BERMÚDEZ, M. and DESAINTE-CATHERINE, M. (2010) Concurrent Constraints Conditional-Branching Timed Interactive Scores. In *Proceedings of Sound and Music Computing*, Barcelona, Spain.
- [19] WHITE, P. (2003) *Basic MIDI*, (London, UK: SMT - Sanctuary Publishing).
- [20] BARRIÈRE, J-B. (1991) *Le Timbre: Métaphore pour la Composition*, (Paris, France: Christian Bourgeois).
- [21] GREY, J. M. (1975) *An Exploration of Musical Timbre*, PhD Thesis, (Stanford, USA: Stanford University).
- [22] HELMHOLTZ, H. (1954) *On the Sensations of Tone as a Physiological Basis for the Theory of Music*, (New York, USA: Dover Publications).
- [23] SCHOUTEN, J. F. (1968) The Perception of Timbre. In *Reports of the 6th International Congress on Acoustics*, 76: 10.
- [24] SIEDENBURG, K., FUJINAGA, I. and MCADAMS, S. (2016) A Comparison of Approaches to Timbre Descriptors in Music Information Retrieval and Music Psychology, *Journal of New Music Research* 45(1): 27-41.
- [25] ZWICKER, E. and FASTL, H. (2013) *Psychoacoustics: Facts and Models*, (London, UK: Springer).
- [26] ALLURI, V. and TOIVAINEN, P. (2010) Exploring perceptual and acoustical correlates of polyphonic timbre, *Music Perception: An Interdisciplinary Journal* 27(3): 223-242.
- [27] AUCOUTURIER, J-J. (2006) *Ten experiments on the modelling of polyphonic timbre*, PhD Thesis, (Paris, France: University of Paris 6).
- [28] ANTOINE, A. and MIRANDA, E. R. (2017) Computer Generated Orchestration: Towards Using Musical Timbre in Composition. In *Proceedings of the 9th European Music Analysis Conference (EuroMAC 9 – IX CEAM)*, Strasbourg, France.
- [29] ANTOINE, A. and MIRANDA, E. R. (2017) Musical Acoustics, Timbre, and Computer–Aided Orchestration Challenges. In *Proceedings of the 2017 International Symposium on Musical Acoustics (ISMA)*, Montreal, Canada.
- [30] ANTOINE, A., WILLIAMS, D. and MIRANDA, E. R. (2016) Towards a Timbral Classification System for Musical Excerpts. In *Proceedings of the 2nd AES Workshop on Intelligent Music Production*, London, UK.
- [31] DE LA HOGUE, T., BALTAZAR, T., DESAINTE-CATHERINE, M., CHAO, J. and BOSSUT, C. (2014) OSSIA : Open Scenario System for Interactive Applications. In *Proceedings of Journées de l'Informatique Musicale*, Bourges, France.
- [32] ESLING, P. and BOUCHEREAU, A. (2014) *Orchids: Abstract and temporal orchestration software*, (Paris, USA: IRCAM).
- [33] TARDIEU, D., CARPENTIER, G. and RODET, X. (2007) Computer-Aided Orchestration Based on Probabilistic Instruments Models and Genetic Exploration. In *Proceedings of International Computer Music Conference (ICMC2007)*, Copenhagen, Denmark.
- [34] HILLENBRAND, J. and HOUDE, R. A. (1996) Acoustic Correlates of Breathly Vocal Quality/dysphonic Voices and Continuous Speech, *Journal of Speech, Language, and Hearing Research* 39(2): 311-321.
- [35] LANDERA, M. A. and SHRIVASTAV, R. (2006) Effects of Spectral Slope on Perceived Breathiness in Vowels, *The Journal of the Acoustical Society of America* 119(5): 3339-3340.
- [36] SCHUBERT, E., WOLFE, J. and TARNOPOLSKY, A. (2004) Spectral Centroid and Timbre in Complex, Multiple Instrumental Textures. In *Proceedings of the International Conference on Music Perception and Cognition*, Evanston, USA, 112-116.
- [37] DISLEY, A. C., HOWARD, D. M. and HUNT, A. D. (2006) Timbral Description of Musical Instruments. In *Proceedings of International Conference on Music Perception and Cognition*, Bologna, Italy, 61-68.
- [38] BOLGER, D. B. (2004) *Computational Models of Musical Timbre and the Analysis of its Structure in Melody*, PhD Thesis, (Limerick, Ireland: University of Limerick).
- [39] TERHARDT, E. (1974) On the Perception of Periodic Sound Fluctuations (Roughness), *Acta Acustica* 304: 201-213.
- [40] AURES, v. W. (1985) A Procedure for Calculating Auditory Roughness, *Acustica* 58(5): 268-281.
- [41] FASTL, H. and ZWICKER, E. (2007) *Psychoacoustics: Facts and Models*, (London, UK: Springer).
- [42] PRATT, R. and DOAK, P. (1976) A Subjective Rating Scale for Timbre, *Journal of Sound and Vibration* 45(3): 317-328.
- [43] CORTES, C. and VAPNIK, V. (1995) Exploring Perceptual and Acoustical Correlates of Polyphonic Timbre, *Machine Learning* 20(3): 273-297.
- [44] OSUNA, E., FREUND, R. and GIROSIT, F. (1997) Training Support Vector Machines: an Application to Face Detection. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 130-136.
- [45] WAN, V. and CAMPBELL, W. M. (2000) Support Vector Machines for Speaker Verification and Identification. In *Proceedings of IEEE Workshop on Neural Networks for Signal Processing X*, IEEE, 775-784.
- [46] PEDREGOSA, F., VAROQUAUX, G., GRAMFORT, A., MICHEL, V., THIRION, B., GRISEL, O., BLONDEL, M., PRETTENHOFER, P., WEISS, R., DUBOURG, V. and VANDERPLAS, J. (2011) Scikit-learn: Machine Learning in Python, *Journal of Machine Learning Research* 12: 2825-2830.
- [47] ALLOMBERT, A., ASSAYAG, G. and DESAINTE-CATHERINE, M. (2007) A System of Interactive Scores Based on Petri Nets. In *Proceeding of the 4th Sound and Music Conference*, Lefkada, Greece, 158-165.