

# REAL-TIME VISUALISATION OF FUGUE PLAYED BY A STRING QUARTET

Olivier Lartillot <sup>1</sup>

<sup>1</sup> RITMO Centre for Interdisciplinary Studies  
in Rhythm, Time and Motion  
University of Oslo  
olivier.lartillot@imv.uio.no

Carlos Cancino-Chacón <sup>2 1</sup>

<sup>2</sup> Austrian Research Institute for  
Artificial Intelligence  
carlos.cancino@ofai.at

Charles Brazier

Institute of Computational Perception  
Johannes Kepler University Linz  
charles.brazier@jku.at

## ABSTRACT

We present a new system for real-time visualisation of music performance, focused for the moment on a fugue played by a string quartet. The basic principle is to offer a visual guide to better understand music using strategies that should be as engaging, accessible and effective as possible. The pitch curves related to the separate voices are drawn on a space whose temporal axis is normalised with respect to metrical positions, and aligned vertically with respect to their thematic and motivic classification. Aspects related to tonality are represented as well. We describe the underlying technologies we have developed and the technical setting. In particular, the rhythmical and structural representation of the piece relies on real-time polyphonic audio-to-score alignment using online dynamic time warping. The visualisation will be presented at a concert of the Danish String Quartet, performing the last piece of *The Art of Fugue* by Johann Sebastian Bach.

## 1. INTRODUCTION

In this article, we present a new project of real-time visualisation of music performance. It is part of a broad objective, by the first author and especially within the context of his *MIRAGE* project <sup>1</sup>, to design tools to make music more easy to understand and more engaging, especially for non-expert listeners.

The presented project is focused on a String Quartet interpretation of the last piece of *The Art of Fugue* by Johann Sebastian Bach. In order to provide a very rich musicological visualisation of the piece, we decided to prepare the analysis in advance and to synchronise the live performance with the help of score-following technologies.

The paper is organised as follows. Section 2 gives a brief overview of previous works, including ours. Section 3 formalises the underlying principles founding our proposed approach. Section 4 describes what the visualisation is

about, concretely. Section 5 details the underlying technical aspects and especially the score-following approach.

In the following, a distinction needs to be made between pitch tracking (detecting pitch height from audio, frame by frame, as described in Section 5.1) and pitch tracing (drawing the pitch line in the visualisation, as described in Section 4.1).

## 2. PREVIOUS WORKS IN MUSIC VISUALISATION

### 2.1 Brief State of the Art

There is a very large number of approaches that have been proposed for video visualisation of music as it evolves along time. For instance, there has been discussion about the application of music-colour synaesthetic mappings in live ensemble performance [1]. A system architecture has been proposed to build real-time music visualisations rapidly, with the view to provide a sense of musical experience to the users [2].

### 2.2 Our Previous Works in Music Visualisation

This interest in music visualisation stems from the first author's work in audio and music feature extraction and especially the design of *MIRtoolbox* [3]. This toolbox was initially conceived with the initial aim to better understand the affective and emotional impact of music [4]. First video visualisations were made in 2012 <sup>2</sup>, so far focusing on timbral and dynamical aspects of audio recordings. Another visualisation <sup>3</sup> was aimed at highlighting audio and musical features contributing to emotions dynamically expressed by music [5]. Recently, we designed a new type of visualisation with transformation of video shootings based on audio analysis, and applied to the production of a music video <sup>4</sup>.

## 3. GENERAL PHILOSOPHY

This visualisation project is guided by a few general desiderata, presented below, that may look incompatible at first sight, due to conflicting constraints. An additional principle addresses this conundrum.

<sup>1</sup> <https://www.uio.no/ritmo/english/projects/mirage>

Copyright: © 2020 Olivier Lartillot et al. This is an open-access article distributed under the terms of the [Creative Commons Attribution 3.0 Unported License](https://creativecommons.org/licenses/by/3.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

<sup>2</sup> [https://www.youtube.com/watch?v=H\\_SpFh2SPmg](https://www.youtube.com/watch?v=H_SpFh2SPmg)

<sup>3</sup> <https://www.youtube.com/watch?v=8S7LyEkzWGE>

<sup>4</sup> <https://www.youtube.com/watch?v=L1FaxT74hTk>

### 3.1 Desiderata

#### 3.1.1 Accessible and Self-Explanatory

The visuals should be understandable by a largest public. Hence no music expertise should be required. In particular, music score representations should not be displayed. The visuals should also be self-explanatory. No verbal or text explanation should be necessary for the viewers to grasp the logic.

#### 3.1.2 Expressive while Thrifty

The visuals should be as impactful as possible while at the same time being minimalist with regard to the use of visual strategies. For instance, there is no need to use colors, forms or textures if they are not motivated by a particular need to convey a particular musical characteristic.

#### 3.1.3 Conveying the Richness of Music

One aspiration is to offer a visual equivalent of the way listeners — and especially experienced music lovers — perceive, grasp and appreciate music through their ears. Music is generally a rich experience consisting of a flow of events appearing together or successively and building rich interactions. Listeners are immersed by this flow of information, inciting to participate and move with the music. One objective is to convey this experience visually.

### 3.2 Addressing Conflicts in the Desiderata

#### 3.2.1 A Visual Complexity Assimilable Gradually

Clearly, describing the whole musical experience visually may lead to a very complex representation that might be unfathomable at first sight. Although complex stimuli might appeal to the spectators' interest, in order to correctly address desideratum 3.1.1, there should be a guiding thread within the flow of stimuli that the spectator can rely on to progressively grasp what is happening and discover little by little the logic of the whole visualisation.

## 4. VISUALISATION SCENARIO

The proposed scenario exploits the particular characteristics of string quartet fugues, where each of the four instruments plays a monody and where the subtle changes in pitch, dynamics and timbre play an important role. In a fugue, a *subject* — a musical theme — is introduced at the beginning in imitation — i.e., repeated successively on each voice at different registers — and reappears frequently in the course of the composition. The particular piece under study, *Contrapunctus XIV*, is actually structured into three sections, each introducing a new subject.

### 4.1 Polyphonic Pitch Tracing

Each instrument plays a single monodic line where successive notes are clearly separated. The slow tempo and the rare use of small rhythmic values such as sixteenth notes also contribute to the clarity of the melodic lines, and to a focus on pitch, dynamic and timbral shaping of each successive notes. As such, it seemed relevant to simply

display pitch curves, resulting directly from the temporal tracking of the fundamental frequency. The pitch curves are progressively drawn from left to right as time goes by. Each instrument generates its own pitch curve.

#### 4.1.1 Rhythm Quantisation

As we will see in Section 4.2, the pitch curves will be superposed so that similar motives could be compared. For that reason, in order to allow their temporal alignments, the horizontal axis directing the pitch curves needs to be temporally quantised, in the sense that it should be graduated in rhythmic values, as indicated in the score. The rhythmic quantisation of the live performance is made possible thanks to score following capabilities, described in section 5.2.

The fact that the tempo is not constant — and actually fluctuates a lot throughout the piece — raises a technical issue: How to progressively draw a pitch curve along this quantised axis if there is uncertainty concerning the quantised position of each successive instant?

For each note, a hypothetical tempo — i.e. a duration in seconds for that note — is inferred. In one strategy,  $S_1$ , the tempo is assumed to be constant, and the hypothetical tempo is identified with the tempo related to the previous note.

Another strategy, not considered here, would consist in measuring the tempo variation throughout the piece for different performances of the same musicians and use it as a guide for the tempo change expectation. If the tempo changes are guessed correctly, they would first appear in the visualisation, and it would look as if the performers follow the visuals, and in the other case, the behavior of the visualisation might appear random, which might be undesirable.

Let  $t$  be the performance time and  $x$  the time axis in the visualisation. Let's consider a note starting at time  $t_i$  and ending at time  $t_{i+1}$ , corresponding to the metrical positions located in abscissae  $x_i$  and  $x_{i+1}$ . Let's suppose that the chosen tempo for the visualisation is such that the note is expected to end at time  $t'_{i+1}$ . Three cases need to be considered:

- If the expected tempo is exact — i.e.,  $t'_{i+1} = t_{i+1}$  —, the curve is drawn in a simple way, by using the simple mapping  $m_t$  defined, for  $t \in [t_i, t_{i+1}]$  as

$$m_t : \theta \in [t_i, t] \mapsto x_i + \frac{\theta - t_i}{t'_{i+1} - t_i}(x_{i+1} - x_i), \quad (1)$$

which means that the part of the curve already drawn does not modify along time afterwards. It remains constant afterward, i.e., for  $t \geq t_{i+1}$ ,

$$m'_t : \theta \in [t_i, t_{i+1}] \mapsto x_i + \frac{\theta - t_i}{t_{i+1} - t_i}(x_{i+1} - x_i), \quad (2)$$

- If the actual tempo is slower — i.e.,  $t'_{i+1} < t_{i+1}$  —, the curve is drawn in the same way until reaching the expected end of note  $t'_{i+1}$ , so equation 1 is valid for  $t \in [t_i, t'_{i+1}]$ . On the other hand, once point  $x_{i+1}$  has

been reached, the mapping  $m_t$  for  $t \in [t'_{i+1}, t_{i+1}]$  becomes

$$m''_t : \theta \in [t_i, t] \mapsto x_i + \frac{\theta - t_i}{t - t_i}(x_{i+1} - x_i), \quad (3)$$

which means that curve is compressed to accommodate the rest of the note in the same visual space. Once time  $t_{i+1}$  is reached, the mapping stabilises to equation 2.

- If on the contrary the actual tempo is faster — i.e.,  $t'_{i+1} > t_{i+1}$  —, the curve is drawn in the same way until reaching the unexpected end of note  $t_{i+1}$ , so equation 1 is valid for  $t \in [t_i, t_{i+1}[$ . But as soon as the note is ended, the mapping is corrected and from that point the mapping switches to equation 2. For the sake of clarity following desideratum 3.1.1 and in order to avoid emphasizing excessively this disruption, an abrupt change in the graphical representation should be avoided, and the transition between mappings  $m_t$  and  $m'_t$  is shown progressively, in accordance with “fluid” design principles.

#### 4.1.2 Pitch Height Ordering

The disposition of pitch curves on the Y-axis is less of a problem, as a simple conversion of frequencies on a logarithmic scale gives an adequate representation. Besides, accurately representing the relative location in pitch of each instrument enables to convey the registral differences. Moreover, the rare cases of swapping of register order between voices (the second violin crossing for instance the first violin and playing the highest pitch for a few notes) can be immediately noticed in the visualisation. By associating a different color with each instrument (for instance red for cello, black for alto, blue for the second violin and orange for the first violin), this makes this rare disruption in the natural ordering in pitch more noticeable.

Vibrato, portamento and any other form of pitch deviation related to performance are directly exhibited by the pitch curve, and additional treatment to make them accessible to the viewers do not seem necessary in our case.

The actual set of pitch heights used as implicit reference for the pitch space defines both diatonic and chromatic scales. This would lead in theory to a graduation of the pitch axis through a large set of horizontal lines. Following desiderata 3.1.1, the display of the whole diatonic or chromatic scale is avoided. Only pitches that are actualized are indicated with horizontal lines. The tonal structuration of this set of lines is further discussed in section 4.3.

#### 4.1.3 Dynamics and Timbre

Amplitude for each pitch curve is represented by controlling the width of the curve. Timbral aspects will also be depicted in future works.

## 4.2 Dynamic Paradigmatic Analysis

So far the position of the pitch curves on the time axis (as discussed in section 4.1.1) was expressed with respect to

the metrical grid given by the notated music. It was implicitly assumed that the X-axis in the visualisation represented the overall temporal evolution of the piece of music from beginning to end.

Yet such strategy would only show how the music sounds like at each individual instant, without any structured reference with what has been played, apart from the mere display of the left part of the curve before the given time. Music listening is a lot about inferring associations from what we are currently hearing and we previously heard. This leads to the inference of a structural understanding of the piece, and in particular of a thematic and motivic lexicon, where music sequences of various sizes are considered as repetitions of themes and motives. Following desideratum 3.1.3, this structural information should be depicted in the visualisation.

One solution commonly used in music analysis is to add annotations on top of the score — or in our case, the pitch curve —, such as boxes around sequences of notes that are repeated. Boxes would be distinguished with respect to color, width, line style, etc., in order to denote particular theme or motive. Figure 1 gives an idea of the type of motivic structure that can be found in the beginning of *Contrapunctus XIV*.

The problem with that formalisation is that it makes the visualisation more complicated, more technical, and would demand a lot of additional attention from the spectators, which would contradict desideratum 3.1.2. In particular, they would need to find by themselves the position in the score (or the overall set of pitch lines) of each type of boxes, which would quickly become tedious. Besides, it does not seem feasible to display the whole pitch curve on one single screen.

A fruitful solution to this problem can be found in the notion of paradigmatic analysis. It was initially a method designed by Levi-Strauss in order to structure texts, and it has then been introduced in music analysis by Nicolas Ruwet [6] (cf. an example in Figure 2), and used extensively in music semiotics [7]. The idea is simply to align vertically musical sequences that share a same motivic or thematic identity. By compiling the successive occurrences one below each other, the whole score can be read from left to right and from top to bottom.

We partly follow this strategy by cutting the pitch lines into segments and aligning those corresponding to repetitions of a same motif or theme vertically. Contrary to the paradigmatic analysis, however, we do not use the vertical axis as a way to distribute the sequences over time, because the Y-axis is already used for representing the whole 4-voice polyphony.

Since our visualisation is dynamic, evolving over time, we do not need to care about the readability of the whole representation once everything is drawn, but only about the understandability of the dynamic progression of the visualisation. As such, it is not a problem to superpose the repetitions of a same pattern one on top of each other. If they are transposed (which is generally the case in fugues), the corresponding vertical translation will be clearly shown. Besides, the music operation of inversion (very typical to



Figure 1. Motivic analysis of the beginning of *Contrapunctus XIV*. The first subject is shown in red and its corresponding countersubject in blue.

figures as well) will lead to the vertical inversion of particular occurrences as well.

Since the pitch curves are not drawn simply from left to right anymore but may appear at any place in the screen, the part of the curve being extended at a given time should be clearly visible. For that purpose, we use the parameter of line width to also control the recency of each part of the pitch curves. The part of the curve most recently appears thicker while the older sections become thinner and thinner. Besides, the rightmost point of the pitch curve corresponding to the current instant is highlighted with a clearly visible pointer, which has the color related to the instrument, as mentioned in section 4.1.2.

Figure 3 presents the progressive construction of the paradigmatic visualisation of the beginning of *Contrapunctus XIV*. The first 5 bars correspond to the first subject, while the rest of the excerpt corresponds to its counter-subject. Only the first occurrence of the counter-subject is shown

here. We notice that the successive 1-bar repetition in the bass line in bars 8—9 and 9—10 (purple box in figure 1) is folded, the second occurrence on top of the first one. Similarly, for the tenor part, bar 12 is an extra bar extending further bar 11, and is therefore folded in the visualisation.

### 4.3 Tonal Analysis

Concerning the tonal dimension of the music, it is necessary to highlight hierarchies between pitches. As the key of the piece is D minor, it seems reasonable to highlight the main triad D-F-A, which is shown in the visualisation in the form of a horizontal grid. All the lines corresponding to a same pitch class are displayed with a same color (for instance<sup>5</sup>, red for D, cyan for F and green for A). Only lines corresponding to pitches that have been actu-

<sup>5</sup> Examples of colors are indicated in this paper. They will be tweaked so as to optimise the clarity of the representation and the hierarchy of importance.



Figure 2. Paradigmatic analysis of a *Geisslerlied* [6].

ally played are shown, so they progressively “light up” as the music goes.

Attentive spectators will notice that while many entries of the fugue’s first subject go from D (red) to A (green) and back to D, other entries go from A to D and back to A, which has a slight shorter ambitus (fourth instead of fifth, shown respectively by the red and black lines in Figure 3.). The visualisation thus explicitly reveals the dichotomy between *subject* and *answer* in a fugue.

To make explicit the concept of pitch classes, which plays an important role in tonal analysis, as we will discuss in the next paragraphs, a circular representation of pitch classes is shown in addition to the pitch curve representation. This is illustrated in Figure 3<sup>6</sup>. When the pitch classes D, A and F successively appear at the very beginning of the fugue, they appear at the same time in the circle, using the same color code as previously defined. Here also, this core triad defining the main key of the piece will remain displayed during the whole piece.

One core characteristic of tonality is that the music can be considered as a succession of triads (simple triads or seventh chords). These successive triads are shown as triangles on the circle. Only the current triad is shown. Seventh chords are shown with a triangle representing the main triad, plus a second little triangle between the fifth, seventh and root, as shown in Figure 3, panels 1, 2 and 4.

For any given triad, if the root does not belong to the D-F-A triad, it is indicated in orange both in the circle and as a horizontal line. Same for the fifth but with a distinction: it is shown in dark green if it is a perfect fifth, and grey in the other cases (not illustrated in the figure). The colors is also used for the side of the triangle linking the root and the fifth. Similar principles are used for the third, but with a distinction between major and minor thirds. Similar principles are defined for the seventh degree.

As soon as the triad has been replaced by a subsequent one, the previous triad is not displayed anymore, neither in the circle, nor as horizontal lines.

The last concept of importance that is visualised is the de-

<sup>6</sup> The tonal visualization has not been implemented yet, and is therefore manually superimposed onto the screenshots.

scending fifth transition, so important in tonal music for its role in cadences but more generally in the “circle of progression”. Descending-fifth transitions are shown through an actual rotation of the triangle in the circle. At the same time, pitch classes that remain constant through the transition are highlighted as such, while voice leading movements are shown through translations, leading tone going one degree up and the seventh one degree down. Finally, when the descending fifth corresponds to a cadence, arriving to an important tonic chord, the rotation ends with a slightly flamboyant animation.

These tonal considerations will surely appear very arcane at first sight for a non-expert spectator. Following principle 3.2.1, these more complex consideration should be displayed in a rather unobtrusive way, available to curious viewers who already understand the rest of the visualisation.

#### 4.4 General Structure of the Piece

*Contrapunctus XIV*, is a fugue with three different subjects, defining three successive parts, each starting when the corresponding subject first appears. It should be noted that when the second, or third, subject appears, the previous one(s) can still occur from time to time.

At the beginning, the cello (*c*) plays the first subject *A* alone, which is displayed on the leftmost part of the screen. Then the alto (*a*) plays the second entry ( $A_2$ ) of *A* while *c* plays the countersubject of *A*. It should be noted that in this fugue the countersubjects are not exact repetitions of each other, but can be rather characterised as looser passages structured by specific motivic cells. Hence these countersubjects are not represented as single curves but as a succession of segments related to the different motivic cells. For instance in this first countersubject, related to  $A_1$ , there is a successive repetition of a small pattern with transposition, which is shown on the right of panel 4 of Figure 3.

The first section continues similarly with the successive entries  $A_i$  of *A* while the other voices continue their entries and develop new motivic materials. All entries of subjects remain visible (while still being progressively dimmed as indicated in section 4.2). On the contrary, to simplify the visualisation, some of the entries’ continuations and developments completely fade away after some point if they are not deeply related to parts that will appear later in the piece.

The first entry by the second violon ( $v_2$ ) of the second subject *B* appears, at measure 114, as one of the many developments within the dense network related to subject *A*. But as  $v_2$  continues its long entry solo, it progressively occupies the middle part of the screen, and all the material related to *A* fades away, so that this new *B* entry suddenly becomes alone.

It should be noted that *B* is far longer than the other two subjects, but includes internal motivic repetitions, which are represented as such, through vertical superposition of small lines.

Then the first violin ( $v_1$ ) plays  $B_2$  while  $v_2$  plays the countersubject of  $B_1$ , and so on. Then at measure 149,

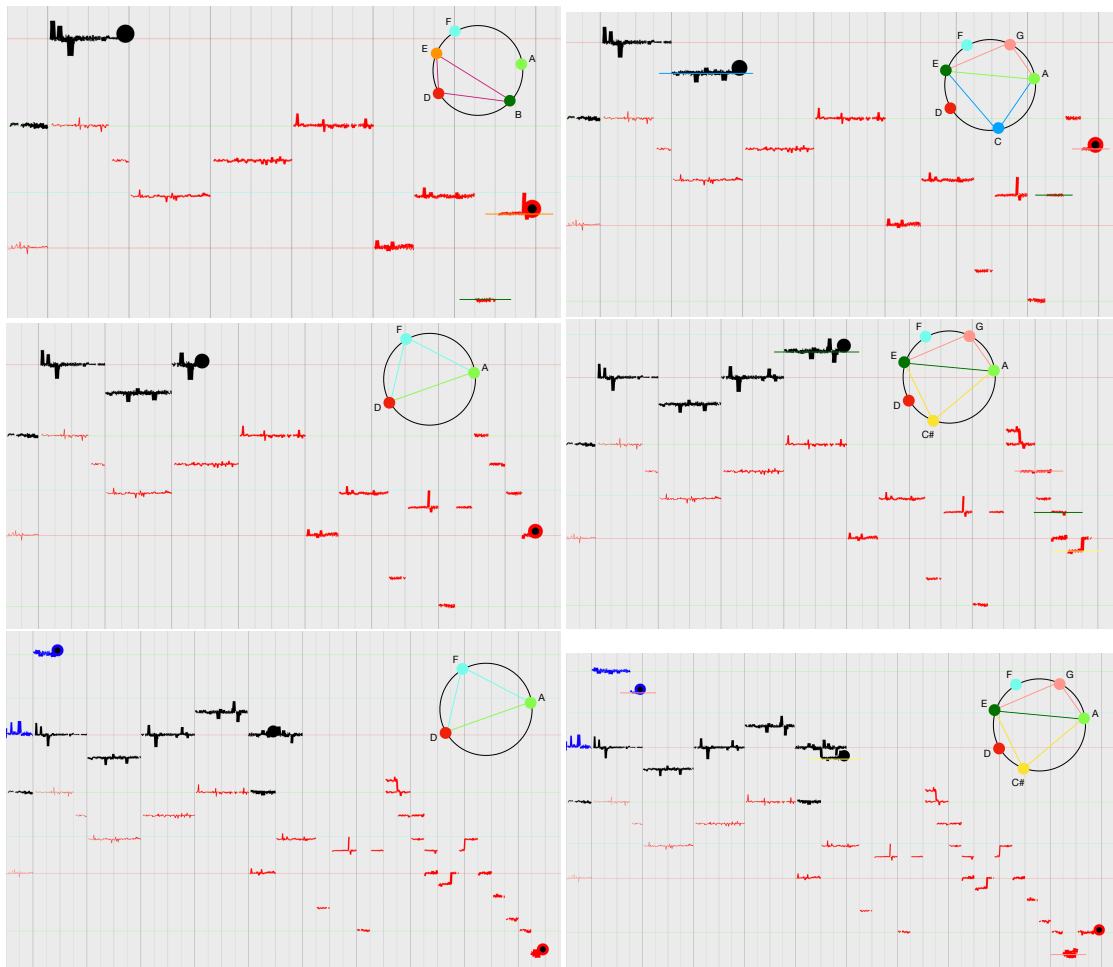


Figure 3. Screenshots of the current version of the visualisation showing various moments at the beginning of *Contrapunctus XIV*. On the top right of each image is added the circular tonal representation, with the main triad D-F-A. Successively from top to bottom, left to right: end of bars 7, 8, middle of bar 9, end of bar 10 and middle and end of bar 12.

$c$  plays a new entry of  $A$  (leading to a progressive visual reactivation of the network of lines related to  $A$ ) while  $v_1$  plays an entry of  $B$  and so on.

There are also rare cases of multiple-voice patterns, where several voices play a pattern that is repeated several times successively. This is a rare case where several pitch curves are superposed not because of their thematic identity but because of their simultaneity.

At measure 193, the third subject  $C$  is introduced by  $a$ . This theme, displayed on the right part of the screen, is known as the “BACH” theme as it starts with  $B\flat$ ,  $A$ ,  $C$  and  $B\sharp$ . So those letters are displayed next to the pitch curve (as well as on the pitch class circle), which also gives some dramatic effect to the visual.

In measure 233,  $v_2$  plays a final entry of  $B$  while  $c$  starts a final entry of  $A$ . These two final entries are particularly highlighted in the visualisation. On measure 239,  $v_2$  and  $c$  end abruptly, while  $a$  continues a short imitation of the ending of  $B$  by  $v_2$  for just one more measure, and everything stops abruptly, as the composition of the piece is un-

finished.

The whole visualisation then fades away with only the four letters  $B A C H$  remaining.

## 5. TECHNICAL SOLUTION

The data-flow diagram of the proposed solution developed for this project is shown in Figure 4.

Each of the four instruments is captured using one of four DPA 4060 close-up microphones. The audio is digitized and mixed into a single 4-channel audio stream using a Behringer UMC404HD audio interface (to be confirmed), which is plugged via USB to a MacBook Pro.

The software part is integrated into a single Mac app developed in Swift. Each of the four channel is fed to a Pitch Tracking module, which sends information about pitch, amplitude and timbre for each successive frame in each instrument to the Visualisation module. In parallel, the four-channel audio stream is also summed into one mono channel and fed to the Score Following module, which sends

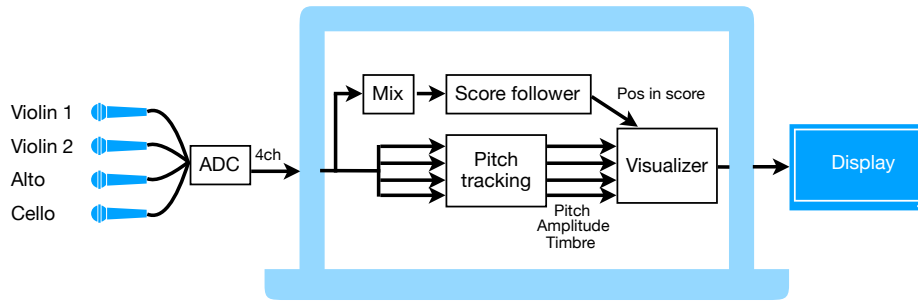


Figure 4. Overall architecture of the proposed solution.

information about the current position in the score to the Visualisation module as well.

### 5.1 Pitch Tracking

The audio signal is decomposed into successive frames of length 50 ms and of hop 25 ms. For each successive audio frame, pitch has been so far estimated by computing the autocorrelation function. To speed up the computation, autocorrelation function is computed using Fast Fourier Transform (FFT), using the vDSP library of Apple’s *Accelerate* framework. The frequency corresponding to the global maximum of the autocorrelation function is selected as the fundamental frequency.

In the first prototype, the analysis is carried out on synthesized rendering of MIDI files. Future works will focus on the particular characteristics of string instruments. Mechanisms will also be developed to avoid incorrect pitch detection due to leakage between instruments.

Other signal processing processes are under consideration to extract timbral descriptions of the sound.

### 5.2 Score Following

Score following refers to the task and technologies for aligning performances of musical pieces to their score in real-time [8,9]. The *dynamic time warping* (DTW) algorithm is the method of choice for performing audio-to-audio alignments [10]. DTW is a general dynamic programming technique for finding the best alignment between two sequences of observations. For the system presented in this paper, we use the *on-line time warping* (OLTW) algorithm [9, 11]. This algorithm adapts DTW to be used in real-time scenarios by computing the alignment incrementally.

We perform audio-to-score alignment by (manually) annotating a live recording of *Contrapunctus XIV* performed by DSQ, which we will refer to in the following as the *reference*. For this recording, the positions of the beats were annotated using Sonic Visualiser<sup>7</sup> by two annotators with 10+ years of formal musical education. These beat annotations are then averaged across annotators. As spectral features to compute the input and reference sequences we used the modified MFCC features described in [10]. These features have been successful in aligning opera in real time [12].

<sup>7</sup> <https://sonicvisualiser.org>

The score follower works as follows:

1. *Before the beginning of the performance:* Extract spectral features for the annotated reference recording and set the current position of the score follower to 0.
2. *During the performance:*
  - (a) Mix the the individual tracks into a single (mono) track.
  - (b) For each new audio frame from the mixed input stream, compute the spectral features.
  - (c) Update the current position of the score follower using OLTW.

### 5.3 Visualizer

The module draws the pitch curves — and their transformations based on amplitude and timbre — related to each instrument into successive segments, and places them in the screen to form the paradigmatic analysis, as discussed in section 4.2. Under the pitch curves are also displayed horizontal lines corresponding to important pitch levels related to the tonal analysis, as explained in section 4.3. The module also draws the tonal circle and the successions of triads and seventh chords. It also displays some particular events such as the pitches B A C H and the particular emphasis on the last entry of subjects *A* and *B*.

Before the real-time process, the paradigmatic analysis is loaded into a particular data structure, containing for each voice (i.e., each instrument) the list of successive segments.

Each segment is represented by a triplet  $(t_1, t_2, x_1)$ , where:

- $t_1$  and  $t_2$  are the time positions of the start and end of the segment, expressed as positions in the score, and
- $x_1$  is the abscissa of the start of the segment in the visual representation.

The tonal analysis is loaded into another data structure, containing the list of successive chords, with also indications of the falls of fifths. Another list indicates when each horizontal lines corresponding to a pitch from the main triad D, F and A lights up.

#### 5.4 Control Interface

A user’s interface is also being designed to allow external control of the system. This will enable to control a number of parameters before and during the concert, related in particular to the amplitude of the input signals or the tuning. This will also allow to manually assign the position in the score in case of problems with the score following or of unplanned performance issues from the musicians’ side.

#### 6. FUTURE WORKS

The system is currently under development. The visualisation will be premiered at a Danish String Quartet concert, which had to be postponed to May 22, 2021, at the Concert Hall of the Royal Danish Academy of Music in Frederiksberg. This leaves a comfortable amount of time to improve the visualisation, in particular related to pitch tracking, and to improve clarity, richness and aesthetic appeal.

The audience will be invited to fill out a questionnaire, to evaluate the fulfillment of desideratum 3.1.1, i.e., that the visualisation should be accessible and self-explanatory to a largest public. The visualisation will also be published online, with a corresponding online questionnaire.

In future works, we plan to generalise the system to broader musical contexts, through the integration of polyphony, of various musical forms. We plan also to integrate a polyphonic pitch tracker that would not require the use of multi-channel instrument pickups, but simply a single mono or stereo pickup. The score following might be used as a guiding point for the pitch tracking.

The musicological analysis, here the paradigmatic and tonal analyses, are for the moment performed manually. In longer term we plan to integrate our systems of automated motivic and structural analysis. By integrating also automated transcription systems, the whole process could be carried out without the need of a score and therefore of a score follower.

#### Acknowledgments

This work was partially supported by the Research Council of Norway through its Centers of Excellence scheme, project number 262762 and the MIRAGE project, grant number 287152; the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme under grant agreement No. 670035 (project Con Espressione<sup>8</sup>); and the Marie Skłodowska-Curie grant agreement number 765068, MIP-Frontiers.

This visualisation project is part of a larger project, the MusicLab5 Copenhagen 2020 event, under the initiative of Simon Hffding, which has been postponed to 2021. We would like to warmly thank the musicians of the Danish String Quartet and Simon Hffding for offering us this fantastic opportunity to carry out this project.

The score in Figure 1 was initially typeset by Tim Knigge, and made available through the Mutopia project<sup>8</sup>.

<sup>8</sup> <http://www.mutopia.org>.

#### 7. REFERENCES

- [1] K. Ng, J. Armitage, and A. McLean, “The colour of music: Real-time music visualisation with synaesthetic sound-colour mapping,” in *Proceedings of the Electronic Visualisation and the Arts (EVA)*, 2014.
- [2] S. C. Nanayakkara, E. Taylor, L. Wyse, and S. Ong, “Towards building an experiential music visualizer,” in *Proceedings of the 6th International Conference on Information, Communications & Signal Processing*, 2007.
- [3] O. Lartillot and P. Toiviainen, “A Matlab toolbox for musical feature extraction from audio,” in *International conference on digital audio effects (DAFx)*, 2007.
- [4] T. Eerola, O. Lartillot, and P. Toiviainen, “Prediction of multidimensional emotional ratings in music from audio using multivariate regression models,” in *Proceedings of the 10th International Society for Music Information Retrieval Conference (ISMIR)*, 2009.
- [5] O. Lartillot, K. Eliard, D. Cereghetti, and D. Grandjean, “Video visualization of predictors of emotions dynamically expressed by music,” in *Proceedings of the 5th Humaine Association Conference on Affective Computing and Intelligent Interaction*, 2013.
- [6] N. Ruwet, “Methods of analysis in musicology,” *Music Analysis*, vol. 6, no. 1-2, pp. 4–39, 1987, translated from *Mthodes danalyse en musicologie*, *Revue belge de musicologie*, 20, pp. 65-90, 1966.
- [7] J.-J. Nattiez, *Music and Discourse: Towards a Semiology of Music*. Princeton University Press, 1990.
- [8] Müller, Meinard, *Fundamentals of Music Processing*. Heidelberg: Springer International Publishing Switzerland, 2015.
- [9] A. Arzt, “Flexible and Robust Music Tracking,” Ph.D. dissertation, Johannes Kepler University Linz, Linz, Austria, November 2016.
- [10] T. Gadermaier and G. Widmer, “A Study of Annotation and Alignment Accuracy for Performance Comparison in Complex Orchestral Music,” in *Proceedings of the 20th International Society for Music Information Retrieval Conference (ISMIR)*, Delft, The Netherlands, 2019.
- [11] S. Dixon, “Live Tracking of Musical Performances using On-line Time Warping,” in *Proceedings of the 8th International Conference on Digital Audio Effects (DAFx’05)*, Madrid, Spain, 2005.
- [12] C. Brazier and G. Widmer, “Towards Reliable Real-time Opera Tracking: Combining Alignment with Audio Event Detectors to Increase Robustness,” in *Proceedings of the 17th Sound and Music Computing Conference (SMC)*, Torino, Italy, 2020.