PHENICX

Off-line music visualization technology

Grant Agreement nr	601166
Project title	Performances as Highly Enriched aNd Interactive Concert eXperiences
Project acronym	PHENICX
Start date of project (dur.)	Feb 1st, 2013 (3 years)
Document reference	PHENICX-D-WP6-150115-D6.1-OfflineMusicVisualizationTechnology-v1
Report availability	PU - Public
Document due Date	Jan 31st 2015
Actual date of delivery	Jan 28th 2015
Leader	UPF
Reply to	Agustín Martorell (agustin.martorell@upf.edu)
Additional main contributors (authors name / partner acr.)	Mark Melenhorst (TUD) Emilia Gómez (UPF) Oscar Mayor (UPF) Gerhard Widmer (OFAI)
Document status	Final

Project funded by ICT-7th Framework Program from the European Commission

SEVENTH FRAMEWORK PROGRAMME

Programme co-funded by the EUROPEAN UNION

PHENICX

Table of Contents

EXECUTIVE SUMMARY

This deliverable summarizes the efforts that PHENICX has carried out until M24 to provide the user with relevant visualizations of symphonic music, intended to enhance the musical experience by audiovisual means (WP6).

This document is focused on the visualisation of information related to the musical piece (WP3) in ”off-line” scenarios, that is, situations in which the visualizations form part of the experience before or after the concert.

Section 1 introduces the goals and context of this deliverable with respect to related tasks and deliverables. Particular attention is given to its relation with D6.3, which focuses on visualisation of performance aspects.

Section 2 introduces the challenge, and sets up the goals and scope of this deliverable.

Section 3 describes a set of selected visualization strategies, motivated by prior analysis of information needs, with focus on the specificities of the symphonic repertoire. Both short-term (local) and large-scale (structural) musical properties, and users with different musical backgrounds are targeted.

Short-term visualisations include regular score information and several simplified abstractions. Piano-roll representations are proposed for melodic tracking. Instrumentation and sound source (spatial) location are simultaneously conveyed by a orchestral layout visualisation.

Large-scale visualisations include multi-scale structural descriptions of the piece, involving several layers of abstraction, as well as instrumentation over time. Simplified tonal content descriptions of the whole piece are also proposed as temporal multi-scale visualisations.

All of the visualisations are aligned in time with an actual symphonic performance (Beethoven’s Eroica, by RCO), in order to explore the mutual information relations when experienced over time.

Five of the proposed visualizations are evaluated by means of a set of user studies, presented in Section 4, intended to identify relevant information needs. Two focus groups with different musical backgrounds have been subjected to the study.

The evaluation of the results confirmed our intuitions about information needs, pointing to the main strengths and weaknesses of the proposed visualisations. Aside the musical information, the actual visual design proved to be particularly relevant for the experience.

The joint effort (visualisation design and user study) has resulted in a proposal of design recommendations. These involves both visual design concerns, as well as the combination of certain musical facets in the same user interface (e.g. joint score and structural information).

Our results provide guidance for the consortium to move forward, and establishes design and integration targets with respect to visualisation strategies.

1 BACKGROUND

WP6 intends to provide users with tools allowing personalized exploration of the information related to a musical concert. In particular, Task 6.1. intends to provide meaningful visualizations of musical pieces and its performance in real time and off-line. This information should be related to different musical facets (WP3 and 4), such as musical structure, tonality and timbre/instrumentation. Information should be extracted from different information modalities (e.g. audio signals or score), and tailored to different user profiles, as defined in WP5.

This document summarizes PHENICX work until M24 on the visualization of different characteristics of musical pieces. Following the definitions and use cases defined in the context of WP2 D2.2 ("Use cases") document, this work focuses on symphonic repertoire and provides input to prototypes in WP7. We work under the use case "**Overseeing the music**": While a concert performance is attended or an audiovisual recorded performance of it is watched, more insight is given with respect to the temporal development during the performance. Different visualizations can be related to features of the piece (e.g. structure or instrumentation), and range from simple abstractions for non-expert audience to more technical (musical) information for people with musical expertise. Alternatively, visualizations can also reflect expressive performance characteristics.

This document focuses on "off-line" scenarios, where the visualizations are designed to be seen "before" or "after" the concert. We exploit the descriptors from audio and score developed in WP3 D3.4 ("Methods to compute music content descriptors"), as well as the score alignment technology developed in WP4 and reported in M12: D4.2 ("Methods for automatic alignment of performance to a score representation"), which permit manual corrections and edition of the alignment. This deliverable is closely related to D6.3 ("Performance visualization technology"). Most of the visualisations discussed in this document have been subjected to user studies. This document covers such studies, as well as the evaluation results. The analysis of these results have resulted in the adaptation of some of the visualisation strategies for its integration in the PHENICX demonstrator, phase II D7.2.

1.1 D6.1 in relation with D6.3

The purpose of this section is to situate deliverables D6.1 and D6.3 with respect to each other. According to the DOW, work around Deliverable D6.1 ("Off-line Music Visualisation Technology") should focus on visualisations of a **piece or composition** itself, and for that purpose work with music descriptors developed in WP3 ("Multimodal Musical Piece Analysis"). D6.3, on the other hand, should develop methods for visualising aspects of a specific **performance** of a piece, that is, the way a piece is played and interpreted and expressively shaped by performers. To that end, it should rely on performance descriptors developed in WP4 ("Multimodal Musical Performance Analysis").

As it turned out, these two dimensions – composition vs. performance – are not always clearly separable. Several of the visualisations to be described can be seen to relate to both structural and performance aspects, particularly when they are shown in synchrony with a specific performance or recording. Some of the visualisations described in D6.1 and D6.3 can be computed from a score alone, or from an audio recording alone, others require both data sources and an alignment between score and recording (see Deliverable D4.2). That is, they are partly computed from the same material (score, audio, alignments) and are related to both piece and performance to varying degrees, but the information is exploited in different ways and with different goals. It was thus decided to assign the developed methods to the deliverables according to the main **purpose** or **aim** of a visualisation:

- D6.1 collects visualisations whose main goal is to to help users **understand how a piece**

is constructed, covering information about the score of the composition, its structural relations, its tonal and melodic content, and its instrumentation. Some of them also involve (inherently) certain performance specificities, as the informational basis of these visualisations come from specific performances. Such is the case of the spatial distribution of the players in the orchestra, or the individual dynamics from each instrumental section.

- D6.3 collects methods geared towards helping users **appreciate a specific performance**, and in particular to **compare performances** to learn about commonalities and differences. That includes representations of expressive timing and dynamics, but also an animated score display synchronised with a given performance. The latter is a main part of D6.1, but it is also presented in D6.3 as a tool for performance appreciation.

2 INTRODUCTION

2.1 Main objectives and goals

There has been a large amount of research within the Music Information Retrieval (MIR) field intended to extract meaningful descriptions from music in audio format or score representation, to compute similarity between music pieces and to classify them according to semantic concepts such as mood, style or preference. However, less effort has been devoted to investigate which are the best strategies to present, in a visual way, this information to users with different profiles (e.g. expert musicians and people with no theoretical musical knowledge) and in different contexts (e.g. enjoyment or education). The main challenges are to provide intuitive visualizations of music pieces, to present information related to different temporal scales (from real-time to global descriptors), and to combine descriptions related to different musical facets.

This document overviews PHENICX approaches to music visualization, where visualizations are key to enrich music concert performances in classical music. We combine descriptors extracted from audio signals and music scores. We study state-of-the-art approaches for music visualization, and adapt them to our particular repertoire. We then integrate them and evaluate them by means of user studies.

The document is organised as follows. Section 3 describes the set of visualisation strategies targetting off-line scenarios. This involves a motivation in terms of information needs and existing solutions. Each visualisation considers the main problems which are unique to (or particularly prominent with) symphonic music, and describes the proposed solution. The different strategies are explained in two blocks, short-term and large-scale respectively, according to the temporal span covered by the conveyed information. Section 4 discusses the evaluation results of a set of user studies, focused on assessing the information needs of users with different musical backgrounds. These user studies covered five of the discussed visualisations in this document.

2.2 Convention

We consider the following writing convention:

- *italics* to newly introduced terms.
- underline to refer to other deliverables of the project.
- **bold** to stress important terms.

3 VISUALIZATION STRATEGIES

Symphonic music is characterised by a particular abundance and sophistication of interrelated musical parameters of interest. While this constitutes a relevant aspect of the listening and aesthetic experience, it may also constitute an overwhelming amount of information for listeners who are not familiar with this kind of repertoires. In this case, the very richness of the music may become a barrier for its enjoyment. Many strategies have been used for bringing (or explaining) unfamiliar music to listeners, specially in educational contexts (e.g. music appreciation courses). Some of such strategies rely upon visual information, which can be designed for conveying or highlighting distinct musical parameters, in order to enhance the listener's awareness about those parameters.

In the PHENICX context, we address some needs/problems which are specific to symphonic music. By specific we mean the musical characteristics and derived signal's properties which are unique to (or particularly prominent with) symphonic music. As discussed in D3.4, such characteristics are mostly related with the large-scale formats of the symphonic compositions and performance practice settings.

Performing arts in general, and music in particular, require the passing of time for their appreciation. Music has to be listened (or imagined) *over time* in order to be experienced. On the other hand, the temporal organisation of the musical *objects* constitutes one of the main compositional and perceptual aspects of music. This calls for visualisation strategies in which the time dimension plays a main role. We conceive music visualisation as a process in which the users play the fundamental (active) role, rather than being mere passive observers. Most of the descriptors discussed in D3.4, capturing musical properties over time, were devised with visualisation in mind. However, a general distinction can be done among the visualisation strategies, according with the temporal scope of the conveyed information.

3.1 Short-term visualisations

By short-term visualisations, we mean the visualisation of musical information *as it sounds*, that is, focusing the viewer's attention towards the musical *now*. This is of course only a convenient way of speaking. Clearly, a physical *instant* cannot convey any musical information, as vibration requires temporal duration. What is visualised is information about a certain temporal (local) segment. This locality depends on the type of information we are interested in, and on the limitations of the human visual and cognitive abilities. The type of visual information in relation with the musical properties that it conveys, calls for considering a diverse set of time-scales and visualisation strategies.

For instance, a simple but useful information is related with the presence/non-presence of individual instruments (or sections) in the orchestra. The perception and recognition of the individual instruments are among the first musical skills learnt in music appreciation courses. The presence/non-presence of a given instrument is a kind of information that can be conveyed at a very high frame rate, for instance, it can be computed from a single analysis frame of few miliseconds. However, human vision abilities are limited with respect to time. Standard visualisation frame rates range from 24 to 100 fps (frames per second). Moreover, instrumentation is not expected to change at such high rates.

On the other hand, a musical *now* can last much longer. Such is the case of a melody, in which the notes change over time, but the sense of coherence or unity is kept all along the melodic discourse. Similar considerations can be argued about larger structures (e.g. thematic groups), although they would generally require higher attentive and musical skills.

3.1.1 Score Follower

The tool by excellence for representing music in graphic formats is related with the concept of *musical score*. The usage of graphical representations for preserving and communicating music compositions is ubiquitous, and it has evolved through many centuries into what is known today as a *score*. Musical scores are highly developed and sophisticated tools, which encode many concepts and abstractions related with the music they represent. A score, thus, requires a sophisticated interpreter, knowledgeable in the specific symbols, musical language, style, interpretation conventions, aesthetics, and so on. A successful decoding of the immense amount of information conveyed by a symphonic score, allows the performers for its interpretation (reconstruction) and its instantiation as actual sound (performance). The score is also one of the main tools for music analysts, as it also encodes many music theoretical and analytical keys to understanding the composition. Users without musical literacy can also find useful information from the score, even if they cannot *read* them, specially when provided with a proper temporal guidance.

It is important to mention that there is not such a thing as *the score* of a given composition. There are as many different scores as potential usages and editors: manuscripts/facsimiles, composing, performing, arrangements/transcriptions, *Urtext*, study or critical scores, are just but some possibilities. As in any information encoding problem, each type of score highlights some specific aspects, while hides or minimizes others. Among the performing scores, the different editorial choices promote distinct interpretative (performance) practices of the same music. Some are very explicit and full of performance indications/suggestions, while others avoid to represent almost anything but the bare notes. Most (if not all) of the actual performances are based in some form of *conducting score* (rarely accessible, often just in the conductor's mind), customised by the conductors to tailor their personal interpretation of the composition. In any case, a graphical score always serves as a useful *proxy* to the composition.

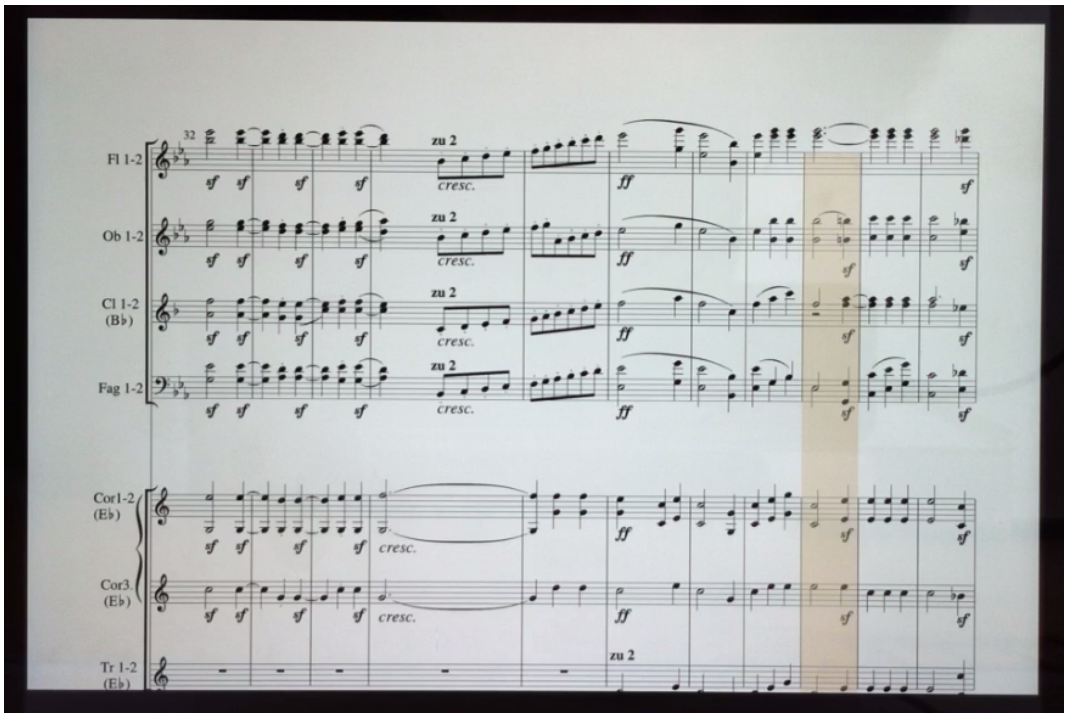

Figure 1: Screenshot of the *Score follower* visualisation case.

Being the score such a rich source of information, a simple visualisation strategy is to *follow* the position of the music in the score, as it sounds. This type of visualisations have achieved a considerable popularity since the last years (e.g., specialised channels and projects devoted

to score following in You Tube, with thousands of followers and positive comments about its usefulness). In the context of PHENICX, this visualisation is referred to as *Score Follower*. Figure 1 depicts a screenshot of the Score Follower visualisation case, as used in the focus group studies (Section 4). A link to the video clip used in the study is available in Appendix 8.1. The visualisation was prepared by OFAI, and it is directly related with D4.2. It consists of an image of the score (Beethoven's Eroica, free edition by the Center for Computer Assisted Research in the Humanities, CCARH), in which the current bar is highlighted with a background colour. According with a temporal alignment at the bar level, the cursor is displaced to the next bar in synchrony with the bar changes in the audio. Pages are also turned accordingly. The method for creating this visualisation is described in D4.2. Visualising a symphonic score presents some problems, as a large number of voices/staves requires large (physical) space to be depicted. The size constraints imposed by most portable devices calls for considering interactive aspects in the visualisation (e.g. zooming in/out, scrolling). Some of these aspects are featured by the video clip used in the study. It is important to notice that Beethoven's Eroica may not be the most appropriate composition for testing this issue, as it is scored for relatively modest instrumental forces (13 staves), compared with the mainstream symphonic repertoire (often above 25 staves).

3.1.2 Reduced piano roll

Among the alternative visualisations to notated scores, the so-called *piano-roll* representations are ubiquitous. Almost any professional Digital Audio Workstation (DAW) uses piano rolls in their standard toolkits, and integrate them in the regular workflow when dealing with music notation. A piano roll simplifies the musical notation, as it only considers pitch (y-axis) and time (x-axis). Using piano rolls for representing symphonic music is challenging, though. The large number of voices, together with the considerable overlapping between them (depending on the orchestration), result in overloaded images in which is difficult to track individual voices, even with the help of colouring strategies. For that, we proposed a simplified variant of the full piano roll, in which only the main melodic line is depicted. For this study, a manual editing of the score was done for selecting the instrumentation of the main melody. This editing was performed by a musicologist, considering few double voicings when required for phrasing or harmonic consistency. The simplified information was then aligned in time with RCO's performance of the Eroica, using a manual alignment at the beat level (as described in D3.4). Bar-based alignments were discarded, because they resulted in annoying (unpredictable) misalignments between the audio and the image (similar perceptual effect than audio/video de-synchronisation in movies). The same beat-level manual alignment was thus used for all the visualisations requiring a precise note-to-note correspondence between audio and video.

Figure 2 depicts a screenshot of the *Reduced Piano Roll* case, as used in the focus groups user studies (Section 4). A link to the video clip used in the study is available in Appendix 8.1. The central part of the image, represents the main melodic line, coloured according to the predominant instrument(s) playing it. Textual/coloured labels account for the involved instruments. A colour legend at the right side serves as a legend to the overall instrumentation. The top side of the image depicts a large-scale view of the same information, covering a duration of few minutes. A vertical white cursor indicates the current time of the performance. In the main (central) piano roll, the cursor is static in the screen, and the piano roll information is gradually displaced to the left. In the long-term visualisation (top), is the cursor that moves over a static image.

3.1.3 Instrumentation and physical space: orchestral layout

A very relevant factor of symphonic music has to do with the physical space it requires to be performed and experienced. Large orchestras require large stages and large halls for projecting the sound to audiences of considerable size. This expands the sonic possibilities, which are richly exploited by both composers and conductors. The spatialisation of the sound is a

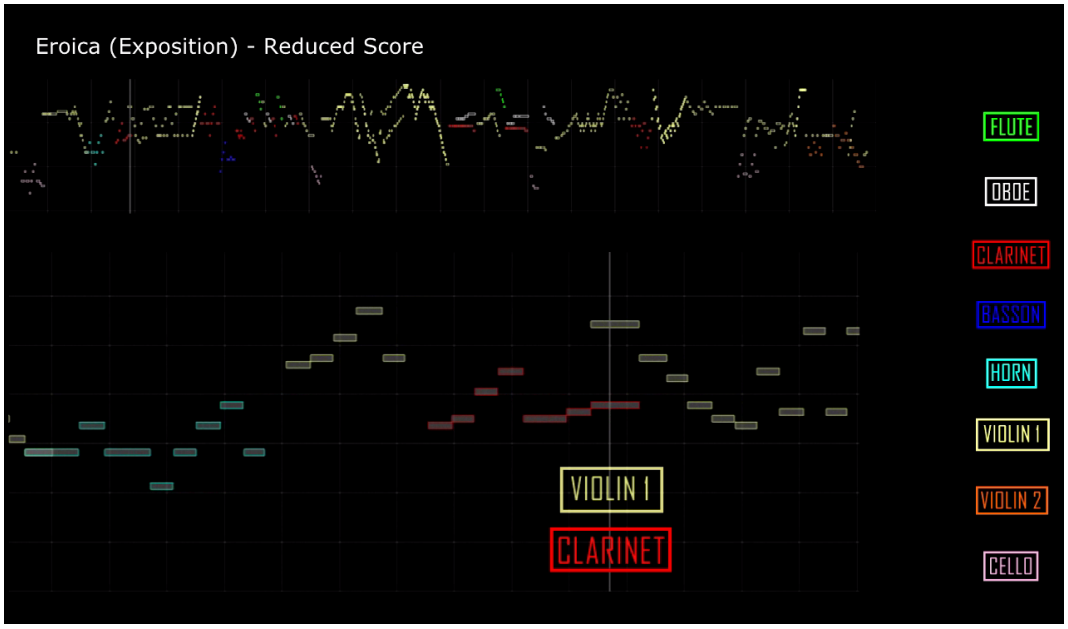

Figure 2: Screenshot of the *Reduced piano roll* visualisation case.

key element of the performance, as it affects the way listeners perceive the different sonorities as coming from different locations. This is also a main concern for producers when mixing symphonic recordings. As discussed for the *Instrumentation* case (Section 3.2.2 below), visualisation can be used for explaining which instruments are playing at each time. A complementary visualisation can extend this information, by explaining from "where" the different instrumental sounds arrive. That is, a spatial localisation of the instruments in the stage. This is referred to as the *Orchestral Layout* visualisation. The base descriptor is the same as for the instrumentation: the activation of every different instrumental section over time. Additionally, it also uses the individual dynamics from each instrumental section, as obtained by means of source separation (D3.6). Figure 3 depicts a screenshot of the Orchestral Layout visualisation case, corresponding to Eroica's performance by RCO, as used in the focus groups user studies on visualisation (Section 4). A link to the video clip used in the study is available in Appendix 8.1.

The figure represents a schematic view of the orchestral layout, as featured during the actual performance. This information was rendered from a very high-angle (almost bird's-eye view) picture of the concert hall. This perspective, as used in photography and cinema, is purposeful for presenting the individual subjects (here, individual players, including the conductor) as part of a wider context, not giving protagonism to anyone in particular. In contrast with the lower-angle perspective in the concert hall, as the audience observes the stage, this extreme angle is intended to induce the viewers with a sense of situational *omniscience*. On top of this layout, different colours are highlighted for representing the activity of each of the musicians. The intensity of the colours follow the individual dynamics of the section. The visualisation, thus, integrates three different musical parameters, related with the composition (instrumentation), the physical space (layout) and the performance (individual dynamics). The visualisation also represents a colour legend with textual labels naming the instruments. Tentatively, the visualisation also depicts simplified pitch information for each instrumental section, represented by the vertical position of an horizontal bar below each instrument's legend.

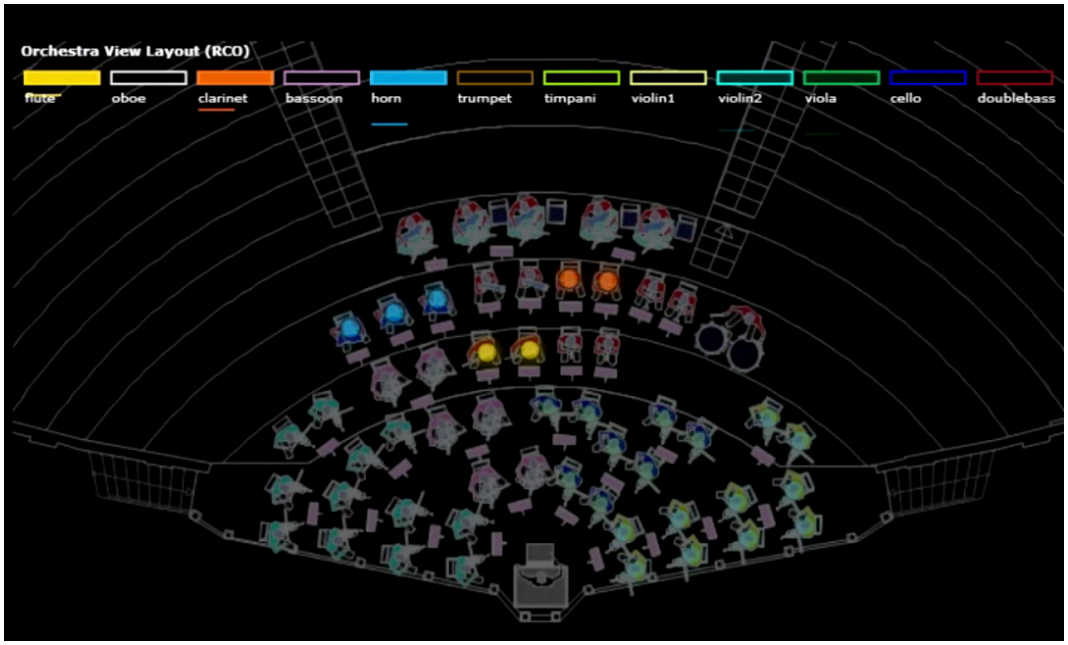

Figure 3: Screenshot of the *Orchestral layout* visualisation case (Eroica).

3.2 Large-scale structural information

Musical structure is a fundamental element in music appreciation. In much symphonic music, the problem of the musical *form* has received particular attention from composers, musicologists/analysts, critics and music lovers alike. The concept of musical form is a very high-level abstraction, related with the *narrative* power of the listening experience. This form can serve as a temporal scheme about the musical discourse, and there are a number of *standard* forms which are expected to appear in many symphonies. The first movement of a prototypical symphony, for instance, is generally expected to follow the so-called *sonata-allegro form*, although the actual compositions are extraordinarily diverse in this respect. The musical form, in the case the music attaches to some of the conventional ones, can (or cannot) be revealed by the actual musical *structure*. Form and structure are related, but not equivalent terms. The musical form is generally a more abstract concept, often involving aesthetic interpretations and/or intentions from the listeners beyond the actual composition's structure. Musical structure is thus prone to interpretation, and it actually constitutes one of the main challenges/pleasures for analysis-oriented listeners. On the other hand, the large-scale symphonic contexts require the usage of large-scale narrative resources, which often call for long and complex structures. Trying to find one's way within such complex musical discourses, however, requires a considerable (active) effort from the listener. This musical richness, on the other hand, may be too challenging for audiences non familiar with the symphonic repertoire, who may get lost in a music difficult to *understand*.

3.2.1 Multi-scale structure

The access to the musical structure can be facilitated by visual means. In the context of PHENICX, the term *structure* is referred to as a temporal scheme in which the music is segmented in chunks and (somehow) *explained*. As a *sense of structure* can be conveyed by different compositional resources, different visual representations can be designed accordingly. For instance, different sections can be distinguished by distinct predominant rhythm patterns, by means of a different tonality, by using different melodic themes, by changes of instrumen-

tation, by inducing distinct musical characters/moods, etc. Most often, it is a combination of parameters what creates the contrast. In some cases, it may be convenient (or at least reasonable) to assign *labels* to those segments. Those labels can inform about the properties of the music in the segment, as well as about its relation with other segments. A wise choice of segments/labels, thus, can help the listener not to get lost in the music. Structural descriptions can be designed at different degrees of detail. A gross segmentation, in terms of movements or large-scale sections, may be enough for many listeners (this is the kind of information usually *described* in most concert's program notes). Listeners with musical training may prefer detailed structures explained in technical terms. However, too detailed information may result in *spoiling* the listening experience. The goal of the structure visualisations is to help the listeners and improve their listening experience (not to substitute it), so a balance has to be found.

In the context of PHENICX, structural information is conceived as a visual index to the music. This is achieved by creating an image, which represents the whole composition. The image is comprised of several segments, which can be labelled by text strings and/or coloured according to different criteria. The current time in the music is highlighted in the structure by a cursor, which is synchronised with the music in order to be located accordingly. In the offline scenario (before and after the concert cases), the structural visualisation is proposed to be used as an interactive *navigation* index as well.

Figure 4 depicts a screenshot of the *Structure* visualisation case, as used in the focus group user studies (Section 4). A link to the video clip used in the study is available in Appendix 8.1. The visualisation, corresponding to the exposition of the Eroica, was prepared by UPF. The segmentation and labelling was done manually upon the score information. This information was then aligned in time with the performance by RCO. The Structure visualisation was not proved problematic when using a bar-level audio to score alignment. However, we used the more precise manual alignment at the beat level, as it was already available (D3.4). In the study, several information layers were tested in a two-tier approach. The top structure depicts a short-scale segmentation, labelled in terms of a proposed set of thematic material. These *themes* are divided in subsections, labelled by abstract characters, following standard (paradigmatic) music analysis terminology. An additional information layer is superimposed by means of a colour code, which indicates the two main tonal regions and a short tonicisation. The bottom structure, depicts the segmentation of the complete first movement of the Eroica, and it is labelled according to the sonata-allegro standard terminology. The choice of the segments, labels and levels of abstraction was decided by a musicologist. On top of each structure (short- and long-term), a running cursor indicates the temporal position of the music as it sounds.

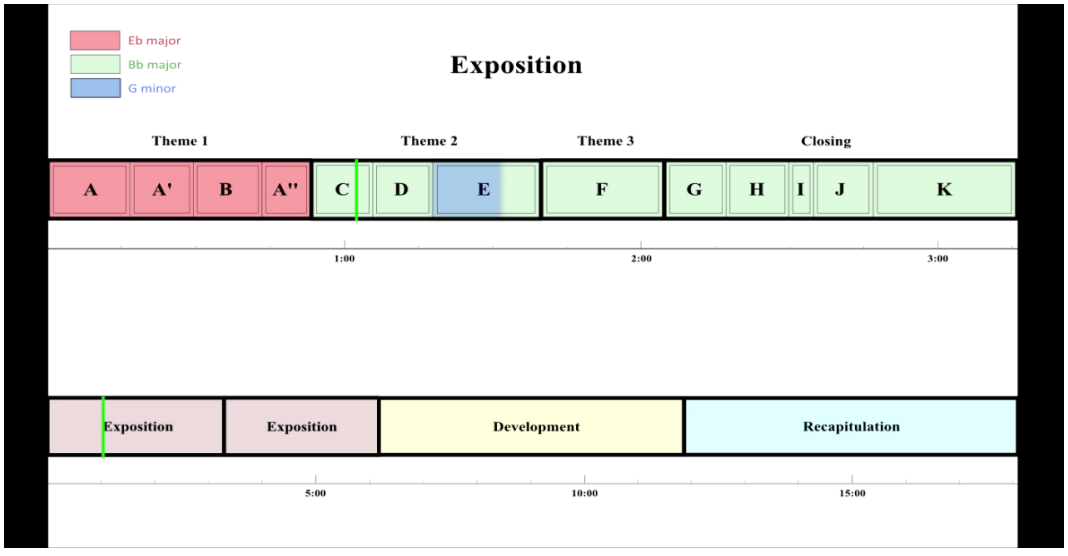

Figure 4: Screenshot of the *Structure* visualisation case.

3.2.2 Instrumentation

A useful simplification from the full score can be achieved by representing only the activation of the different instruments over time. This constitutes a considerable information reduction, which frees a substantial visual space in the device's display. This allows the *compression* of the time dimension, so that a complete symphonic movement can be visualised in a single image. The representation, thus, may serve as an index or navigation map to the composition in terms of its instrumentation. In D3.4, we described previous research about the *instrumentation descriptor*. Instrumentation is among the most important resource for clarifying the musical structure in symphonic music, often informing about main structural boundaries. Similar visual representations are common in professional DAWs and music scoring applications (e.g. Sibelius), as tools for improving the composing/arranging workflow. Figure 5 depicts the instrumentation descriptor computed from the Eroica's exposition, as described in D3.4. Figure 6 depicts a screenshot of the visualisation designed for the focus groups studies (Section 4). A link to the video clip used in the study is available in Appendix 8.1. It presents the same essential information as the descriptor, but including explicit distinction between each individual instrumental sound, by means of a different colour and a textual label. The instrumental families (woodwinds, brass, percussion and strings) are also labelled. Similarly to the rest of the visualisation, this information is aligned in time with RCO's performance of the Eroica (beat-level manual alignment), and the position of a running cursor (green vertical line) is updated along the music.

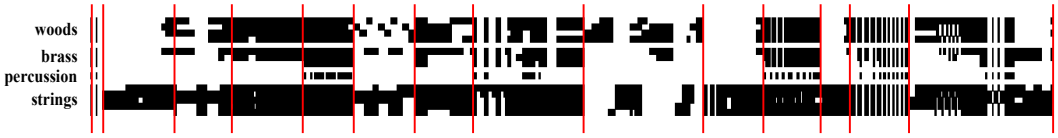

Figure 5: *Instrumentation* descriptor, as discussed in D3.4.

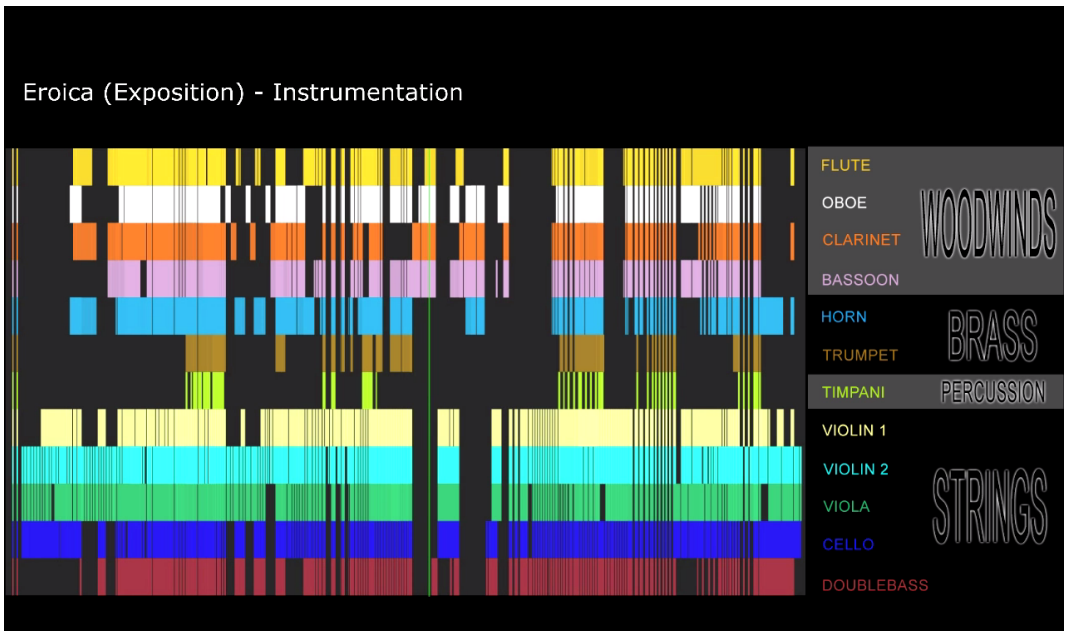

Figure 6: Screenshot of the *Instrumentation* visualisation case.

3.2.3 Tonal structure: simplified keyscape and modescape

As described in D3.4, several music-related features can be computed from both the audio and the score alike. Among the tonal descriptors, the *keyscape* accounts for the tone centre estimation over time, computed at many time-scales (see D3.4 for an overview of the general method). We conducted informal evaluations of the keyscapes as visualisation aids for understanding symphonic music. Aside their potential for revealing structural boundaries (see D3.4), we investigate on the reliability of the tonal information itself: what a keyscape really says in terms of the actual tonal content of the music. As it has been widely documented (Martorell, 2013 for a critical review), the state-of-the-art tonal estimation methods are still limited, when it comes to analyse tonally *complex* music. The case of the Eroica is prototypical: it is clearly a tonal piece, but it features continuous short- and long-term modulations, many of them with a considerable degree of ambiguity and uncertainty for its estimation. This results in keyscapes patched with many colours, which are difficult to interpret even with the help of a colour legend. Put simply, there are *too many* keys/tonicisations in a single movement of a symphony, and no current key description method is free of estimation errors. We investigated different summarisation and colouring strategies to address these issues.

The problem of visualising tonal information pertaining to many classes, introduced in (Martorell, 2013) and extended in (Martorell and Gómez, 2015), was addressed by adopting a relative solution. Instead of using an absolute (one-to-one) key-to-colour mapping, we use a gradual colourmap that represents a single dimension. This dimension is referred to as the tonal distance between the analysed segment and a given reference. This distance is computed from a well-established (perceptually validated) space of interkey distances (Krumhansl, 1990), adapted for dealing with ambiguity (Martorell, 2013). In our case, the tonal reference is the tonic of the movement. The resulting simplified keyscape represents the closeness of every segment of the music to the tonic of the movement. Figure 7 (central pane) depicts this information, computed from the audio signal corresponding to the first movement of the Eroica (RCO). The darkest areas correspond with the closest segments to the tonic, so it is easy to observe that the main body of the movement stays away of the composition's tonal reference (without specifying that this reference is Eb major). The lighter areas, on the contrary, represent the farthest tonal regions used in the piece. These far-reaching tonal excursions are expected to abound in the development section of a sonata-allegro, which is the case of the Eroica.

An alternative tonal simplification is related with only the modal sonority. This is motivated by the assumption that many listeners would distinguish the difference between major and minor, while only those with absolute pitch would tell the *name* of the tonic. For that, we do not consider the specific tone center, but only the estimation of the mode (major or minor). We also compute a measure of the confidence (how strongly major or minor is the segment). The visual representation, named a *modescape*, requires a different type of colouring, usually referred in colorimetry to as *divergent* colourmaps. Divergent colourmaps are designed for improving the gradual distinction between two *opposing* classes (typically, positive and negative), while being able to discriminate small differences in between unambiguously. Figure 7 (bottom pane) depicts the modescape computed from the same audio signal as for the keyscape, in which the bluish colours represent major mode (the darker, the *more major)*, while the reddish represent minor mode in similar terms. This simplified tonal representation alleviates many of the misestimation problems of keyscapes, however, like any automatic tonal description, it is not completely free of misleading information.

One of the lessons learnt from the focus groups user studies (Section 4) is related with the usage of colours in the visualisations. The usage of many different colours has been proved unappealing for some subjects. Accordingly, for the tonal visualisation research, we investigate further about perceptual colourmaps in the colorimetric literature (CIE, 2004), in relation with the properties of the underlying data to be represented. For that, we also test image processing techniques (smoothing) that minimize the introduction of misleading visual artefacts. Among the many tested colourmaps, we have focused on those more robust to individual (perceptual) differences. This includes some colourmaps specifically designed for users affected with the

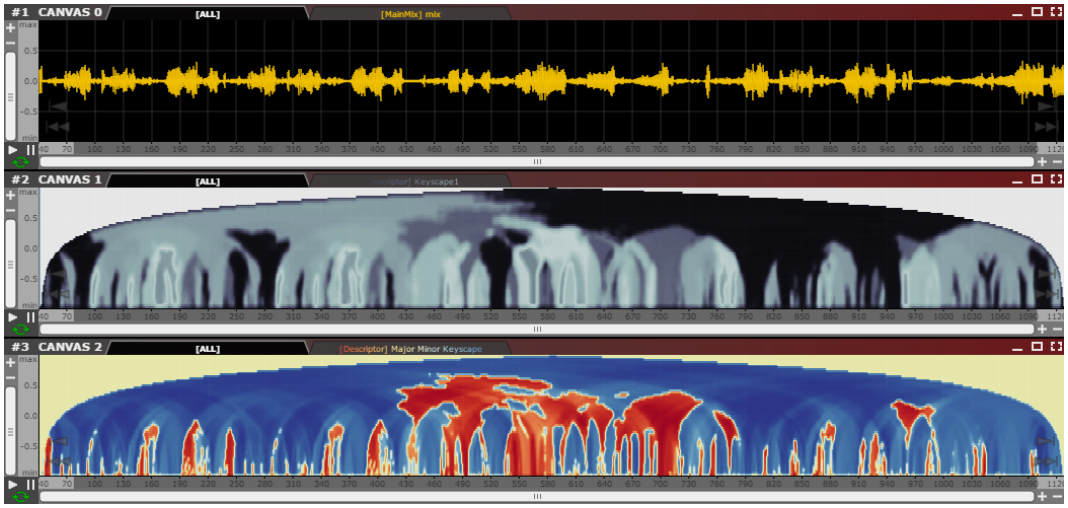

Figure 7: Visualization along listening tests: *keyscape* and *modescape*

most extended types of daltonism. We also conducted informal evaluations, taking advantage of UPF's repoVizz framework, as it permits to synchronise any of the visualisations with the audio performance. Figure 7 depicts a screenshot of the evaluation framework, which includes the audio signal, the keyscape, and the modescape, corresponding to RCO's performance of the Eroica. These visualisations have been proposed for its integration and testing in the PHENICX demonstrator, phase II (D7.2).

4 EVALUATION

In the previous section, a set of visualisation strategies were proposed under the assumption that the resulting implementations support users to cope with the complexity of symphonic music. In this section we report on two focus groups we have carried out to address this assumption.

4.1 Research questions

- What are the main challenges when experts and non-experts are confronted with a piece of music they are unfamiliar with?
- To what extent can information be useful for experts and non-experts to make these challenges easier? Which information is preferred in terms of modality, timing (before, during, after the concert), and content?
- How do experts and non-experts evaluate different implementations of visualisation strategies to help them understand the musical structure of a piece?

4.2 Methodology

4.2.1 General approach

Two focus groups were organized. One group was organized in Amsterdam, with predominantly casual consumers as well as some heavy consumers (see D2.8). The second group was organized in Barcelona, which involved conservatoire (higher music education) students.

Five different visualization concepts were shown to the participants with the purpose of eliciting information needs and design requirements. The concepts were meant as probes (Boehner *et al.*, 2007) to stimulate discussion rather than as early prototypes of working applications.

4.2.2 Participants

Participant requirements:

- Participants must be interested in new ICT developments.
- Participants must be able to reflect on their needs with regard to information visualization and on how ICT can support their needs.

Recruitment in Amsterdam: Participants are recruited via RCO's concert friends. Participants were recruited via e-mail. An invitation was sent to the casual consumer groups that have been involved in earlier studies.

Recruitment in Barcelona: Conservatoire students (majors in performing and composing) were recruited via ESMUC. Participants were recruited from the regular courses and via e-mail.

4.2.3 Materials

Five video clips were composed, each proposing a different visualisation idea, that would address the information needs of users who are unfamiliar with (symphonic) music works. The analysed visualisation concepts (from Section 3) were: Score Follower, Reduced Piano Roll,

Orchestral Layout, Structure, and Instrumentation. An music excerpt was selected with relevant elements that would benefit from visualization. For this purpose, the exposition of Beethoven's Eroica (about 3-minute long, as performed by RCO) was used. Links to all the video clips, as used in the study, are available in Appendix 8.1.

A questionnaire was composed to evaluate the visualisation ideas, containing questions about the perceived added value, the support they offer to understand the music, and the distraction impact when used during a concert. A brief demographics and background survey questionnaire was also used. The questionnaires, as used in the study, are available in Appendix 8.2.

4.2.4 Procedure

1. Signing informed consent forms.
2. Introductions and introduction to the project.
3. Discussion about challenges to listen to new music (RQ1).
4. Discussion about opportunities for suppport (RQ2).
5. Presentation of visualizations and filling out questionnaires.
6. Ranking the visualizations in order of preference.
7. Discussion about rankings.

As the length of the session was restricted to 75 min. in Barcelona, steps 3 and 4 had to be left out in that discussion. The ranking discussion was set up in order to have participants think about why they prefer one visualization over the other in terms of the information they present.

4.2.5 Data analysis

The discussion parts of the focus group were transcribed and segmented into episodes. A thematic analysis was conducted. Thematic analysis is a method for identifying, analysing, and reporting patterns (themes) within (qualitative) data (Braun and Clark, 2006; Fereday and Muir-Cochrane, 2006). Data were coded and themes were identified with the purpose of deriving support needs, design principles for the visualisations and in-depth feedback on the visualisation concepts.

Survey responses were averaged and standard deviations were computed. Given the small numbers of participants, the data should be treated with care.

4.3 User needs with regard to visualisations

4.3.1 Types of needs

Across the two focus groups, two classes of needs can be identified:

- The need for specific information, possibly on a specific moment in time
- The need to have control over the information that is displayed and/or the sound that can be heard.

Fig. 8 summarises this information. Please note that these information needs are not mutually exclusive. For instance, support in focusing attention can be done by pointing out the structural elements of a piece of music.

Need	Description	Example
Structure	Need for information that explains the structure of a piece of music.	[f1#p7] "If you have a very good hearing, you can pick up on the structure. However, most people don't. This is therefore much richer to explain what you are hearing."
• Themes	• Need for information that helps listeners to discover the themes in the music	[f1#p9] "You don't know where themes overlap and you can't relate them to the music you hear"
• Score	• Need to have access to the music's score	[f1#p3] "You can find out exactly what you like and where you pay attention to"
Instrumentation	Need for information on what instrument is playing, and possibly also when.	[fg1#9] "I like to have a seat on the balcony, so that I can see where the sound comes from. I want to understand it and I also investigate it to understand how it sounds with other groups."
Focal points	Need for information on what elements in the music the listener should pay attention to	[fg1#3] "Could also help me to hear different things in how the piece is structured. Becomes nicer if you know what to pay attention to" [fg1#6] "It doesn't really tell what to pay attention to. You can see that the timpani let loose, but you don't see it a second before ('look what's going to happen')"
Background information	Need for background information about the piece of music and about the musicians	[fg1#p2] "This type of information makes the concert far more lively. You often sit down listening for 2 to 3 hours. In those cases I'd like to know more about the musicians on stage. Where did they study? What do they do? Why are they here?"
Score	Need for having the score in view while listening	
Status information	The need for information about the current part that is being played.	[fg2#1] "It is also useful if you get bored, to know that there is still 30 minutes to go [collective laugh]" [fg1#1]. "Suppose you get lost in the paper program. That could happen to me" (e.g.: the lyrics)
Preview	The need for information about upcoming parts	[fg2#4] "It would be better if you can see what comes next, because expectation makes you to get caugt."
Support in recognizing themes	The need for information that helps the listener recognize recurring patterns in the music	[fg2#5] "I would find it interesting if it shows the main themes or motives, for instance, if something will be repeated. Maybe the structure information is not that important during the concert."
Support in recognizing non-obvious elements	Support in the recognition of parts of the piece that do not immediately attract the attention of the (non-expert) listener.	[fg1#7] "At some point I have observed a melodic line I haven't heard before, because I don't know the score (…) Pick out the concealed voices. To pick out the sounds that are latent."

Figure 8: Classes of needs.

Factors were:

- The quantity of the information presented (fg1#9, “I found this one quite distracting. It made me hyper. An enormous amount of information. A lot of things flash, and you see lines going up and down”, orchestral layout).
- Interestingness (fg1#3, “This one was a bit boring for during the concert”, structure).
- Required effort (fg2#1, “I think it is very restricted: or you see too many things at a time, or else you have to interact continuously with the application, and this distracts yourself from the music”, score follower).

The focus groups were consistent with earlier studies in the sense that participants were concerned about the impact of visualizations during the concert on the public as a whole (fg1#6, “If you don’t feel like it yourself at that particular moment, but the one next to you is continuously using it”, concluding discussion).

4.3.2 Effect of expertise

The thematic analysis of the focus group results provided indications on the effect of musical expertise on the need visualization information. However, the relationship proved to be more fine-grained than one might initially suspect.

One might suspect that the score follower is more suitable for those with some level of formal music education in order to be able to read the score. However, some participants in the first focus group who indicated that they couldn’t read the score were still interested, because the score allowed them to follow the main melodic line (fg1#4, “I can’t read the sore, but I don’t miss it either. The way notes vary already indicates what is going to happen”).

Other participants did confirm the hypothesis that score information is primarily interesting for experts (fg1#9, “If you can’t read the score, then it becomes more difficult to do something with it”).

Participants expressed that instrumentation information was only interesting for non-experts who just got started with listening to classical music. (fg1#2, “Not so interesting for experts, but people who just got started might benefit from it”, “Hey, a flute is starting to play”, instrumentation).

The structure visualizations appealed to both expert and non-expert participants. For experts, structure visualizations have added value for contemporary pieces whose structure is difficult to understand (fg2#5, “I think it will be nice for difficult works, those for you don’t understand most of it, and they are many (laugh), and specially for some contemporary music, in which you are a bit lost”, structure).

In the first focus group, participants expressed a general, genre-independent need for structure information (fg1#7, “If your hearing is very good, then you can understand the structure, but that doesnt apply to most people. This [visualization] is therefore much richer in terms of explaining what you hear”, structure).

Even though the data analysis has revealed some patterns in the relationship between musical expertise and information needs, a more detailed insight is needed in order to personalize the information that is offered to users. The conclusions will elaborate on our plans for future studies.

4.3.3 Use of the visualisations before, during and after the concert

In PHENICX we perceive the concert experience to include the concert-goer’s preparation for the concert and the reliving of the concert afterwards. In this study, we therefore also addressed

the contribution of music visualizations to these three different stages of the concert experience: before, during, and after the concert.

We asked the participant whether the visualizations were suitable for the concert experience phases. The results are displayed in Fig. 9. The numbers refer to the number of participants that expressed the intention to use the visualization during the corresponding phase when it were available as a working application.

Visualization	Group	N	Phase		
			Before	During	After
Orchestra layout	Casual consumers	9	5	1	4
	Conservatoire students	6	4	6	3
Instrumentation	Casual consumers	9	5	2	5
	Conservatoire students	6	5	5	3
Structure	Casual consumers	9	4	3	5
	Conservatoire students	6	5	5	5
Score follower	Casual consumers	9	4	4	7
	Conservatoire students	6	4	5	4
Reduced score	Casual consumers	9	5	3	4
	Conservatoire students	6	4	6	5

Figure 9: Suitability for concert experience phases.

The results demonstrate that the Barcelona group was more inclined to use the visualizations during the concert than the Amsterdam group. The latter result is in line with earlier studies in which we found that concert-goers are very hesitant to use technology during the concert. While the type of information that is offered is appreciated, the level of distraction from the performance is the main factor that causes participants hesitation. For example, in this study a participant said about using the orchestra layout during the concert (fg1#6, “Rather not, you can see it in front of you (and also prettier) on the stage”, orchestral layout).

The professional interest of the Barcelona group is likely to influence the participants’ attitude towards using visualizations during the concert. Participants expressed that they considered the visualizations a useful learning tool (fg2#3, “learning about the meaning of the work, and the relations between structure and harmony”, structure). However, their responses are likely to include an estimation of the added value for non-experts (fg2#6, “learn who play the main melody, better control of your listening. Only for amateurs, as musicians are not interested only

in the melody", reduced piano roll).

4.4 User perceptions of the visualisations

For each of the visualizations, user perceptions were assessed with a brief questionnaire using Likert scales. The questionnaire can be found in Appendix 8.2. User perceptions were assessed with respect to:

- The support the visualization provides for the discovery and comprehension of the music
- Perceived added value and impact of the visualizations.

Fig. 10 summarises this information. The visualizations each contained different types of information to support users in listening to the music and discovering new aspects they might not have heard otherwise. The results indicate that the intended contribution to the fulfilment of the users' information needs was recognized by the participants.

Visualization	Group	Support			
		Discovery of new elements	Understanding the structure	Recognition of instruments	Recognition of themes
Orchestra layout	Casual consumers	3.3 (1.5)	3.2 (1.6)	4.1 (1.3)	2.7 (1.3)
	Conservatoire students	4.2 (.4)	3.5 (1.4)	4.8 (.4)	3.5 (.5)
Instrumentation	Casual consumers	3.7 (1.1)	3.9 (.7)	4.2 (1.3)	2.1 (.9)
	Conservatoire students	3.2 (1.0)	2.8 (1.5)	4.2 (.8)	2.0 (1.0)
Structure	Casual consumers	3,8 (.8)	4,1 (.9)	1,8 (1.3)	4,0 (1.1)
	Conservatoire students	3,8 (.8)	4,7 (.5)	1,2 (.5)	4,2 (.8)
Score follower	Casual consumers	4.4 (.7)	4.2 (.7)	4.2 (1.3)	3.9 (1.4)
	Conservatoire students	3.5 (.8)	3.2 (.4)	4.7 (.5)	3.7 (1.2)
Reduced score	Casual consumers	3.4 (1.3)	3.4 (1.5)	4.3 (.5)	3.0 (1.6)
	Conservatoire students	3.6 (0.5)	2.8 (.9)	4.6 (.6)	4.2 (.8)

Figure 10: Users' perception.

Overall, the score follower contributed to most to the understanding and discovery of the music. This is not surprising, given that having access to the score allows for the discovery of the whole

composition: the structure of the piece, instrumentation, and other elements a listener might not discover while listening.

However, the positive evaluation of the score follower by non-experts was surprising. The participants' explanations showed that a listener's ability to read the score did not detract from the value of the score follower: according to the participants it still allows for, for instance, the discovery of the main melodic lines in the music.

Apart from the user's assessment of the support to the discovery of the music, we also assessed the perceived added value for the concert experience. The questions included the distraction that might occur when the concept was developed as an application to be used during the concert. These questions were introduced, because participants consistently reported this as a reason for their reluctance to introduce technology in the concert hall during earlier focus groups.

The results in Fig. 11 show that participants were most positive about the added value of the score follower. The relatively low perceived added value can be the result of the prototypical nature of the designs: the videos were designed to explain the concept, without paying much attention to the interaction design.

Visualization	Group	User perceptions				
		Added value[a]	Comprehension of information[a]	Distraction of oneself	Distraction of others	Motivation to listen
Orchestra layout	Casual consumers	2.2 (1.6)	4.9 (.3)	4.3 (1.0)	4.6 (.7)	1.9 (1.1)
	Conservatoire students	2.8 (.8)	5 (.0)	2.7 (1.0)	3 (.9)	4.2 (.4)
Instrumentation	Casual consumers	2.9 (.9)	3.8 (.4)	4.7 (.7)	4.8 (.4)	2.8 (1.4)
	Conservatoire students	2.0 (1.0)	5.0 (.0)	2.8 (.5)	3.2 (.9)	3.6 (.8)
Structure	Casual consumers	3.8 (1.1)	3.2 (1.0)	3.4 (1.5)	3.8 (1.4)	2.9 (1.0)
	Conservatoire students	2.4 (1.0)	4.0 (.0)	2.8 (1.0)	2.8 (.8)	3.6 (1.4)
Score follower	Casual consumers	4.1 (1.1)	4.9 (.4)	4.1 (1.5)	4.0 (1.2)	3.0 (1.0)
	Conservatoire students	3.7 (.8)	5.0 (0)	2.7 (1.4)	3.0 (1.3)	3.6 (.8)
Reduced score	Casual consumers	2.0 (1.7)	2.8 (1.3)	4.4 (.9)	4.2 (1.1)	2.6 (1.2)
	Conservatoire students	3.0 (1.0)	5.0 (.0)	2.4 (.6)	3.0 (.9)	3.6 (.5)

Figure 11: Added value and distraction.

As a result, the visual stimulation from the designs is low, which is expected to be an important factor for the concert experience and the perceived added value. Before introducing the visualizations into the demonstrator, careful interaction design will take place. Additionally, in upcoming studies, we therefore explicitly measure the hedonic quality (Hassenzahl, 2004) participants gain from using the visualizations.

For the casual consumers, the structure visualization was also positively evaluated, which is different from the perception of the conservatoire students. This is likely to be caused by their level of expertise: these participants expressed that they were able to pick up the structural elements in the music without the support of the visualization.

Apart from the reduced piano roll (casual consumer group), participants did not experience any issues in understanding the information that was offered in the visualizations.

The distraction levels differ between the casual consumers and the conservatoire students. While casual consumers do not listen to concerts professionally, conservatoire students have a professional interest in music. Therefore, they might be more inclined to actively engage with and study the music. As a result, the motivations for conservatoire students that relate to escapism and being away from their everyday's life, are expected to be less important than for the casual consumers. Consequently, they consider the level of distraction to be smaller.

4.5 Design recommendations

4.5.1 General design principles

The visualization concepts were discussed in the two focus groups with the purpose of understanding user needs as well as eliciting general design principles for music visualizations in a symphonic concert setting. In this section we summarize the tentative design principles that were derived from the focus group discussions. Fig. 12 summarises this information.

An important element for future studies is the surprise factor. Participants in focus group 2 have demonstrated the need to be surprised during the music: (fg2#6, "I voted this the worst. It is related with what I said before, as a musician, sometimes I don't want to see what comes next. It's kind of spoiling a movie. Well, sometimes I would like to. But this does not give much information", structure).

However, on the other hand they expressed a need to get an overview of what is coming up (fg2#4, "It would be better if you can see what comes next, because expectation makes you to get catched"). For this participant, engagement is a function of his awareness of what comes next. In that case, there is a risk of a spoiler effect, as argued by fg2#6.

A trade-off between attracting attention towards specific elements in the piece versus overstimulation was identified, for which timing seems to play a major role: (fg#6, "It doesn't tell to what you need to pay attention. You can see that the kettledrum goes wild, but not a second in advance", "Look what's about to happen", orchestra layout), (fg#7, "There is only one thing there, so that's what you start to pay attention to. It's the pink elephant effect", reduced piano roll).

4.5.2 Design recommendations for the visualisations

In addition to the general design principles for classical music visualizations, the data were also analysed with the purpose of generating design recommendations for the visualizations that were discussed. In Fig. 13 and Fig. 14, an overview of the design recommendations is provided, structured according to each visualization.

Design principle	Source	Time[1]	Explanation
Minimal effort expectancy	1	B	Participants (fg1) expressed a need for condensed, tailored information that can be consumed in a brief amount of time, with a preference of video over textual information.
Minimalistic design	1, 2	B, D, A	Designing visualizations in such a way that it helps user to focus their attention, but refrains from overloading them with visual clutter, was considered important by all participants.
Unobtrusiveness	1, 2	D	Participants argued that the main focus should be on the music. Apps can add to the experience, but without removing too much attention from the music itself. Designing applications for within the concert hall should be done with great care.
Appealling design	1		Participants frequently commented on the visual design of the visualization concepts. Visual design seemed to influence their perception of the added value of the concept.
Adaptability	1, 2	B, D, A	Participants need to be able to change what is displayed according to their personal information needs. The focus groups have demonstrated that prior knowledge is the main factor determining the information need with respect to score, structure, and instrumentation. Adaptability was assumed to lower the barriers to start listening to classical music.
Surprise	2	B, D, A	Particpants in fg2 have an ambivalent attitude towards surprise. On the one hand, they want to see what is coming. On the other hand, they appreciate being surprised, as surprise makes (new) music interesting. There is a risk of a spoiler effect when an app shows upcoming particularly interesting fragments listeners should pay attention to.

Figure 12: General design principles. B = before a concert; D = during a concert; C = after a concert

Problem type	Issue description	Feedback
Orchestra layout		
Feature request	Switch on and off who's playing	"I hope that you can choose per person who is playing and that you can really interact with it"
Feature request	Time information - combine with Structure (the current position in the piece)	
Feature request	Opportunity to go back in time	
UX issue	Visual clutter	"I became restless and distracted. An enormous amount of information"
Feature request	Show in advance when something is going to happen	"What I miss is that one does not know what will happen." "It would be better if you can see what comes next, because expectation makes you to get catched."
Score follower		
UX issue	Lack of overview - too many staffs displayed at once	"or you see too many things at a time, or else you have to interact continuously with the application, and this distracts yourself from the music" "the problem is how to locate visually which instrument plays what you want to listen to, for instance, if you listen the bassoon, you have to find where the bassoon is in the score, and sometimes it is not even on the screen." "Those people start to sweat when they have to read one port. If you do all eight, then you can go to bed right afterwards."
Feature request	User's should have the opportunity to pick out the instrument of which they want to see the score	
UX issue	Highlighted bar indicating the current position is too distracting.	"The bar was not so nice. It could have been left out (...) I kept trying to look ahead, but the color kept changing. It wasn't smooth. Jumpy."
Feature request	Display of the musicians' notes on the score	"I'd like to see the musician's notes." "You get a very personal perspective on one's interpretation of the piece"

Figure 13: Design recommendations (Orchestral Layout and Score Follower).

	Problem type	Issue description	Feedback
Instrumentation			
	UX issue	Too many colored blocks	"A lot of color", "Very 80ies", "Space invaders"
	Perceived usefulness	Continuously active instruments don't need to be displayed	
	Feature request	Depiction of intensity (like a graph)	
	Feature request	Opportunity to choose which instruments are playing (e.g. influence the audio stream)	
	Feature request	Opportunity to choose which playing instruments are depicted	
	Perceived usefulness	Information whether instrument is playing or not is perceived as limited	"Only 'on' and 'off' don't do justice to an instrument"
	Feature request	Combination with structure labels	
Structure			
	UX issue	Naming themes with characters (A-G) raised comprehension issues.	"I first thought, oh wow, the alphabet, but now I understand and I think that people would find this interesting"
	Feature request	Opportunity to use the structure visualization to select interesing parts to listen to.	
	Feature request	Combination with score follower (e.g. display structure elements like themes below the score)	"I'd like for this concept to be displayed below the score follower. Then you can both read the score and get an explanation of the score"
	Feature request	Instrumentation information is lacking.	
Reduced score			
	UX issue	Comprehension - goal of the visualization not clear	
	Perceived usefulness	Listeners can pick up the pitch themselves	
	UX issue	Comments on visual appeal and comprehension of color use. Colors for different instruments are too similar.	
	UX issue	Highlighted instrument labels are distracting	"The highlighted instrument labels made me restless. It was very Las Vegas."
	Feature request	Opportunity to use the structure visualization to select interesing parts to listen to.	"Maybe it has more sense as an index, rather than for visualization during the concert."

Figure 14: Design recommendations (Instrumentation, Structure, Reduced Score).

5 CONCLUSIONS

This document has reported the efforts done in PHENICX with respect to off-line visualisation strategies, addressed to the specificities of the symphonic repertoire. Five of these visualisations have been subjected to user evaluation, in terms of information needs, and have covered two focus groups with different musical backgrounds.

The analysis of the two focus groups have demonstrated that there seem to be a need for each of the types of information the visualization ideas appeal to, confirming our prior analysis on information visualisation strategies. Even though the study was aimed to the visualisation concepts (in somewhat abstract terms), the specific visual design seemed to be a strong factor for the users. The discussions have enabled us to formulate a list of general design principles for the visualizations. In addition, the focus groups have resulted in a list of visualization-specific design recommendations that enable the consortium to move forward and that increase the level of detail for the requirements. Some of these recommendations, in terms of simplified visual designs and integration of complementary informational sources (e.g. score follower or orchestral layout with the addition of structural information, colorimetric studies), are currently been taking into consideration.

The focus groups have highlighted the importance of musical expertise. The need for specific types of information depends on the level of musical expertise. However, the results have also demonstrated that the relationship is less clear-cut than expected. For example, score information not only appeals to conservatoire students that have had formal training, but also to laymen that cannot read scores. The complexity of this relationship has motivated us to set up a large-scale experiment in which we further explore the relationship between musical expertise and preferences for visualisations.

6 REFERENCES

Boehner, K., Vertesi, J., Sengers, P., and Dourish, P. (2007), “How HCI interprets the probes”, *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, San Jose, CA: ACM, 1077-1086.

Braun, V., Clarke, V. (2006), “Using thematic analysis in psychology”, *Qualitative Research in Psychology*, **3** (2), 77-101.

CIE 015 (2004), *Colorimetry*, Commission Internationale de L’éclairage, 3rd. ed.

Fereday, J. and Muir-Cochrane, E. (2006), “Demonstrating rigor using thematic analysis: a hybrid approach of inductive and deductive coding and theme development”, *International Journal of Qualitative Methods*, **5** (1), 1-11.

Hassenzahl, M. (2004), “The interplay of beauty, goodness, and usability in interactive products”, *Journal of Human-Computer Interaction*, **19** (4), 319-349.

Krumhansl, C. L. (1990), *Cognitive foundations of musical pitch*, New York: Oxford University Press.

Martorell, A., Gómez, E. (2015), “Contextual set-class analysis”, in D. Meredith (ed.), *Computational Music Analysis*, Berlin:Springer (in press).

Martorell, A. (2013), *Modelling tonal context dynamics by temporal multi-scale analysis*, Unpublished PhD diss., Universitat Pompeu Fabra.

7 LIST OF AUTHORS

- Agustín Martorell (UPF)
- Mark Melenhorst (TUD)
- Emilia Gómez (UPF)
- Oscar Mayor (UPF)
- Gerhard Widmer (OFAI)

8 APPENDIX

8.1 Video clips (Evaluation)

The links to the video clips used in the focus group studies follow:

- Score Follower: `https://www.youtube.com/watch?v=v0h6BRXxrEE`
- Reduced Piano Roll: `https://www.youtube.com/watch?v=A1-OKD8BXnU`
- Orchestral Layout: `https://www.youtube.com/watch?v=4IZlUUyQvjA`
- Structure: `https://www.youtube.com/watch?v=VU_ObRIHyD8`
- Instrumentation: `https://www.youtube.com/watch?v=rNSSLB_4RwI`

8.2 Questionnaires (Evaluation)

Background survey Nr _____

Gender
☐ Male ☐ Female

Age _____

How often do you attend classical concerts? Pick the answer that comes closest to your frequency.
☐ Multiple times a month
☐ Once a month
☐ Once every quarter
☐ Once every year

Please describe your taste for music in keywords below.

Which of the following statements is most applicable to you?
☐ 1 I'm usually the last to adopt new technology
☐ 2 I'm usually among the last to adopt new technology
☐ 3 I'm usually in the middle when it comes to adopting new technology
☐ 4 I'm usually among the first to adopt new technology
☐ 5 I'm usaully the first one to adopt new technology

Please mark the technology you make use of in the list below.
☐ iPad of een andere tablet
☐ Smartphone (iPhone, Samsung, etc.)
☐ Interactieve televisie (bijv. uw tv pauzeren of films huren)

Which of the following internet applications do you make use of once a week or more?
☐ YouTube
☐ Facebook
☐ Twitter
☐ Spotify

May we contact you for future studies?

☐ Yes My e-mail address is: _____
☐ No

Figure 15: Background survey questionnaire.

Idea no. _____ **Participant no.** _____

Can you briefly describe what you have seen in the video?

Imagine that each video would be developed into a working application. One can make an application that can be used before, during, and/or after the concert. When would you prefer to use this application? And for what purpose? (multiple answers allowed)

When?	*For what purpose?*
☐ Before the concert	_____
☐ During the concert	_____
☐ After the concert	_____

With what device would you like to use this application?

☐ With a desktop and a browser
☐ With a tablet
☐ With a smartphone
☐ With a smartwatch
☐ Other: _____

Please indictate for each statement below the exent you agree with the statement. Imagine again that the idea you have seen in the video has been developed into a working application. And that the application has been developed for the moment and the device of your preference.

	Fully disagree	Disagree	Neutral	Agree	Fully agree
This application helps me to discover new things in the music.	1	2	3	4	5
This application helps me to learn how the music is structured.	1	2	3	4	5
This application helps me to learn what instruments are playing.	1	2	3	4	5
This application helps me to recognize themes.	1	2	3	4	5
This application has little added value for me.	1	2	3	4	5
I don't understand what is shown in this video.	1	2	3	4	5
If I would have to use this application during the concert, it would distract me too much.	1	2	3	4	5
If I would have to use this application during the concert, I would expect to distract others.	1	2	3	4	5
This application motivates me to keep listening, because I can see what is coming up.	1	2	3	4	5
This application makes listening to this music more attractive.	1	2	3	4	5
I expect it to be easy to learn to use this application.	1	2	3	4	5

Figure 16: User perception questionnaire.