

Visualizing the Semantic Structure in Classical Music Works

Wing-Yi Chan, Huamin Qu, *Member, IEEE*, and Wai-Ho Mak

Abstract—A major obstacle in the appreciation of classical music is that extensive training is required to understand musical structure and compositional techniques toward comprehending the thoughts behind the musical work. In this paper, we propose an innovative visualization solution to reveal the semantic structure in classical orchestral works such that users can gain insights into musical structure and appreciate the beauty of music. We formulate the semantic structure into macrolevel layer interactions, microlevel theme variations, and macro-micro relationships between themes and layers to abstract the complicated construction of a musical composition. The visualization has been applied with success in understanding some classical music works as supported by highly promising user study results with the general audience and very positive feedback from music students and experts, demonstrating its effectiveness in conveying the sophistication and beauty of classical music to novice users with informative and intuitive displays.

Index Terms—Classical music, macro-micro relationship, glyph, typography, threads, weaving, information visualization.

1 INTRODUCTION

WHILE pop music is intended to be friendly to the general public, difficulties arise when we listen to classical music. The actual structure and form of the music are usually only understood by music lovers, who have received extensive training in music theory and history. This high learning curve, thus, makes classical music seemingly less approachable. In this paper, we demonstrate how information visualization techniques can be novelly applied to visualize the semantic structure of music such that users are immediately impressed by the sophistication and beauty of classical music with the vivid visual displays, and then able to visually gain insights into musical structure.

Realizing musical structure is an important step in comprehending music, but it is often difficult for nonexperts due to its abstract and temporal nature. While one can readily identify the concrete visual features in a painting, it is generally challenging for unskilled ears to recognize complicated musical elements from multilayered music. Some believe that the lack of visual equivalence is what makes music different from other art forms, and is also why music appears to be more abstract. Together with its dynamic feature, analysis and comparisons become more laborious.

The most common way to learn music is to study musical scores and technical essays (Fig. 1). Reading the score is, however, extremely demanding. Beginners have to spend considerable time learning the basic notations and technical terms in music theory before mastering a score for in-depth understanding. Another alternative is to read pages of technical essays, but it is certainly not a very

fascinating experience for the general audience as the essays are often very abstract and full of jargon.

To aid the general public in understanding the sophisticated structure in classical music, visualization is one of the promising approaches, due to our strong visual cognition ability. Although numerous efforts have been devoted to the visualization of sonic features, few have attempted to visualize the semantic structure of music instead of the low-level physical properties for perceptual understanding. Moreover, the important layer interactions and theme variations are barely addressed.

The beauty of classical music comes from its elegantly complicated yet beautifully arranged structure. The lowest semantic level is constructed by themes, which are recurring main melodies and principal musical ideas. Composers achieve balance in time by arranging and setting in order, the basic framework which is then clothed with interesting details by applying a variety of musical materials, resulting in different variations [1]. Apart from this horizontal foundation across time, a composition can be viewed from the layer perspective. Different layers in the music collaborate, accompany, blend, and contrast with one another, giving the rich vertical texture and depth of music. In brief, the microrelationship between themes contributes to the breadth of the music, while the macrorelationship among layers gives the depth of music. Furthermore, these two equally important relationships are not independent. A layer contains information about themes and themes are played by various layers; this interaction results in an extra macro-micro relationship between themes and layers, ultimately yielding a piece of sophisticated music.

We pioneer several intuitive design principles and effective visualization prototypes to convey the semantic structure in classical music, which novice users without musical background can gain insights into musical structure rapidly. A native color scheme is proposed to depict orchestral instruments for macrolevel layer information. The microlevel themes are represented by configurable glyphs, with theme variations encoded by a glyph grammar based on typographical concepts for amplified cognition.

- The authors are with the Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong. E-mail: {winnie.chan, null.mak}@family.ust.hk, huamin@cse.ust.hk.

Manuscript received 7 Nov. 2008; revised 7 Feb. 2009; accepted 28 May 2009; published online 4 June 2009.

Recommended for acceptance by P. Rheingans.

For information on obtaining reprints of this article, please send e-mail to: tcvg@computer.org, and reference IEEECS Log Number TVCG-2008-11-0182. Digital Object Identifier no. 10.1109/TVCG.2009.63.




Fig. 1. Traditional ways of learning music with (left) a musical score and (right) a technical essay [2].

With these building blocks, the layer braid and theme fabric prototypes are implemented to visualize the layer interactions and theme repetitions across layers, respectively. By introducing an analogy to plaiting and weaving in textile art, the general audience can rapidly gain insights into complex musical structure with familiar and impressive visualization. They can rely on vivid visual displays, instead of dull technical essays, toward understanding semantic structure, and thus, appreciating the sophistication and beauty of music for a more enjoyable listening experience. Music teachers, music students, and musicians also benefit as the visualization reveals musical structure in an intuitive way for teaching and learning different compositional and musical styles, as well as providing interesting visual clues for musical features that may be otherwise overlooked without the visual aid.

This paper is organized as follows: After introducing the related work in Section 2, we explain musical structure data in Section 3. Section 4 presents the design principles for encoding the important layer and theme components from which the braid and fabric prototypes in Section 5 are constructed to reveal the sophisticated relationships in music. Section 6 outlines the user interactions provided by our system for visual exploration, followed by Section 7, which summarizes the implementation details. Next, we discuss some case studies in Section 8 to illustrate the merit of our techniques. Section 9 describes the user study conducted with the general audience and the user survey feedback received from music students and experts. We conclude the paper in Section 10.

2 RELATED WORK

The visualization of the semantic relationships in musical structure has not been explicitly addressed in the literature. Typical music visualization approaches are highly mathematical and physical, while we emphasize the human interpretation of musical structure. Among the few relevant techniques, Arc Diagrams [3] are the first attempt to visualize repetitions in music using information visualization. The beauty of music is aesthetically revealed by connecting repeated units on the same track with an arc. We adapt the arc layout for the collapsed theme fabric view with extra visual channels for theme repetitions and variations across tracks. Isochords [4] visualize chord structures with animation, which is not the focus of semantic structure. Other works such as ImproViz [5] and Graphic Scores [6] only provide brief guidance for visualizing specific compositions.

Other related works [7], [8], [9], [10] aim to understand music performances in cognitive terms described by melody, rhythm, and phrasing. They examine the expressiveness of music with some basic visualization techniques. Other approaches [11], [12], [13], [14] use 3D views to visualize musical components. However, since we emphasize the insights into musical structure from which users can appreciate the beauty of music, 3D views are not included due to potential occlusion that would hinder music comprehension.

Several commercial products have also been developed to visualize the structure of music. TimeSketch [15] facilitates music teaching and analysis in music schools with simple half-disc shapes. Music Animation Machine [16] contains various effective graphical modules to visualize physical structural units in an MIDI file. Hyper-score [17], [18] allows users to sketch and compose music with freehand drawing. Color, shape, and texture are mapped to high-level musical features like timbre, melodic contour, and harmonic tension.

Since semantic structure is mainly qualitative, general quantitative techniques [19], [20], [21] cannot be deployed directly. Instead, new approaches are needed to address the special qualitative features.

3 THE MUSICAL STRUCTURE DATA

The major musical structural elements discussed in this paper are layers and themes, as defined below:

Layer. A composition essentially comprises various layers, which are musically known as parts or voices. A layer usually refers to an instrumental part, such as the flute, the first violins, or the cellos. It is determined by musical significance rather than physical section. For example, a group of cellos is equally important to a flute in terms of musical function and each is regarded as one layer, whereas the first and second violins have different functional roles and are, thus, split into two layers. The audience may not be aware of how the layers are actually played, though they would have a general idea that groups of layers sound distinct from one another from the macroperspective. The term “layer” is used instead of “track” to provide a generic view of musical parts. For instance, the sound effect “barking dogs” could be a layer in some contemporary music.

Layer roles. A layer may have certain functional roles in the harmony of all layers at a time, which contributes to the overall sounding effect of the musical texture. A layer possibly plays different roles throughout the music and some layers may belong to the same functional group. For example, violins are initially *dominant* and may later have a *contrast* role playing countermelodies, while the flute and the oboe still share the *accompany* role supporting other layers. A layer may also have multiple roles like *accompany* and *contrast* at the same time, when it accompanies with contrasting materials. Sometimes a layer may be playing transitional materials that do not have any specific role. Overall, layers exhibit some macrorelationships among themselves by having different functional roles over time, regardless of what melodies they play at the microlevel.

Theme. Musical themes, also known as motifs and subjects, are main melodies and musical ideas persisting in the music. Composers usually construct the entire piece

with only a few themes since too many ideas would make the music appear to be aimless and lacking in form. Themes are repeated to bring unity to the music and contrasted to introduce variety and interest [1].

Theme variations. Each theme occurrence is usually varied by melodic, rhythmic, instrumental, and transformative means, thus, forming microrelationships. Also, a theme is often played by different layers when it reoccurs, with some solo layers and perhaps other layers as accompaniment and reinforcement, contributing to the macro-micro relationship between layers and themes.

The uniqueness of musical structure and relationships make visualization a challenging task. Layers form the macrocontexts, where they interrelate with different roles. At the microlevel, themes are repeated with different variations along the time line. Furthermore, the macrolevel layer roles and the microlevel theme variations may be possibly related, resulting in a complicated macro-micro relationships. For example, it would be interesting to know the functional role of a particular theme variation. These coordinated relationships with different levels of viewpoints lead to a challenging visualization problem.

3.1 Data Preprocessing

Musical structure data are retrieved manually from descriptive essays, since presently there is no existing algorithm to parse the high-level structural information automatically from MIDI files or raw sound data. This is feasible since there are common, objective agreements for musical structure; the difference between essays is usually in their comprehensiveness but not perspective. It also allows educators to customize the input for teaching purposes, as there are strong human elements in musical structure. Nevertheless, front-end preprocessing can be automated if such algorithms become available, but it is beyond the scope of this paper and would not affect how the back-end visualization is generated.

All the examples in this paper are based on the listening guidelines from an introductory textbook [2]. Without any computer assistance, it took an amateur with high school music background two days to process the layer and theme information of the first movement of Beethoven's Symphony No. 5, which include the beginning and ending beat of an interval, the instruments involved, the layer roles, and the theme variations. The most demanding part is, surprisingly, not making musical judgments but digitizing the data from the score. It is, therefore, estimated that, with a simple graphical interface that allows users to mark on a digital score as they always do on sheet music, preprocessing time can dramatically be reduced to at most a few hours.

4 DESIGN PRINCIPLES

Since the score is already a special visualization of music with demanding notations, we aim to propose an intuitive design, which is not just another functional but overwhelming new language for penetrating musical structure. Challenges lie in how the design can strike a balance between aesthetics and effectiveness. In general, there are four design principles for the visualization of the semantic structure in music as follows:

1. **Consistent:** The encoding scheme should be consistent in all visualization prototypes such that users

would not be confused or misled. Thus, the basic visual channels should carefully be assigned.

2. **Intuitive:** Novice users should be able to comprehend the visual displays rapidly without much training. This could be accomplished by introducing familiar metaphors in the visualization.
3. **Effective:** The underlying musical structure should be shown effectively for novice users to gain insights into the structural features.
4. **Aesthetic:** Being a visual representation of beautiful and sophisticated musical compositions, the display should be of high aesthetic to reflect both the artistry and complexity of music.

The design starts with layers and themes, which are the most important and persistent structural components in the data. Apart from being intuitive and efficient, their encoding schemes should also be highly general and flexible as building blocks of diverse visualization prototypes. For instance, while an instrument icon can naturally represent a layer, it considerably limits the design of visualization tools as an icon needs to take a block of pixels to become discernible, and does not fuse well with other visual channels such as color and shape. An intuitive color scheme is, therefore, used to represent layers in our visualization, which incorporates well with different visual designs. For themes, instead of listing the variations with filled boxes and plain texts, the visual representation should convey the semantic meanings of the variations effectively so as to give a feel of how a theme varies. It should also allow users to easily spot which variations a theme applies and which themes apply the same variations. Therefore, a specific pictorial glyph grammar with typographical concepts is proposed to encode the important theme entity in our framework such that each variation can intuitively be mapped to an iconographical feature of a configurable theme glyph.

4.1 Color Scheme for Layers

The correlation between music and color has been of mutual interest to the researchers from various fields. Despite the many works on music and color, none can directly be applied to the visualization of musical structure data as the semantic context has not been addressed.

Tone color. In music, the color of a tone, also known as *timbre*, refers to the quality of a sound. Even if two instruments play the same musical note at the same amplitude for the same length of time, one can still easily distinguish them since different instruments have unique timbres. Conceptually, tone color is associated with sensation such as "warm" and "mellow" rather than optical colors. Thus, there is no explicit mapping for tone color in the optical spectrum. Also, an instrument can have different tone colors depending on performance styles, making it impractical to represent structural data.

Psychology. In music psychology and psychoacoustics, researchers are particularly interested in how colors are related to music in terms of emotions invoked, such as joy and sadness. They study the human perception of sound, which is again not relevant to the musical structure.

Physics. Sound pitch can be encoded by a physical mapping between sound and color spectrum. For instance, the color of noise is loosely determined by mapping sound frequencies to equivalent light-wave frequencies. Similar mappings exist for people suffering from a form of

Woodwind:	flute	oboe	clarinet	bassoon
Brass:	horn	trumpet	trombone	tuba
Percussion:	triangle	snare drum	timpani	bass drum
Strings:	violin	viola	cello	double bass

Fig. 2. Color scheme for layers. Instrument families are represented by their physical hues, except that percussions are assigned an outlier color. Lightness indicates the pitch range of an instrument for intrafamily contrast.

Chromesthesia called color hearing through which they see consistent colors when they hear the same musical notes. This kind of mapping is rational but lacks perceptual justification in visualization.

For effective visualization, the use of color should be coherent and unambiguous. We, thus, assign colors to layers as traditional orchestral instruments can be divided into four families: strings, woodwinds, brasses, and percussions. This categorical nature is ideally represented by colors of unique hues, with lightness distinguishing instruments in the same family [22]; the lower the pitch range, the darker the color. Although we can randomly assign the hues, it is less desirable since confusing emotions could be invoked by misleading color assignments. For example, if strings are represented in blue, listeners might expect some sad violin sounds even for a cheerful violin solo. Instead, the physical color of the instrument families is used, as it is objective and intuitive that users can easily recognize the instrument without memorizing the hues. It follows that strings are in reddish brown, woodwinds are silver, and brasses are golden (Fig. 2). Percussion instruments, however, do not share a common color like the other families. Hence, some degree of arbitrary assignment like mapping to an outlier purple color is feasible because they are defined as those not belonging to the other three families. Similarly, nonorchestral layers can be assigned some distinct colors based on users' preferences.

4.2 Theme Glyph

ThemeRiver [24] proposed a river metaphor to depict thematic variations in document collections, but it is less intuitive in representing musical themes since the semantic context of musical variations is lost. On the contrary, we use an iconographic glyph metaphor to represent each theme occurrence. However, existing glyph encoding schemes [7], [25], [26] and icon designs [27], [28] are not appropriate as the qualitative meanings of musical variations cannot be conveyed effectively with these quantitative methods.

In this paper, we propose a musical symbol glyph for themes to deepen the connection with music. Our musical symbol glyph is inspired by typography (Fig. 3), which is the study and process of typefaces, including how to select, arrange, size, and modify type glyphs, such that type is integrated with other design elements for a complete composition. Fig. 3 shows some basic concepts used in typography such as x-height, ascender, and descender. The musical symbol glyph consists of three basic components, namely a flag, a straight stem, and a filled-in note head (Fig. 4). Users may also define what indicative decorations are to be added around the glyph. Naturally, we can treat the flag in



Fig. 3. Basic type anatomy in typography [23].

the musical symbol glyph as the ascender in the typography, the note head as the descender, and the core stem attached with the note head as the x-height part.

We then establish a glyph grammar with typography so that the glyph can vary easily with sound design and cognitive foundation for an informative representation. Specifically, we propose the following glyph grammar for five basic types of theme variations (Fig. 4).

1. **Theme** (e.g., theme 1, theme 2): **Duplicate the flag of the glyph.** As the number of themes is usually limited, it is feasible to duplicate the flag to indicate which theme it is. A glyph with one flag indicates the first theme while a glyph with two flags indicates the second theme, and so on.
2. **Form** (e.g., form A, form B, and form C): **Change the note head shape.** A theme may be further classified into different forms. This categorical attribute is encoded by changing the shape of the filled-in note head part such that the overall glyph appearance is still preserved.
3. **Transformation** (e.g., augmented, diminished, and inverted): **Transform the glyph/changes the note head size.** A theme may structurally be transformed into other shapes. For instance, a theme is augmented by enriching it with extra musical notes and patterns, instead of simply elongating it. We can change the note head size to encode augmented and diminished transformations and invert the glyph to indicate an inverted theme transformation.
4. **Duration** (e.g., lengthened, shortened): **Adjust the length of the flag.** The duration of a theme may be altered due to the changes in the rhythmic pattern or partial appearance. It is denoted by adjusting the length of the flag to reflect this one-dimensional variation.
5. **Singleton** (e.g., echoed, transposed, in different key): **Add indicative decorations.** Some special variations are more like tags and they do not actually vary with




Fig. 4. The complete glyph variation board for the musical symbol glyph.

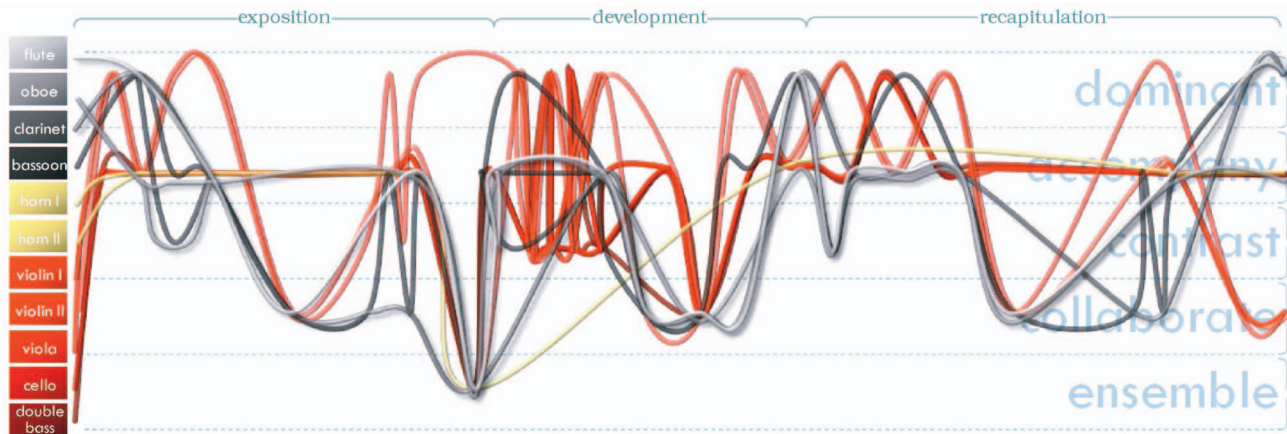


Fig. 5. Layer braid for the first movement of Mozart's Symphony No. 40. Layer threads are plaited according to their functional roles over time. The threads are blended ($\alpha = 0.7$) to better reveal the braid pattern.

the theme. These singleton variations are represented by indicative decorations that can cope with more information, such as the relationship with the new key or the degree of transpositions.

In our design, only mutually exclusive variations, like lengthen and shorten, are assigned to the same visual channel such that all possible combinations of theme variations are encoded effectively. As a result, various music note glyphs can be generated with the proposed grammar to represent different sets of variations.

5 VISUALIZATION PROTOTYPES

With the careful encoding schemes for layers and themes, many visualization prototypes can be built with these two basic structural units to reveal the macrorelationships among layers, the microrelationships of theme occurrences, and the macro-micro relationships between layer roles and theme variations. An ideal prototype should be intuitive and effective in revealing musical structure, as well as being able to arouse aesthetic sensation to reflect the beauty and sophistication of music. We propose two modules based on a textile art metaphor, namely the layer braid and theme fabric, which satisfy the above criteria. We observe that the vertical structure formed by layers and the horizontal foundation contributed by themes share some common features with weaving and knitting. A piece of delicate cloth is generated by interlacing and intertwining threads along both the x and y -axis, while weaving patterns are still clearly perceived for examining textile structure in this sophisticated art piece. Moreover, this visualization metaphor enables novice users to comprehend the displays easily without being cognitively overloaded by complex encoding schemes and unfamiliar visual layouts.

5.1 Layer Braid

Inspired by plaiting, layer braid uses a braid analogy to visualize how different threads, representing different layers by the color scheme discussed in Section 4.1, interact and correlate with each other over the horizontal time line (Fig. 5). Initially, the threads are laid down one by one following the musical score order of the woodwind, brass, percussion, and string sections. Next, the interweaving position is determined by its layer role over time. For

example, if cellos and double basses both stay as *collaborate* layers for a period of time, we will have two control points at the beginning and ending point accordingly in the *collaborate* span (Fig. 6). When a layer has more than one functional role at the same time, such as clarinet in this example, the thread is spread between two streams at the previous control point of *accompany*, which plait at the respective locations in the *dominant* and *contrast* spans, and rejoin at the next control point. The intertwining process then begins by interpolating the threads at these interweaving positions, generating different braid patterns for each unique macrostructure. Since the x -position corresponds to the time and the y -position refers to the current layer roles of the thread, layers belonging to the same functional group, like violins and violas in Fig. 6, would be plaited together to create informative braids. The order of roles can further impose some semantics of the layer relationships. *Dominant* and *ensemble* can be placed aside, encapsulating the supportive *accompany*, *contrast*, and *collaborate* roles in between for better visual impact.

More interactive techniques can also be introduced to the basic braid layout. Segments not having a specific role can be darkened upon user request to eliminate confusion caused by threads crossing some roles they do not belong to during interpolation. Users may specify the opacity value of each thread to reduce visual clutter and highlight some specific thread (Fig. 7). Various focus+context approaches

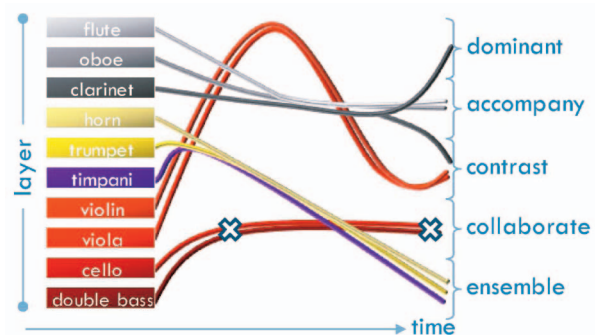


Fig. 6. Layer braid prototype. Layer threads intertwine to show how layers interact at the macrolevel. Thread clusters of various functional groups are seen over time. Control points for interpolating the cello thread are marked.




Fig. 7. The string threads are highlighted by adjusting the thread opacity values. The visual clutter in Fig. 5 can be reduced.

can also be deployed to the long braids. For example, Fig. 8 shows the same visualization as in Figs. 5 and 7 using the bifocal display [29]. A smaller version of the same display allows us to observe the overall braid pattern without knowing how the threads are actually plaited for a quick overview. Most importantly, the dynamic nature of music can be visualized by animating the braids. Simple highlights of the threads can indicate the current playing intervals or the braids may appear in accordance to the music to simulate a flow of musical threads illustrating how the layers interact in real time.

5.2 Theme Fabric

While the layer braid visualizes the macrointeractions of layers with intervening threads, theme fabric aims to reveal the microrelationships between theme occurrences across layers. A theme not only varies, but is also played by different layers when it reappears. As a result, these microrelationships are found along both the x -axis over time and the y -axis for different layers.

The theme fabric is first constructed by overlaying the theme glyphs, which are encoded by the grammar explained in Section 4.2, on a conventional musical score layout according to which layers play the theme and when the theme is played (Fig. 9). If a theme occurrence is played by more than one layer, the glyph is duplicated at every layer that plays the theme like the two-flagged glyph in this example. To avoid glyphs from overlapping, the space between the glyphs is scaled nonlinearly such that a minimum glyph gap is guaranteed. This is applicable since the relative gap proportion is more useful than the actual gap distance in understanding thematic structure when we listen to the music. Typography also suggests that letterspacing would affect the legibility of the text because it would be easier to distinguish one letter from the next with sufficient space. Therefore, slight adjustment to the spacing is made to ensure that every glyph is equally visible while retaining a relative gap proportion.

Although themes having the same variations are represented by the same glyph, the repetitions are still not obvious in a large graph without strong visual clues to




Fig. 8. Deploying some focus+context technique like bifocal display [29] to the long layer braid for better navigation.




Fig. 9. Theme fabric prototype. The glyphs are first placed in a 2D plane according to when and where the theme appears at the microlevel. The glyphs of the same themes can then be connected aesthetically in various weaving styles.

connect the glyphs together. Originating from the textile metaphor, glyphs can be treated as stitches connected by threads in various styles, weaving a sheet of theme fabric that effectively reveals how the glyphs are organized on a 2D plane. We include three types of stitch in our framework, namely plain style, bundled style, and collapsed style, to provide decreasing levels of detail of the microrelationships of theme variations.

5.2.1 Plain Style

The most straightforward way is to connect repeating glyphs in an orthogonal direction only. That is, identical glyphs of a layer are chained horizontally, and those played by different layers within a short time interval are linked vertically. This basic layout is an analogy to plain weaving in textile art, where two sets of threads, known as warp and weft, are interlaced to create a fabric (Fig. 10a). The horizontal weft threads are woven under and over, back and forth through the vertical parallel warp threads. Consequently, glyphs across layers are connected with straight yarns in our plain weaving metaphor, and those appearing at the same layer are woven by curvy yarns interwoven with the vertical lines (Fig. 11). Moreover, each set of identical glyphs is stitched with threads of unique lightness for a clearer weaving pattern. The plain theme fabric serves as the most fundamental and expanded view of the thematic structure at the microlevel as every detail on theme variations and repetitions across layers are literally visualized in a plain weaving style.

5.2.2 Bundled Style

As a theme is usually played by various layers at the same time, the glyphs are duplicated at these layers and connected by replicated horizontal threads. The plain style may, therefore, be enhanced by bundling all horizontal weft threads originating from glyphs initially connected by vertical warp threads, by adding a control point in between two repetitions for interpolating and, thus, bundling the




Fig. 10. Theme fabric styles in textile art. (a) Plain weaving. (b) Butterfly stitch in knitting that bundles yarns. (c) Open strips produced by drop-stitch technique.




Fig. 11. Theme fabric with plain style weaving.

thread at the midpoint (Fig. 12). A less cluttered view is created while maintaining the betweenness centrality of the theme glyphs (Fig. 13). For example, the first column of glyphs in Fig. 12 is obviously connected in a same bundle although vertical threads are absent. While thread bundles make groups across other glyphs clearer, they are not applied on successive glyphs that are not repeated elsewhere, such as the four double-lined groups at the end of this example, since they are already shown effectively without bundles. These special groups can be further outlined by a double dashed line pattern. Alpha blending of threads is also included to tackle the overlapping of thread bundles and knots so that orange groups in Fig. 12 remain visible behind the blue ones.

The bundled style is, in fact, derived from the butterfly stitch in knitting (Fig. 10b). It is an abstract form of the plain style since the weft threads are bundled and the warp threads are removed. Although repetitions at the same layer are less obvious, the bundled view effectively reveals the




Fig. 13. Generating bundled style from plain style theme fabric. Vertical threads (warp) are removed and horizontal threads (weft) are bundled, retaining the betweenness centrality of the theme glyphs with a concise layout.

more high-level thematic structure by visually clustering all identical glyphs across layers. The actual musical structure is further highlighted by thread bundles and knots, which figuratively exhibit the sophistication of theme and variations at the microlevel.

5.2.3 Collapsed Style

Conceptually, the collapsed style is very similar to drop-stitch knitting that can produce any number of open stripes by dropping stitches intentionally (Fig. 10c). Apart from clustering similar threads as in the bundled view, the numerous redundant glyphs can be further eliminated by collapsing identical glyphs at different layers into one single glyph only, resulting in some Arc Diagrams [3] layout (Fig. 14). That is, a theme occurrence is represented by one glyph only, despite the number of layers it is played by; and repeated glyphs are, hence, connected by one arc only. As the original design cannot accommodate multiple tracks in one display, we propose to encode the layer information by rainbow arcs divided into color bands based on the layer color scheme. A band is colored if the corresponding layer contributes to the current theme occurrence. The arc is split into two quarters to show the layer details of the former and latter occurrences, respectively. In Fig. 15, the augmented




Fig. 12. Theme fabric in bundled style for the first movement of Mozart's Symphony No. 40. Each theme occurrence is represented by a musical symbol glyph encoding its variation. Identical glyphs are connected by bundled threads.



Fig. 14. Collapsing a theme fabric by removing duplicated glyphs. Layer information is now encoded by rainbow arc.

form *A* at the start of development section is first played by silver woodwinds and golden brasses, and later repeated by violins only, resulting in an arc split accordingly. Also, in contrast to the original Arc Diagrams where the arc height equals the duration between two repetitions, this visual channel is further exploited by assigning a specific arc height to every unique glyph such that users can locate all repetitions of the same theme glyph efficiently.

The arc height can randomly be assigned for solely showing the thematic structure or seamlessly combined with the spatial order of layer roles in the layer braid to reveal the interesting macro-micro relationships between layers and themes. For example, the echoed form *A* at the beginning is often associated with a *collaborate* role. As a result, theme glyphs that are correlated with certain layer roles would have their arcs frequently appear in the roles' allocated partition. Similarly, layers strongly related to a certain functional role would have their respective colors often filled in the rainbows at a particular height. If a layer has multiple roles at the same time, duplicated arcs at different heights may be drawn.

In summary, the three theme fabric styles provide different levels of detail on the microrelationships by connecting the theme glyphs with reference to several textile art techniques, to aesthetically express the sophisticated nature of theme and variations in music. The plain style contains the expanded structure, while the bundled view highlights groups of identical glyphs by thread bundles. The collapsed style is the most abstract view of

the thematic structure with compressed layer information. By mapping layer roles to arc heights, it can further reveal the unique macro-micro relationships on how themes and layers correlate in a compact display.

6 USER INTERACTIONS

Various user interactions are implemented to help users explore and navigate among different visualization prototypes [30]. The system takes an MIDI file and its processed layers and themes information in a predefined text format as input. The layer braid and theme fabric, on which users can further interact to reveal the sophisticated musical structure, are then automatically rendered. Besides typical user interactions such as selection, filtering, zooming, and panning, the following features are also provided.

- **Customization:** While the proposed color scheme is used as default, users may configure their own layer colors via a graphical interface.
- **Linking and brushing:** Users can interactively explore different modules with simple brushing. Parts selected in one layout will be brushed correspondingly in other modules. Moreover, the prototypes can be overlaid given their similar layouts. The layer braid can be fused with theme glyphs or embedded to the arc height stack representing layer roles in the collapsed theme fabric.
- **Abstraction:** The theme fabric is abstracted by rendering in different styles or by retaining only a sketchy outline of the glyphs. For smooth transition, fading is added when the theme fabric style is changed. Abstraction of the layer braid is performed at the data level, where short intervals not musically significant can be aggregated to eliminate abrupt zig-zag leaps during thread interpolation.
- **Focus+context:** Focus+Context is critical as users often wish to maintain the context of the overall



Fig. 15. Collapsed theme fabric for a segment of the first movement of Beethoven's Symphony No. 5. Macro-micro relationships are shown by mapping arc heights to layer roles. In this example, no theme occurrence is associated with the *dominant* or *accompany* layer roles. Pale reflection of the arcs is added for visual reinforcement.

structure while examining some interesting details in lengthy musical compositions. Bifocal display [29], [31] is currently employed since it is similar to browsing sheet music. The focused area will be shifted according to the music.

- **Details on demand:** Apart from showing the details of a selected visual entity, more information is also provided if users zoom in on the displays. For example, instrument icons, lists of variations, or the actual musical score can be shown on demand.
- **Synchronization with music:** The visualization is synchronized with music in real time. Users can set the playing position and range through a progress bar. Moreover, the animated layer braid may flow with the music concurrently; likewise the theme bundles can appear according to the music. For the collapsed theme fabric, not only the rainbow arcs emerge along with the music, but the current arcs may also be rippling to reinforce the thematic structure.

7 IMPLEMENTATION

The visualization system is developed with Java and the prefuse toolkit and installed on a ThinkPad T61p notebook with an Intel Core 2 Duo 2.20 GHz CPU, 2 GB memory, and an Nvidia Quadro FX 570 M graphics card. After the data are loaded into the memory, the layer braid and the plain theme fabric are rendered in two separate windows in real time. All curves in the system, namely the layer braids and the threads in theme fabric, are rendered as cardinal spline to ensure that the curves pass through the control points. The implementation of the collapsed theme fabric is based on the original arc diagrams, except that quarter-arcs are drawn instead at different offsets to generate split arcs of different color bands. All variations of the glyphs are generated by our system according to the theme attributes with the basic glyph defined in a vector format. For example, the musical note glyph is composed of three primitive shapes: an oval as the note head, a vertical line as the stem, and a curved line as the flag. Different variations of the glyph are then generated by applying affine transformations on individual parts. Specifically, fixed point scaling at the joint with other parts is required so that the glyph is varied properly.

8 CASE STUDIES

In this section, we demonstrate how the proposed visualization prototypes, the layer braid, and the theme fabric, allow the general audience to quickly gain insights into the semantic structure without going through the overwhelming technical essays. The analysis tends to be more general to avoid specific knowledge of classical music, but already provides many interesting observations on musical structure and compositional styles. Classical symphonies are examined as their clear-cut forms are relatively easier to understand. Nevertheless, our system is also capable of handling nonorchestral music.

The analysis is based on two famous symphonies, Mozart's Symphony No. 40 in G minor, K.550 and Beethoven's Symphony No. 5 in C minor, Op. 67. We focus

on the first movement since it is commonly regarded as the signature of a symphony. A normal performance of both examples takes around 10 minutes, which is fairly standard in the classical music. The first movement is often built up in *sonata form* including exposition, development, and recapitulation. The *exposition* and its repetition are found at the beginning to expose basic musical ideas; more complicated relationships are seen in the *development* section, and the music eventually returns to some familiar materials in *recapitulation*. It definitely requires some effort for untrained ears to recognize sonata form with plain text. With the proposed visualization, novice users can also understand the semantic structure easily.

8.1 Macrorelationship

The macrorelationship of layers for Mozart's Symphony No. 40 is first studied using Fig. 5. Remarkably, the development section exhibits highly complicated relationships with rapid exchanges of roles. Overall, ensembles are only heard toward the end of exposition. After modifying the thread opacity values and highlighting the string threads (Fig. 7), we observe that lower strings mostly play the accompaniment when the violins are dominant. The roles played by other instruments can be examined similarly.

The layer braid figuratively illustrates Mozart's talents in handling an orchestra at the apex of the mature Classical era, including the frequent interchanging relationship between instrument groups and the careful sharing of musical interests across layers [32]. This compact layer braid display intuitively portrays the rich musical texture with a familiar braid analogy to novice users.

8.2 Theme and Variations

The horizontal framework of repeating themes provides an alternative perspective for understanding the microstructure of melodies as shown by the theme fabric (Fig. 12). The overall textile patterns in bundled style give a lasting impression of thematic structure, with the thread bundles depicting how the themes are structured in the entire piece. The sonata form is usually framed by two subjects in contrasting styles. The bluish thread bundles indicate that the principal subject is mostly played by the strings and occasionally heard in echo among the clarinets, bassoons, and violins following the second subject stitched with orange yarns. Its normal form without variations clearly marks the beginning of exposition, repeated exposition, and recapitulation, suggesting how the music is restated in sonata form. Vivid variations that never repeat are seen in development and recapitulation, such as shortened, augmented, or in different keys. These orphan variations with double-lined threads imply how the composer built up the musical variety and tension toward the ending climax.

Overall, the strings are the primary theme players in this Mozart symphony. The woodwind section is less significant; in particular, the oboe contributes to only one theme occurrence. Although the score also calls for horns as seen from the layer braid, they are not involved in the thematic structure at all. It aligns with the historical development of the orchestra that the string section was the foundation of a typical Classical orchestra, and the brasses had only become

increasingly significant in the early 19th century under Beethoven's influence.

With the informative layer braid and theme fabric, abstract compositional styles become visible: while Mozart's melodies and forms are mostly graceful and elegant, the complexity of his music perhaps comes from his genius sensation of instruments that makes every simple micro-module an essential contribution to the macrotexture, producing these beautiful pieces of music.

8.3 Macro-Micro Confluence

The interesting macro-micro confluence of layers and themes is analyzed with the collapsed theme fabric (Fig. 15). A segment of the first movement of Beethoven's Symphony No. 5 is visualized, with arc heights indicating the respective layer roles of the theme. The display focuses on the famous "Fate knocking at the door" motif; the lyrical second theme is, thus, not connected with arcs. In addition, orphan theme variations not repeated in this segment are highlighted in yellow.

Surprisingly, the fate motif never acts as a typical principal melody that dominates the music as in other Classical symphonies; it mostly serves as a contrasting or echoing component. Briefly, there is no sole melodic line in this example, which is one of Beethoven's pivotal compositional styles. Unlike his predecessors Haydn and Mozart, Beethoven's motif is usually some forceful rhythmic patterns rather than conventional melodic figures. On the contrary, Mozart's music is characterized by his elegant melodic themes with a clean and balanced structure as revealed from the bundled theme fabric.

Detailed orchestration is compressed into splendid colors in this abstract view. The rainbow arc patterns imply that the woodwinds and brasses usually collaborate at the macrolevel by playing the augmented fate motif at the microlevel, whereas the inverted, shortened motif in the lower strings is used to contrast the music. Generally speaking, the functional roles in this Beethoven's example are more aligned with the motif rather than the high-level layer interactions, which follows that the boldness of Beethoven's music is attributed mainly to his extensive development of musical materials.

9 EVALUATION

Our system was evaluated by both the general audience and music experts. We performed a user study with the general audience to quantitatively evaluate the proposed visualization. A user survey was also conducted with a group of music students and experts to collect their qualitative feedback on the visualization framework.

9.1 User Study

A total of 20 college students without prior musical background participated in the user study, which aims to validate the effectiveness and efficiency of the proposed encoding schemes in conveying musical structure.

9.1.1 Tasks

After a five-minute introduction, the subjects were asked 12 questions regarding musical structure similar to the

Prototypes	Accuracy		Time (min)	
	Mean	SD	Mean	SD
Layer Braid	87%	12%	4.52	1.32
Theme Fabric (Bundled)	94%	8%	5.03	0.98
Theme Fabric (Collapsed)	89%	12%	4.93	1.47

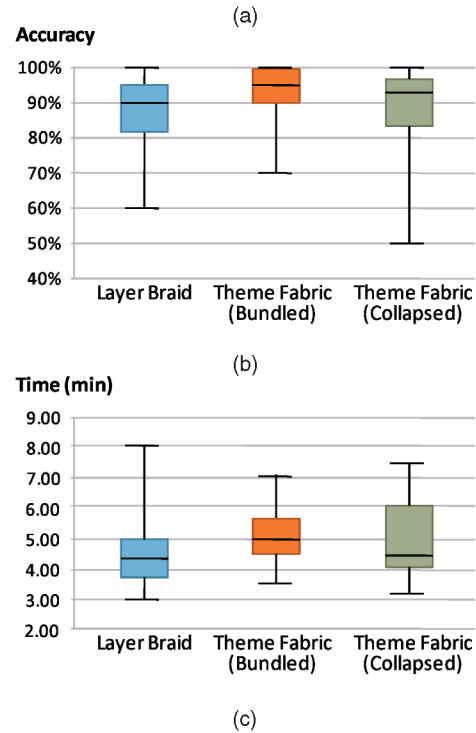


Fig. 16. User study result with the general audience to evaluate the effectiveness of the proposed visualization prototypes in revealing musical structure.

exercises from music textbooks. For the layer braid, they were questioned about the layer roles over time and the relationships between functional groups, while for the theme fabric, they were required to answer how, when, and where the theme appears. The macro-micro interactions between layer roles and theme variations were also tested with the collapsed theme fabric.

9.1.2 Results

The accuracy and response time are shown in Fig. 16. One-way analysis of variance (ANOVA) shows that the variation among the means of the three displays is not significantly greater than expected by chance with $p > 0.05$ for both accuracy and response time. Generally, the average response time for the bundled theme fabric is the longest since frequent user interaction was required for such an intensive display, but the noticeably high accuracy, with three quartiles above 90 percent, proves that the thematic structure is delivered effectively despite a slightly cluttered layout. Both the layer braid and the collapsed theme fabric are fairly efficient due to the simpler visual designs. The layer braid, however, shows a statistically significant difference with other modules in response time; the lowest quartile yields a remarkable range of three minutes. It is suggested that some participants found the macrolevel

layer role relationships highly abstract, and thus, took a longer time to familiarize themselves with the visual layout. Moreover, the accuracies of the layer braid and collapsed theme fabric are comparatively lower as some information is condensed and detailed analysis is demanded instead of simply spotting the visual patterns as in the bundled view. The statistically diverse accuracy of the collapsed style further indicates that it is more challenging for some users to understand the macro-micro interactions between layers and themes. Nevertheless, the participants were able to make better judgment when further training was provided.

9.1.3 Comparisons and Discussions

To further validate our visualization, we conducted an experiment with another 10 college students without prior musical background. They were provided with the standard learning tools, including the full score of the music, the supplementary listening guide, and the technical essays. After a similar five-minute introduction on the provided materials, they were then tested on the same set of questions. Unfortunately, the students encountered tremendous problems and left most questions unanswered. The participants could generally answer questions discussed explicitly in the essays. However, most questions required them to identify components from the scores, which was proven to be challenging even they could listen to the music at the same time. For example, although the score segment of the original theme was given in the listening guide, the participants had difficulties in classifying what variations were applied in other occurrences by listening to the music with the score. They were, thus, unable to answer such kind of questions. A music educator suggested that the lack of musical background was the main reason that the participants could not finish some questions with the materials provided, as reading the score, for instance, definitely requires some degree of musical background. Therefore, in contrast to learning music traditionally with scores and essays, our visualization prototypes are effective and efficient for novice users to gain insights into musical structure promptly without prior extensive training and guidance. While musical training is always required for analyzing the music by ears, the visual modules provide a friendlier pictorial guidance for music learning and appreciation by symbolizing the sophistication and beauty of the music, thus, making the music listening experience more enjoyable and informative.

9.2 User Survey

In addition to the user study with the general audiences, we also conducted an in-depth survey with a group of 15 domain experts consisting of 9 music students, 5 music educators, and one musician for their expert feedback. They are all graduates in music (i.e., with ABRSM Grade 8 music theory or equivalent) and specialized in at least one instrument (i.e., with DipABRSM or equivalent).

9.2.1 Results

The survey results are summarized in Fig. 17. The participants show a very positive attitude to the visualization prototypes. All modules are equally valuable in terms of usefulness, effectiveness, and aesthetics. The theme fabric in bundled style is the most useful display as users are more concerned about thematic structure, whereas the collapsed




Fig. 17. Survey results with music experts on a 5-point scale about the usefulness, effectiveness, and aesthetics of each prototype in visualizing musical structure.

style derived from the Arc Diagrams layout is the most aesthetic one. In particular, they find the rainbow arcs with reflections strongly appealing.

9.2.2 Discussions

Comments from the music experts were also very encouraging. Overall, they appreciate the familiar analogy to textile art, which intuitively expresses the sophistication and aesthetics of a composition with vivid displays. They believe that braid and fabric are excellent metaphors of music since its physical structural patterns and perceptual artistic aspects are both revealed effectively.

Among all the prototypes, the participants appreciate the collapsed theme fabric most because it is aesthetic and less complicated than the bundled style, while retaining sufficient information on layers and themes at a glance. They agree that the layer braid is extremely useful in understanding how different instruments interact in a huge orchestra, as the layer role concept might be less definite to novice users without the visual aid. The participants having substantial musical background find the bundled theme fabric effective in presenting complicated microrelationships, but less experienced students consider it slightly complex and preferred the collapsed style for different levels of detail. This implies that the proposed prototypes can cooperate with or complement one another in suiting a variety of needs.

The music educators like the intuitive layer colors and theme glyphs very much, as they believe that reading pages of words would possibly diminish a student's interest in learning a composition. It is also easier for the students to follow the music with such a real-time system, which is especially beneficial for visual learners. Moreover, they appreciate that instrumentation, which has the multilayered property, is addressed in our framework. They find this layer information particularly important, which is, however, not captured by the TimeSketch program they currently use for learning and presentation. The music theory teacher specifically states that the displays are useful for teaching children and beginners how to analyze a piece of music work. They believe that the theme fabric is a helpful tool for learning themes and variations as a compositional technique. While the advantages may seem less obvious for performers, the high school teacher suggests that school orchestra members without much musical background can better understand their individual functions with respect to the entire orchestra in the music with the layer braid visualization. The collapsed theme fabric can also shed light on the overall structure when the players start to learn an unfamiliar work.

The music students also appreciate the layer feature because differentiating instruments in an orchestral work is difficult. By dissecting a piece into various layers, it can boost the enjoyment of orchestral works and facilitate a sensible presentation to music audience. They find that the visualization prototypes are especially useful in comparing compositions of distinctive characteristics from different musical periods, which help learners identify the relevant period and its notable features that are otherwise abstract without the visual aid. Moreover, they like the idea of representing a musical theme with a semantic glyph, as the structurally important thematic changes can be understood and analyzed more easily by projecting the linkages and variations of themes on an intuitive diagram like a mind-map. They find it interesting and exciting to see how portions of thematic information are connected in a graphical layout, which could be very different from the pure imagination stimulated by listening to the music.

While the musician holds the philosophy that music should speak for itself, he agrees that these visual displays truly assist music learning and analysis. Classical music can, thus, be publicized with such an impressive visual guidance that exhibits the sophistication and beauty of the music from which the general audience will find classical music more approachable and eventually be better amused by the music itself. He also looks forward to attaching the visualization to the concert programme so that audiences can better understand the performance and appreciate the genius ideas behind the music. By gaining insights into musical structure, listeners can advance from solely emotional appreciation to more conscious learning and deeper understanding, resulting in new enjoyment, excitement, and satisfaction.

One major concern, as stated by the educators, is that it could be problematic to visualize musical works without a clear form since the prototypes focus on revealing the structure, and also for handling exotic instrumentations not taking the conventional orchestra in contemporary pieces. Given the variety of musical styles, it is, in fact, impossible for a visualization framework to be capable of handling compositions of any style. Our design tries to strike a balance among simplicity, scalability, effectiveness, comprehensiveness, and depth. Specifically, the layer concept we introduced is not limited to instrumental layers but can refer to any recognizable track.

The participants also identified some possible improvements. While the proposed prototypes are well designed and comprehensive, the students are also interested in having more attributes such as the change in dynamics, tempo, and tone color in the resulting displays. This can be achieved by exploiting more visual channels like size and shading in the current prototypes or by developing new visual modules based on the flexible layer color scheme and theme glyph grammar. Moreover, the educators believe that adding some descriptive texts can guide students in discovering the aesthetic part of the music, which is currently not included due to copyright issues. Finally, the educators conclude that they are eager to apply the visualization in music schools.

10 CONCLUSION

In this paper, we proposed an innovative visualization solution to show the semantic structure in classical music works, including macrolevel layer relationships, microlevel theme variations, and macro-micro interactions between layers and themes. We first discussed several design principles for encoding the important structural units, layers, and themes, with an intuitive color scheme and a sound glyph grammar founded on typography, respectively, from which visualization prototypes can then be constructed flexibly. Instead of developing some general displays, we specifically identified the similarities between music and textile art in their structural and aesthetic aspects. The layer braid and theme fabric prototypes were initiated to visualize the semantic relationships in music with an analogy to common weaving and knitting techniques. Apart from revealing the structure effectively, the vivid displays also intuitively portray the sophistication and beauty of the music. While the prototypes essentially form a comprehensive framework, they are complementary to one another to provide different levels of detail. Several representative case studies were then presented in which complicated semantic relationships are efficiently depicted, providing visual clues to novice users about the abstract musical styles. Its merit was further reflected by the user study that proves the effectiveness and efficiency of the encoding schemes, and by the user survey, where highly positive feedback from music students and experts has been received.

In summary, we demonstrated the potential of applying information visualization to another discipline that dramatically enhances the understanding of a new data type, the musical structure data. The strong human factors present in its semantic relationships distinguish this work from other information visualization approaches that mainly focus on the quantitative aspects of the data. The qualitative and subjective elements in musical structure pose special challenges to the visualization task that not only the underlying structure should be conveyed efficiently, but the artistic sensation aroused from its sophistication should also be expressed intuitively. In our framework, this perceptual issue was addressed using familiar metaphors from other art fields, including typography in graphical design and weaving in textile art, to generate visual displays for cognitive pleasure. Well-established design principles can be easily adopted to create informative and aesthetic displays once their similarities are identified. It also indicates that we may further explore other art and design analogies for applications that involve considerable perceptual concerns, such as casual information visualization and semantic data visualization. The prototypes presented in this paper are, thus, not limited to revealing musical structure, but can be generalized to visualize other data types like news and storytelling entities.

There are several avenues for future work. First, more diverse compositions could be visualized to fully realize the potential of the system. Different types of music will be continuously added to the database. Second, a simple user interface allowing educators to directly mark on a digital score could dramatically reduce data input time, and thus, generalize the usage of the visualization. Finally, we seek the

possibility of applying the visualization as a teaching tool for some introductory music courses to reach a wider audience.

ACKNOWLEDGMENTS


The authors thank Yee-Wing Chan, Jonathan Soon, Andrew Horner, and David Rossiter for providing valuable feedback, and all the participants of the user studies. They are thankful to the anonymous reviewers for their constructive comments. This work is partially supported by HK RGC grants CERG 618705 and 618706.

REFERENCES

- [1] R. Bennett, *Cambridge Assignments in Music Series*. Cambridge Univ. Press, 1999-2001.
- [2] C. Hoffer, *Music Listening Today*, third ed. Schirmer Books, 2006.
- [3] M. Wattenberg, "Arc Diagrams: Visualizing Structure in Strings," *Proc. IEEE Symp. Information Visualization '02*, pp. 110-116, 2002.
- [4] T. Bergstrom, K. Karahalios, and J.C. Hart, "Isochords: Visualizing Structure in Music," *Proc. Int'l Conf. Graphics Interface*, pp. 297-304, 2007.
- [5] J. Snyder and M. Hearst, "ImprovViz: Visual Explorations of Jazz Improvisations," *Proc. Conf. Extended Abstracts on Human Factors in Computing Systems*, pp. 1805-1808, 2005.
- [6] B. Evans, "The Graphic Design of Musical Structure: Scores for Listeners," *Proc. Electroacoustic Music Studies Conf.*, 2005.
- [7] R. Hiraga, "Case Study: A Look of Performance Expression," *Proc. IEEE Conf. Visualization*, pp. 501-504, 2002.
- [8] R. Hiraga and N. Matsuda, "Graphical Expression of the Mood of Music," *Proc. Seventh IEEE Int'l Conf. Multimedia and Expo*, pp. 2035-2038, 2004.
- [9] R. Hiraga, R. Mizaki, and I. Fujishiro, "Performance Visualization: A New Challenge to Music through Visualization," *Proc. ACM Int'l Conf. Multimedia*, pp. 239-242, 2002.
- [10] R. Miyazaki, I. Fujishiro, and R. Hiraga, "Exploring MIDI Data Sets," *Proc. ACM SIGGRAPH '03 Sketches & Applications*, 2003.
- [11] R. Miyazaki and I. Fujishiro, "comp-i: 3D Visualization of MIDI Data Set," *Proc. IEEE Conf. Visualization*, 2002.
- [12] J. Ox, "Two Performances in the 21st. Century Virtual Color Organ: GridJam and Im Januar am Nil," *Proc. Seventh Int'l Conf. Virtual Systems and Multimedia*, p. 580, 2001.
- [13] S.M. Smith and G.N. Williams, "A Visualization of Music," *Proc. IEEE Conf. Visualization '97*, pp. 499-503, 1997.
- [14] L. Valbom and A. Marcos, "An Immersive Musical Instrument Prototype," *IEEE Computer Graphics and Applications*, vol. 27, no. 4, pp. 14-19, July/Aug. 2007.
- [15] "Timesketch," Electronic Courseware Systems, Inc., 2000.
- [16] S. Malinowski, *Music Animation Machine*, 2001.
- [17] M.M. Farbood, E. Pasztor, and K. Jennings, "Hyperscore: A Graphical Sketchpad for Novice Composers," *IEEE Computer Graphics and Applications*, vol. 24, no. 1, pp. 50-54, Jan./Feb. 2004.
- [18] Hyperscore, Harmony Line, 2006.
- [19] C. Chen, *Information Visualization: Beyond the Horizon*, second ed. Springer, 2004.
- [20] C. Ware, *Information Visualization: Perception for Design*, second ed. Morgan Kaufmann, 2004.
- [21] W.S. Cleveland, *The Elements of Graphing Data*, second ed. Hobart Press, 1994.
- [22] D. Borland and R.M. Taylor, "Rainbow Color Map (Still) Considered Harmful," *IEEE Computer Graphics and Applications*, vol. 27, no. 2, pp. 14-17, Mar./Apr. 2007.
- [23] "A Brief Introduction to Typography," Dept. of Art, Western Illinois Univ., 2002.
- [24] S. Havre, E. Hetzler, P. Whitney, and L. Nowell, "ThemeRiver: Visualizing Thematic Changes in Large Document Collections," *IEEE Trans. Visualization and Computer Graphics*, vol. 8, no. 1, pp. 9-20, Jan.-Mar. 2002.
- [25] H. Chernoff, "The Use of Faces to Represent Points in k-Dimensional Space Graphically," *J. Am. Statistical Assoc.*, vol. 68, no. 342, pp. 361-368, 1973.
- [26] G. Kindlmann, "Superquadric Tensor Glyphs," *Proc. Joint Eurographics-IEEE TCVG Symp. Visualization*, pp. 147-154, 2004.
- [27] J.P. Lewis, R. Rosenholtz, N. Fong, and U. Neumann, "VisualIDs: Automatic Distinctive Icons for Desktop Interfaces," *Proc. ACM SIGGRAPH '04*, pp. 416-423, 2004.
- [28] P. Kolhoff, J. Preub, and J. Loviscach, "Music Icons: Procedural Glyphs for Audio Files," *Proc. 19th Brazilian Symp. Computer Graphics and Image Processing*, pp. 289-296, 2006.
- [29] M.D. Apperley, I. Tzavaras, and R. Spence, "A Bifocal Display Technique for Data Presentation," *Proc. Conf. Eurographics*, pp. 27-43, 1982.
- [30] S.K. Card, J.D. Mackinlay, and B. Shneiderman, *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann, 1999.
- [31] J.D. Mackinlay, G.G. Robertson, and S.K. Card, "The Perspective Wall: Detail and Context Smoothly Integrated," *Proc. ACM Conf. Human Factors in Computing Systems*, pp. 173-179, 1991.
- [32] H. Abert and S. Spencer, *W.A. Mozart*, C. Eisen, ed. Yale Univ. Press, 2007.



Wing-Yi Chan received the BEng degree (highest honors) in computer science and the MPhil degree in computer science and engineering from the Hong Kong University of Science and Technology (HKUST) in 2007 and 2009, respectively. Her research interests include information visualization and visual analytics.



Huamin Qu received the BS degree in mathematics from Xi'an Jiaotong University, China, and the MS and PhD degrees in computer science from the Stony Brook University, New York. He is an assistant professor in the Department of Computer Science and Engineering at the Hong Kong University of Science and Technology (HKUST). His main research interests are in visualization and computer graphics. He has conducted a wide range of research on scientific visualization, information visualization, real-time graphics systems, virtual reality, and medical imaging. He is a member of the IEEE.



Wai-Ho Mak received the BEng degree (first-class honors) in computer science and the MPhil degree in computer science and engineering from the Hong Kong University of Science and Technology (HKUST) in 2007 and 2009, respectively. His research interests include scientific visualization and information visualization.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.