

Designing Visual Guides for Casual Listeners of Live Orchestral Music

Catherine Solis*

Fahimeh Rajabiyazdi†

Fanny Chevalier‡

University of Toronto

ABSTRACT

The experience of attending live orchestra performances is rich in cultural heritage and can be emotionally moving; however, for those unfamiliar with classical music, it can be intimidating. In this work, we explore the use of visual listening guides to supplement live performances with information that supports the casual listener's increased engagement. We employ human-centred design practices to evaluate a currently implemented guide with users, from which we extracted design requirements. We then identify dimensions of a music piece that may be visualized and created sample guide designs. Finally, we presented these designs to experts of visualization and music theory. Feedback from the two evaluations informs design implications to consider when creating visual guides of classical music for casual listeners.

Keywords: Music visualization, visual listening guides, orchestra.

1 INTRODUCTION & BACKGROUND

With an aging and declining audienceship at classical music concerts [5], fewer people are experiencing the rich beauty and mastery of a live orchestral performance. Lack of exposure in popular media and cuts to music education in schools have been cited as contributors to this phenomenon [13]. Classical music, compared to modern mainstream music, is complex and elaborate, often performed by large instrumental ensembles playing long pieces that consist of many parts (movements). For the lay person with a limited background in music, these barriers can be intimidating to overcome, particularly given that the use of devices is prohibited during most performances, leaving them with no guiding information beyond what is printed in a provided booklet. As a result, many miss out on the opportunity to participate in this long-standing tradition of the orchestra, to deeply engage with music whose renown persists in our time (through the likes of Bach, Mozart, Vivaldi, Tchaikovsky and more) and even reap benefits from the experience [2].

There is a wealth of visual resources that are used to perform, understand or appreciate music pieces and their technical elements, but they do not accommodate a lack of music expertise. Aside from traditional music notation—which requires training to read—artistic visualizations called graphic scores express music in a visually engaging manner [11]. Graphic scores require a significant amount of performer interpretation, making them unsuitable for supporting listeners' understanding of the piece. The resources that exist for non-music specialists, meanwhile, tend to be academic in nature. Listening guides for music appreciation students integrate text, musical notation, and diagrammatic elements to provide detail about a piece [12]. However, these guides still require some basic technical knowledge, and are intended for study rather than enjoyment. In



Figure 1: Excerpted graphics from the Visual Listening Guides [3]. Reproduced with permission.

contrast, the orchestra concert attendee is provided with written program notes, which are often historical and contextual and provide only an overview of a piece. There is still a gap for a source of information that is engaging and requires no music expertise.

Orchestras have been eager to develop new strategies to broaden engagement by catering to the lay person [13]; market research identified a significant potential in attracting people who are not trained in classical music, yet are interested in the educational and artistic value of concerts [14]. This gives rise to an opportunity to provide guidance and information to enhance their listening experience, comparable to an audio-guided tour at a museum. The Toronto Symphony Orchestra (TSO) took a novel approach to this by developing Visual Listening Guides printed in program booklets (see Fig. 1). The guides have the stated purpose of “enhancing the understanding of classical music through engaging visual design” [3]. While orchestra patrons and the community at large generally received the guides positively [15], no formal evaluation was done to validate the effectiveness of the design, nor to incorporate user feedback.

We present an initial exploration of visualizations that present information with the aim of enhancing the live classical music experience for casual listeners, following a human-centred design approach that involves evaluation of current and novel designs. We gathered user feedback through a design probe to learn about casual listeners' constraints and requirements (§2). We created diverging sample designs (§3) and presented them to experts of visualization, and of music theory and musicology (§4). This allowed us to infer implications to the design of these visual artifacts (§5). This represents a first research step from which further user studies can be undertaken.

We chose to investigate these visual guides in paper form to align with the current culture of screen-free concerts, thus avoiding the potential distractions caused by dynamic visualizations and inviting their wider adoption by orchestras into the common program booklet. The major contribution we present is a set of guidelines about the perceptual (auditory and visual) cues which support or hinder a novice listener's understanding (§2.2.2); as well as important considerations that must be made with respect to the environmental constraints of a live concert, and the implementation limitations set by the nature of different orchestral pieces (§5). This work can inform further development of effective visual guides, with the future potential of wider adoption of them, and the creation of (semi-)automated methods to support the creative authoring of listening guides more broadly.

2 UNDERSTANDING CASUAL LISTENERS' REQUIREMENTS

As an initial probe into how casual listeners' experience may be improved through supplementary information, we conducted an evaluation of the TSO listening guides. For the purposes of this work, we define casual listeners as those not trained in music theory, nor in appreciation of classical music works. This yielded insight into the effectiveness of the TSO's use of visuals, the value of the information that was visualized, and suggestions from the casual listener audience for further design ideas to explore.

* cat.solis@mail.utoronto.ca

† fahimeh.rajabiyazadi@mail.utoronto.ca

‡ fanny@cs.toronto.edu

2.1 Methods of User Study

We recruited 17 participants (Ps) to undergo a simulated concert experience; demographic questionnaires confirmed limited exposure to music theory and orchestra performances. We provided them with a paper booklet containing the TSO visual guide for the piece, written program notes, and a “How to Read the Visual Listening Guide” page that they could consult at will in a well-lit room. This was done to mimic the real concert experience of waiting for the performance to begin. Lights were then dimmed and a six-minute single-shot video was shown of the third movement of Tchaikovsky’s fifth symphony. The task instruction, that we intentionally kept open-ended, was to “listen to the piece and try to follow along with the guide”, with no obligations to do so accurately or consistently. During the listening activity, we captured behavioural observations, including glancing direction and posture.

We administered the NASA Task Load Index (TLX) [6] post-study to assess subjective workload associated with following the guide (see results in Fig. 2), as well as an engagement questionnaire [9] (see results in Fig. 3), followed by a semi-structured interview in which we asked Ps about their strategies for following the guide, challenges to understanding the visuals, and reasoning behind observed behaviours and responses to the workload questionnaire. We transcribed interviews and analyzed data through affinity diagramming, for which there were two coders.

2.2 Design Requirements for Developing Visual Guides

Data from observations, questionnaires and interviews yielded helpful information about the overall impact of having a visual listening guide, and about perceptual (auditory and visual) cues from a performance which support or hinder a novice listener’s understanding. General feedback, as well as comments specific to the TSO guides, contributed to laying out more general design implications for the creation of visual listening guides (§5).

2.2.1 Overall Impact of Visual Listening Guides

Though the overall reception of the visual guide was mixed, the potential benefit of a visualization in this scope was evident through the interviews. For instance, 7 Ps stated that they gained knowledge about the piece’s instruments, keys and structure. Four Ps also commented that the guide helped them pay attention to the music. These observations indicate that the guide enhanced the participants’ listening experience, serving effectively as an information source and as an anchor to the music. In contrast, 9 Ps stated that they did not connect what they read in the text-based notes to their listening. These results motivate further exploration of a use of a visual guide.

The dominant contributors to negative feedback from participants came from challenges to understanding the visualization. All subjects encountered at least some difficulty, and most reported high mental demand (see Fig. 2), while following the visual elements. This suggests that improving the implementation of the visuals themselves will serve to reduce these negative impacts. Details of this feedback are described in §2.2.2.

The engagement questionnaire results show that most participants were engaged with the music, if not neutral (Fig. 3); despite the challenges faced with the guide, it was not detrimental to the listening experience. The identified barriers to understanding the guide provide strong motivation for a redesign, with promise that the positive impacts observed can be further increased.

2.2.2 Design Constraints

Affinity diagramming of the semi-structured interview transcripts revealed that user requirements primarily related to the technical constraints, i.e. the perceptual capabilities, of the casual listener.

Technical requirements. The most frequent observations related to the participants’ ability to perceive elements of the music and performance are described below:

- *Transitions:* Three Ps attributed difficulty in following the guide to an inability to hear when the music transitions to a new section.
- *Instruments:* While many found the instruments’ icons (5 Ps) and names (5) helpful cues, several had challenges with the icons (7), notably with distinguishing between instruments visually (5). Two said that they could not tell instruments apart by listening.
- *Melody:* Six Ps expressed difficulty in following the melody visual, with three of them expressing that hearing the melody itself was a challenge.
- *Key:* Three Ps said that the key indication was not helpful because they cannot tell the key from listening. However, 3 Ps said that hearing key changes helped them follow the guide.

Usability. When asked about the reported increased workload (see Fig. 2), 4 Ps expressed the lack of familiarity with the guide due to first time exposure as one of the contributors.

In summary, the main challenges to understanding the TSO guide faced by our participants primarily involved visual elements that closely followed traditional music notation, suggesting that lack of experience with it was a barrier. It is apparent from the aggregated feedback that there was no single piece of visual information that was a dominant cue for participants’ effective following, suggesting that an effective guide would visualize multiple cues that can be understood with no knowledge of music notation. Non-intuitive visual elements also created difficulty, even with instructional content, suggesting that a guide tailored to casual listeners must leverage their innate perceptual capabilities.

3 EXTRACTING DIMENSIONS OF MUSIC

In order to provide a structure for discussion around the design of visual guide to support the listening experience, we distill a set of different dimensions that may be visualized, derived from insights gained through our observations and interviews (Table 1). Each dimension represents a characteristic of a (performed) musical passage that can be individually described. To illustrate, we can apply this list to an excerpt of the TSO guide for Mozart’s




Figure 2: NASA-TLX raw responses from user study.

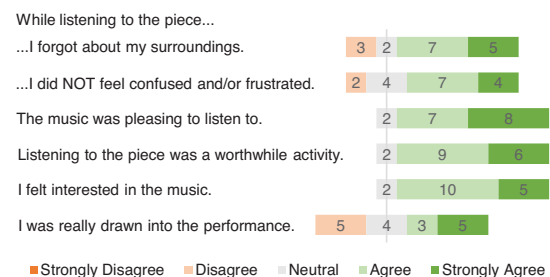


Figure 3: Engagement questionnaire responses from user study.

“Jupiter” (Figure 1, left) and observe the discretization of a fanfare motif (dimension #7, theme) in a coloured pill, with the melody (#1) encoded by dots and lines, the instruments playing it (#2) denoted by words and icons, and loudness (#3) indicated by traditional music notation (“f”).

We consulted literature from the interdisciplinary field of music information retrieval, which seeks to identify and categorize the information content within music for various applications involving computer-based parsing. In particular, Lesaffre et al. proposed a framework for user-oriented description levels of music that classifies them by their level of abstraction, spanning both lower-level, non-contextual information (physical properties of the sound), and higher-level, contextual information (affective, expressive properties) [8]. We leverage this framework not only to ensure completeness in consideration of the information potential of the music, but to also distinguish contextual from non-contextual information. The dimensions in Table 1 serve as a starting point to further investigate the larger design space of a visual listening guide and facilitate discussions around our problem space.

4 EXPLORING ALTERNATIVE DESIGNS

In order to initiate discussions with our experts about the opportunities and challenges associated with a design of a visual guide, we developed three sample design alternatives that we used as a probe, utilizing the dimensions identified above. Our goal was to further understand what information and what presentation methods have the most potential to support casual listeners.

4.1 Development of Sample Designs

We developed three sample design alternatives, building on our initial findings (Fig. 4). Diverging alternatives were devised using the morphological chart method [4], comparing the dimensions in Table 1 with different visualization metaphors. The range of solutions was narrowed by incorporating environmental constraints of the concert, feedback from the user study, and good practices from visualization literature. This divergence is necessary as visualizing all the dimensions of interest at once lends itself to information overload. Key characteristics of each design are highlighted below.

The “Waves” sample (Fig. 4a) is designed to emphasize the affective properties of the piece. We use expressive lines [10] as the prominent visual feature: smooth, wide curves conveying moments of gracefulness are contrasted with tight zig-zags which represent faster-paced, agitated melodies. Lines are encapsulated in pill shapes that differentiate melodic themes, and are colour-coded redundantly to further highlight changes in emotional character. To direct the listener’s attention to the instruments playing, the orchestra layout schematic is included alongside the pills, shaded based on volume.

The “Staff” sample (Fig. 4b) is designed to illustrate the structure of the piece. We prioritize representation of repeating themes and sections by separating them into designated rows for

Table 1. Dimensions of music, extracted from user study.

Dimension	Description
1 Melody	the relative pitch of the prominent sounds
2 Instrument	the instrument(s) playing, or playing the melody
3 Loudness	the overall volume
4 Key	the group of pitches which the music comprises
5 Time	the duration of the piece and its subdivisions; time markers
6 Spatial layout	where musicians of each instrument are sitting on the stage
7 Theme	the subdivision of the piece into sections, comprising themes, motifs etc.
8 Context	key emotions, inspiration, or ideas expressed in the piece

easier navigation. Discrete pills represent individual melodies that are repeated, and icons depict single instruments or sections that carry each melody.

The “Tunnel” sample (Fig. 4c) utilizes the spatial layout of the orchestra to visualize the passing of the sound among orchestra sections. This allows listeners to relate the sound at each moment to the instruments producing it. We use a diagonal timeline to simulate the visuals progressing in depth, with markers indicating changes in the instruments playing. To further orient the listener, speaker icons highlight particularly loud and emphatic portions of the piece (and are also included in the Staff design).

Based on the visibility limitations of the concert setting, colour was deemed suitable only for data with few categories; we used it to express three distinct emotional characters, deriving affective properties from existing program notes for the piece and selecting corresponding colours as recommended by Bartram et al. [1].

4.2 Expert Evaluation

We evaluated our prototypes with six experts in either information visualization or music theory/musicology in our institution. We considered this step necessary before further user evaluation in order to efficiently narrow the space of potentially effective designs, and to confirm that the information value of visual guides is in alignment with music experts’ conception. Three experts are instructors from the Faculty of Music, and three are PhD students researching Human-Computer Interaction in the Department of Computer Science. None had prior experience with our project before participating. Each interview followed the format below:

- Providing an overview of the project, and stating the scope and target audience for the visual guides,
- Presenting each design sequentially (in a different order for each expert), having experts think aloud as they made sense of each one, and identify any usability issues they encountered,
- Asking experts to select the design they considered best for the stated context and explaining their choice,
- Playing a recording of the visualized piece, allowing the experts to peruse any or all of the designs as they listened,
- Debriefing, asking whether/how the experience caused them to re-evaluate their assessment of the designs.




Figure 4: Excerpted graphics from the three sample designs used in expert evaluation: (a) “Waves”, (b) “Staff” and (c) “Tunnel”

From our music experts we also inquired about the potential scalability of these guides, as well as their experience teaching music to non-musically trained students. Interviews were audio recorded, and written notes were used to aggregate and analyze insights.

5 DESIGN IMPLICATIONS

We aggregated feedback from the experts (§4) to elucidate design implications to consider when creating visual listening guides. Key results are presented below, synthesized with supporting results from the initial user study (§3).

5.1 Present High-Level Information with Interpretability

All three designs incorporated a combination of low- and high-level information. This led some of the experts to try and ascribe technical meaning to visuals that were meant to be abstract, such as the expressive lines (Fig. 4a). This highlighted the need to *be clear with audiences the abstraction level* of the presented information.

While contextual information was appreciated by experts, with all gaining an idea about the character of the music from the colours used to convey emotions (Fig. 4), some were concerned about the prescriptive-ness of the visual. They advised to leave room for interpretability, and have the visual guide *set up only limited/ partial expectations that allow for discovery* during listening. Specifically, two experts recommended to remove the text labels on the colours. This highlights the value of the visual medium, which carries its own interpretability; studies have shown that semantically precise language-based descriptions, such as those found in program notes, decrease listeners' enjoyment of classical music [7]. These findings match our user study observations, in which 2 Ps experienced being "taken away" from the music by the specific and detailed visuals in the TSO guide. Despite these concerns, there was a consensus that providing high-level, con-textual information (specifically, the affective characteristics) is beneficial to support users' understanding of what they are hearing.

One element of the visual guides that had a mixed reception was the visualization of time. After a strong consensus from the user study that time would have been helpful for following, we included time markers in our presented designs. However, some experts stated that this was unnecessary, and the music experts noted the variability in timing based on the conductor's personal artistic interpretation. What experts proposed instead was to identify "landmarks" in the piece, such as solos, and to use changes in the music rather than time markers as anchors. Doing so would allow for listeners to follow the guides less rigidly and give designers greater interpretability of a piece. Future studies should seek to better define what constitutes an easily perceptible landmark for a casual listener.

5.2 Supplement Low-Level Details with Meta-Information

A concern expressed by four experts involved the use of technical terminology ("jargon"). They anticipated confusion from including the section names (e.g. "Waltz", "Trio", "Coda") which are used in music theory, and descriptive terms such as "runs". Counterintuitively from the usability heuristic of minimizing text in visualizations, some experts suggested brief descriptions, or a glossary for these terms. In a similar vein, they also advised adding text labels for elements which we attempted to convey solely through icons, specifically the instrument sections (Fig. 4b). These comments suggest that the visual representations of these *lower-level details should be supported with more text, or "meta-information"*, rather than relying on having users attempt to grasp novel concepts through listening alone. This aligns with the comments from casual listeners in our first study (6 Ps), who

noted that they ignored—and thus gained no value from—unfamiliar visual descriptors in the TSO guide. Additional meta-information would also support the interest that some users (3 Ps) expressed in general information about the orchestra (for example, the difference between "first" and "second" instrument).

5.3 Facilitate Different Levels of Reading

When asked to choose the guide they thought best fit the stated purposes, most experts gave different answers depending on when they intended on engaging with the design. For the use case of looking at the guide before the performance, they preferred the Waves design (Fig. 4a) for its simplicity and ability to effectively give an overview of the piece. However, when exposed to the designs while listening to the piece, experts noted that they were engaged with the details in the Staff (Fig. 4b) and Tunnel (Fig. 4c) designs, gaining insight from the visuals where they had not earlier. Based on these observations, we recommend *facilitating under-standing at both high and low levels of detail*, so the guide can be effectively used both before and during the performance. While our work was inconclusive regarding the casual listeners' booklet-perusing preference, we acknowledge that further studies of this behaviour would be valuable for programme designers.

5.4 Feasibility and Scalability of Implementation

Considering the visual listening guides' demonstrated potential, we looked to the music experts for an idea of the feasibility of creating them for different classical music pieces, to facilitate broad implementation. One stated that the symphonic works that comprise the "bread and butter" of a standard orchestra repertoire would be appropriate to visualize with these guides; this is reflected by the regularity with which the TSO was able to produce guides for their pieces throughout two seasons [3]. Experts also acknowledged, however, that many contemporary styles would not fit these design templates; there was specific mention of 20th and 21st century works. When presented with visual guides that identified which instruments were carrying the theme melody, music experts noted that this would be challenging or messy to visualize when a melody was being passed around very quickly, or for works in which all instruments play throughout. Additionally, experts were concerned about how longer pieces would scale to fit the limited space available in program booklets; however, this challenge may prove irrelevant for future design iterations that reflect experts' feedback about strict timelines being less important.

6 CONCLUSION

This work is, to the best of our knowledge, the first investigation of the role of visual guides in live orchestral performances. The orchestra environment is an interesting one for visualization design, given the imposed environmental constraints of low lighting and the restricted use of screens, posing unique design challenges. While an initial exploration, our study sheds light on important design considerations for improving experience of casual listeners. This work can inform the design of guides whose effectiveness can be demonstrated in further studies, and support development of (semi-)automatic computerized authoring methods for classical music concerts and beyond.

ACKNOWLEDGMENTS

The authors wish to thank Hannah Chan-Hartley for rich discussions on her work, all participants for their time and insights, the reviewers for their feedback, and Anastasia Bezerianos for her valuable input on the methodology.

REFERENCES

- [1] L. Bartram, A. Patra, and M. Stone. Affective Color in Visualization. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*, pp. 1364–1374. ACM Press, New York, New York, USA, 2017. doi: 10.1145/3025453.3026041
- [2] S. Chafin, M. Roy, W. Gerin, and N. Christenfeld. Music can facilitate blood pressure recovery from stress. *British Journal of Health Psychology*, 9(3):393–403, sep 2004. doi: 10.1348/1359107041557020
- [3] H. Chan-Hartley. About - Symphony Graphique. <https://www.symphonygraphique.com/about>.
- [4] N. Cross. Generating Alternatives. In *Engineering Design Methods: Strategies for Product Design*, chap. 9, pp. 123–138. Wiley, Chichester, 3rd ed., 2000.
- [5] E. Grannis. Orchestras Fight Hard Times Through Bankruptcy Seeking New Model. <https://www.bloomberg.com/news/articles/2012-08-21/orchestras-fight-hard-times-through-bankruptcy-seeking-new-model>, Aug. 2012.
- [6] S. G. Hart and L. E. Staveland. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In P. A. Hancock and N. Meshkati, eds., *Human Mental Workload*, vol. 52 of *Advances in Psychology*, pp. 139–183. North-Holland, 1988. doi: 10.1016/S0166-4115(08)62386-9
- [7] E. Hellmuth Margulis. When program notes don't help: Music descriptions and enjoyment. *Psychology of Music*, 38(3):285–302, July 2010. doi: 10.1177/0305735609351921
- [8] M. Lesaffre, M. Leman, and J.-P. Martens. A User-Oriented Approach to Music Information Retrieval. In T. Crawford and R. C. Veltkamp, eds., *Content-Based Retrieval*, number 06171 in Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, Dagstuhl, Germany, 2006.
- [9] H. L. O'Brien, P. Cairns, and M. Hall. A practical approach to measuring user engagement with the refined user engagement scale (UES) and new UES short form. *International Journal of Human-Computer Studies*, 112:28–39, Apr. 2018. doi: 10.1016/j.ijhcs.2018.01.004
- [10] A. T. Poffenberger and B. E. Barrows. The Feeling Value of Lines. *Journal of Applied Psychology*, 8(2):187–205, 1924. doi: 10.1037/h0073513
- [11] T. Sauer. *Notations 21*. Mark Batty Publisher, New York, 2009.
- [12] R. L. Todd. Discovering Music - Listening Maps. <http://www.oxfordpresents.com/ms/todd/listening-maps/>.
- [13] M. Vanhoenacker. Requiem: Classical music in America is dead. <https://slate.com/culture/2014/01/classical-music-sales-decline-is-classical-on-deaths-door.html>, Jan. 2012.
- [14] Classical Music Consumer Segmentation Study 2002 [United States]. Technical report, Knight Foundation, 2015. doi: 10.3886/ICPSR35535.v1
- [15] Toronto Symphony Orchestra listening guide - Information is Beautiful Awards. <https://www.informationisbeautifulawards.com/showcase/1322-toronto-symphony-orchestra-listening-guide>, 2016.