PATTERNVIEWER: AN APPLICATION FOR EXPLORING REPETITIVE AND TONAL STRUCTURE

Ali Nikrang
Dep. of Computational Perception
Johannes Kepler University Linz
an@musicresearch.eu

Tom Collins
Faculty of Techonology
De Montfort University
tom.collins@dmu.ac.uk

Gerhard Widmer
Dep. of Computational Perception
Johannes Kepler University Linz
gerhard.widmer@jku.at

ABSTRACT

The purpose of the PatternViewer application is to offer non-expert listeners new insights into the inner workings of pieces of music. This is achieved by presenting interactive representations of repetitive and tonal structure as a piece plays. [1] Our submission to ISMIR will have a demonstration component, allowing listeners to interact with recordings made by the Royal Concertgebouw Orchestra. The recordings will be of Ludwig van Beethoven's (1770-1827) symphonies, and interaction will occur via the interface shown in Figure 1 overleaf. Beethoven's symphonies are ideal candidates for the PatternViewer demo, because while they are central to the Western music canon, their size (both in terms of large orchestral forces and duration of individual movements) makes them an imposing and impenetrable prospect for many non-expert listeners. By visualizing their inner workings, the PatternViewer makes these pieces less imposing and more accessible.

The layout of the PatternViewer will be described below, but in short it combines music informatics research on calculation of chroma vectors [1, 2, 7, 8] with extraction of motivic, thematic, and sectional content [3, 6]. Thus, as a piece establishes a key and/or modulates, and as motives appear and disappear, perhaps forming part of larger themes/sections, all of this activity becomes evident to the user. Our objective, then, is for the PatternViewer to emerge as a salient example of the application of music informatics research to general music listening.

This research is built in turn on experimental findings and models from music psychology, where empirically derived chroma vectors have been used to model tonal cognition [5], and hierarchical graphs have been used to model processing of repetitive structure [4]. The question arises of whether, by exposing non-experts to structural representations derived from the habits of more expert listeners, it is possible for the PatternViewer to improve a non-expert listener's music appraisal skills. We are partway through an investigation into this matter, and would welcome the opportunity to present at the ISMIR demo, and the additional feedback on the PatternViewer that this would provide.

Lastly in this abstract, the layout of the PatternViewer is described in more detail. The application visualizes symbolic music data synchronized to a corresponding audio file. There are five elements to the display in Figure 1, which have been labeled with white indicator boxes. Each indicator box refers to the following:

1. **Media control bar** in the top-right corner, which shows the current playing time in the audio and can be clicked to navigate to different time points;
2. **Global-scale piano-roll representation** of the piece in the top panel, with different colors indicating changes in key estimates (again this can be clicked for navigation purposes);
3. **Local-scale piano-roll representation** of active (currently playing) patterns in the bottom-right corner;
4. **Colored circle of fifths** in the bottom-left corner, which provides a reference for key estimates shown by the differing colors in the global-scale piano-roll representation (e.g., currently the piece in Figure 1 is in C major);
5. **Pendular graph** in the top-left hand corner (so-called because the nodes and edges resemble pendulums), where each node represents a note collection that repeats. Nodes for active note collections appear red and their note content appears in the local-scale piano-roll representation.

The pendular graph shows all occurrences of a particular repeated pattern in the same row. For instance, $C_1, C_2, \ldots, C_6$ in the pendular graph of Figure 1 are all occurrences of the main theme from the first movement of Beethoven's Symphony no.1. Each note in C_1 belongs also to the large repeated section A_1, and so an edge connects the two nodes. Nodes in a pendular graph are ordered vertically by pattern duration, so shorter motives/themes tend to appear lower down in the graph, and may connect to longer themes/repeated sections, which tend to appear nearer the top of the graph. Nodes in the PatternViewer are clickable, bringing the audio to the beginning of the clicked pattern occurrence, and showing its note content in the local-scale piano roll representation.

[1] Video demonstration at http://tinyurl.com/l3ppov9

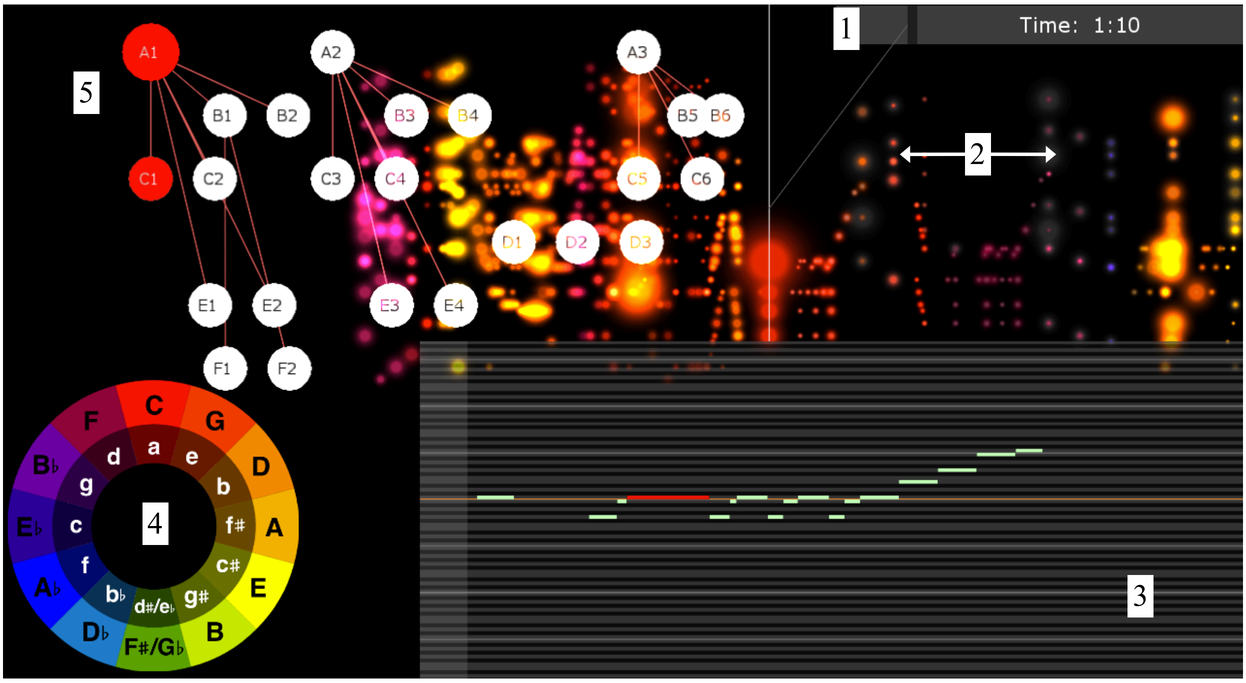

Figure 1. The PatternViewer visualization of the first movement from Beethoven's Symphony no.1 in C major op.21. Top right (indicator box 1): a media control bar shows the current playing time, and can be clicked/dragged for navigation purposes. Top panel (indicator box 2): a global-scale point-set representation, colored by the current key of the music. A white vertical bar indicates the position within the audio, which can be moved back or forward by clicking to the left or right. Bottom right (indicator box 3): a local-scale piano-roll representation, showing the contents of a pattern occurrence from the top-left graph (currently the main theme C_1). Bottom left (indicator box 4): a colored circle of fifths, which can be used to identify the key color shown in the top panel (currently C major). Top left (indicator box 5): a pendular graph representing the hierarchical, repetitive structure of the current piece. For instance, C_1 is the first occurrence of the main theme, whose notes are a subset of the large repeated section A_1 known as the exposition.

1. ACKNOWLEDGMENTS

This paper benefited from the use of the MuseData repository. The work is supported by the Austrian Science Fund (FWF) under project number Z159 (Wittgenstein Grant) and the European Union under the FP7 Framework (Project PHENICX, grant no. 601166).

2. REFERENCES

[1] J. P. Bello and J. Pickens. A robust mid-level representation for harmonic content in music signals. In *Proceedings of the International Symposium on Music Information Retrieval*, pages 304–311, 2005.

[2] W. Chai and B. Vercoe. Detection of key change in classical piano music. In *Proceedings of the International Symposium on Music Information Retrieval*, pages 468–473, 2005.

[3] T. Collins, A. Arzt, S. Flossmann, and G. Widmer. SIARCT-CFP: Improving precision and the discovery of inexact musical patterns in point-set representations. In *Proceedings of the International Symposium on Music Information Retrieval*, pages 549–554, 2013.

[4] D. Deutsch and J. Feroe. The internal representation of pitch sequences in tonal music. *Psychological Review*, 88(6):503–522, 1981.

[5] C. L. Krumhansl and E. J. Kessler. Tracing the dynamic changes in perceived tonal organization in a spatial representation of musical keys. *Psychological Review*, 89(4):334–368, 1982.

[6] D. Meredith, K. Lemström, and G. A. Wiggins. Algorithms for discovering repeated patterns in multidimensional representations of polyphonic music. *Journal of New Music Research*, 31(4):321–345, 2002.

[7] C. S. Sapp. Visual hierarchical key analysis. *ACM Computers in Entertainment*, 3(4):1–19, 2005.

[8] J. Serrà, E. Gomez, P. Herrera, and X. Serra. Statistical analysis of chroma features in western music predicts human judgments of tonality. *Journal of New Music Research*, 37(4):299–309, 2008.