Music IR for Music Theory

Eric J. Isaacson
Indiana University School of Music
1201 East Third Street
Bloomington, IN 47405 USA
+1-812-855-0296
isaacso@indiana.edu

ABSTRACT

"Academic musicians," students and faculty in schools of music, have music IR needs that differ from those of the mass-market consumer. This paper describes the characteristics of this user group, the types of musical information they use and how they use them, and the kinds of IR tasks they need to perform. A section describes the special needs of the music theorist. Finally, the implications of this group's needs for a music IR testbed are outlined.

1. USERS

I represent academic music institutions, mainly schools and departments of Music, in U.S. and Canadian colleges and universities. Activities carried out in these institutions include teaching and learning, performance, and scholarly research,. Three groups of people are involved in these activities: (1) undergraduate and graduate music majors, who are training to be music teachers, professional musicians, scholars, arts administrators, and music technologists; (2) non-majors taking music courses (such as music appreciation, rock or jazz history, music and culture) on an elective basis; (3) and faculty, who teach the students, engage in performance and other creative activities, and carry out research in areas such as musicology, music theory, music perception and cognition, and music education. Within this latter group, I am interested primarily in the music theorist, who as a faculty member (or as a graduate student), is involved in both teaching and research, about which I will have more to say below. I will refer these three groups of people collectively as "academic musicians."

Academic musicians are an important audience for music IR research. Though smaller in number than the mass-market consumer, they are a highly knowledgeable user-base, with needs that, though not altogether different, are differently focused. The musical repertoire of interest to academic musicians is largely dissimilar to that of the mass-market. The kinds of IR tasks they require are generally broader, and the types of information they need to access are different.

2. TYPES AND USES OF MUSICAL INFORMATION

The teaching and performance activities of North American music schools tend to focus on western art music (typically referred to, somewhat imprecisely, as "classical" music), particularly that written in the period 1650-1900. Increasingly, (art) music of the twentieth century is being performed, and that century's music may is the focus of what may now be the majority of music theoretical research. In addition, greater attention is being given to jazz, popular song, film music, rock and pop, and world musics.

I emphasize the pedagogical uses of musical information because music students take a large number of courses involving music theory and music literature, and the pedagogy of music theory is an important part of what music theorists think about. In the course of teaching, music in various formats is used. (In this section, I will focus on institutionally maintained musical information—those things that are stored in some way by the institution, normally through a library, which then provides access to users. I will not be concerned with things that students purchase for themselves.) Printed scores are arguably the most important kind of musical information used in this environment. They are used by students looking for new music to learn and perform, and in their courses in music theory and music history. And aside from their instructional uses, they are used by music theorists for whom the analysis of compositional procedures is often a central research activity.

Recorded music, which was a rarity until the advent of the LP, is now widely used. Student use recordings to listen to expert performances of music they are studying or to become familiar with new pieces. Many courses include listening exams that require students to recognize and describe features in recorded music. Researchers are beginning to use recorded music for the study of performance nuance (see Dodson [2002] for a recent, and online, example). Though used less often, video is also used, especially to study performances of opera and ballet.

Although these uses typically involve physical media (printed scores, CDs, VHS tapes), the potential benefits of having access to digital versions of these media are many. (The Indiana University *Variations2* project, http://variations2.indiana.edu/, is based on this premise.)

Symbolically encoded scores are a third important type of musical information. The encoding of scores dates back to at least the 1960s. Except for MIDI files, the use of encoded scores remains largely localized, with collections developed by an individual scholar for that scholar's own research. In music instruction, encoded scores could be useful because it would be possible to modify existing works to illustrate alternative compositional choices, or to create composition assignments in which a student is given an excerpt from a piece and asked to complete it in the same style. For researchers, the availability of encoded scores provides a way to do certain kinds of analysis over a large collection of examples. The best collection of encoded scores right now is the MuseData database, under development at the Center for Computer-Assisted Research in the Humanities (http://musedata.stanford.edu/).

3. RETRIEVAL TASKS

In this section I will describe some of the retrieval tasks that students and faculty in music departments either currently perform, or would like to perform if the mechanisms existed to support such tasks. Then in section 4 I will describe some of the music-analytic tasks carried out by music theorists. Finally, I will discuss the implications of these tasks on the design of a testbed for Music IR research in section 5.

3.1 Metadata-based search

Music students and faculty looking for scores or recordings often do so knowing information *about* the item or kind of item they are looking for that is not related to contents of the item. A metadata-based search is appropriate for these situations. Such searches are most often done based on one or more of the following: composer, title, performer. Other search criteria might include genre (which may not appear as part of the title—tone poem, for instance); date or place of composition, publication, or recording; other personnel, such as performers, editors, librettists, or choreographers; key; instrumentation; format (full-score, piano-vocal reduction); language; and so on. To be effective, a search engine must be sensitive to variant spellings.

Here are some hypothetic searches of this kind:

- Looking for sound recording for C.P.E. Bach Piano sonata, W55 no 2 (H130)
- Looking for a video of *Don Pasquale* by Donizetti
- Looking for the score to Byrd, "Come, woeful Orpheus"
- Looking for recordings of cello music by Beethoven, Martinu, and Brahms
- Looking for settings of poems by Goethe

3.2 Content-based search

Sometimes faculty or students want to find scores or recordings of one or more works that meet certain criteria. Such searches may add content-based criteria to a metadata-based search, or may be based strictly on content. Following are some examples. All of these search examples must currently be carried out in large part manually, perhaps after an initial metadata-based screening process.

- A contralto preparing a recital wants to find songs in French or Italian for alto and piano in which the voice part goes no higher than C5 and no lower than F#3
- A saxophonist learning to play "high tones" wants to find all works in the library that include an alto saxophone part that uses the altissimo register (above concert pitch A5)
- A musicologist studying the use of quotation wants to find several examples of works that contains the B-A-C-H motive
- A music history professor teaching a course in Renaissance music wants to find masses based on a particular plainchant melody.
- A music theory instructor teaching a course covering chromatic harmony wants to find a piece in a major key that includes a Neapolitan chord (=II) in root position
- A music theorist teaching a form-and-analysis course wants to find pieces in sonata-allegro form in which the themes recapitulate in the reverse order
- A student in an eighteenth-century counterpoint course wants to find a fugue or invention in which the opening theme is subsequently stated in inversion.
- A music theorist writing a book on form wants to find a five-part rondo (ABACA) in an 18th-century symphony not by Mozart or Haydn.

Such searches are carried out every day by music students and faculty, but they are done by manually by browsing shelves of music in the library and thumbing through volumes of music. While such time-consuming searching can sometimes be valuable, it more often involves simply the inefficient use of time.

4. INFORMATION ANALYSIS TASKS

In this section I will describe some of the tasks carried out by music theorists. After a general overview, I will provide a representative sampling of activities that have been done using computer-based techniques.

As a professional scholarly discipline, music theory is concerned with the following: studying the abstract properties of temporal and pitch spaces and developing theories to describe those spaces; studying groups of compositions (unified by such features as composer, genre, era, geographical region, or culture) to discover stylistic traits, compositional procedures, or other organizational features; developing theories to describe the discovered characteristics; using those theories to analyze the structure of other works, both as a way of testing or extending the theories, and to discover idiosyncratic features of the other work; studying relationships between music and literature, other arts, science, and mathematics; researching human cognitive and perceptual abilities and responses to various aspects of music; applying educational principles to the teaching of music theory and to developing instructional materials; study of aesthetics and the philosophy of music; and researching the history of music theory.

Much of the time the study of music is centered on the musical score, though score-based analysis often involves at least an implicit understanding of how the piece "sounds," even if no specific performance is in mind. Less often, the actual sound—whether the details of a particular performance or the raw audio signal—forms the basis of a study. In such cases there is often an explicit understanding of the underlying symbolic "content" of the music (the pitches, rhythm, musical lines, etc.).

It should be noted that, although the term "music theory" connotes some type of scientific rigor, much of the activity carried out by music theorists is better described as "art." This is because, although relationships involving basic musical materials can be codified and even "proven" in a mathematical sense, the judicious application of those theoretical principles in the analysis of music requires creativity, in large part because music exists as the result of a creative process.

I will now describe a small sampling of research that has, or would have, benefited from the existence of some kind of a music IR database (usually a small, local one), appropriate retrieval mechanisms, or tools for the analysis of musical information.

Brinkman (1979) studied the works in J. S. Bach's *Orgelbüchlein* to find traces of chorale melodies in the preludes on which they are based. Brinkman used a database containing the melodic contours of all the works found in Bach's collection.

Gjerdingen (1991, and in a series of follow-up articles) uses an Adaptive Resonance Theory (ART 2) neural network (Carpenter and Grossberg, 1987) to study recurring chord and voice-leading progressions, as well as the role of meter in the perception of harmonic patterns, in early Mozart piano pieces. These works are "encoded" as a series of input patterns to a neural network. (The study of such prototypical patterns or, more accurately, cognitive schemata, formed the basis of a more extensive, though not computer-based, study [Gjerdingen 1988]).

Based on work by Meyer (1956), and deeply indebted to principles from Gestalt psychology, Narmour (1990) develops a theory of melody based on cognitive principles involving expectation and realization. Narmour's work has been the subject of psychological verification studies by Krumhansl (e.g., 1997), and computational testing by Von Hippel and Huron. Von Hippel and Huron (2000) tested one of Narmour's assumptions by analyzing a large collection of encoded melodies against the predictions of Narmour's theory. Huron and members of his group have published other studies examining some of music theory's conventional beliefs (e.g., Huron 2001a, Huron and Royal 1996) and methodological practices (Huron 2001b), many of which draw on a large database of symbolically represented music. This work, which represents arguably the most interesting of current computer-based theory and analysis activities, is carried out largely using the Humdrum toolkit developed by Huron (1999). (A list of over 350 Humdrum tasks, which include a number of interesting music-analytical queries, is given at http://www.music-cog.ohio-state.edu/Humdrum/problems.html.)

The ability to digitally encode musical scores and audio recordings has led to new visualization techniques for music. Brinkman and Mesiti (1991) describe tools that display new kinds of graphical representations of music, derived from the score. Cogan (1984, 1998) on the other hand uses spectrograms of recorded performances to study various large-scale aspects of the music that are particularly apparent visually, such as the role of musical parameters such as dynamics, register, and textural density in shaping musical form.

The possibilities for additional computer-based study in music theory and analysis are substantial. New visualization techniques, the testing of existing theories against large bodies of works, and the discovery of previously unrecognized musical patterns through data-mining techniques are among the promises in this area.

5. IMPLICATIONS FOR A MUSIC-IR TESTBED

What would be the benefits of research using a Music IR testbed for the academic musician? For students, it would provide the ability to more effectively locate pieces that fit criteria they were interested in, and to be able to listen to and view the music. For the teacher, it would provide the ability to do content-based search for music containing certain features, and to provide synchronized playback/display of audio and score. For researchers, it would provide the ability to analyze large amounts of data for previously unknown patterns. The existence of a testbed would also facilitate such research by ultimately eliminating the arduous task of encoding large amounts of music. In addition, it would give different researchers an opportunity to tackle a common repertory from multiple perspectives.

What should such a testbed look like? The Music IR testbed should include music in multiple formats, including audio, scanned score images, and symbolically represented music. Having retrieved a work in one format, it should be possible to retrieve instances of the same work in other formats. The testbed should provide a rich metadata structure that permits searches for items using a wide range of criteria pertaining both to a work and particular instances of it (see http://dml.indiana.edu/pdf/DML-DataModel-V1.pdf).

The testbed should support a structural metadata scheme that enables synchronization between different representations of the same work, if possible to the level of the measure, and should support research that would lead toward the automated linking of audio and symbolic data. This is an interesting and challenging task in itself, but is has the important practical benefit of addressing the problem of scale in creating synchronization data for a potentially huge number of works.

The symbolic data format used in the testbed should include sufficient detail to support the display and editing of common music notation (see Byrd and Isaacson, 2002, for a possible feature list for such a representation). The symbolic format should support the ability to trace music both horizontally (harmonic) and vertically (melodic).

Audio in the testbed should be of high-resolution (near CD-quality). Support for multiple levels of audio quality, though helpful for distance learning applications, is not essential, nor is it desirable for face-to-face teaching.

Scanned score images should be scalable without losing the integrity of lines in the notation (staff lines and stems). The DjVu format works very well for this (http://www.lizardtech.com).

The testbed itself should include music of a wide variety styles and genres within the Western art music tradition. It should include data in all three formats (audio, scanned score images, encoded scores) for as many of the works it contains as possible. Including items from jazz, world musics, and popular music is more challenging because of the problems with notation, but these repertories are less important. Ideally, it should contain at least two sets of "complete" works that would permit researchers to develop analytic approaches to a composer's entire output or a complete segment of that output. I would encourage the music IR community to identify a handful of music theorists would who agree on an appropriate "complete" collection to use as the basis for several types of analysis, all derived from the same collection.

6. ACKNOWLEDGMENTS

Some of the material in this paper is based upon work supported by the National Science Foundation under Grant no. 9909068. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the view of the National Science Foundation.

7. REFERENCES

[1] Brinkman, A. R. 1980. "The Melodic Process in Johann Sebastian Bach's *Orgelbüchlein*." *Music Theory Spectrum* 2: 46-73.

[2] Brinkman, A. R., and Mesiti, M. P. 1991. “Graphic Modeling of Musical Structure.” *Computers in Music Research* 3: 1-42.

[3] Byrd, D., and Isaacson, E. 2002. “Music Representation Requirement Specification for *Variations2*.” Unpublished manuscript.

[4] Carpenter, G. A., and Grossberg, S. 1987. “ART 2: Self-organization of Stable Category Recognition Codes for Analog Input Patterns.” *Applied Optics* 26: 4919-30.

[5] Cogan, R. 1984. *New Images of Musical Sound*. Cambridge, Mass., Harvard University Press.

[6] Cogan, R. 1998. *Music Seen, Music Heard: A Picture Book of Musical Design*. Cambridge, Mass.: Publication Contact International.

[7] Dodson, A. 2002. “Performance and Hypermetric Transformation: An Extension of the Lerdahl-Jackendoff Theory.” *Music Theory Online* 8.1 <http://societymusictheory.org/mto/issues/mto.02.8.1/toc.8.1.html>.

[8] Gjerdingen, R. O. 1988. *A Classic Turn of Phrase: Music and the Psychology of Convention*. Philadelphia: University of Pennsylvania Press.

[9] Gjerdingen, R. O. 1991. “Using Connectionist Models to Explore Complex Musical Patterns.” In *Music and Connectionism*, edited by Peter M. Todd and D. Gareth Loy, pp. 138-149. Cambridge, Mass.: MIT Press.

[10] Huron, D. 1999. “Humdrum User’s Guide.” <http://www.music-cog.ohio-state.edu/Humdrum/guide.toc.html>

[11] Huron, D. 2001a. “Tone and Voice: A Derivation of the Rules of Voice-leading from Perceptual Principles.” *Music Perception* 19: 1-64.

[12] Huron, D. 2001b. “What Is a Musical Feature? Forte's Analysis of Brahms Op. 51, No. 1, Revisited.” *Music Theory Online* 7.4 (July). <http://societymusictheory.org/mto/issues/mto.01.7.4/toc.7.4.html>

[13] Huron, D. and Royal, M. S. 1996. “What is Melodic Accent? Converging Evidence from Musical Practice.” *Music Perception* 13: 489-516.

[14] Krumhansl, C. L. 1997. “Effects of Perceptual Organization and Musical Form on Melodic Expectancies.” In *Music, Gestalt, and Computing: Studies in Cognitive and Systematic Musicology*, pp. 294-320, edited by Marc Leman. Berlin: Springer Verlag.

[15] Meyer, L. B. 1956. *Emotion and Meaning in Music*. Chicago: University of Chicago Press.

[16] Narmour, E. 1990. *The Analysis and Cognition of Basic Melodic Structures: The Implication-realization Model*. Chicago: University of Chicago Press.

[17] Von Hippel, P., and Huron, D. 2000. “Why Do Skips Precede Reversals? The Effect of Tessitura on Melodic Structure.” *Music Perception* 18: 59-85.