# Variations2: Toward Visual Interfaces for Digital Music Libraries

**Mark Notess**
University Information Technology Services
Cook Music Library
Indiana University
Bloomington, IN 47405 USA
mnotess@indiana.edu

**Natalia Minibayeva**
University Information Technology Services
Cook Music Library
Indiana University
Bloomington, IN 47405 USA
nminibay@indiana.edu

**ABSTRACT**
In this paper, we discuss unique challenges of visualizing musical data, address limitations of existing music interfaces and data structures, and propose some approaches to music visualization based on the data model developed by Variations2 digital music library project.

**Keywords**
Digital library, music, information visualization.

**INTRODUCTION**
Visualization of music bibliographic data encounters difficulties on two levels. First, today's music bibliographic data is typically force-fit into a catalog structure that obscures most of its richness and utility. Second—and certainly as a result of the first obstacle, useful visualization tools for music data inquiry and presentation simply do not exist. Variations2, the Indiana University Digital Music Library Project [12], is making a significant effort to address these problems by providing a metadata record structure to capture the richness and relationships inherent in music bibliographic data. Variations2 is also examining the visualization opportunities provided by this new catalog structure. This paper describes the limitations of existing data structures and technologies and offers a glimpse of work in progress on the related issues of catalog structure and visualization.

**UNIQUENESS OF MUSICAL DATA**
Music representation and organization present a number of challenging problems that affect the ways in which visual music interfaces are constructed. First of all, as defined by the Standard Music Description Language (SMDL) [8], music exists in several physical and logical dimensions:

- visual/graphic (music notation)
- performance/gestural (actual sounding and performance)
- logical (logical information required for music realization)
- analytical (concepts involved in musical analysis)

Each domain captures certain music properties yet excludes others.

Second, any given musical work is typically manifested in various performances and printed editions. These various instantiations of the same work often differ drastically from one another, as well as from the more abstract concept of the musical work in terms of their instrumentation, structure, duration, and interpretation. Thus, certain attributes can be shared among the instantiations of a work and some can differ. This presents a challenge in terms of linking instantiations to works, which is necessary for efficient and effective organization, retrieval, and manipulation of music information.

Finally, instantiations of a musical work exist in a variety of formats (score, recording, video, text). Therefore, in addition to the differences resulting from various interpretations of the musical work, there are also variations imposed by the different formats.

Without consideration of these issues specific to the music domain it would be difficult if not impossible to design an effective visual music interface. The needs of music users can be satisfied only if musical objects are adequately described, represented, and organized. This paper focuses on visual interfaces related to music information retrieval, which is concerned mostly with bibliographic music information rather than actual music representation schemas (which present their own fascinating issues).

**CATALOG RECORD STRUCTURE LIMITATIONS**
The existing text-based catalog record structures such as the USMARC [7] provide a foundation for capturing some of the required descriptive music information but also present some significant limitations. The biggest problem is the "flat" nature of the record structure, which does not reflect the complexity of the relationships existing among various musical objects. For example, when a music library user is looking for some recordings and scores of a specific musical work, s/he has to perform several complex searches to collect all the needed records (e.g., "Mozart AND Sonatas AND K332 AND Score", etc.). This is

because MARC records mostly capture container-level information and do not provide a robust mechanism for linking various instantiations of the same musical work. This presents difficulties in searching for music information, as search queries are often very complex and require special user training.

## INQUIRY AND PRESENTATION TOOL LIMITATIONS

Visual interfaces to music bibliographic data are not broadly used today. To explore what could be done with existing tools and data, we carried out some experiments, which we describe in this section.

### Variations and IUCAT

The original Variations project [4, 11], for which Variations2 is named, provides for the on-line storage and retrieval of 9000 hours of audio recordings. This system is in current use in the IU Cook Music Library as a listening reserves system, serving up between hundreds and thousands of audio recordings each day. Users may access recordings via URLs embedded in course reserve lists, or users may search the IU library catalog using IUCAT [6].

IUCAT is a typical web-based library catalog front-end in that it offers a variety of search options and returns a limited list of matches in summary form. This is not very useful data for visualization because it is MARC-based container-oriented data. Details such as date of composition, instrumentation, genre, or key are not always available in a reliable, consistent manner.

### Experiments with Spotfire DecisionSite

Given the issues mentioned above with our existing catalog data, we looked elsewhere for a source of music bibliographic data to use for visualization experiments. The Classical.Net website [3] provided a large database (20,000+) of musical works based on a personal collection. The collection is not representative of the IU music library collection (e.g., it contains very little vocal literature). The Classical.Net data was made available in a Microsoft Access database. The fields in the database used in the visualizations were *composer*, *key*, *date of composition*, and *solo instrument*.

Spotfire® DecisionSite™ is a commercially available visualization tool used in such varied industries as oil and gas exploration and biotechnology [10]. It grew out of dynamic queries and starfield simulation research at the University of Maryland. An early demo of the technology showed a user interface to a film database [1]. Spotfire provides a variety of visualizations, including scatter plots and histograms. Spotfire's particular strength is the "tight coupling" between the user interface controls and the data—as you move a slider, the visualization changes immediately.

Here we give just two examples of the kind of visualizations possible with Spotfire. Figure 1 shows a histogram of musical periods with the bars divided into three parts, giving the proportion of the works written in a major key, a minor key, or unspecified (from top to bottom of each bar). Figure 2 shows a 3D visualization where the axes are as follows: x-date of composition, y-composer, z-solo instrument.

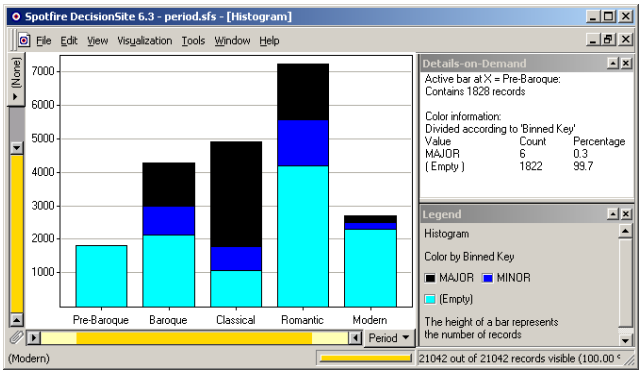


Figure 1. Key by Period

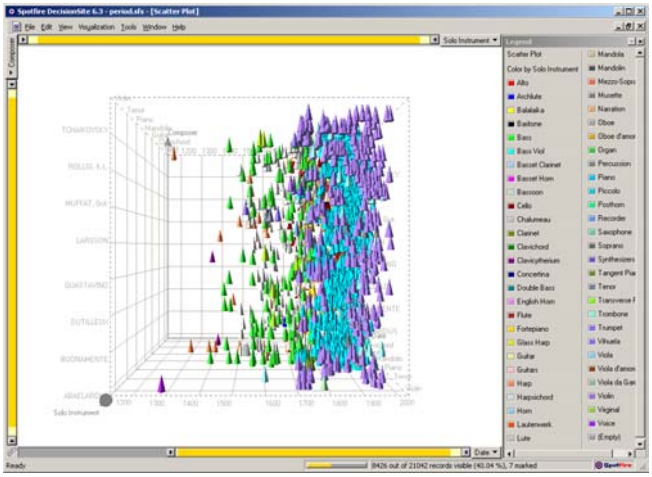

Figure 2. Solo Instrument by Composer and Date

A commercial visualization tool like Spotfire has the advantage of being robust, with a rich set of visualization capabilities. The dynamic queries provide an effective way to zoom in on the data of interest while not losing a sense of the broader context, even with over 20,000 records. Traditional library interfaces provide nothing comparable. However, because Spotfire is an off-the-shelf tool not aimed particularly at meeting library needs, the experiments and examination of technical capabilities yielded the following limitations.

- Spotfire is for expert users. The interface as it exists could not be used by a walk-up on-line catalog user.
- Spotfire works with a flat table of data—it provides no way to handle a hierarchical tree of different record types, e.g., users cannot navigate from a work to an instantiation of a work and then examine information about a performer associated with the specific instantiation. The only relationships Spotfire handles are categorical grouping (e.g., solo instrument) and quantitative (e.g., date of composition).

- Sequential dates can be modeled as numerical values; Spotfire has no special datatype for "date". Date intervals cannot be displayed as timelines.
- Spotfire requires the complete database to be available on the local system.
- Spotfire is expensive—there is no inexpensive pricing model for a stripped-down run-time version.

Although experiments with Spotfire are helpful in figuring out what sorts of visualizations may be useful for music information seeking, it is unlikely that off-the-shelf visualization tools developed for other markets will be useful in visualizing music data.

### VARIATIONS2 Project

We now give a brief overview of our current project, Variations2 [12], and focus particularly on the new data model and metadata scheme.

#### Project Description

Variations2, the Indiana University Digital Music Library Project, is establishing a digital music library testbed system, as a foundation for digital library research in the areas of instruction, usability, human-computer interaction, and intellectual property rights. The initial Variations project focused primarily on audio playback. Variations2 will provide synchronized playback of multiple music media, including audio, scores, and encoded music notation. In addition, Variations2 will use a new data model and metadata scheme that expresses the richness and complexity of musical bibliographic data. A search interface will provide access to and retrieval from the collection. Variations2 will be used by students, faculty, librarians and library patrons at Indiana University as well as by users at several remote locations in the U.S. and overseas. Remote locations include: University of Illinois at Urbana-Champaign, University of Massachusetts at Amherst, Northwestern University, King's College-London, Loughborough University, Oxford University, and Waseda University. For information on other research and development activities related to Variations2, see [12].

#### Data Model and Metadata Scheme

Variations2 overcomes the limitations of traditional catalog structures, captures the complexity of music information, and satisfies the special needs of music library users. The Variations2 data model accommodates the relationships existing among various types of musical objects by introducing five related entities, each represented by descriptive, administrative, and structural metadata elements:

- **Work:** Represents the abstract concept of a musical piece or set of pieces. For example, Beethoven's "Symphonies, no. 7, op. 92, A major" is a work.
- **Instantiation:** Represents a manifestation of a work as a performance or a score. An actual performance of Beethoven's 7th by the London Symphony Orchestra is an instantiation of the abstract work.
- **Container:** Represents the physical item or set of item(s) on which instantiations of works can be found. If a publisher releases a CD recording of the London Symphony's recording of Beethoven, along with recordings of the rest of his symphonies, this CD set is a container. A published score of Beethoven's 7th or a video of a performance are also containers.
- **Media Object:** Represents a piece of digital media content, such as a sound file or score image. For example, a CD set may be stored on-line as multiple audio media objects.
- **Contributor:** Represents persons or groups that contribute to a Work, Instantiation, or Container. Common categories of contribution include composer, performer, arranger, conductor, and producer, although there are many other roles as well.

Although catalogers must concern themselves with all of these entities, end users search for and interact primarily with the work, the contributor, or the container.

This data model allows for a more effective and efficient navigation, browsing, and search of different digitized music formats. It provides a foundation for designing visual interfaces that would utilize the wealth of music attributes and relationships and assist the user in successful retrieving and manipulating of musical data.

### VISUALIZATION EXAMPLES

With the new data model and metadata scheme, we can provide useful visualizations. Consider the following two examples.

#### Example 1—Cello Music by Specific Composers

A music reference librarian today might be told: "I'm looking for cello music by Beethoven, Martinu and Brahms." Many of the occasions on which patrons ask the librarian to help locate a resource are provoked by the difficulty of figuring out the right way to query the catalog or to find the desired item amidst the profusion of imprecise search results. With Variations2, the patron could enter the following query:

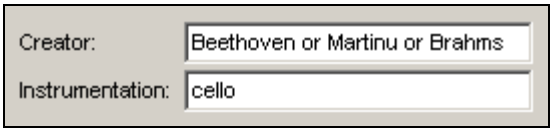


The response to such a query could be a visualization such as the one shown in Figure 3.

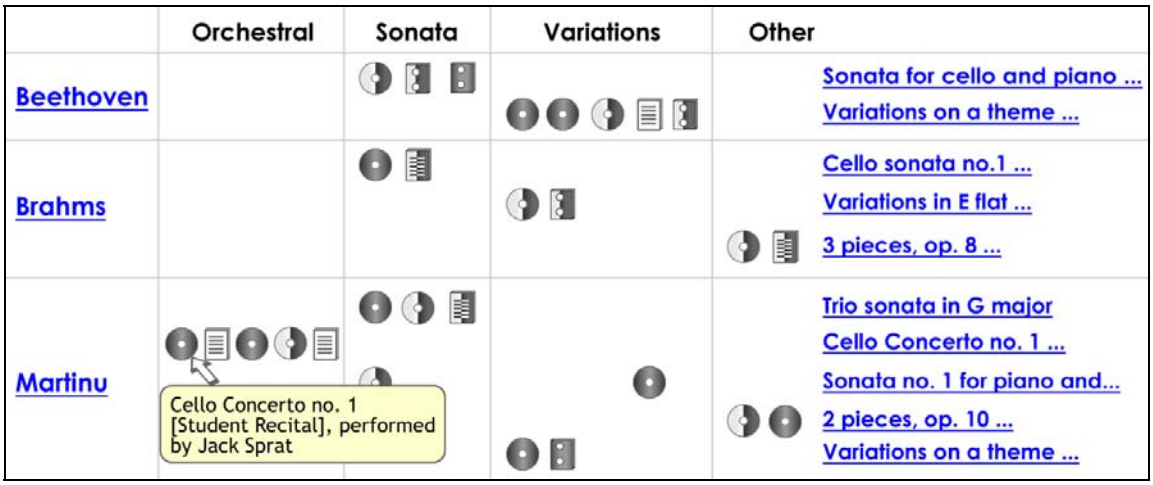


Figure 3. Visualization of Search Results

In this figure, we see a grid-based visualization that uses icon shape to represent media type (audio, score, or video), color to represent the performer, and position within the grid to indicate both work genre (x-axis) and composer/work (y-axis). Hovering over an object gives details on performer, etc. One can imaging using techniques such as Table Lens [9] or GRIDL [5] to explore a large dataset without losing the larger context. The visualization gives immediate feedback on the relative frequency of different works in the collection and makes it easy to browse—clicking on the item's icon invokes the player or viewer window appropriate to the data type. Even though users may simply click on the icon to access the digital media, some users may wish to see further information about the work. While users could get more information by simply hovering over the icon, this information is transient and brief. An alternative is to click on the name of the work on the right-hand side of the display. This hyperlink brings up a list of instantiations (Figure 4), which is itself another grid-based visualization.


Figure 4. Work Window

The purpose of the instantiation display is to show all the on-line media holdings for the specified work. This visualization also uses icon shape to indicate media type, but it provides additional information about the completeness of the instantiation (e.g., a CD might hold only one movement from a longer work) and uses the y-axis to give the publication date context. Hovering over, e.g., a recording, provides such instantiation details as title (which may differ from the work title), performer(s), instrumentation, language, and performance date. As with Figure 3, clicking on the icon invokes the player or viewer.

The other navigation available from the search results display in Figure 3 is the composer name hyperlink. Clicking on this link displays the composer overview window (Figure 5).

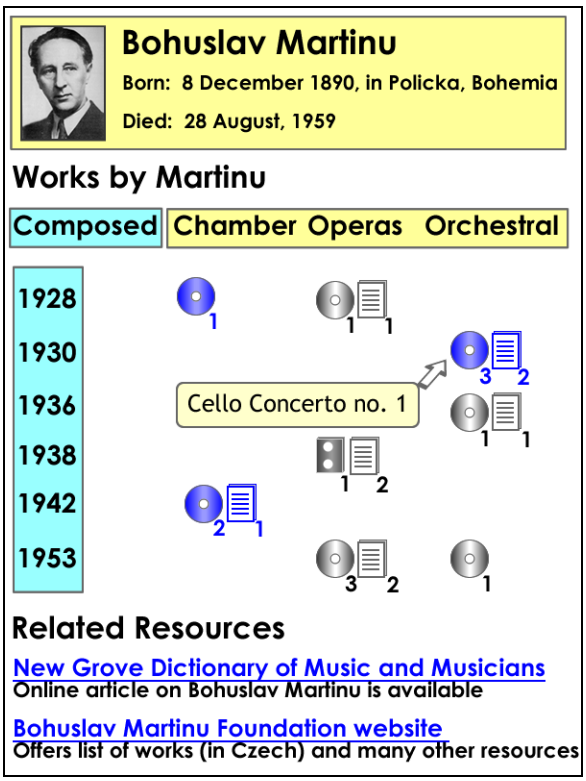


Figure 5. Composer Window

This visualization shows all of the digital library holdings by media type, for each work. The x-axis shows the musical genre and the y-axis the date of composition. The icons in this window can be solitary or compound. If, for example, the library holds a particular orchestral work of Martinu in the form of three recordings and two scores, a compound icon showing both the recording icon and the score icon, with a subscripted "3" and "2," are displayed. Icon color in this example highlights the works that were included in the previous display's search results. The gray icons show other works so that the searcher can see the search results in the broader context of the composer's output. Hovering over an icon (or compound icon) will show the work's title. Clicking on the icon will lead to the instantiation list (Figure 4).

**Example 2—Visual Reserve List**

Instructors provide music students with listening lists for the semester. Today, these lists are, at best, textual HTML pages with links to the containers (CDs, scores) which include the pieces of interest. In addition, instructors provide a syllabus, either on-line or on paper, which shows which pieces the student needs to listen to by week, sometimes with some additional information about the composer or piece. These documents are tedious to create and keep synchronized with each other.

With Variations2, instructors could provide a graphical timeline of listening assignments, showing composers, dates, nationalities and other details. The timeline view is of particular importance to music history classes. By clicking on the pieces, students can access pre-selected media or media excerpts (Figure 6).

The pieces can be shown by week in the semester. Weeks can be expanded or collapsed. When a week is expanded, students see icons representing the media type. In this example, some of the media is synchronized for listening and score viewing. The x-axis position of the item shows the date of composition. An interval bar shows the composer's lifespan. Color used to indicate the composer's nationality.

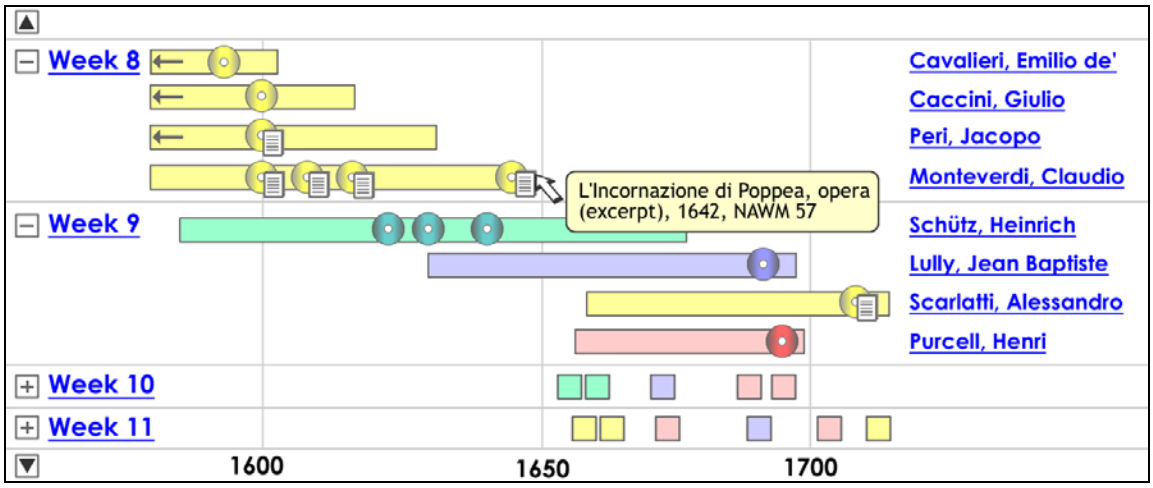


Figure 6. Timeline Visualization

## FUTURE WORK

Many visualization tools develop somewhat generically, based on technical capabilities and broad concepts. Our goal with Variations2 visualization is to base visualization tools on observed user needs. While we have some data from reference desk question logs and field observations of students and faculty, we have much more work to do to figure out the highest value visualization tools to provide. A second—but equally important—concern is to provide the tools in such a way that the casual user is not overwhelmed with options and extraneous capabilities. Basing tool design on field data should enable us to provide a minimal but useful set of visualization tools that will not require extensive user training.

Another future activity will be to experiment with additional existing visualization tools such as GRIDL.

An issue raised by these investigations is the contrast between a given library's holdings and the possible universe of data. For example, when we display a work in the context of a composer's other works (Figure 5), we are limited to displaying works listed in our catalog, which at best represents our total collection. However, far more comprehensive lists of composers, works, and containers exist. Ideally, we could display an item in the total context of a composer's output, a performer's discography, etc., calling out those items in our own collection. The FirstSearch service of OCLC's WorldCat [13] is a text-based step in this direction, letting users search for an item, view results from a union of collections, and see which ones are held by a particular library. A future direction of research would be to provide visualizations of collection results in the context of the universe of data.

As Variations2 progresses, we will be able to work with real data in the new data model. Our current prototypes suggest that we will be able to provide a variety of useful visualizations for users of digital music libraries.

## ACKNOWLEDGEMENTS

This material 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 views of the National Science Foundation. The authors also wish to thank Spotfire for access to their visualization software and Classical.Net for providing a large database of music data.

## REFERENCES

1. Ahlberg, C. And Shneiderman, B., Visual information seeking: Tight coupling of dynamic query filters with starfield displays. In *Proc. of CHI '94*, pp. 313-317.
2. Card, S.K., Mackinley, J.D. and Shneiderman, B. (eds.). *Readings in information visualization: using vision to think.* Morgan Kaufmann Publishers, San Francisco, 1999.
3. Classical.Net. See http://www.classical.net.
4. Dunn, J.W., and Mayer, C.A. VARIATIONS: A digital music library system at Indiana University. In *DL '99: Proceedings of the Fourth ACM Conference on Digital Libraries*, Berkeley, California, August 1999, pp. 12-19.
5. GRIDL: GRaphical Interface for Digital Libraries. See http://www.cs.umd.edu/hcil/west-legal/gridl/.
6. IUCAT. See http://www.iucat.iu.edu/.
7. Library of Congress. MARC Standards Website. See http://www.loc.gov/marc/.
8. Mounce, S.R. A brief discussion of Standard Music Description Language. ISO/IEC Draft International Standard 10743. See http://www.techno.com/smdl.htm.
9. Rao, R., and Card, S.K. The table lens: Merging graphical and symbolic representations in an interactive Focus+Context visualization for tabular information. In *Proc. of CHI'94*, pp. 318-322 and 481-482.
10. Spotfire. See http://www.spotfire.com.
11. Variations. See http://www.dlib.indiana.edu/variations/
12. Variations2: the Indiana University Digital Music Library Project. See http://dml.indiana.edu/.
13. WorldCat. See http://www.oclc.org/worldcat/