

Music Information Retrieval Annotated Bibliography Website Project, Phase I

J. Stephen Downie

Graduate School of Library and Information Science
University of Illinois at Urbana-Champaign

501 East Daniel St.
Champaign, IL 61820

jdownie@uiuc.edu

ABSTRACT

Music information retrieval (MIR) as a nascent discipline is blessed with a multi-disciplinary group of people endeavoring to bring their respective knowledge-bases and research paradigms to bear on MIR problems. Communication difficulties across disciplinary boundaries, however, threaten to impede the maturation of MIR into a full-fledge discipline. The principal causes of the communications breakdown among members of the MIR community are a) the lack of bibliographic control of the MIR literature; and, b) the use of discipline-specific languages and methodologies throughout that literature. This poster abstract reports upon the background, framework, goals and ongoing development of the *MIR Annotated Bibliography Website Project*. This project is being undertaken to specifically address and overcome these bibliographic control and communications issues.

1. INTRODUCTION

The problems associated with the creation, deployment, and evaluation of robust, large-scale, and content-based (i.e., music queries framed musically) music information retrieval (MIR) systems are far from trivial. Music information is inherently multi-faceted, multi-representational (i.e., can be represented in many different ways), multi-modal (i.e., experienced in many different ways), and multi-cultural. The complex interaction of Pitch, Temporal, Harmonic, Timbral, Editorial, Textual, Bibliographic, Representational, Experiential, and Cultural facets makes music information difficult to store, and then retrieve, in any robust, large-scale, and comprehensive manner. Simply put, this dizzyingly complex interaction *is* the “MIR problem”.

Because MIR is such a complex and multi-dimensional research problem, many diverse groups of scholars, researchers, and interested parties have begun to explore MIR issues within the frameworks of their particular disciplines. These groups include music and digital librarians, computer scientists, audio engineers, music publishers and retailers, musicologists, information retrieval specialists, intellectual property lawyers, music hobbyists, music psychologists, educators, Internet content providers, broadcasters, and business managers. Students, representing all the aforementioned disciplines, at levels ranging from undergraduate to post-doctorate, are also seeing MIR issues as fruitful and interesting areas of study.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.

2. THE PROBLEM

A recurring theme brought to the fore by participants at the *International Symposium on Music Information Retrieval* (23-25 October, 2000) was the ignorance many participants felt about MIR work being done in disciplines other than their own. This ignorance had two manifestations. First, participants of discipline *W* bemoaned the fact that they did not know that members of discipline *X* had been working on a given problem *Y* and publishing their findings on *Y* in the *X* literature for years. For example, computer scientists, and others, were not aware of the extensive musicology literature dealing with music representation codes. Second, participants of discipline *X* were distressed by their inability to fully comprehend, and thus evaluate fairly, the contributions being made by members of discipline *W* because the contributions of *W* were so deeply rooted in its discipline-specific language and methods. For example, the music librarians, and others, struggled with the audio engineering presentations because they did not have the educational background needed to evaluate the application of Fast Fourier Transforms and other highly mathematical techniques to a particular MIR problem.

The MIR corpus is scattered willy-nilly across the scholastic landscape with important papers found in the musicology, computer science, information retrieval, information science, and engineering literatures, to name but a few sources. Because of this scattering, it is nowhere uniformly represented in any one of the traditional indexing sources. For example, the musicology-based MIR work is found in various music, arts, and humanities, indexes but not necessarily in the computer science and engineering indexes. Similarly, important engineering-based papers are missing from the arts and humanities indexes, and so on. Since researchers are generally unaware of the differences in scope of the various discipline-based indexes, they tend to focus upon those with which they are most familiar and thus overlook the contributions of those based in other disciplines. Unfamiliarity with the wide-range of vocabularies used by the various disciplines further compounds the communication difficulties by making it problematic for MIR investigators to conduct thorough and comprehensive searches for MIR materials. Until these issues are addressed, MIR will never be in a position to fully realize the benefits that a multi-disciplinary research and development community offers, nor will it be able to develop into a discipline in its own right.

3. PROPOSED SOLUTION

The creation of a Web-based, two-level, collection of annotated bibliographies will overcome many of the communications problems currently plaguing the MIR community (Fig. 1). The first level, or core bibliography, will bring together those items identified as being germane to MIR as a nascent discipline. Thus, the core bibliography will comprise only those papers dealing specifically with some aspect of the MIR problem, such as MIR system development, experimentation, and evaluation, etc. The second level, or periphery bibliographies, will comprise a *set* of discipline-specific bibliographies. Each discipline-specific bibliography in the set will provide access to the discipline-specific background materials necessary for non-expert members of the other disciplines to comprehend and evaluate the papers from each participating discipline. For example, an audio engineering bibliography could be used by music librarians and others to understand the basics of signal processing (i.e., Fast Fourier Transforms, etc.). Another example would be a musicology bibliography that computer scientists could draw upon in an effort to understand the strengths and weaknesses of the various music encoding schemes, and so on. Thus, taken together, the two-levels of the MIR bibliography will provide:

- a) the much needed bibliographic control to the MIR literature; and,
- b) an important a mechanism for members of each discipline to comprehend the contributions of the other disciplines.

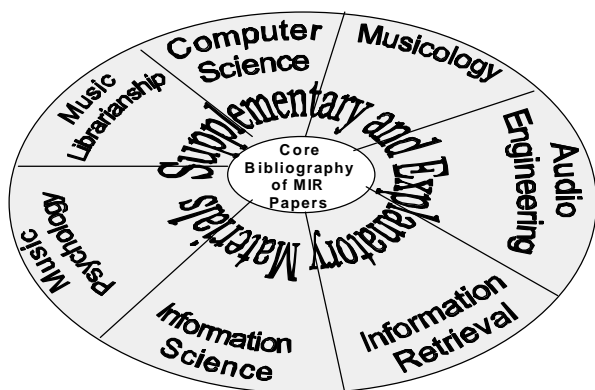


Figure 1. Project Schematic

4. PHASE I COMPONENTS

An important operating principle of the project is the use of non-proprietary formats and software. We are committed to the ideals of the Open Source Initiative [6] and the GNU General Public License [2] and thus intend to make our innovations freely available to others. In keeping with this commitment, we have chosen the *Greenstone Digital Library Software* (GSDL) package [5], the *Apache HTTP* server [1], the PERL scripting language [7], and the Linux operating system [4] to create the basic technological foundation of the project. We have purchased copies of the commercial bibliographic software package, *ProCite* [3] for initial, in-house, data-entry. *ProCite* also provides us with a representative instance of commercially available software that many end-users might utilize in manipulating the records they retrieve from our bibliography.

We have acquired the domain name *music-ir.org* under which access to the bibliography will be located (<http://www.music-ir.org>). At present, there are two central components of project undergoing development and alpha testing:

- a) the bibliographic search and retrieval interface using the GSDL package; and,
- b) the Web-based end-user data entry system.

For both of these, the goal is to create a system that will permit ongoing viability of the bibliography by minimizing—but not necessarily eliminating—the amount of human editorial intervention required. Item A issues being addressed include modifications to the basic GSDL system to permit the importation of specially structured bibliographic records and their subsequent access through a variety of field selection options. Item B is a CGI-based input system that guides end-users through the process of constructing well-structured bibliographic records through a series of step-by-step interactions and the on-the-fly generation of input forms specifically designed to provide the appropriate fields for the various types of bibliographic source materials (i.e., journal articles, conference papers, theses, etc.).

4.1 Next Steps

Now that the general framework for the core bibliography has been laid, we are moving forward on the acquisition of the supplementary and explanatory materials. For these we will be drawing upon the expert advice of those that have graciously signed on as advisors to the project. These advisors will not only lend their disciplinary expertise but will also afford us a very important multinational perspective on the potential uses and growth of the bibliography.

5. ACKNOWLEDGMENTS

I would like to thank Dr. Suzanne Lodato for helping the project obtain its principal financial support from the Andrew W. Mellon Foundation. I would also like to thank Dr. Radha Nandkumar and the National Center for Supercomputing Applications Faculty Fellows Programme for their support. The hard work and insightful advice of Karen Medina, Joe Futrelle, Dr. David Dubin, and the other members of the Information Science Research Laboratory at the Graduate School of Library and Information Science, UIUC, is gratefully appreciated. Those members of the MIR research community who have volunteered to act as project advisors are also thanked.

6. REFERENCES

- [1] Apache Software Foundation. 2001. *Apache project*. Available at: <http://httpd.apache.org/> .
- [2] GNU Project. 2001. *Free software licenses*. Available at: <http://www.gnu.org/licenses/licenses.html>
- [3] ISI Researchsoft. 2001. *ProCite, Version 5*. Available at: <http://www.procite.com> .
- [4] Linux Online. 2001. *The Linux homepage*. Available at: <http://www.linux.org/> .

- [5] New Zealand Digital Library Project. 2001. *About the Greenstone software*. Available at: <http://www.nzdl.org/cgi-bin/library?a=p&p=gsdl> .
- [6] Open Source Initiative. 2001. *The open source definition, version 1.9*. Available at: <http://www.opensource.org/docs/definition.html>
- [7] Perl Mongers. 2001. *Perl Mongers: The PERL advocacy people*. Available at: <http://www.perl.org/> .