

Interdisciplinary Communities and Research Issues in Music Information Retrieval

Joe Futrelle

Graduate School of Library and Information Science
University of Illinois
Champaign, IL 61820
+1 217 265 0296
futrelle@uiuc.edu

J. Stephen Downie

Graduate School of Library and Information Science
University of Illinois
Champaign, IL 61820
+1 217 265 5018
jdownie@uiuc.edu

ABSTRACT

Music Information Retrieval (MIR) is an interdisciplinary research area that has grown out of the need to manage burgeoning collections of music in digital form. Its diverse disciplinary communities have yet to articulate a common research agenda or agree on methodological principles and metrics of success. In order for MIR to succeed, researchers need to work with real user communities and develop research resources such as reference music collections, so that the wide variety of techniques being developed in MIR can be meaningfully compared with one another. Out of these efforts, a common MIR practice can emerge.

1. INTRODUCTION

Music Information Retrieval (MIR) is a rapidly growing interdisciplinary research area encompassing computer science and information retrieval, musicology and music theory, audio engineering and digital signal processing, cognitive science, library science, publishing, and law. Its agenda, roughly, is to develop ways of managing collections of musical material for preservation, access, research, and other uses. In this way it resembles traditional library science, and indeed, libraries have historically led the development of music collections. The idea of applying automatic information retrieval (IR) techniques to music dates back to the 1960's (Kassler 1966). But in particular, MIR has grown recently out of an explosion of interest in networked collections of musical material *in digital form*, precipitated by the development of compression technologies such as mp3, online services such as Napster, advances in optical musical recognition (OMR), and the ever-plummeting costs of digital storage and bandwidth. In this sense MIR is closely related to Digital Libraries.

As in other interdisciplinary fields, discourse in MIR is impeded at disciplinary boundaries by unfamiliar jargon, differing methodology, and even philosophical and ethical differences. To understand the field, it is currently necessary to acquire at least a cursory understanding of each of the disciplines, and MIR researchers are undertaking this as they begin to develop a common practice (Downie 2001).

This paper investigates MIR's interdisciplinary communities and research issues by surveying the proceedings of the International Symposia on Music Information Retrieval (ISMIR 2000 and 2001). The ISMIR series of meetings is an explicit attempt to gather together all of the disciplines and research areas pertinent to MIR. ISMIR can also claim to be the only conference series exclusively devoted to the advancement of MIR research. For

these reasons, we believe the proceedings of ISMIR provide a representative "snapshot" of the major issues comprising MIR research and development. Based on these proceedings and with reference to components of their supporting literatures, this paper will characterize the field, outline the major research communities involved in the field, assess the state of the art in each community, identify coverage gaps, and propose a research agenda aimed at addressing those gaps.

2. WHAT IS MIR RESEARCH?

What are MIR researchers trying to build and what problems are they trying to solve? MIR researchers often characterize their motivations by pointing out that the increasing volume of digital music available necessitates new retrieval techniques (Durey et al. 2001; Hoos et al. 2001; Kornstädt 2001; Yang 2001). However, the lack of effort to assess this volume, its rate of growth, and/or compare its rate of growth against the cost of bandwidth, storage, and processing power, and the relative scarcity of research such as Jang et al. (2001) focusing on scaling existing techniques, indicate that MIR researchers are primarily concerned with larger, more fundamental problems.

MIR researchers understand that the increasing availability of digital music is merely an aggravating factor of a more significant issue: few effective retrieval techniques exist for digital music collections. The problem has existed since music was first encoded digitally, but has become pressing only recently as the cost of storing large digital music collections has dropped to almost nothing and the number of such collections has consequently exploded. For MIR, developing effective retrieval techniques is basic research, which continues to advance on a number of interrelated fronts.

Developing IR techniques for music is challenging because of the wide variety of ways music is produced, represented, and used (Smiraglia 2001). Basic research in MIR can be categorized roughly by the kind of music representation employed.

Table 1 shows some representations and the kinds of MIR research being applied to them.

In addition to the variety of music representations, their complexity presents a problem as well. Like language, music in virtually all of its representations contains difficult-to-extract layers of significance, such as harmony, polyphony, and timbre. Even the most robust representations still require sophisticated processing techniques to extract some of these features, and developing these techniques is an active area of MIR research. This area is often called "content-based" MIR, to distinguish it from more traditional digital and pre-digital approaches based on manually-produced metadata of bibliographic and related varieties.

As basic MIR research begins to produce results, it raises the questions of what kind of MIR systems can be built, what their

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.
© 2002 IRCAM – Centre Pompidou

user interfaces would be, and what their institutional or economic contexts would be. These questions occupy a broad swath of MIR research, ranging from system architecture and testbeds (such as VARIATIONS (Dunn et al. 1999; Dunn 2000; Dunn et al. 2001), MELDEX (Bainbridge et al. 1999; Bainbridge 2000), and the Levy Sheet Music Collection (Choudhury et al. 2000)) to Intellectual Property rights (Levering 2000).

Table 1: Music Representations in MIR

Representation	Description	Research
Symbolic	Notation (scores, charts), Event-based recordings (MIDI), Hybrid representations	Matching, Theme/Melody Extraction, Voice Separation, Musical Analysis
Audio	Recordings, Streaming Audio, Instrument Libraries	Sound/Song Spotting, Transcription, Timbre Classification, Musical Analysis
Visual	Scores	Score Reading (“Optical Music Recognition”)
Metadata	Cataloging, Bibliography, Descriptions	Library Testbeds, Traditional IR, Interoperability

3. MIR COMMUNITIES

Understanding a multi-disciplinary field requires understanding the disciplines involved and the variety of research interests they represent. Table 2 summarizes the research communities involved in MIR along with their typical home institutional settings and typical areas of research.

Table 2: MIR Communities

Community	Type of Institution(s)	Research Areas
Computer Science, Information Retrieval	Academic, Commercial	Representation, Indexing, Retrieval, Machine Learning, User Interface Design
Audio Engineering, Digital Signal Processing	Academic, Commercial	Compression, Feature Detection, Pitch Tracking, Machine Learning, Classification, Musical Analysis
Musicology, Music Theory	Academic	Representation, Musical Analysis
Library Science	Libraries, Academic	Representation, Metadata, User Studies, Classification, Intellectual Property Rights, User Interface Design
Cognitive Science, Psychology, Philosophy	Academic	Representation, Perception, User Studies, Ontology
Law	Government, Legal Profession, Academic	Intellectual Property Rights

Table 3 describes major research areas in MIR, along with basic research questions and exemplar papers in those areas.

Table 3: MIR Research Areas

Research Area	Description
Representation	How should musical material be represented in digital form? What aspects of music are critical to represent for the purpose of building music collections? At what level of granularity can we represent music? What kinds of representation are the most efficient? How can markup languages be applied to music? (Good 2000; Hoos et al. 2001; Lindsay et al. 2001; Maidin et al. 2001)
Indexing	How can database indexing techniques be applied to musical material so it can be retrieved effectively and efficiently? (Downie 1999; Chen 2000; Pickens 2000)
Retrieval	What kinds of queries can we perform on indexed collections of musical material? How can the performance of these queries be evaluated and improved? (Hsu et al. 2001; Lemström et al. 2001)
User Interface Design	How can user interfaces be built which enable users to effectively find and use digital musical material from a collection? (Fernström et al. 2001; Kornstädt 2001; MacMillan et al. 2001)
Compression	How can audio be encoded more efficiently? What are the implications for MIR of various emerging compression technologies? (Lindsay et al. 2001)
Feature Detection	How can distinguishing features of music be detected from audio signals? How can these techniques be applied to MIR systems? (Foote 1999; Ismirli 2000; Logan 2000; Nam et al. 2001; Tzanetakis et al. 2001)
Machine Learning	How can we deduce or induce aggregate musical features of collections, so that they can be organized for retrieval? What are the most efficient and effective ways of representing these aggregate features? (cf. Classification)
Classification	What kinds of classification techniques and schemes can/should be applied to digital music collections? (cf. Machine Learning) (Herrera et al. 2000; Reiss et al. 2001)
Musical Analysis	How is a musical composition organized? How is it similar to, or different from, other pieces of music? How can MIR systems meet the needs of musicologists? (Bonardi 2000; Larson 2000; Cope 2001; Kornstädt 2001)
Metadata	What kinds of descriptive or contextual information about musical material can and/or should be managed by an MIR system, and how should such metadata be represented? (cf. Representation) (Choudhury et al. 2000; Smiraglia 2001)
User Studies	What kinds of MIR capabilities do users need? How do users search for musical material, and why? What would be the ideal MIR system for a given user community? (Itoh 2000; Selfridge-Field 2000; McPherson et al. 2001)

Intellectual Property Rights	Who owns musical material? Under what conditions? Under what arrangement can digital libraries of musical material and owners of IP rights for musical material peacefully coexist? (Levering 2000)
Perception	How do people perceive music? How can music perception inform the design of MIR systems? What is music? How is musical similarity perceived? (cf. Ontology) (Perrot et al. 1999; Huron 2000; Hofmann-Engl 2001)
Epistemology/Ontology	What is music? What is a musical composition? What is the relationship between different representations of “the same” piece? How do improvised aspects of music relate to composed aspects with respect to collections of musical material? (Smiraglia 2001)

3.1 Computer Science, Information Retrieval

Of course, virtually all MIR research employs techniques from computer science. But there is an important subset of MIR research whose origins can be traced back to the Information Retrieval research on bibliographic text retrieval systems in the early 1960’s. This ongoing research emphasizes techniques for locating items in a collection or index which match a query, rather than techniques for analyzing aggregate properties of collections of items (e.g. data mining). For an overview of traditional IR see Baeza-Yates et al. (1999). MIR research based on traditional IR is typically aimed at supporting a scenario in which users know characteristics of the music they desire, and use an MIR system to locate musical material that most closely matches those characteristics. Downie (1999) calls this a “locating” MIR system.

In traditional IR, a query on a collection can be thought of as a fragment or reduced form of the desired item from the collection. For text collections, the query is often a word that occurs in the desired document. MIR researchers have taken some pains to devise MIR systems which fit this model, most notably Downie (1999) who reduces the music in his collection to n-grammed sequences of intervals, which can then be indexed using inverted files. Other research on locating MIR systems uses other traditional IR strategies, such as probabilistic modeling (Pickens 2000) and approximate string matching (Lemström et al. 2001). Some work also addresses IR issues such as relevance and ranking (Uitdenbogerd 2000). Most of this research is based on symbolic music representations, but it has also been applied to audio which is pre-processed and converted to symbolic sequences of audio feature classes as in Aucouturier (2001) and Batlle et al. (2000).

A great deal of attention has been paid to so-called “Query by Humming” (QBH) systems, which retrieve pieces based on melodic fragments sung by the user. QBH systems typically combine melody extraction from an audio query with a locating MIR system to match the melody against a target database. In the ISMIR 2001 proceedings, 10 out of 43 papers, posters, and talks (23%) concerned QBH systems.

3.2 Audio Engineering, Digital Signal Processing

A major category of MIR research concerns audio representations of music (i.e. recordings, audio streams, or live performance). Techniques used in this area grow out of decades of work in Digital Signal Processing (DSP) and speech recognition. A good overview of these techniques is found in Foote (1999). The techniques are applied to an interrelated set of problems:

- What are the most important features of audio representations of music for MIR, and how can they be extracted from audio? (e.g. melodies, harmonies, instrument timbres, etc.)
- Given a set of features extracted from audio, what techniques can be used to understand the relationships between those features in an audio collection?
- How can we use audio to perform structural analyses of music, and how can these be used to improve MIR systems?

A number of audio features have been used in MIR research. Virtually every audio MIR system uses some kind of frequency-domain transformation of the signal, such as the Fast Fourier Transform (FFT) (Brigham 1988) or its more musically-relevant derivative, mel-frequency cepstral coefficients (MFCC) (Logan 2000). A number of other features are used, including time-domain autocorrelation (mostly used for pitch tracking) and wavelet transforms (Tzanetakis et al. 2001), but most of the features employed are statistics computed from the FFT. The reason frequency-domain transformations are so prevalent is the primacy of periodicity in the perception of musical aspects such as pitch, timbre, and rhythm.

Once a set of features is selected and can be reliably extracted from audio, the problem is essentially one of multivariate analysis, in which each piece of music in a collection can be conceptualized as a vector in n-dimensional feature space. Traditional multivariate techniques as well as probabilistic machine-learning techniques such as Hidden Markov Models (HMM) and neural networks can be applied to identify salient features and perform data reduction, often through classification. A good overview of these and other classification techniques for audio can be found in Herrera et al. (2000).

3.3 Musicology, Music Theory

The study of music is an important application area for MIR, and thus drives much of MIR research. Musicology is an ancient and interdisciplinary field which has been transformed by computational techniques (Bel et al. 1993) and promises to be further transformed by ready access to large digital music collections. Musicology-related MIR research ranges from computational music analysis, as in Cope (2001), to MIR systems specialized for musicologists, as in Bonardi (2000) and Kornstädt (2001). An interesting tension exists in MIR between musical analysis which concerns the nature of music *per se* and is often qualitative as in Larson (2000), and approaches that attempt to empirically demonstrate improved retrieval performance and thus rely on quantitative techniques such as statistical analysis and machine learning. Cope’s work with computational musical analysis (Cope 2001) and algorithmic composition (Cope 1992) interestingly bridges these two very different perspectives by using computational techniques to divine aspects of music, such as style, which have traditionally been investigated qualitatively or with exhaustive manual effort as in Van der Merwe (1989). There is clearly some middle ground that remains unexplored, since aspects of music that have currently only been characterized by musicologists may yet prove useful in the design of MIR systems.

3.4 Library and Information Science

Libraries and library scientists are involved in MIR as part of their ongoing effort to cope with ballooning multimedia collections. Libraries face all of the issues raised by MIR, from basic research questions such as how to represent and index musical material, to applied information technology issues such as integrating traditional bibliographic systems with advanced MIR tools, to

policy issues such as how to manage intellectual property rights for the producers and users of music collections. Of particular importance to MIR are a number of testbed projects being undertaken at academic libraries and digital library research facilities, including Indiana's Digital Music Library projects (Dunn et al. 1999; Dunn 2000; Dunn et al. 2001), the University of Waikato's MELDEX digital library (Bainbridge et al. 1999; Bainbridge 2000; McPherson et al. 2001), and the Levy sheet music collection at Johns Hopkins (Choudhury et al. 2000). These testbeds integrate a variety of MIR tools with significant music collections in order to explicitly address issues such as usability, scale, and multi-modal access to musical works (e.g. linking scores with recordings). They also support the application of MIR tools and collections to specialized use cases such as music theory education.

Testbeds are a good way to begin to evaluate who the potential users of MIR systems are and what features they are most interested in. Preliminary user studies such as McPherson et al. (2001) indicate a trend away from speculative user requirements analysis such as that found in, for example, Bonardi (2000) towards empirically grounded approaches and techniques. The user modeling methods put forward by Rolland (2001) suggest that future MIR systems can be tailored to meet the needs of a variety of user communities.

3.5 Cognitive Science, Psychology, Philosophy

A small subset of MIR research concerns the implications of music perception on the design of MIR systems. Research efforts range from models of music perception such as Dannenberg (2001) and Hofmann-Engl (2001) to epistemological analysis of music information such as Smiraglia (2001). Much research has been done on music perception in psychology, music psychology and cognitive science (Deliège et al. 1997; Cook 1999). There also have been notable efforts in both music philosophy (Adorno 1973) and cultural studies (Attali 1985; McClary 1991) to characterize how music is understood as a social and cultural phenomenon. Significantly, however, MIR researchers have so far rarely adopted work in these areas as a basis for MIR studies.

3.6 Law

High-profile cases such as Napster demonstrate that MIR systems are being developed in an uncertain regulatory environment, and legal issues will continue to be important to MIR researchers until this situation changes. Issues such as the Digital Millennium Copyright Act (Levering 2000), intellectual property rights management, and researcher access to music databases (Byrd 2001; Downie 2001) will be important to the field indefinitely. To a large extent, copyright law is a policy issue rather than a technical issue, but legal issues dramatically affect the priorities of commercial and non-commercial agencies funding MIR research, and thus are of critical importance.

4. CRITICAL ANALYSIS OF COVERAGE GAPS IN MIR RESEARCH

MIR's newness and multi-disciplinary constituency make the field strong on innovation and basic research, but weak on evaluation and application to real user communities. The problem is twofold:

1. There are no commonly accepted means of comparing the efficacy of retrieval techniques; and,
2. There have been few if any attempts to study potential users of MIR systems to find out what they need.

The two problem areas are interrelated in that meaningful evaluation of retrieval techniques must be grounded in a significant understanding of user requirements.

In addition to these evaluation-related gaps, there are also areas of basic research that are receiving more and less attention than they should. In particular, the amount of emphasis on QBH systems appears to be unsupportable given doubts about their usefulness (McPherson et al. 2001) and scalability (Sorsa et al. 2001). Research on recommendation systems, common in the DL and commercial communities, is inexplicably rare in the MIR community. User interface research, now undertaken most often as an afterthought to research into retrieval techniques, is clearly under-emphasized, especially since retrieval interfaces may have to incorporate complex audio strategies such as those explored by Fernström et al. (2001). And finally, MIR research as a whole has failed to significantly address music outside of the common-practice Western music canon.

4.1 Difficult To Compare Techniques

Research into MIR techniques rarely presents results that can be compared with other research. Some studies such as Spevak et al. (2001) and Rolland (2001) do not present evaluation results at all. Others present results based on very small sample sets. For example, the particular technique of Mazzoni et al. (2001) is evaluated on a database with less than 20 pieces of music. Other papers report overall results without reference to any common measure of significance: for example Yang (2001) reports "90% retrieval accuracy" without explaining what constitutes "accurate" retrieval, Nishimura et al. (2001) report a "search rate" computed by averaging precision and recall together, etc.

We believe this inconsistency largely arises from MIR's interdisciplinary nature. Evaluation metrics that are well understood in one field, such as precision and recall in traditional IR, are new and unfamiliar to other fields such as audio engineering. In addition, there are no community-wide music collections against which researchers can cross-evaluate a wide variety of different techniques, a problem which the community is eager to address (Byrd 2001; Downie 2001). So far, the most sophisticated attempts to rigorously compare a variety of MIR techniques are being done in limited domains as in Hsu et al. (2001) and Uitenbogerd (2000).

4.2 Few Attempts to Assess User Requirements

Jef Raskin's talk at ISMIR 2001 about how to make computer systems more usable was notable in that it suggested using theoretical models of users to guide user interface design choices rather than involving users in the design process (Raskin 2001). This emphasis on basic research over application to, and involvement with, users is common in MIR research, and may result from the influence of the computer science and audio engineering communities.

Already, MIR is beginning to emphasize certain areas of research without having identified user communities and evaluated whether the techniques developed will meet the needs of those communities. As mentioned before, QBH systems are being intensively developed, but there is virtually no evidence cited that users prefer these systems, and even some that suggests that they do not (McPherson et al. 2001). QBH papers typically begin with speculations that such a system would be useful, such as "singing is naturally used as input" (Haus et al. 2001), or anecdotal evidence such as

The potential utility of such systems is attested to by music librarians, who report that library patrons often hum or whistle a phrase of music and ask them to identify the corresponding musical work. (Smith et al. 2001).

Even MIR research focusing on usability rarely involves user studies. For instance, Kornstädt (2001) presents a graphical user interface apparently tailored for the needs of musicologists, but cites neither research into what kinds of tools musicologists need nor any evaluation of the system by musicologists. Research into genre classification from audio (e.g., Tzanetakis et al. (2001)) repeatedly cites Perrot et al. (1999) as a basis for its assumption that users can make effective genre judgments based on short musical examples, but does not evaluate the performance of genre classification techniques against real users' genre judgments. One exception to the dearth of user studies is Fernström et al. (2001) which at least attempts a preliminary user evaluation of a new user interface design based on a perceptual phenomenon known as the "cocktail party effect".

To some extent, this is a chicken-and-egg problem; MIR researchers cannot evaluate techniques that have not yet been developed. Even worse, it is difficult to meaningfully study user behavior without users having access to large, relatively comprehensive collections of music with which they can spontaneously interact. As testbed projects continue to develop, they will be in the best position to analyze their own users (e.g., McPherson et al. (2001) and Dunn (2000)), which should provide valuable guidance to basic researchers. However, it still remains to be seen whether or not real user needs match the interests or technological capabilities of the many disciplines currently involved in MIR research.

4.3 Undue Emphasis on Western Music

Another significant coverage gap in MIR concerns non-Western music. The music used in MIR studies is predominantly common-practice Western music. This is primarily a problem with symbolic MIR systems, which tend to use representations derived from common-practice Western music notation. Notable exceptions include Linardis et al. (2001) who describe a retrieval system based on Byzantine neumatic notation. Audio MIR systems are presumably more flexible than symbolic MIR systems because the audio representations and features they employ are presumably more culturally neutral, but no audio MIR research has specifically investigated this hypothesis.

Addressing the undue emphasis on common-practice Western music in MIR research requires, finally, a radical rethinking of MIR research practice. Assumptions commonly made by MIR researchers about music – that it has melodies, that its rhythm is metrical, and that it can be treated as re-contextualizable information objects – must be replaced by provisional sets of assumptions resonant with the cultural milieu of real user communities.

5. FIRST PRINCIPLES AND A MIR RESEARCH AGENDA

To best overcome the gaps in current MIR research, we believe that MIR research must embrace as *sine qua non* the following three principles:

1. MIR systems are developed to serve the needs of particular user communities.
2. MIR techniques are evaluated according to how well they meet the needs of user communities.
3. MIR techniques are evaluated according to agreed-upon measures against agreed-upon collections of data, so that meaningful comparisons can be made between different research efforts.

To realize these first principles, we believe the following steps must be taken:

First, MIR research should begin by assessing existing MIR systems (broadly defined to include both digital and traditional formats), including libraries, music retailers, on-line media merchants, and individual collections. Existing practice should be evaluated to establish baselines of usability upon which new MIR systems must improve. These evaluations should be systematic and empirical and involve the participation of both users and maintainers of existing systems, rather than being based on the opinions and speculations of MIR researchers. These same techniques can then be used to evaluate new MIR systems, and the results can be compared. It is imperative that studies of existing music systems include non-Western resources and their use in non-Western contexts.

Second, the investigation of existing MIR practice should identify distinct user communities and investigate what they need from MIR systems. As MIR researchers have already pointed out, the musicological community has, for example, quite different requirements for MIR systems than other communities. Future MIR research should explicitly identify which community's needs it is attempting to address.

Third, MIR research programs should also agree upon evaluation measures. Retrieval accuracy and system effectiveness should be measured using clearly delineated, agreed-upon methodologies and reported consistently across studies. To this end, MIR researchers should share music collections, so that a variety of techniques can be applied to the same collection and results replicated or refuted by independent research teams. The development of a set of "universal" test collections as called for by the ISMIR 2001 "resolution" (see <http://music-ir.org/mirbib2/resolution>) is an important step in this direction.

Fourth, and finally, MIR researchers should develop under-represented research areas such as recommendation, browsing, and user interface design. Advances into these areas should be grounded in existing knowledge of user requirements and music perception.

If the MIR research community embraces these principles and this research agenda, we believe that future MIR systems will better provide real users with the tools, features, and ease-of-use they need to get the most out of rich and comprehensive collections of music in digital form.

6. REFERENCES

- Adorno, T. W. (1973). Philosophy of Modern Music. New York, Seabury Press.
- Attali, J. (1985). Noise: the political economy of music. Minneapolis, University of Minnesota Press.
- Aucouturier, J.-J. and M. Sandler (2001). Using Long-Term Structure to Retrieve Music: Representation and Matching. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 1-2.
- Baeza-Yates, R. and B. d. A. N. Ribeiro (1999). Modern information retrieval. New York, ACM Press.
- Bainbridge, D. (2000). The Role of Music IR in the New Zealand Digital Library Project. International Symposium on Music Information Retrieval.
- Bainbridge, D., C. G. Nevill-Manning, et al. (1999). Towards a Digital Library of Popular Music. Fourth ACM Conference on Digital Libraries, Berkeley, California.
- Batlle, E. and P. Cano (2000). Automatic Segmentation for Music Classification using Competitive Hidden Markov Models. International Symposium on Music Information Retrieval.

Interdisciplinary Communities and Research Issues in Music Information Retrieval

- Bel, B. and B. Vecchione (1993). "Introduction." Computers and the Humanities 27(1 -- special issue on Computational Musicology).
- Bonardi, A. (2000). IR for Contemporary Music: What the Musicologist Needs. International Symposium on Music Information Retrieval.
- Brigham, E. O. (1988). The Fast Fourier Transform and Its Applications. Englewood Cliffs, New Jersey, Prentice-Hall.
- Byrd, D. (2001). Music-Notation Searching and Digital Libraries. ACM/IEEE Joint Conference on Digital Libraries, Roanoke, VA: 239-246.
- Chen, A. L. P. (2000). Music Representation, Indexing and Retrieval at NTHU. International Symposium on Music Information Retrieval.
- Choudhury, G. S., T. DiLauro, et al. (2000). Optical Music Recognition System within a Large-Scale Digitization Project. International Symposium on Music Information Retrieval.
- Cook, P. R. (1999). Music, cognition, and computerized sound : an introduction to psychoacoustics. Cambridge, Mass., MIT Press.
- Cope, D. (1992). "Computer modeling of musical intelligence in EMI." Computer Music Journal 16(2): 69-83.
- Cope, D. (2001). Computer Analysis of Musical Allusions. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 83-84.
- Dannenbergh, R. B. (2001). Music Information Retrieval as Music Understanding. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 139-142.
- Deliège, I. and J. A. Sloboda (1997). Perception and cognition of music. Hove, East Sussex, Psychology Press.
- Downie, J. S. (1999). Evaluating a Simple Approach to Music Information Retrieval: Conceiving Melodic N-grams as Text. Graduate Program in Library and Information Science. London, Ontario, University of Western Ontario: 179.
- Downie, J. S. (2001). Music Information Retrieval Annotated Bibliography Website Project, Phase I. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 5-7.
- Downie, J. S. (2001). Whither Music Information Retrieval: Ten Suggestions to Strengthen the MIR Research Community. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 219-222.
- Dunn, J. W. (2000). Beyond VARIATIONS: Creating a Digital Music Library. International Symposium on Music Information Retrieval.
- Dunn, J. W., M. W. Davidson, et al. (2001). Indiana University Digital Music Library Project: An Update. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 137-138.
- Dunn, J. W. and C. A. Mayer (1999). VARIATIONS: A Digital Music Library System at Indiana University. Fourth ACM Conference on Digital Libraries, Berkeley, California.
- Durey, A. S. and M. A. Clements (2001). Melody Spotting Using Hidden Markov Models. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 109-117.
- Fernström, M. and D. Ó. Maidín (2001). Computer-supported Browsing for MIR. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 9-10.
- Foote, J. (1999). "An Overview of Audio Information Retrieval." ACM Multimedia Systems 7(1): 2-11.
- Good, M. (2000). Representing Music Using XML. International Symposium on Music Information Retrieval.
- Haus, G. and E. Pollastri (2001). An Audio Front End for Query-by-Humming Systems. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 65-72.
- Herrera, P., X. Amatriain, et al. (2000). Towards instrument segmentation for music content description: a critical review of instrument classification techniques. International Symposium on Music Information Retrieval.
- Hofmann-Engl, L. (2001). Towards a Cognitive Model of Melodic Similarity. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 143-151.
- Hoos, H. H., K. Renz, et al. (2001). GUIDO/MIR -- an Experimental Musical Information Retrieval System based on GUIDO Music Notation. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 41-50.
- Hsu, J.-L. and A. L. P. Chen (2001). Building a Platform for Performance Study of Various Music Information Retrieval Approaches. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 153-162.
- Huron, D. (2000). Perceptual and Cognitive Applications in Music Information Retrieval. International Symposium on Music Information Retrieval.
- Ismirli, O. (2000). Using a Spectral Flatness Based Feature for Audio Segmentation and Retrieval. International Symposium on Music Information Retrieval.
- Itoh, M. (2000). Subject search for music: Quantitative analysis of access point selection. International Symposium on Music Information Retrieval.
- Jang, J.-S. R., J.-C. Chen, et al. (2001). MIRACLE: A Music Information Retrieval System with Clustered Computing Engine. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 11-12.
- Kassler, M. (1966). "Toward Musical Information Retrieval." Perspectives of New Music 4(2): 59-67.
- Kornstädt, A. (2001). The JRing System for Computer-Assisted Musicological Analysis. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 93-98.
- Larson, S. (2000). Searching for Meaning: Melodic Patterns, Combinations, and Embellishments. International Symposium on Music Information Retrieval.
- Lemström, K., G. A. Wiggins, et al. (2001). A Three-Layer Approach for Music Retrieval in Large Databases. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 13-14.
- Levering, M. (2000). Intellectual Property Rights in Musical Works. International Symposium on Music Information Retrieval.
- Linardis, P., D. Politis, et al. (2001). Musical Information Retrieval for Delta and Neumatic Systems. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 23-24.
- Lindsay, A. and Y. Kim (2001). Adventures in Standardization, or, how we learned to stop worrying and love MPEG-7. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 195-196.
- Logan, B. (2000). Mel Frequency Cepstral Coefficients for Music Modelling. International Symposium on Music Information Retrieval.
- MacMillan, K., M. Droettboom, et al. (2001). Gamma: a Structured Document Recognition Application

Interdisciplinary Communities and Research Issues in Music Information Retrieval

- Development Environment. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 15-16.
- Maidín, D. Ó. and M. Cahill (2001). Score Processing for MIR. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 59-64.
- Mazzoni, D. and R. B. Dannenberg (2001). Melody Matching Directly from Audio. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 17-18.
- McClary, S. (1991). Feminine endings: music, gender, and sexuality. Minneapolis, University of Minnesota Press.
- McPherson, J. R. and D. Bainbridge (2001). Usage of the MELDEX Digital Music Library. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 19-20.
- Nam, U. and J. Berger (2001). Addressing the "Same but different - different but similar" Problem in Automatic Music Classification. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 21-22.
- Nishimura, T., H. Hashiguchi, et al. (2001). Music Signal Spotting Retrieval by a Humming Query Using Start Frame Feature Dependent Continuous Dynamic Programming. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 211-218.
- Perrot, D. and R. O. Gjerdingen (1999). Scanning the Dial: an Exploration of the Factors in the Identification of Musical Style. Society for Music Perception and Cognition, Evanston, IL: 88.
- Pickens, J. (2000). A Comparison of Language Modeling and Probabilistic Text Information Retrieval Approaches to Monophonic Music Retrieval. International Symposium on Music Information Retrieval.
- Raskin, J. (2001). Making Machines Palatable. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 39-40.
- Reiss, J., J.-J. Aucouturier, et al. (2001). Efficient Multidimensional Searching Routines for Music Information Retrieval. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 163-171.
- Rolland, P.-Y. (2001). Adaptive User Modeling in a Content-Based Music Retrieval System. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 27-30.
- Selfridge-Field, E. (2000). What Motivates a Musical Query? International Symposium on Music Information Retrieval.
- Smiraglia, R. P. (2001). Musical Works as Information Retrieval Entities: Epistemological Perspectives. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 85-91.
- Smith, L. and R. Medina (2001). Discovering Themes by Exact Pattern Matching. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 31-32.
- Sorsa, T. and J. Huopaniemi (2001). Melodic Resolution in Music Retrieval. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 33-34.
- Spevak, C. and R. Polfreman (2001). Sound Spotting -- A Frame-Based Approach. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 35-36.
- Tzanetakis, G., G. Essl, et al. (2001). Automatic Musical Genre Classification of Audio Signals. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 205-210.
- Uitdenbogerd, A. L. (2000). Music IR: Past, present, and future. International Symposium on Music Information Retrieval.
- Van der Merwe, P. (1989). Origins of the popular style : the antecedents of twentieth-century popular music. Oxford Oxfordshire, Clarendon Press.
- Yang, C. (2001). Music Database Retrieval Based on Spectral Similarity. International Symposium on Music Information Retrieval, Bloomington, IN, USA: 37-38.