

Toward a Theory of Music Information Retrieval Queries: System Design Implications

J. Stephen Downie

Graduate School of Library and Information Science
University of Illinois at Urbana-Champaign
Room 103 LIS Building
Champaign, Illinois, USA
Phone: +1(217) 351-5037

jdownie@uiuc.edu

Sally Jo Cunningham

Department of Computer Science
University of Waikato
Private Bag 3105
Hamilton, New Zealand
Phone: 64-7-838-4402

sallyjo@cs.waikato.ac.nz

ABSTRACT

This paper analyzes a set of 161 music-related information requests posted to the *rec.music.country.old-time* newsgroup. These postings are categorized by the types of detail used to characterize the poster's information need, the type of music information requested, the intended use for the information, and additional social and contextual elements present in the postings. The results of this analysis suggest that similar studies of 'native' music information requests can be used to inform the design of effective, usable music information retrieval interfaces.

1. INTRODUCTION

Interest in the development of content-based music information retrieval (MIR) systems is growing rapidly. The MIR research community consists of a multidisciplinary amalgam of librarians, digital librarians, information scientists, computer scientists, musicologists, audio engineers, lawyers and business persons. This multidisciplinary approach has given rise to significant technological advancements in retrieval algorithms, audio interfaces and data representation schemes.

Notwithstanding these technological advancements, MIR research is currently a systems-centered research domain. For a variety of reasons—including intellectual property law, limited access to substantial, multi-genre, multi-format collections and a lack of a historical user-base—MIR research has hitherto been unable to develop and exploit data concerning the nature of real-world user needs and use of music information. Although [2] and [1] have both commented on the paucity of research involving real users of MIR systems, there have been limited attempts to garner information about and from users. The two most notable examinations of MIR system usage in the research literature are a partial analysis of the OPAC transaction logs from a music library [4] and a log analysis of text and audio queries to a testbed collection of bibliographic and MIDI files [6]. The insights gained in these studies are necessarily limited to the systems that generated the usage data—that is, the studies provided valuable hints as to the usability of features implemented in the two MIR systems, but the data could not provide insight into what additional search facilities or output types that users desire. Moreover, the testbed collection whose usage was studied in [6] is an artificial collection created to test the facilities of the MELDEX MIR software [5]—and so it is not

clear that this log analysis describes 'real' information needs.

MIR research is being hindered by its lack of user-centered theoretic or empirical groundings for its systems design decisions. It is the purpose of this paper to address this obvious need by laying out an empirically derived theoretic framework for the categorization of MIR queries through the context of the needs and uses made manifest in a collection of unsolicited and spontaneous real-world music queries. We will also demonstrate the possible design implications that emerge from our query analyses and classifications.

2. DATA GATHERING

This study analyzed 161 music queries posted to the *rec.music.country.old-time* Usenet newsgroup [3] during the period July 2001 – Jan 2002. During this period, approximately 760 'threads' are represented. Each thread was manually examined, and only the 161 postings which began a thread and which contained a music-related request were retained. This newsgroup was selected as a source of queries for several reasons: it is an active newsgroup (by February 2002, the group's archive contained over 32,800 discussion 'threads'); the group focuses tightly on discussions of music and music events; and the discussions are relatively 'serious'—that is, the participants focus on the genre, rather than on personalities or gossip. Satisfying an information need related to old-time music is not a simple matter of browsing in a CD shop or running a quick search on a peer-to-peer file-sharing service

Why analyze postings to a newsgroup? Given the dearth of query data available for analysis from formal MIR systems, we sought a source of authentic music information requests. This newsgroup provides such a source of queries, and in practice functions as a human-based MIR 'system': it is effectively an old-time music reference desk, run by volunteers keen to share their expertise in and passion for old-time music. The 'interface' to this people-powered system is flexible and usable: the users can post natural language requests, expressing their information needs in their own words and unrestricted by the artificial constraints of a search syntax. These requests contain rich contextual details and background information, permitting the researcher a greater insight into search motivations than is generally available in, for example, a transaction log from a computer-based IR system. Further, information requests to a formal IR system are often constrained by the user's pre-conception of what types of information or document formats are available—the user tailors requests (consciously or unconsciously) to what s/he thinks could be retrieved from that system. In contrast, newsgroup readers recognize that posted requests to a newsgroup could literally retrieve anything—a desired fact, an opportunity to purchase a much-desired album, a

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

pointer to an MP3 file, and so forth. The text-based format of newsgroup postings limits this study to primarily text-based queries, and so we cannot examine in any depth the audio or symbolic techniques that people may wish to use when describing a music information need. Even in the sparse, ASCII newsgroup interface—newsgroup posters include transcriptions and links to audio files in their queries.

3. ANALYTIC CATEGORIES AND DATA

Four principal analytic categories emerged from our reading of the queries:

- *Information need* (see Table 1)
- *Desired outcomes* (see Table 2)
- *Intended uses of information* (see Table 3)
- *Social and contextual elements*

The data from the first three analytic categories are easily quantified and are presented in relatively self-evident summary form below (unfortunately, space precludes more detailed explication of the categories). The fourth category, however, defies *prima facie* quantitative summarization as it involves the qualitative description of extra-musical information that users presented to better contextualize their queries. For example, an associative or environmental context of the desired music-related information was mentioned in 18% (30) of the postings analyzed. By this, we mean that the posting mentioned a social connection that brought the music to the requestor’s attention, or an emotional/social memory associated with the music, or that the music is associated with an event other than simply listening to a CD or LP (“I’ve heard it at a couple of jam sessions”, “ex-friend of mine ... used to sing a song when bowling”, “a mandolin class”).

In some cases, this associative context may be helpful in satisfying the information need; for example, the detail that a song was heard by the requestor “Last week at the Fraley festival” is potentially useful to responders, as the responders may have also attended the festival, may have access to the festival programme, or may use background knowledge about the type of music included in that festival to identify the song in question.

Table 1. Features used to describe the information need.

Information need description	Percentage	Occurrences
BIBLIOGRAPHIC	75.2%	121
LYRICS	14.3%	23
GENRE	9.9%	16
SIMILAR WORKS	9.9%	16
AFFECT	7.5%	12
LYRIC STORY	6.8%	11
TEMPO	2.5%	4
EXAMPLE	1.8%	3

Table 2. Characterizations of desired information

Outcome type	Percentage	Occurrences
BIBLIOGRAPHIC	35.4%	57
LYRICS	29.8%	48
RECORDING	16.8%	27
NOTATION	13.0%	21
BACKGROUND	13.0%	21
RESOURCE	5.0%	8
OTHER	2.5%	4

Table 3. Intended uses for requested music information.

Category	Percentage	Occurrences
LOCATE	49.7%	80
RESEARCH	19.3%	31
PERFORM	18.6%	30
COLLECTION BUILDING	18.0%	29
LISTEN	6.8%	11

4. SUMMARY AND FUTURE RESEARCH

Our study analyzed a set of music information queries across several facets: we examined the types of details presented to describe the information need, the types of information that are desired to satisfy the information need, the explicitly described intended uses for the information, and the social/environmental/associative contexts mentioned in the request. This study suggests that findings from studies of music information requests can be used to inform the development of effective, usable MIR system interfaces, as well as to indicate the types of documents or document representations required to support specific information needs.

However, these queries are drawn from a single source and describe music information needs based around a single genre. Further studies from a variety of sources of music information requests, across a variety of musical genres drawn from diverse cultures, are required to validate the categories presented in music groups are a rich source of music query data, as are the multitudinous email discussion lists and website dedicated to particular types of music or groups of music appreciators. It is also important to move beyond the limitations of text-based query description, to capture the multi-modal richness of real-world, spontaneous “music” inputs—for example, to record and analyze queries including hummed and sung snippets of songs. These latter types of queries are more difficult to obtain; we plan, for example, to negotiate with producers of a popular radio programme to gain access to the musical queries phoned in to their on-air experts.

5. REFERENCES

- [1] Byrd, D., and Crawford, T. Problems of music information retrieval in the real world. *Information Processing & Management* 38(2), 249-272.
- [2] Downie, J.S. Music information retrieval. To appear in *Annual Review of Information Science and Technology* 37. Learned Information, Medford NJ, (in press).
- [3] FAQ of rec.music.country.old-time. Available at <URL:http://home.earthlink.net/~stevesag/otfaq.html>, viewed Feb. 21 2002.
- [4] Itoh, M. Subject Search for Music: Quantitative Analysis of Access Point Selection. In *Proceedings of the 1st Int. Symposium on Music Information Retrieval*. Center for Intelligent Information Retrieval, Amherst MA, 2000. Available at <URL:http://ciir.cs.umass.edu/music2000/>.
- [5] McNab R.J., Smith L.A., Witten I. H. and Henderson C. L. Tune retrieval in the multimedia library. In *Multimedia-Tools and Applications* 10, 113-132. Kluwer Academic Publishers, 2000.
- [6] McPherson, J.R., and Bainbridge, D. Usage of the MELDEX digital music library. *Proceedings of the 2nd Int. Symposium on Music Information Retrieval* (Bloomington IN, USA; October 15-17, 2001), 19-20. 2001.