

ABSTRACT

COMMUNITY-BASED CONTENT CONTROL

Harriette Hemmasi

In the Web environment, the meaning of content necessarily encompasses much more than the basic document, artifact, or file. Communities must provide context for their content that includes descriptive, structural, and technical metadata, the schema that defines the metadata, as well as the ontology that provides an explicit framework for interpreting the data and metadata. For effective online search and retrieval, content must be controlled through the development and application of community-specific rules, logic, and labels or vocabulary. This paper documents how the music community, like many other domain-specific communities, is attempting to improve the definition and management of its content and provides some comments on what additional steps might be taken. While many members of the community still cling to outgrown traditions, progress is being made in music information retrieval and the formalization of a domain-specific ontology that will potentially benefit academia, the Web, and the commercial and entertainment sectors. Ultimately, the ability to share and reuse parts of formalized domain-specific knowledge is vital to the management, integration, distribution, and preservation of broad-scale knowledge. It is the quality of both content and control that must guide community efforts and that will bind our commitment to the future in these areas.

COMMUNITY-BASED CONTENT CONTROL

Harriette Hemmasi

INTRODUCTION

Whether one asks by humming, drumming, nodding, knocking, blinking, speaking, or keying, the asking is for content. Content, whether it is melody, rhythm, color, concept, shape, date, formula, frame, scene, or pixel, begs for context. Context provides a framework for understanding and interpreting content and is shaped by individual and community knowledge. This knowledge can be further refined through a set of imposed controls or tracking mechanisms. These controls are embodied in community-defined rules, a shared logic for these rules, and a common set of labels, or vocabulary, that must be communicable across humans and computers. Web-based search and retrieval processes demand the engagement of total content, i.e., data, metadata, schema, and ontology. Community-specific content control exercised at the points of describing, organizing, and searching helps put content into context and ultimately improves retrieval. The purpose of this article is to demonstrate how the music community, illustrative of other domain-specific communities, is attempting to improve the definition and management of its content as well as what additional steps might be needed to accomplish these goals. Many of the challenges faced by the music community are common to other communities. It is hoped that this discussion will not only lead to more self-awareness among the music community but may also serve as a useful example or basis of comparison for those communities embarking on or already involved in the process of community-based content control.

COMMUNITY-BASED ONTOLOGY

By their very nature communities loosely or strictly generate either implicitly or explicitly a common set of rules, a shared logic for applying those rules, and a collective vocabulary or set of labels derived from that content and its contextual environment. Together, these rules, logic, and labels comprise the fundamentals of a community-specific knowledge base or ontology. While the emergence of a knowledge base can be thought of as a natural evolution within an established group, domain, or discipline, a knowledge base is not usually recognized as such and it is rarely documented, even informally. To improve communication both within a given community and among other communities in the Web environment, it is critical to formalize domain-specific knowledge into a domain-specific ontology.

A domain- or community-based ontology serves as the conceptual backbone for providing, accessing, and structuring information in a comprehensible and comprehensive fashion. Building ontologies is a social process where different stakeholders need to agree on shared classes, terminology, relationships, and constraints. In the online environment, ontologies provide the foundation for processing resources based on the meaningful interpretation of their content rather than simply their physical structure or metadata surrogates. Ontologies define a common vocabulary for researchers who need to share information within and beyond a given domain. In addition, ontologies include machine-interpretable definitions of basic concepts within the domain, established properties and restrictions of each concept, and delineated relationships

among the concepts. Ontologies facilitate the understanding of the structure of information among a community and software agents; they enable the reuse of domain knowledge, document and explicate assumptions shared by a community, separate domain knowledge from operational knowledge, and promote the analysis of domain knowledge.¹

By definition, ontologies contain rules, logic, and labels that are communicable across people and computers. Because of the integral role that ontologies play in interpreting content, they become an essential and ultimately inseparable part of content. As it is collectively the medium *and* the message, so, too, ontologies must be both content and context that comprise and exercise the complete digital resource.²

THE MEANING OF CONTENT

Expressions of intellectual and artistic content could hardly be more multifarious than what we find in today's commercial, educational, and research markets. Simplicity is replaced by multiplicity on every level. In addition to "just text," there are objects of art and architecture, printed and recorded works of music, endless statistics and geospatial representations, videos, photographs, files, slides, and all manner of learning and commercial objects that are presented as viable candidates for description, identification, retrieval, and preservation. Users expect that if "it" exists (or existed) either online or in print, they should be able to find it, any part of it, and everything about it precisely where and how they search. As a result, information specialists are expected to assemble, save, and serve mass quantities of content as well as tiny bytes of information, ad infinitum.

The amount and variation of content and the fact that searchable content may extend beyond the traditional cataloging records or metadata to include the artifact create major challenges for information technologists.³

Therefore, the meaning of content, i.e., what is (or was) and what is sought, can no longer be limited to the descriptive surrogate or even the surrogate plus object. Instead, a more complete definition of content encompasses the data (i.e., document, artifact, etc.), accompanied by its descriptive, structural, and technical metadata, the schema that defines the metadata, and also the ontology that provides an explicit framework for interpreting the data and metadata. An even more comprehensive statement of content, particularly when considering long term preservation, would include not only the multi-layered artifact-metadata-schema-ontology compound but also the original software design and operating system used to support the object's transmission, retrieval, and analysis. For purposes of this paper, the author defines content as the combination of artifact, metadata, schema, and ontology. Together, these four components comprise layer upon layer of intrinsic, potentially identifiable, and ultimately preservable content. Each element representing this expanded definition of content is an expression of the community's rules, logic, and labels.

COMMUNITY-BASED CONTENT: THE MUSIC DOMAIN

Works of Music

The music domain's primary content is comprised of works of music, i.e., compositions that may be manifested on the printed page, in recorded sound, or in live performance,

and works about music, meaning text-based and multimedia works written about music, musicians, and related topics. As the music community is attempting to increase accessibility of this content, it is also attempting improve the definition and management of music-related content. Deceptively simple, the most fundamental “physical” elements of music that are transmitted and understood are sound and its production. This is like saying that visual art is nothing more than color and shape, or literature nothing more than words. The richness of music, as with other arts, is its intricate interweaving of these basic elements into a myriad of patterns and variations. The results are an artistic complexity worth pursuing and worth preserving.

Organizing sound into patterns results in musical composition, with sound occurring either singularly (melody) or simultaneously (harmony) and according to a measured or unmeasured beat (rhythm). Compositions may be written or unwritten and thus may or may not be ascribed to a particular composer, time, or place. Musical performance is the “sounding” of organized patterns, either precisely or imprecisely at the will and/or ability of the performer. Lastly, there are multiple physical manifestations of compositions and their performances in various forms of notated scores, recordings, and live performances. Together, these four broad attributes -- composer, composition, performer, and manifestation -- represent what are commonly known as “bibliographic” elements. Taken together with the physical elements of sound and sound production (such as melody, harmony, rhythm, instrumentation, musical form, etc.) they generate a complex array of multidimensional and deeply interrelated content for which the community must develop explicit rules, logic, and labels.

All aspects and pursuits in Western music, whether artistic, academic, or commercial, center on these basic bibliographic and physical attributes. They serve as the touchstones for upcoming and experienced performers, musicologists, regardless of area of expertise, and the commercial and entertainment industries. Composer, composition, performer, and manifestation are the most common attributes provided in a surrogate describing the “contents” of music, while melody, harmony, and rhythm are the most common elements found in actual musical content. Despite the frequency with which these elements are used, there is little consistency in how various members of the music community, much less the general public, name these attributes and relate them within a given composition or among groups of compositions.⁴ Examples of such inconsistencies were documented in an article in the *New York Times* about Apple computer’s new music service, *iTunes*.⁵ The author states that *iTunes* has “10 different listings for Tchaikovsky, from Piotr Ilych to just plain Peter;” the title for Saint-Saen’s First Cello Concerto is listed as Violincello [sic] in A major, Op. 33, No. 1; and the only “artist” (i.e., performer) cited for this selection is the cellist, Mstislav Rostropovich (totally disregarding the participation of the symphony orchestra and its conductor). Disparity in naming conventions and incomplete or inadequate identification of roles and relationships among the composer, work, and performers continue to be major barriers to successful music information retrieval.

Works about Music

In addition to establishing improved access to works, performances, and manifestations of music there are also community concerns about improving access to text-based and

multimedia works *about* music. While writings about music may concentrate heavily on musical elements and ideas, these texts are likely to include musical examples (e.g., score, sound, video), and they are often highly interdisciplinary in nature. The literature of Western historical musicology, theory and analysis, performance practice, ethnomusicology, etc. addresses not only purely musical ideas but also the rich cultural, historical, and political environment from which music evolved and in which it existed and influenced its surroundings. Likewise, writings and multimedia works emanating from other disciplines often allude to and at times focus on musical matters. These multimedia, multicultural, and multidisciplinary works about music present complex representational and retrieval problems for music information retrieval (MIR), for bibliographic description, and for the general searcher. For effective retrieval in the Web environment, works about music also need an explicit set of rules, logic, and labels.

COMMUNITY EFFORTS TOWARD CONTENT CONTROL

Music information retrieval is a growing research community involving audio engineers, musicologists, music theorists, computer scientists, lawyers, librarians, and others who are concerned with content identification and control.⁶ Music representation languages are in the early stages of facilitating the isolation and tracking of both individual and combined physical elements of a musical work, such as melody, harmony, rhythm, instrumentation, and text (when present), and addressing the rules, logic, and labels that apply to these elements. While much progress has been made, an enormous amount of work clearly remains to be done in developing effective MIR systems.⁷

Stephen Downie, professor in the Graduate School of Library and Information Science at the University of Illinois at Urbana-Champaign and a leader in music information retrieval, highlights two major concerns in MIR research: system evaluation and user studies. According to Downie in his 2003 overview of music information retrieval, “Each contributing discipline brings to the MIR community its own set of goals, accepted practices, valid research questions, and generalizable evaluation paradigms.”⁸

Communicating across disciplines can be problematic and Downie believes that lack of standardized evaluations is one of the biggest obstacles. He and others have called for the creation of standardized test collections, queries, and relevance judgments to help stabilize the evaluation process and advance MIR research.

Closely related to system evaluation is the concern for usability studies, or community analysis, both online and offline. Among the ten central questions for MIR research, Downie cites eight in his overview that deal with system capability and two that represent fundamental questions about usability: (1) What do “real” users of MIR systems actually want the systems to do, and (2) how will “real” users actually interact with MIR systems? The continuing need to answer these rudimentary questions about what users want demonstrates the lack of explicit rules and logic about how the music community conceives of and connects musical elements, whether bibliographic or physical. It is not uncommon for researchers and practitioners to launch system development before gaining a clear understanding of what users want. Additionally, researchers and practitioners may prematurely establish what they believe to be a satisfactory manner of anticipating users’ searching behavior. This approach is typical of many research

projects and creates a major stumbling block to the successful communication and retrieval of information. The music community has the opportunity and the responsibility to lay a solid foundation in user studies that will enable future sharing and reuse of information about music. To construct this framework, the community must analyze not only its content and context, but also its constituents.

A major MIR research project is *Variations2*, the Indiana University Digital Music Library project funded by the National Science Foundation and National Endowment of the Humanities through the Digital Libraries Phase 2 (DLI2) program.⁹ This project is comprised of an interdisciplinary research team working together to coordinate innovations in system design, usability, metadata, intellectual property, music instruction, and networking. The team has made advancements in establishing a digital music library test bed system supporting multiple formats: audio, video, score images, and score notation. During each phase of the project's development, special care has been taken to gather and respond to users' reactions. Based on user feedback, the *Variations2* team has established explicit relationships among composer, work, performer, and manifestation in its data model, similar to Functional Requirements of Bibliographic Records (FRBR).¹⁰⁻¹¹ The *Variations2* data model is work-centered (a change from the container-centered, static object representation of traditional library cataloging records), strengthening the links between individual works and their associated properties. The data model clearly links and identifies the role of each contributor (composer, performer, editor, etc.) to each work and facilitates both the collocation of and distinction among the many manifestations or versions of individual works within the database. The data model is

comprised of a set of music-specific metadata encoded in XML format for scores and sound recordings and is supported by a robust search and retrieval system developed for digitized works of music with the capability of synchronized playback and display of sound and scores files. Efforts are also underway to establish a representative music vocabulary, particularly for form or genre and musical instruments, based in part on terminology found in Library of Congress Subject Headings (LCSH).¹² The long term value of *Variations2* will be measured by user response and by the extensibility of the project beyond its current use at Indiana University and the handful of satellite sites.

COMMUNITY FACING CHANGES

The music community, like other specialized communities, has an uneven understanding and acceptance of emerging technologies and the enormous impact these technologies have on music resources as well as the search and retrieval processes. Some community members continue to cling to outgrown but longstanding traditions. Many of these traditions are still at least partially functional, and some are deeply embedded in and prized by both the community and society at large. Among such traditions are: widespread inconsistencies in naming conventions that would be difficult, if not impossible to codify (e.g., multiple spellings of Tchaikovsky); use of stale, imprecise, and incomplete “controlled” vocabulary and classification schemes such as those contained in LCSH, Library of Congress Classification, Dewey Decimal Classification (e.g., limited representation of many styles of music, particularly in the areas of popular and ethnic music). Also problematic is the lack of explicit relationships between composer, composition, performer, and manifestation, mentioned above, as propagated

by the Anglo American Cataloging Rules (AACR2), the MARC record, online public access catalogs, commercial databases, and the Web (see discussion above on *Variations2* data model) .¹³⁻¹⁴ The acceptance, and possibly even the advancement of technology is impeded by the continuing use of print publications, such as thematic indexes, for processes that could be much more efficiently accomplished via a search engine. By contrast, over-reliance on what is available on the Web, a trusting acceptance of the Web's limited search strategies, and the assumption that going to Google is an adequate replacement for going to the library also represent barriers to sophisticated music information retrieval. As in other disciplines, there is a sense of community pride in and protection of the expertise held by certain individuals or groups within the music community (e.g., "only the finest Beethoven scholar can answer your question;" or "only expert catalogers are able to establish the correct form of a name") that can inhibit the willingness to entrust an equivalent level of knowledge or skill to computer programs. In addition, the academic music community strongly supports individualism: finding one's own answers the best way possible, even if it is the hardest way; pursuing one's own research project without sufficient regard to how the project might meet broader needs or fit into the community's more comprehensive research agenda.

Even when the need for change is clear, such as the need for metadata beyond the traditional bibliographic descriptions of scores and sound recordings supported by AACR2 description and MARC encoding, reaching consensus and obtaining community endorsement is a tedious and time-consuming process. The Music Library Association, for instance, has only tentatively begun to examine the possibility of expanding metadata

for digitized musical scores and sound recordings beyond the confines of the MARC bibliographic record.¹⁵ Yet it is precisely this group, with its extensive experience and knowledge of describing and accessing works of and about music, who should be among the leaders of music metadata reform. Instead, individual research projects, commercial companies, and Web search engine designers are setting the agenda for the next generation of music metadata and music retrieval.

For many years the music community has clearly expressed the need for improved subject access to works of and about music.¹⁶ Several attempts, both within the US and internationally, to create a music thesaurus have fallen by the wayside due to lack of community coordination and adequate funding.¹⁷ Gathering, documenting, organizing, and sharing domain-specific terminology (descriptors, names, and works) would be an enormous contribution to the music community, as well as to other communities. No less significant than the development of the *Art and Architecture Thesaurus* which flourished under the J. Paul Getty Trust sponsorship, identifying and arranging the vast array of scholarly and popular musical terms, both foreign and English, would be invaluable.¹⁸ The potential impact of providing and maintaining a standardized searching and indexing vocabulary that could be applied across academia, the Web, and the commercial and entertainment sectors argues for a more concerted community effort and more serious consideration by potential funding agencies.

CROSS-COMMUNITY COMMITMENT

In order to facilitate large-scale knowledge integration, ontologies need to be viewed from a highly interdisciplinary perspective. The ability to share and reuse parts of formalized bodies of knowledge is vital to the management and preservation of knowledge. While the ability to transparently handle variations in content and the meanings of content is needed, it is neither realistic nor practical to seek interoperability through the adoption of a single standard or a single vocabulary. Instead, developing domain-specific schemas and vocabularies are essential first steps toward improving cross-community knowledge sharing. Harmonization of metadata vocabularies through the use of schemas that consist of data elements drawn from one or more namespaces seems to offer an optimal approach.¹⁹ Assessing the cost of building and maintaining such systems must be weighed against the value of the content. It is the quality of content *and* control that must guide community efforts and that will bind our commitment to the future.

CONCLUSION

Music, like other specialized communities, has a well-developed but largely undocumented knowledge base. Even though the basic physical and bibliographic attributes of music are seemingly apparent and may be considered to be widely known and recognizable, they lack the necessary formalization of explicit rules, logic, and labels. Reliance on informal rules, implicit logic, and variant names forms for the same concept, object, work, or person results in insurmountable gaps in communication and accessibility, particularly as the amount of information and breadth of distribution

increase. The many variations and intricate interweaving of musical elements generate a set of complexities that require an expanded view of content in order to improve information retrieval both within the music community and beyond. For successful music information retrieval, the texts and artifacts (data) must be wrapped in community-developed metadata, a community-endorsed metadata schema, and a community-specific ontology. This total framework for content will ensure that it is communicable across people and computers.

The MIR research community is making strides in isolating, identifying, and presenting basic elements of musical content for search and retrieval. Meanwhile, there are still fundamental questions about what users want and need from a music information retrieval system. The music community has yet to do a thorough self-analysis, to lay a solid foundation in user studies. While the outcome of such studies is unlikely to be one of consensus, user input, however varied, is critical to system development and to documenting the community's rules, logic, and labels.

Even as segments of the music community are eager to make advancements in information retrieval, there remain pockets within the community who are either satisfied or willing to live with the current state of affairs, no matter how inadequate. Widespread and persistent education about the need for change and the advantages of collaboration and cooperation in forming a common set of rules, logic, and labels cannot be overestimated. Building a common understanding of the goals and projecting the

potential results of achieving these goals are important steps toward involvement, endorsement, and ultimate success.

FOOTNOTES

1. Noy, Natalya F. and Deborah L. McGuinness. "Ontology Development 101: A Guide to Creating your First Ontology."
http://protege.stanford.edu/publications/ontology_development/ontology101-noy-mcguinness.html (accessed 9/21/03)
2. McLuhan, Marshall. *The Gutenberg Galaxy: The Making of Typographic Man*. Toronto: University of Toronto Press, 1962.
3. Lagoze, Carl. "From Static to Dynamic Surrogates: Resource Discovery in the Digital Age," *D-Lib Magazine*, June 1997.
<http://www.dlib.org/dlib/june97/06lagoze.html> (accessed 9/21/03)
4. Smiraglia, R. "Musical Works as Information Retrieval Entities: Epistemological Perspectives," *Proceedings of the 2nd Annual International Symposium on Music Information Retrieval (ISMIR 2001)*; 85-92, 2001.
<http://ismir2001.indiana.edu/pdf/smiraglia.pdf> (accessed 9/21/03)
5. "Adventures in Downloading Haydn," Arts Section. *The New York Times*, June 15, 2003.
6. Music information retrieval. <http://music-ir.org> (accessed 9/21/03)
7. Downie, J. Stephen. "Music Information Retrieval." (Chapter 7) *Annual Review of Information Science and Technology*, 37. ed. Blaise Cronin, 295-340. Medford, NJ: Information Today, 2003. http://music-ir.org/downie_mir_arist37.pdf (accessed 9/21/03)
8. Ibid.
9. Variations2: Indiana University Digital Music Library.
<http://www.dml.indiana.edu/index.html> (accessed 9/21/03)
10. Hemmasi, Harriette. "Why Not MARC?" *Proceedings of the 3rd Annual International Conference on Music Information Retrieval (ISMIR 2002)*; 242-248, 2002. <http://ismir2002.ismir.net/Proceedings\02-FP08-3.pdf> (accessed 9/21/03)
11. Functional Requirements of Bibliographic Records.
<http://www.ifa.org/VII/s13/frbr/frbr.pdf> (accessed 9/21/03)
12. Library of Congress Subject Headings, part of the Library of Congress Program for Cooperative Cataloging. <http://www.loc.gov/catdir/pcc/> (accessed 9/21/03)

13. Library of Congress Classification. <http://lcweb.loc.gov/catdir/cpsolcco/lcco.html> and Dewey Decimal Classification. <http://www.oclc.org/dewey/> (both URLs accessed 9/21/03)
14. *Anglo-American Cataloging Rules, Second Edition, 1998 revision*. Prepared under the direction of the Joint Steering Committee for Revision of AACR, a committee of the American Library Association. Chicago: American Library Association, 1999- and, Library of Congress. MARC Standards web site. <http://www.loc.gov/marc/> (accessed 9/21/03)
15. The Music Library Association Metadata Standards Working Group was charged and appointed in September, 2003. [author is a member of this group]
16. McKnight, Mark and others. "Improving Access to Music: A Report of the Music Thesaurus Project Working Group," *NOTES*; 45:714, 1989.
17. Hemmasi, Harriette. "The Music Thesaurus: Function and Foundations," *NOTES*; 50 (4): 875-882, 1994.
18. Art & Architecture Thesaurus Online. <http://www.getty.edu/research/tools/vocabulary/aat/> (accessed 9/21/03)
19. Lagoze, Carl and Jane Hunter. "The ABC Ontology and Model," *Journal of Digital Information*; 2(2). 2001. <http://jodi.ecs.soton.ac.uk/Articles/v02/i02/Lagoze/lagoze-final.pdf> (accessed 9/21/03)