This paper is a pre-print version presented at the ISKO UK 2009 conference, 22-23 June, prior to peer review and editing. For published proceedings see special issue of Aslib Proceedings journal.

Organising music for movies

Charles Inskip and Andy MacFarlane
City University London

Pauline Rafferty
University of Aberystwyth

Abstract

Purpose - The purpose of this paper is to examine and discuss the classification of commercial popular music when large digital collections are organised for use in films.

Design/methodology/approach - A range of systems are investigated and their organization is discussed, focusing on an analysis of the metadata used by the systems and choices given to the end-user to construct a query. The indexing of the music is compared to a checklist of music facets which has been derived from recent musicological literature on semiotic analysis of popular music. These facets include aspects of communication, cultural and musical expression, codes and competences.

Findings -In addition to bibliographic detail, descriptive metadata is used to organise music in these systems. Genre, subject and mood are used widely; some musical facets also appear. The extent to which attempts are being made to reflect these facets in the organization of these systems is discussed. A number of recommendations are made which may help to improve this process.

Originality/value - This paper discusses an area of creative music search which has not previously been investigated in any depth and makes recommendations based on findings and the literature which may be used in the development of commercial systems as well as making a contribution to the literature.

Keywords Music information retrieval, Indexing, Domain analysis

Paper type Research paper

1. Introduction

Record companies and music publishers (rights holders or music owners) control rights of recordings and compositions respectively. When music is used to accompany moving images (music synchronization), in a film or TV commercial, the film maker or advertising agency (music user) is encouraged by the rights holders to use their bespoke proprietary systems to perform searches for suitable material from their catalogues. Specialist intermediaries also perform searches on behalf of users based on written, verbal and visual queries. Production music libraries, which control music made specifically to accompany moving images, organise their music by genre, subject and activity rather than by artist and title. Some of these approaches have been adopted by owners of pre-existing commercial recordings, whose primary purpose is not for synchronization, in order to exploit their material further.

Previous research (Inskip *et al*., 2008a, b) discusses the difficulties involved in meeting this type of information need. There are a number of stakeholders, and a range of factors which affect the decision (including non-musical factors such as budget and availability). In addition, it is difficult to frame music queries as

descriptions can be highly subjective and conveying musical meaning adequately relies on shared codes and competences (Inskip *et al.*, 2008c, Middleton, 1990; Stefani, 1987; Tagg, 1999). An expert intermediary (independent or music owner staff) acts in an advisory role, performing a search of their catalogue on behalf of the user by interpreting their query (known as a "brief") in a way that matches the organization of the catalogue. Many of these intermediaries have previously worked in film and advertising and have insight into the codes and competences of the users. Most large rights holders have made an attempt to deal with these issues in disintermediation by developing bespoke proprietary search tools for their catalogues.

The purpose of this paper is to investigate a range of bespoke music Information Retrieval (IR) systems used in the music industry. The methodology, which is designed to facilitate analysis of the metadata used by the systems and the choices given to the end-user to construct a query, is discussed in Section 2. This leads on to a discussion in Section 3 of the way commercial music controlled by these music owners is organised for the purpose of music synchronization, referring to some examples of real queries used in this process. Section 4 then refers to Knowledge Organization and semiotics literature with a view to offering some possible solutions to the problems of applying descriptive metadata consistently to music. Section 5 summarises the paper, presents a selection of recommendations regarding possible improvements to the process, and suggests further research objectives for the future based on these findings.

2. Methodology

This investigation discusses search engines operated by five multinational music publishing companies (Warner Chappell, EMI, Sony ATV, Chrysalis, and Universal) and a service run by an intermediary (Ricall) which searches a collection of other music catalogues operated by a number of smaller rights holders including EMI Records and Sony Music. Although the size of these catalogues ranges from around 70,000 works (or compositions) to more than 4 million recordings (some compositions are recorded many times) the works or recordings are not all digitised and available online. Although some are held back for legal reasons, digitization strategies prioritize higher earning or more "synch-able" material, leading to the owners "narrowing down" the material on behalf of the users.

All of the descriptive terms on each application's user interface were copied into tables for analysis. This removed them from their visual context, allowing a clearer comparison to be made between the ranges of facets being presented to the users. These types of terms were ranked according to frequency (Table I) to find any commonalities and notable differences in their organisation. After the terms had been entered into tables these were each imported into NVivo8 software. A word count was performed on each set. This generated a list for each search engine of the top 45 words of three letters or more (including "pop" but excluding "an"), ranked by frequency. Tag clouds were also produced for each set, enabling a visual analysis of the collections.

Representatives of the companies operating the six services were approached directly with interview requests for more detailed information surrounding their organisation and use. Four agreed to interview, two did not reply within the deadlines of this paper

submission. The subject of these semi-structured interviews, included: detail of the classification schemas, information around the usage of the services, and identification of problems of organising music. The interviews were transcribed and used to enhance and illustrate the findings of the search engine text analysis.

3. Organization: facets

Table I represents the top fifteen ranked musical facets derived from the study of the six music IR search engines:

Take in Table I. Top 15 ranked musical facets

It can be seen that all six of the search engines present a mixture of bibliographic and descriptive metadata to the user, allowing them to search by writer, year, artist and title (bibliographic) as well as by genre, tempo and keyword (descriptive). Other descriptive metadata included mood, subject, lyrics and instrumental/vocal mix and bibliographic terms included album title, chart position, one stop (ease of copyright clearance) and originating territory. These facets of the musical document, then, are deemed by the creators of these search engines to be key in organising music when it is being searched for synchronisation purposes.

The bibliographic metadata is unchanging over time and is mainly set when the work or recording is created. Some of this information may be used by the organisation for a wide range of administrative purposes other than retrieval, including rights clearances and royalties payments. Descriptive metadata, on the other hand, can change over time, can be subjective, can vary between performances of the same work, and can be more domain-related. Traditional music cataloguing approaches do not focus on classifying musical works or performances according to their mood or the subject of their lyrics. Determining mood or subject requires expert input related to the reasons behind applying this high-level type of classifier. In the scenario under investigation the objective of the application is to retrieve a manageable selection of pieces of music that meet the requirements of a film or commercial maker who wishes to match a piece of music to a section of film.

3.1 Unknown item

In this type of searching there are two types of query: known item and unknown item. With a search for a known item, the user knows exactly the piece of music they want to use, and need to negotiate this use with the rights holder. This "clearance" is dealt with initially by the synchronisation intermediary who subsequently passes the request on to legal and business affairs for contractual purposes. Often, however, the user is unsure about their choice of music. The owners organise their interfaces in an attempt to meet this particular need, in a wide variety of approaches which are based on socio-cultural elements of music and by musical descriptions. For example an advertising agency may state:

> We need to avoid the spot from going too 'mushy' or 'schmaltzy', and therefore wants to be upbeat whilst charming. It doesn't want to be melancholic, down beat, over the

> top, over powering or dark, but optimistic, light hearted, contemporary and charming. (Music brief, 2009a, excerpt)

Or:

> We are looking for a great, proper love song. Something unashamedly romantic so don't be afraid to be sentimental or warm. It should lift the heart. Instinctively we think that the track we need is probably from the 50's or 60's, maybe the 70's. […] We want a big crescendo at the end, when the main characters kiss and everybody wakes up, so songs that build up to swelling, string-soaked choruses are especially welcome. Lyrics and choruses involving sleep, eyes, waking, dreaming, touch, or some other direct link to the narrative, would be great. (Music brief, 2009b, excerpt)

3.2 Description

It is the job of the expert intermediary or the search tool to match these types of unknown item queries to relevant items in the catalogue. It is problematic that terms such as "mushy", "schmaltzy", "charming" and "warm" can be interpreted in many different ways by users and owners alike. Some applications offer a controlled vocabulary of moods or feelings to deal with this, and these types of non-specific emotive terms are described variously as styles, moods, subjects, topics or keywords. Offering a range of choices in a controlled vocabulary improves consistency in indexing and acts as a guide to the user when searching. This relies on the chosen terms reflecting the users' language choices and, from various comments in the interviews, it appears that the principle of user warrant (NISO, 2005) is generally applied when building descriptive categories, while literary warrant is applied to bibliographic terms.

3.2.1 Genre

Genres are semiotic codes which are agreed by the community (Fabbri, 1981). Musical events within a piece of music combine with social uses and interpretations of the piece, which may shift over time. They may be very broad (art music, folk music, pop music) or highly specific. They may also vary according to the codes and competences of the listener (Middleton, 1990). Despite their fluid nature they are widely used in arranging music collections. In record stores genre is frequently the sole category used to organize music. This may mean the music being sought by a user may not be classified in the way they may expect (Radiohead under "Pop" or "Rock"? Simon and Garfunkel under "Folk" or "Pop"?), but their all-pervasive nature indicates a strong acceptance by the community that genre is a suitable way to organize music collections even though inadequate genre definitions may lead to mis-classification (Abrahamson, 2003).

All search engines in the study presented a range of genres to choose from. These ranged in specificity from a small selection of ten broad genres (rock, R&B/soul, jazz/easy listening, country/blues/folk, rap/hip-hop, pop/dance/electronica, alternative, world/reggae/latin, new age, gospel/Christian (EMI, 2009)) to a selection of 32 genres and 1,295 sub genres (Universal, 2009). These genres are offered to the user as a controlled vocabulary.

A word frequency query was run on the entire collection of terms. This gave rise to the tag cloud in Figure 1.

Take in Figure 1. Total word frequencies tag cloud - all search engines combined

This shows that a very wide range of genre terms are the most used facets, the highest in all of the collected data being the word “pop”, followed by “rock” and “soul”. On closer examination it can be seen that some of the terms are subject based, such as “adult”, “love”, “fun” and “desire”, or mood based (“dark”, “dramatic”, “reflective”). Musical features are also included (“slow”, “soft”, “fast”, “instrumental”).

The frequency tables for each individual search term collection indicate that although each of them focuses mainly on genre terms, especially “rock” and “pop”, there are variations in depth of cataloguing and in focus. For example, Table II shows the variety in the frequency of use of “pop” and “rock” in the search tools.

Take in Table II. Frequency of “pop” and “rock”

While most of these references can be found in the expected genre category, a small number appear in subject classifications and are used to clarify; for example, “worship” music – allowing the user to choose between Christian rap and Christian rock (Warner Chappell, 2009).

3.2.2 Lyrics/subject

Four of the applications allowed the user to perform a search through lyrics of the catalogue, again by free text input. This often reflects the purpose of the search, which is to find music with lyrics that in some way enhance the footage it accompanies. In the second example above (3.1) lyrics which link to the narrative are seen to be important. Using the search terms “*sleep, eyes, waking, dreaming*”, etc. helps to focus the results on relevant material. Unfortunately it is not possible to narrow down this set further by incorporating other criteria (“*70s*”, “*crescendo*”, etc.). The user is forced to choose which criteria are most important to the search and then narrow down by reading through a list and listening to the songs to determine their relevance. The subject of the lyrics, however, was searchable, mainly by controlled vocabulary, while one collection was searchable by free text – the owners had already determined the meaning or subject of the lyrics on behalf of the users. The categories included Actions, Aesthetic, Enjoyment, Time/Travel (Warner-Chappell, 2009) and Travel, Party, People, Action (EMI, 2009). They are illustrated in the tag cloud in Figure 2.

Take in Figure 2. Tag cloud - subjects - all search engines

In advertising and film, the words in the lyrics, or their subject, can be used to communicate messages from the film maker to the viewer. This is not to say the lyrics are always used to convey the meaning intended by the songwriter. In TMobile’s recent use of Vashti Bunyan’s “Diamond Day”, for example, a surface interpretation of the lyric indicates that it celebrates a beautiful day in the countryside. The visuals for this advert show young people using mobile phones in urban settings featuring surfaces which are unnaturally flexible and bear no direct relation to the lyric. Barclaycard’s use of the Bellamy Brothers’ “Let Your Love Flow” in their waterslide advert, however, does have some lyrical relevance inasmuch as the protagonist *flows*

down the slide in the water for the duration of the commercial. The original brief for this commercial described the visuals and stated:

> Any lyrics should relate loosely to the story of the ad, which conveys a positive journey. (Music brief, 2009c, excerpt)

These examples indicate there is a high level of creativity in the choices behind these music uses, which cannot easily be broken down into genre/subject/mood categories. This is reflected in the briefs examples above, which are non-specific and subjective.

3.2.3 Mood

In an attempt to regularise the subjective nature of mood descriptions, four applications gave a controlled set of options to the user, ranging from a selection of 8 (aggressive, brooding, happy, mellow, romantic, sad, sentimental, upbeat (EMI, 2009)) expanded to a further 28, to Universal's 41 "mood groups" and 612 sub moods.

Mood is not universally used in these search engines (although entering mood categories into keyword boxes where a drop down choice does not exist does generate a set of results) and it is subjective although it is widely used in briefs. The mood of a piece is determined by whether the listener is considering the music, or the lyrics, or their combination. Many 1960s Motown songs, for example, have uptempo major key melodies, indicating a "happy" or "positive" meaning, while the lyrical content is often to do with breaking up relationships or general hardship. This creative counterpoint is not reflected in the search tools under investigation.

The mood descriptors that are offered include both positive ("high spirited", "passionate", "dynamic") and negative ("sad", "aggressive", "angry") emotions. Listening to selected pieces in these categories indicates they are indexed by a combination of musical and lyrical content. A song can have multiple mood categories.

Take in Figure 3. Tag clouds - moods - all search engines

3.2.4 Keyword

All search engines offered a choice to the user of entering a keyword. While most of them allowed this search to be by free text, one offered a dropdown choice of 271 selections, which included anger, breakfast, family, heartbreak, suicide, and youth. Using natural language indexing this links the users chosen keyword(s) to words in the title or lyric of the song, or to descriptive metadata applied by the owner.

3.2.5 Musical features

Three services allowed the user to narrow down their search by the specific musical facet of vocal mix (female / male / chorus / duet) and a selection of featured instruments, while all six offered choices in Tempo, which ranged from a selection of three (Fast, Medium, Slow (EMI, 2009)) to a more sophisticated choice of eleven, using Beats Per Minute (BPM) as a descriptor along with a text explanation (e.g., Slow - (71-90) BPM (Warner-Chappell, 2009). Other types of musical features

included instruments and beat but detailed features normally associated with musical analysis such as key, texture, tonality, range, accentuation, harmonies, etc. are not offered as choices to the user.

The emphasis on these search tools is to offer a highly detailed selection of features which are not "musically" specific but are more general and descriptive. This suggests that the users do not need to be musically trained, and that there are certain facets which they find more important than others when they are talking about music in this type of searching.

4. Discussion

4.1 Classification and knowledge organization

Music can be organized by bibliographic or descriptive means. Some of these are specific to music rather than text. Vellucci recommends that:

> understanding the various ways in which music resources are used will allow the metadata creator to resolve the practical problems when determining the type of metadata information required to meet the search and retrieval needs of musicians (2004, p. 39).

She discusses how the problems of multiple authors, performances, languages and manifestations of works can cause problems with music classification that are not dealt with in text classification methods. There have been efforts to deal with these problems and the International Federation of Library Associations (IFLA) recommend a focus on the Functional Requirements of the Bibliographic Record (FRBR) outlined states of: work, expression, manifestation and item. These are reflected by the bibliographic nature of MARC and AACR2R schemes which deal particularly effectively with notated music (Vellucci, 2004). Dublin Core (DC), however, is widely used by digital libraries for digital audio and is recommended by International Association of Sound and Audiovisual Archives (IASA, 2009) while MPEG-7 has been found to be comprehensive, wide-ranging and valid in the music information retrieval community (Corthaut *et al*., 2008).

These methods and approaches are integral to the organization of music in a comprehensive sense, and are indeed applied by some of the organizations in this study. Bibliographic terms such as title, writer, publisher, artist, and album feature prominently. However they do not directly deal with the difficulties in the situation of searching for music in large collections to match the unknown item queries found in the area of music synchronization. They either focus on bibliographic metadata, which specifically deals with known item searching, or their musically specific formats fall outside the natural language queries of the users, who are experts in their field (film-making, advertising) but not all are experienced musicians or music analysts.

Abrahamson discusses three types of musical analysis: structural, sociological, and semiotic. Structural analysis focuses on inner musical elements and style, while sociological analysis concentrates on how music is affected by social factors such as its production and consumption (Abrahamson, 2003, pp.156-157). Semiotics of popular music attempt to analyse the meaning of music. It is clear that these

organizations have taken a domain analytic approach to the development of these applications.

Hjørland (2008, p. 86) raises the issue of the importance of extracting the meaning of documents when organizing them, and notes how this may not be as objective as the use of standards may imply. He discusses six approaches to KO: traditional, facet-analysis, information retrieval (IR), user-oriented, bibliometric, domain analytic and others (including semiotic, discourse-analytic, etc.). He recommends that classifying objects should be determined not by "*trivial or naïve description*" but by "*broader meaning-producing concepts*" (Hjørland and Nissen Pedersen, 2005, p. 593) and that a pragmatic, rather than a positivist approach be taken. In the domain analytic approach documents are classified by their purpose, which reflects the paradigm in which they sit and the needs of the intended users ((Hjørland and Nissen Pedersen, 2005).

Specialized communities, such as the music and film industries or the even more specialized synchronization community, develop their own vocabularies and discourses over time (Buckland *et al.*, 2001). Analysing these discourses (including briefs, websites, promotional material and interviews) can inform a domain analytic approach as they will give insights into the special language and attributed meanings used within this discourse community. Understanding what elements of musical meaning are important to the community will help to identify the most relevant facets for organizing the music.

The inclusion of specialised descriptive metadata such as Subject and Mood especially has been developed in order to reflect the wording of briefs and other types of query. These facets seem to be a response to the problem of matching musical and filmic meaning. Taking some examples from Figure 2 we can see how songs about numbers", "sports", "memories" or "drinking" could be used to enhance the visual message of commercials or scenes in a film featuring mobile phones, training shoes, cameras or liquor. Similarly, from Figure 3, moods such as "dramatic", "passionate" and "sentimental" could equally be applied to describing scenes in a film or be associated with products.

The language and organization appears to reflect some of the special needs of this community (catalogue exploitation for the owners, and synchronization for the users) and the aim of these applications is to meet a particular need. This seems to relate to the use of music to enhance particular cultural meanings intended by the creators and other stakeholders to be carried to the viewer by combining it with moving images. The problem which arises frequently in interview is that there are no "musical" rules to determine how to apply these descriptors. Although the textual meaning of a song lyric may be interpreted in a similar way by a range of listeners (and this is certainly not as likely as a unified interpretation of the meaning of, say, an academic paper), the musical mood can vary amongst listeners and even one listener can change their interpretation over time.

If this is the case then it may be suitable to use ideas from the field of popular music semiotics in analysing and indexing music for these purposes.

4.2 Semiotics of popular music

Tagg derives a checklist of "Parameters of Musical Expression" (1999, pp. 29-31), stating:

> It is also vital that those parameters and the musical structures they create are related to the world outside the music, i.e. to the social and cultural position, intentions, motivations of those producing and using the music as well as to the functions and acoustic context of the music.(1999, p. 28)

This seems to reflect the domain analytic approach recommended by Hjørland (2002). In his checklist, Tagg (1999) recommends analysis should consider the communication of the parties involved in the process, their interests and motivations, attitudes and situations as well as detailed examination of non-musical texts (sleeve notes, venues, movements, visuals) around the music and musical parameters such as instrumentation (including timbre, performance techniques), composition (including texture and polyphony), temporal parameters (duration, tempo, metre), tone (pitch and harmony) and dynamics (soft/loud). Notably, otherwise all the analyst would end up with would be a list of facets, he also recommends analysing how the communication and the non-musical elements relate to the music itself. This holistic approach to musical analysis, if framed within the paradigm of music synchronization informed by a domain analytic approach, could be used to derive a comprehensive and useful analysis of music which would include the facets already discussed (such as genre, subject and mood) and benefit the process of music search in this particular context. It should be noted that the level of detail of analysis, of course, should be informed by the findings of a deeper analysis of the discourses of this community. Although some of the search tools in this investigation do not give significant weight to "musical" facets such as timbre, phrasing, compositional techniques and detailed harmonic parameters, intermediaries discussed how they used these types of factors to inform their choices when searching, and pointed out how these facets contribute to genre and mood definitions. A number are formalising this approach in new developments of these tools which are being designed for internal use and index by "thickness" and "key" as well as those made available to the specialist user by these tools.

5. Summary and recommendations

We have shown that in order to exploit the material they control, record companies and music publishers have made an effort to develop systems that enable outside users to search areas of their catalogues when looking for music to accompany moving images. These music owners have developed categorisation techniques that are outside the traditional paradigm of music cataloguing and are specific to their domain. Although they index using bibliographic terms, they also use descriptive terms which are more important when users are searching for unknown items. We have shown the key descriptive elements to be Genre, Subject and Mood and how some musical terms such as Tempo are also important.

From the analysis of these search tools it has been shown that the music owners are partly taking a domain analytic approach to their music classification, making them appropriate for a particular use and user in terms of language and presentation. It is possible that ideas from the semiotics of popular music could inform more comprehensive classification. This may enhance these services which are reportedly underused and inadequate. There are some important issues that may be raised here:

1. This subjective indexing needs to be consistent and reliable – although controlled vocabularies help this, they are inflexible and adding a new term requires the whole catalogue to be re-indexed.
2. The music is indexed by humans who are not always experts – they do not always share the same codes and competences (members of other departments within a music owner may have different interests and motivations, and may not be experienced in the language of synchronization).
3. It is clear that none of these tools search through the whole of their catalogue, which has either not been fully digitized, not been fully indexed, or has been narrowed down to exclude non-priority material.
4. Excluding musical facets such as texture, harmony, key, timbre, etc., forces the user to search with subjective descriptions and then narrow down by listening.
5. Interviewees frequently refer to the clichéd query "*quirky and upbeat with a bit of a build*". None of the search tools here made it clear that they indexed facets of the type of *"a bit of a build"* (although "build" did appear in some free text keyword searching).
6. Although these applications were developed for the user they do not appear to be widely used by them. Some owners use them internally but are not positive about them either because they do not offer useful results sets or because they feel the human approach is more suitable to creative search.
7. There are frequent occurrences of mis-applied keywords, mis-spellings, incomplete metadata, low precision and recall, poor stemming, inappropriate ranking, and difficulties in narrowing down or expanding results sets.
8. Frequently users employ similarity metrics to guide the intermediary. Although this type of approach is restricted in these types of tools because they can only refer to music they control for proprietary and copyright reasons. Ricall (2009), however, is able to search a wide range of catalogues and does employ this as a query.

The existence of these tools indicates a commitment on the part of the music owners to a certain amount of disintermediation in this search process. Two of the companies stated they were currently Beta-testing new versions of these tools which were more comprehensive and flexible; however these are currently internal and not yet accessible for research purposes. From the insight given in interview (but not for detailed disclosure) these tools will be more fit-for-purpose than currently.

5.1 Further research

This research is part of a wider project investigating the communication, meaning and retrieval processes surrounding music use within the music industry, focusing on use with moving images. We plan to further evaluate these systems and analyse the discourses around the process in more detail in order to investigate the domain of music synchronization in the context of LIS theory and practice.

Acknowledgements

We would like to thank all the anonymous participants who have taken part in this research, who have kindly given up valuable time and knowledge and supplied a range of briefs for analysis. The corresponding author also gratefully acknowledges PhD studentship support from Arts and Humanities Research Council.

References

Abrahamson, K.T. (2003), "Indexing of musical genres: an epistemological perspective", *Knowledge Organization*, Vol. 3/4, pp. 144-169.

Buckland, M., Jiang, H., Kim, Y. and Petras, V. (2001), "Domain-based indexes: indexing for communities of users", in *3e Congrès du Chapitre français de L'ISKO, 506 juillet 2001. Filtrage et résumé informatique de l'information sur les réseaux, Université Nanterre Paris X, Paris*, pp. 181-185.

Chrysalis (2009), Song Search available at: <http://www.uk.chrysalismusic.co.uk/core/song_search.cfm> (accessed 21 April 2009).

Corthaut, N., Govaerts, S., Verbert, K. and Duval, E. (2008), "Connecting the dots: music metadata generation, schemas and applications", *Proceedings of the Ninth International Conference on Music Information Retrieval*, Philadelphia, PA, 14-18 September 2008.

EMI (2009), Song search and license available at: <http://www.emimusicpub.com/worldwide/around_the_world/united-kingdom_home.html> (accessed 21 April 2009).

Fabbri, F. (1981), "A theory of musical genres: two applications" in Horn, P. and Tagg, P. (Eds), *Popular Music Perspectives*, International Association for the Study of Popular Music, Goteborg and Exeter, pp. 52-81.

Hjørland, B. (2008), "What is Knowledge Organization (KO)?", *Knowledge Organization*, Vol. 35 No. 2/3, pp. 86-101.

Hjørland, B. and Nissen Pedersen, K. (2005), "A substantive theory of classification for information retrieval", *Journal of Documentation*, Vol. 61 No. 5, pp. 582-597.

IASA (2009), *IASA TC-04 Guidelines on the Production and Preservation of Digital Audio Objects*, 2nd edition, edited by Bradley, K., IASA, Auckland Park, South Africa.

Inskip, C., MacFarlane, A. and Rafferty, P. (2008a), "Music, movies and meaning: communication in film-makers' search for pre-existing music, and the implications for music information retrieval", in Bello, J., Chew, E. and Turnbull, D. (Eds), *Proceedings of the Ninth International Conference on Music Information Retrieval*, 14-18 September 2008, Philadelphia, PA.

Inskip, C., Macfarlane, A. and Raffety, P. (2008b), "Content or context? Searching for musical meaning in task-based interactive information retrieval", in Borlund, P., Schneider, J., Lalmas, M., Tombros, A., Feather, J., Kelly, D., de Vries, A. and Azzopardi, L. (Eds), *Proceedings of the Second IIiX Symposium on Information Interaction in Context*, 14-17 October 2008, London.

Inskip, C., Macfarlane, A. and Rafferty, P. (2008c), "Meaning, communication, music: towards a revised communication model", *Journal of Documentation*, Vol. 64 No. 5, pp. 687-706.

Middleton, R. (1990), *Studying Popular Music*, Open University Press, Buckingham.

Music Brief (2009a, b, c), (Personal communications, March 2009).

NISO (2005), Guidelines for the Construction, Format, and Management of Monolingual Controlled Vocabularies ANSI/NISO Z39.19-2005, NISO Press, Bethesda, MD, available at: <www.niso.org> (accessed 27 April 2009).

Ricall (2009), Ricall Music Licensing, available at: <http://www.ricall.com/home/home.seam> (accessed 21 April 2009).

Sony ATV (2009), Sony / ATV Music Publishing Music Search, available at:<http://www.sonyatv.com/en-uk/> (accessed 21 April 2009).

Stefani, G. (1987), "A Theory of Musical Competence", *Semiotica*, Vol. 66 No. 1-3, pp. 7-22.

Tagg, P. (1999), *Introductory Notes to the Semiotics of Music, Version 3*, unpublished, available at: http://www.tagg.org/xpdfs/semiotug.pdf (accessed 7 December 2006).

Universal (2009), Universal Music Publishing Group SyncExpress, available at: <http://www.synchexpress.com> (accessed 21 April 2009).

Vellucci, S. (2004), "Music Metadata", in Gorman, G.E. (Ed.), *International Yearbook of Library and Information Management 2003-2004 Metadata Applications and Management*, Facet, London.

Warner Chappell (2009), Warner Chappell Music Song Search, available at: <http://www.warnerchappell.co.uk/wcmuk/song_search/search_menu/songsearch.jsp> (accessed 21 April 2009).

Received
Revised
Accepted

This paper is a pre-print version presented at the ISKO UK 2009 conference, 22-23 June, prior to peer review and editing. For published proceedings see special issue of Aslib Proceedings journal.

Take in Table I. Top 15 ranked musical facets

FACET	QUANTITY	CONTROLLED / FREE TEXT
Bibliographic		
Artist	6	Free text
Song title	6	Free text
Writer	6	Free text
Year	6	Controlled
Album title	3	Free text
Chart position	3	Controlled
One stop	3	Binary
Originating territory	3	Controlled
Descriptive		
Genre	6	Controlled
Keyword	6	Controlled / Free text
Tempo	6	Controlled
Lyrics	4	Free text
Mood	4	Controlled
Subject	4	Controlled
Vocal mix / instrumental	3	Controlled

Take in Table II. Frequency of “pop” and “rock”

Music Owner	Pop	Rock
Warner Chappell	6	8
EMI	3	10
Sony ATV	1	1
Chrysalis	1	9
Universal	302	209
Ricall	18	12

This paper is a pre-print version presented at the ISKO UK 2009 conference, 22-23 June, prior to peer review and editing. For published proceedings see special issue of Aslib Proceedings journal.

Take in Figure 1. Total word frequencies tag cloud - all search engines combined

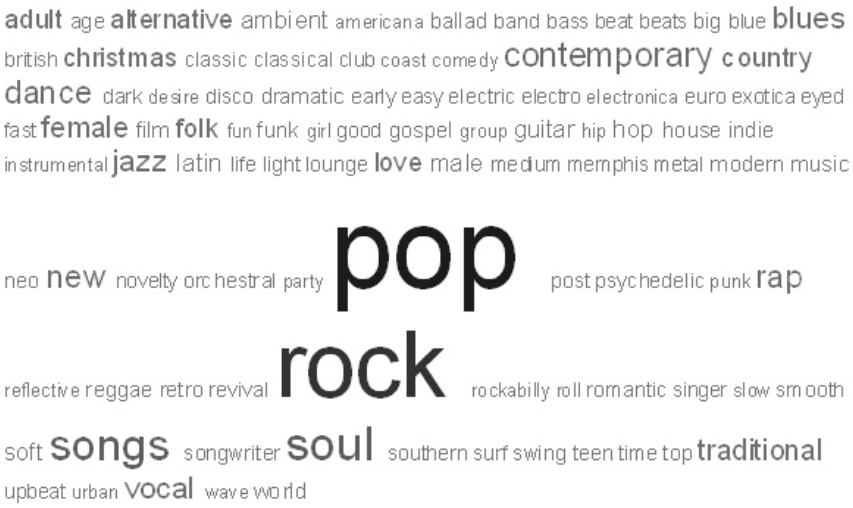

Take in Figure 2. Tag cloud - subjects - all search engines

Figure 3. Tag clouds - moods - all search engines