

Abbey E. Thompson. Playing Tag: An Analysis of Vocabulary Patterns and Relationships Within a Popular Music Folksonomy. A Master's Paper for the M.S. in L.S. degree. April, 2008. 33 pages. Advisor: Jane Greenberg.

The cataloging of musical materials has always posed challenges for librarians, requiring special treatment for classification and organization. Accurate description is critical for achieving high-recall retrieval and access by patrons. This study considers the challenge of music description through content analysis of the popular music websites AllMusic.com and Last.fm. The goal of the research was to gain a better understanding of how users describe their music collections. The findings in this study illustrate that hierarchical vocabulary structures are clearly evident within the Last.fm folksonomy. The findings also show that tagging data is more reliable in representing musical genre/subject than previously speculated, indicating that with proper analysis and coding, social tag data could be harvested to provide genre-level metadata for popular music titles. The work presented here contributes a methodology for further study of this topic, specific to music folksonomies and vocabularies, which may also be useful for other disciplines.

Headings:

Bibliographic control -- Music

Controlled vocabularies

Folksonomies

Music information retrieval

Popular music

PLAYING TAG : AN ANALYSIS OF VOCABULARY
PATTERNS AND RELATIONSHIPS WITHIN
A POPULAR MUSIC FOLKSONOMY

Abbey E. Thompson

A Master's paper submitted to the faculty
of the School of Information and Library Science
of the University of North Carolina at Chapel Hill
in partial fulfillment of the requirements
for the degree of Master of Science in
Library Science.

Chapel Hill, North Carolina
April 2008

Approved by:

Jane Greenberg

ACKNOWLEDGEMENTS

To my advisor, Jane, for her tireless encouragement of my efforts: You have been a source of constant reassurance, heart-warming pep-talks, and of course, chocolate from your secret stash. Without your help and mentorship, I would not have come this far.

To my mother and father, Marna and Jerry, for their love and enthusiasm: You have always been there for me, cheering me on at every step. Without your guidance and support, I would not be who I am today.

And to Nick, just because.

TABLE OF CONTENTS

Introduction	4
Introduction to Music Information Retrieval.	4
The Problem with Popular Music	5
Folksonomies: To Tag or Not to Tag.	6
Tagging Popular Music: Last.fm.	7
Research Objectives.	8
Background and Literature Review	10
Taxonomy vs. Folksonomy	10
Tagging of Specific Formats	11
Music Tagging Studies	13
The Role of Controlled Vocabularies	14
Methodology	17
Method of Analysis	17
AllMusic.com.	17
Last.fm.	19
Coding, Method 1	20
Coding, Method 2	21
Findings/Results	23
Implications and Conclusions	27
Bibliography	29

LIST OF TABLES

Table 1: Billboard Top-Five Singles.	20
Table 2: AllMusic Facet Analysis.	23
Table 3: AllMusic Match Analysis.	24
Table 4: Description of Music Folksonomy Facets	24
Table 5: Music Folksonomy Facet Analysis.	25

INTRODUCTION

Introduction to Music Information Retrieval

The indexing, organizing, and collocation of musical materials has always posed challenges for librarians and scholars alike. Unlike textual resources, music documents express ideas and concepts that go beyond the bounds of common linguistic communication, therefore requiring special treatment by those seeking to classify or organize them. The number of possible access points for any one musical object is far above the number for the typical textual work. The careful and deliberate description of musical works situated within information retrieval systems is critical for achieving high-recall retrieval and access by patrons, since they may be searching for that work from any number of possible angles. As we collectively move further into a world of technological cooperation and interoperability, it is important for library catalogs and bibliographic/metadata records to function and interact with knowledge organization systems beyond the bounds of the library itself in order to fully satisfy user demand and expectations.

When computerized bibliographic records began to enter the marketplace in the mid-twentieth century, libraries promoted and encouraged greater adherence to standardization and cooperation in their cataloging. This allowed the individual records they created to function properly in another library across the country, or even the world.

This shift in the cataloging community promoted even more rigorous standards for classification and description than in years past, such as the *Anglo American Cataloging Rules* (AACR) and its format-standard counterpart, *MARC* (MAchine Readable Cataloging). Even then, the codification of rules for cataloging music documents was difficult to use; for nearly every rule, there would be an exception. As more and more audio recordings were collected by libraries and archival institutions, these problems only became more glaring, as audio content is – by virtue of its format – even harder to quantify in text than written music.

While some cataloging metadata for audio works is usually fairly easy to procure (such as artist name, composer and title of the work), once one moves beyond the author/publisher supplied textual data, the boundaries become a lot less clear. Subject description is particularly problematic: how does one translate the essence of an audio file into text? How does one distinguish between the *genre* of a piece of music (classical, popular, jazz, etc.), the *instrumentation* (choral, solo piano, string quartet, etc.), and the *subject* of the work (love, death, heartache, etc.)? Traditional cataloging for both printed and audio musical works tends to focus more heavily on the first two types of “subject” associations, genre and instrumentation. It is the latter type of subject access – theme, concept, author intention, or sometimes lyrics – which is severely lacking from most formalized cataloging systems both within and outside the library realm.

The Problem with Popular Music

Classifying and categorizing audio recordings for patron access is already a difficult process. Musical genre is notoriously hard to pinpoint, and the concept itself is difficult to

define (Aucoutrier & Pachet; McKay & Fujinaga). For library collections which consist of primarily “classical” music, or that of the Western art music tradition, there is a substantial amount of literature and guidance to aid in their classification and categorization, primarily because this class of music is heavily studied in academic circles and garners the most attention within the academy. Support for the study of popular music within academia has been growing for decades, yet is still underrepresented. Because of this, for the collectors of popular music there is far less cataloging and classification assistance available, if any. Lacking the hierarchical structures of classification schemes and vocabularies of the classical music sphere, popular music is still somewhat the “red-headed stepchild” of the music bibliographic universe. The ever-changing, ever-growing field of popular music can be hard to quantify due to its fluidity and likely also because of our lack of distance from the subject at hand. Trends come and go in an instant, and it takes time to develop the kinds of formalized vocabulary structures that librarians and scholars are familiar with.

Folksonomies: To Tag or Not To Tag?

The catchphrase “Web 2.0” and its subsequent library-themed derivative, “Library 2.0,” are now quite well-known in the scholarly community (Miller, 2005). The phenomenon in question relates to harnessing the power of savvy and opinionated Web users to submit their thoughts, rankings, ratings, tags, or other input to digital content. The bigger concepts at play in this 2.0 trend are those of collaboration, the sharing of information, and social networking of online peers. Websites such as Flickr, Del.icio.us and others have received much scholarly attention in recent years for their use of social

tagging, or the collective annotation, indexing and collocation of online objects through user-defined vocabularies. Users of these types of websites can label (or “tag”) objects as they see fit, as well as view and browse the labels applied by other users. The more times an item is labeled with a certain tag, the more that tag is assumed to be relevant. The type of collection of tags eventually became known as a *folksonomy*. This term was coined in 2004 by Thomas Vander Wal via a listserv post, in which he described a folksonomy as “user-created bottom-up categorical structure development with an emergent thesaurus” (Vander Wal, 2007). Since then, folksonomy and social tagging has only continued to expand into new genres of digital content.

Tagging Popular Music: Last.fm

Last.fm is a popular internet music networking site where users are encouraged to tag the artists and tracks in their personal music collections. Last.fm then uses this community-based tagging data to collocate like artists and tracks, performing recommendations based on the listening habits of its user base. The Last.fm FAQ describes their tags as

keywords or labels. You can assign as many tags as you like to any track, album, or artist in your musical profile. Tags are a great way to label items by genre...but the possibilities are endless. Think a song would sound great on the highway? Tag it with ‘driving’. If you find an album very relaxing, tag it with ‘chillout’ or ‘relax’ (or both). Keep a list of artists to tell a friend about by creating a tagset called ‘singers Sarah would like’. (Last.fm Ltd., 2008)

In this way, Last.fm can be seen as actively encouraging users to think outside the bounds of traditional music description, to use tagging as a method of personal music organization. In contrast to AllMusic.com’s hierarchical, editor-driven vocabulary, the tags found on Last.fm run the gamut from codified genre names to personal statements, to

strings of seemingly random characters whose meaning is only known to their creator. How, one might ask, can this data be used for any kind of reliable music description purposes? Detractors to folksonomy use often ask this very question, arguing that the amount of tag-debris lowers the overall effectiveness of the created vocabulary; because the folksonomy vocabulary is flat in structure without any sort of defined semantics, it cannot be trusted. However, this conjecture is merely a set of assumptions. Numerous studies have been performed on other social tagging websites (such as Flickr and Del.icio.us) to analyze the vocabulary structures within that particular tagging community, yet to this date, very little analysis has been done regarding the tagging of popular music and the Last.fm community.

Research Objectives

The purpose of this research study is to explore the popular music tagging community of Last.fm and perform analysis on tagging data collected from their website. Through analysis of this tagging data, we can better understand how users describe their music collections, and also begin to form a more consistent methodology for tag analysis in the genre of popular music. The following two objectives guided the work presented in this paper:

1. *To explore AllMusic.com's controlled popular music vocabulary and compare it to the tagging vocabulary of Last.fm?*
2. *To develop a better system of facets or categories for analysis of popular music tagging data*

A review of relevant literature on this topic supports these objectives and the analysis of the study's findings. Following the literature review is a discussion of methodology, limitations of the study, data analysis, the study results/findings, and suggestions for next steps within this research area.

BACKGROUND AND LITERATURE REVIEW

Taxonomy vs. Folksonomy

There has been intense debate over the usefulness of folksonomies for classification, cataloging, metadata generation, and traditional search and retrieval activities (Dye, 2006; Gruber, 2007; Mathes, 2005; Peterson, 2006; Shirky 2005). Folksonomies are, according to their critics, notoriously imprecise, full of tags that are “often ambiguous, overly personalized and inexact,” resulting in “an uncontrolled and chaotic set of tagging terms that do not support searching as effectively as more controlled vocabularies do” (Guy & Tonkin, 2006). Elaine Peterson goes so far as to state that the very principle of folksonomies has deep philosophical flaws:

A folksonomy universe allows both true and false statements to coexist. Because tags are relativized, personal, idiosyncratic views can coexist and thrive in the form of tags, in spite of their inconsistencies...Even should all interpretations be of equal worth, if users can continuously add tags to articles, at some point it is likely that the whole system will become unusable. A folksonomic system threatens to undermine its own usefulness. (Peterson, 2006)

Despite the lack of command and authority that traditional controlled vocabularies wield, it has been proven that not all folksonomy structures are without their own hierarchical relationships and structure. Recent studies, such as Kome (2005) argue that because categorization is a “fundamental human cognitive activity,” semantic relationships do in fact exist within much folksonomy metadata (p. 2). Kome’s study of tags from

Del.icio.us discovered that hierarchical relationships were common, found in nearly 45% of all tested data, and as high as 90% in certain areas (p. 17).

If we look at folksonomies as more of a social phenomenon rather than an effort to replace top-down hierarchical vocabularies, more interesting possibilities arise: Jessica Dye, in her 2006 article “Folksonomy: A game of high-tech (and high-stakes) tag,” notes that search engines “are constantly trying to think like people: how they search and how they say what they mean. Human-generated metadata, when applied correctly, can be more valuable than that generated by a robot” (p. 43). Gathering data on user perspective and description of library resources is not a new phenomenon: user studies have been used for decades to learn more about how patrons think and describe resources. Using folksonomy and tagging data is merely a new approach for this type of research.

Tagging of Specific Formats

Research of folksonomy structures and tagging behavior has grown considerably in the scholarly and research community, and one genre in particular which is showing interesting developments is in the tagging of images. The indexing and description of images shares many of the same complications as music indexing, since they are both genres which are based on non-textual documents.

Joan Beaudoin (2007) performed a study of tags on Flickr to look for underlying patterns within the tags themselves, and then use those patterns “to alleviate some of the problems associating with tagging” (p. 26). She gathered the top ten tags of fourteen randomly-chosen images from Flickr and subsequently applied conceptual labels to each

of the tags, eventually agreeing on eighteen labels or facets. Beaudoin's methods for producing this faceted system was used as a basis for this study of Last.fm tagging data.

Marlow et al (2006) also examined tagging data from Flickr in their development a tagging system model. Their view is much more focused on the social networking aspect of tagging than Beaudoin's study. Rather than creating a system of labels or facets to place already-developed tags, they instead conceptualize ways in which tagging systems could be improved for future data-mining, and then apply this concept to Flickr. Their intent was to describe "key dimensions of tagging systems' design that may have immediate and considerable effect on the content and usefulness of tags generated by the system" (p. 34).

Language and general vocabulary clearly play a very strong role in tagging systems for non-textual items such as images and musical objects. Elaine Ménard, in the October/November 2007 edition of *JASIST*, discussed collaborative tagging as the "latest trend" in image indexing. Ménard performed a partial analysis on data gathered comparing the indexing of terms in both controlled and uncontrolled (i.e., folksonomy) vocabularies. Her results so far indicate a propensity in uncontrolled vocabularies for terms "referring to size, color, texture, gender or trademarks, contrary to controlled vocabularies which have a tendency to be less graphic and, in many cases, less detailed or descriptive" (p. 24). A more extensive examination of her data is still in progress, but at this time she recommends moving forward with collaborative tagging systems, and that if used in conjunction with controlled vocabularies "the two approaches may co-exist and be very helpful" (p.25).

Music Tagging Studies

To date, a few studies have examined music social tagging systems. The study performed by Geleijnse, Schedl, and Knees (2007), like the study presented in this paper, examined tagging data from Last.fm. Their focus, however, was on artist description rather than general music description, to see whether “Last.fm data can be used to generate a ground truth to describing musical artists” (p. 525).¹ Their findings showed that similar artists share a large number of tags, and that Last.fm data proved to be reliable in generating the ground truth sought by the researchers. The authors conclude the paper by proposing more research involving community based tagging data for music information retrieval, asserting that the very fact that a large community of users determines the data (rather than a small community of experts) results in data that is much richer and more robust.

Levy and Sandler (2007) chose to examine music social tags from the Last.fm and MyStrands web services as a potential high-volume source of semantic metadata. Unlike the other studies discussed here, the harvested and studied tags were all from individual tracks. Some 45,000 tags were incorporated into their dataset, helping to show that “tags define a vector space with highly attractive properties for music retrieval, and which appears to have genuine semantics” (p. 411). The authors conclude that despite the fact that many music tags are discursive in nature, they “appear to capture sensible attributes grounded in individual tracks, defining a well-behaved similarity space with an effective dimensionality of around 10^2 ” (p. 416).

¹ The term “ground truth,” as defined by the *Oxford English Dictionary*, refers to “information obtained by direct observation of a real system, as opposed to a model or simulation; a set of data that is considered to be accurate and reliable, and is used to calibrate a model, algorithm, procedure, etc.”

Hu and Downie (2007) chose to focus on music “mood metadata.” Mood, they assert, is a music access feature which remains poorly understood by traditional music information retrieval efforts. “Since mood is a very subjective notion,” they state, “there has yet to emerge a generally accepted mood taxonomy that is used within the MIR research and development community” (p. 67). The aims of their study were to explore the relationships between musical mood and genre, artist, and usage metadata in order to develop automated mood access techniques. Comparisons were made across three metadata collections, AllMusicGuide.com, eopinions.com, and Last.fm. The results of their study were somewhat inconclusive: the relationship between mood and usage statistics were “stable enough to warrant further consideration,” while the mood-genre and mood-artist relationships “show great promise” (p. 72). Hu and Downie note that the major obstacle to further research is the sheer volume of possible mood vocabulary terms, as many mood terms in use are highly synonymous.

The Role of Controlled Vocabularies

At this time there are a great deal more formal structures in place for describing images than for music, such as the *Art and Architecture Thesaurus* (AAT), the *Thesaurus for Graphic Materials I and II* (TGM I and TGM II), and the *ICONCLASS* vocabulary (Ménard, 2007). The field of music still lacks such important vocabulary tools for either the classical or popular music sphere. The *Library of Congress Subject Headings* (LCSH), which is currently used for subject access and cataloging of musical works in most music libraries, shows a distinct bias towards Western art music and is insufficient in many ways for popular music subject access. Though many new stylistic terms for

popular music have been added to LCSH in recent years, their use by catalogers remains minimal and inconsistent, most popular music recordings receiving only a generic heading such as “Popular music” or “Rock music” with the occasional chronological subdivision (Cronquist, 2004, p. 4).

The Music Thesaurus Project was originally conceived in 1989 by a working group of the Music Library Association. The goal of the *Music Thesaurus* was:

a system that supports both pre-coordinate indexing (in which terms are combined by the indexer) and post-coordinate searching (in which users combine their own desired descriptors from a list of single terms) [that] would more clearly define such multielement works and provide for more accurate retrieval of music materials. (McKnight, Griscom, & Young, 1989, p. 715)

Several papers were published in the 1990s concerning the project and its progress, but since 2000, no new results have been published and progress on the *Thesaurus* seems to have stagnated (Cronquist, 2004, p. 9). In recent years, an effort has been made by scholars at Brigham Young University and IAML (the International Association of Music Libraries) to develop an international music thesaurus, but as of yet, nothing concrete has been produced (Spilker, 2005).

The nearest development to a controlled vocabulary for popular music at this point in time can be found on the website AllMusic.com. Launched in 1995, AllMusic.com was intended as “a place for music fans to indulge their passion.” With the aid of an editorial staff and a team of expert contributors, AllMusic bills itself as “the most comprehensive music reference source on the planet” (AllMedia Guide, 2008). A variety of data about artists and musical works, both popular and classical, is available free of charge to the public, along with what AllMusic describes as “relational content,” or details about artists and pieces of music to illuminate stylistic, genre, and other

similarities. Social tagging and Web 2.0 concepts are not currently a part of their business model; AllMusic oversees and controls classification and description through their own editorial staff, which is why their hierarchical model is much closer to a controlled vocabulary model than a folksonomy. Besides searching their database for artist, album, or song titles, users have the option to browse their database via genres/styles of music, specific instrument types, geographical regions, as well as musical *moods* (such as “happy,” “sentimental” and “dramatic”) and *themes* (such as “Christmas Party,” “Hanging Out” and “Road Trip”). The AllMusic vocabulary structure is discussed in more detail in the Methodology section. However, it should be noted that the *moods* and *themes* browsing capability represents a true shift in the formal description of musical items, away from the types of safe, tangible data favored by traditional cataloging and categorization methods towards the less-stable – but very fulfilling! – realm of opinion and emotion. It is these types of subjective, natural-language descriptions that are found in abundance in folksonomy structures. Because of the emotional attachment that listeners have to their music collections and favorite artists, the concept of mood, emotion, and other types of descriptive metadata for popular music is worth investigating further. By examining how users describe their own music collections and tastes, we can better understand how to classify and organize them for the sake of music information retrieval. Analysis of the folksonomy and tagging patterns of popular music social networking website users is a way to achieve this goal.

METHODOLOGY

Method of Analysis

This study employs content analysis method to obtain information about popular music tagging and vocabulary patterns. The content analysis was performed on vocabulary data from two websites, AllMusic.com and Last.fm for purposes of comparison between a controlled vocabulary and a broad folksonomy. This comparison was then used to develop a more rigorous system of facets or categories for analyzing popular music folksonomy data in future studies. Content analysis was chosen as the best method to examine the research objectives, and due to the availability of data for analysis. This method was also desirable because it is unobtrusive; vocabulary and tagging decisions could be observed without any interaction or interference from the researcher.

AllMusic.com

AllMusic.com was selected to represent the controlled popular music vocabulary. Originally known as the *All Music Guide*, a print resource, the company was founded in 1991 with the intent “to help consumers navigate the increasingly complex world of recorded music and discover the very best recordings”(AllMusic Guide, 2008). The website itself does not define its music categorization system as a “controlled

vocabulary,” but the faceted structure of genres, themes, moods, and other musical descriptors and categorizations shown on the website is collected, organized, and maintained by AllMusic’s editorial staff, giving the vocabulary a sense of legitimacy and authenticity. The navigation tabs at the top of the page allow users to browse the vocabulary facets, including a collection of larger genres (such as “Rock,” “Jazz,” “R&B,” etc.), “Instruments,” “Countries,” “Moods,” and “Themes.” AllMusic shows distinct differences when compared to the more traditional categorization system in the *Library of Congress Subject Headings*. LCSH primarily focuses on distinctions of genre, instrumentation, and to some degree, geographical information. AllMusic has expanded their vocabulary to include subjective information via their “Moods” and “Themes” facets. The “Moods” facet is comprised primarily of adjectives related to the emotional experience of music-listening. Some examples from this category are: “Bittersweet,” “Cheerful,” “Hypnotic,” and “Soothing.” The other subjective facet of the AllMusic vocabulary is “Themes,” a collection of extra-musical associations that listeners would possibly make with a track or artist. Examples include: “Background Music,” “Dinner Ambiance,” “Patriotic,” and “Slow Dance.” Through examination and analysis of the AllMusic vocabulary, four main categories or facets emerged from their somewhat cumbersome tab-navigation browsing system:

1. *Genres: Vocabulary terms which describe the genre of the musical work or artist.*
2. *Audio Attributes: Vocabulary terms which describe the prominent instrument of the musical work or artist.*

3. *Geographic: Vocabulary terms which associate the musical work or artist with a country, region, or geographic identifier.*
4. *Opinions: Vocabulary terms which represent subjective or opinion data of the AllMusic editors (including both “Moods” and “Themes”)*

This four-facet system was used as a basis of measurement for two sets of tagging data downloaded from Last.fm.

Last.fm

In order to achieve a broad view of tagging activity, two sets of tag data were downloaded from the Last.fm website. Data Set #1 (DS-1) consists of the one hundred most frequently-used tags within the Last.fm database. The information was collected using the *Audioscrobbler* system, described as “a database that tracks listening habits and does wonderful things with statistics” (Audioscrobbler Ltd., 2008). Realtime access to *Audioscrobbler* data is available online through specific browser URLs. DS-1 was acquired on March 11, 2008 via the following URL:

<http://ws.audioscrobbler.com/1.0/user/RJ/tags.txt>. These top 100 tags were then migrated into an Excel spreadsheet for coding and analysis. Included with each tag was the total number of times the tag had been used and a static URL. For example, the most popular tag in the Last.fm database is “rock,” which has been used 1,742,353 times and can be found at <http://www.last.fm/tag/rock>.

In order to achieve a more complete view of tagging activities on Last.fm, a second data set was needed to examine tagging on an individual-track level. Five single popular music tracks were selected for tag analysis in Data Set #2 (DS-2). The sampling process

for selecting these tracks was purposive. The *Billboard Music Charts* have been serving the information needs of the entertainment business since 1894 and are widely regarded as an authoritative source for statistical popular music information (Nielson Business Media, Inc, 2008). Five separate singles charts were selected to represent five different segments of the popular music world: *Adult Contemporary*, *Hot Country*, *Hot Dance Club Play*, *Hot R&B/Hip Hop*, and *Modern Rock*. Data was gathered from the March 8, 2008 issue of *Billboard* magazine (No. 1 on the Charts, 2008, p. 3). For each of these five charts, the number one single was chosen for tag analysis.

Table 1: Billboard Top-Five Singles

Chart	Title	Artist
Adult Contemporary	<i>Bubbly</i>	Colbie Caillat
Hot Country	<i>Cleaning This Gun (Come on Boy)</i>	Rodney Atkins
Hot Dance Club Play	<i>Piece of Me</i>	Britney Spears
Hot R&B/Hip Hop	<i>I Remember</i>	Keyshia Cole
Modern Rock	<i>Long Road to Ruin</i>	Foo Fighters

All tags assigned to each of these five tracks were transcribed into an Excel spreadsheet for coding and analysis.

Coding, Method 1

The first coding method involved comparing both data sets to the AllMusic four-facet vocabulary system. Each tag from both sets was analyzed and placed into one of the AllMusic facets (Genres, Audio Attributes, Geographic, Opinions). If none of the facets seemed to fit the tag in question, it was placed in the Other/None category. The only tags excluded from analysis were those in foreign languages, as this study is limited to English-language tags. In total twelve tags were excluded from DS-2. The tags were also

compared to the AllMusic vocabulary for similarity. One of three labels was attached to each tag to represent its degree of match:

1. *(Y)es*: The Last.fm tag appears within the AllMusic vocabulary (or within the artist, album, or title metadata) exactly as written.
2. *(P)artial*: The Last.fm tag is a partial match to an AllMusic vocabulary term.

Partial matches include variant spellings, variant concatenation of multiple words, synonyms and other similarities. For each partial match, the AllMusic vocabulary term or terms to which the Last.fm tag matches is listed.

3. *(N)o*: The Last.fm tag does not appear within the AllMusic vocabulary.

Coding, Method 2

The purpose of the second coding method was to develop a more detailed facet system for analyzing popular music folksonomy tags. The four-facet system presented by AllMusic was not sufficient for fully analyzing the underlying patterns in the tagging data; it is too crude of a tool. Therefore, an approach was taken similar to that of Beaudoin with Flickr tags, applying labels to the Last.fm tags in an iterative fashion, going through both lists of tags looking for patterns and similarities. This process produced a total of eleven facets. Like in Beaudoin, tags were allowed to be assigned to multiple categories as needed, since many of the facets developed were not mutually exclusive. The eleven facets are as follows: *Personal Meaning*, *Extra-musical Associations*, *Superlative Expressions*, *Descriptors*, *Charts*, *Genres*, *Geographic*, *Audio Attributes*, *People*, *Song ID*, and *Time*. The facets are described in detail in the subsequent Findings/Results section. Drawing from the research of Al-Khalifa and Davis

(2007), these facets were grouped into three broad groups: Personal, Subjective, and Factual tags:

(P)ersonal tags: These are tags that have an intended audience. They are often used to organize a user's own resources, and can be roughly classified into: self-reference tags, task and time management tags and others.

(S)ubjective tags: These are tags that express people's opinions on the bookmarked web resource.

(F)actual tags: These are tags which identify 'facts' about the described web resource such as people, places, or concepts. (p. 165)

One facet, *Charts*, was placed in both the **Subjective** and **Factual** groups. It was often unclear whether a tag was referring to a published chart (such as *Billboard*) or one's personal chart choices, so it was necessary to place this facet within both groups.

FINDINGS/RESULTS

The two sets of tag data collected were very different in content. DS-1, the one hundred most frequently-used tags, contained primarily genre data. The other four facets each received less than ten percent of the overall total. Results of facet distribution in Data Set #2, tags from the selected top five single tracks, were slightly more evenly distributed. **Table 2** shows the distribution patterns of DS-1 and DS-2 within the AllMusic four-facet system:

Table 2: AllMusic Facet Analysis

<i>Facet</i>	DS-1	DS-2
<i>Genres</i>	68 68.0%	49 23.8%
<i>Audio</i>	6	22
<i>Attributes</i>	6.0%	10.7%
<i>Geographic</i>	8 8.0%	8 3.9%
<i>Opinions</i>	9 9.0%	80 38.8%
<i>Other/None</i>	9 9.0%	47 22.8%

The propensity towards genre tags in DS-1 shows that Last.fm users are concerned, even if unconsciously, with genre identification as a primary function of tagging. When looking at the spread of tags across a single track, however, there is much more variety. In the five single tracks examined for this study, nearly 39% of tags were placed in the *Opinion* category. The disparity between the percentage of *Opinion* tags in the two data

sets is likely due to the high degree of vocabulary variation in how users state their personal opinions. *Genres*, in contrast, tend to conform more readily to community standards; “classic rock” is a more standard term than “awesome” or “catchy.” **Table 3** shows the degree to which the tags in each data set conform to the AllMusic controlled vocabulary:

Table 3: AllMusic Match Analysis

<i>Degree of match</i>	DS-1	DS-2
(Y)es	38 38.0%	37 18.0%
(P)artial	35 35.0%	39 18.9%
(N)o	27 27.0%	130 63.1%

The high amount of tags in DS-2 which could only be classified in the *Other/None* category (22.8%) justified the development of this study’s new eleven-facet system for tag analysis. In order to clearly see the distribution of tag types and patterns, and to facilitate further discussion of popular music folksonomies, a system must be in place to analyze them. The eleven facets which resulted from analysis of the two data sets are described below:

Table 4: Description of Music Folksonomy Facets

Group	<i>Facet</i>	Description
Personal	<i>Personal Meaning</i>	A tag which has an intended audience or is of use only to the author
Personal	<i>Extra-musical Associations</i>	A tag which associates the item with a non-musical concept, theme, idea, or entity
Subjective	<i>Superlative Expressions</i>	A tag in which the author implies a hierarchy or personal system of rank
Subjective	<i>Descriptors</i>	A tag which can be said to describe the item in some way

Subjective	<i>Charts</i>	A tag which associates the track with some kind of chart, either official (<i>Billboard</i> , etc.) or unofficial
Factual	<i>Genres</i>	A tag which describes musical genre
Factual	<i>Geographic</i>	A tag in which a geographical location is identified or associated with the item
Factual	<i>Audio Attributes</i>	A tag which describes an audio attribute of the item, such as instrumentation or vocal type
Factual	<i>People</i>	tag which names an individual or musical group
Factual	<i>Song ID</i>	A tag which incorporates the name of the track
Factual	<i>Time</i>	A tag which associates the item with some denomination of time

When these new facets were applied to the two data sets, new patterns emerged. Though the tags from DS-1 still reflected a dominance of genre identifiers (68%), the tags from DS-2 now show a much different distribution than via the more imprecise AllMusic facet system. **Table 5** shows the results from this analysis:

Table 5: Music Folksonomy Facet Analysis

<i>Facet</i>	DS-1	DS-2
<i>Personal Meaning</i>	4 4.0%	35 17.0%
<i>Extra-musical Associations</i>	4 4.0%	20 9.7%
<i>Superlative Expressions</i>	2 2.0%	17 8.3%
<i>Descriptors</i>	9 9.0%	62 30.1%
<i>Charts</i>	0 0.0%	11 5.3%
<i>Genres</i>	68 68.0%	50 24.3%
<i>Geographic</i>	11 11.0%	10 4.9%
<i>Audio Attributes</i>	8 8.0%	25 12.1%
<i>People</i>	0 0.0%	9 4.4%
<i>Song ID</i>	0 0.0%	2 1.0%

<i>Time</i>	5	14
	5.0%	6.8%
<i>Other</i>	0	11
	0.0%	5.3%

Though many critics of folksonomy use for classification note the excess of “personal meaning” tags, or those which have an intended audience or is of use only to the author, it should be noted that these types of tags make up only 17% of DS-2 and a mere 4% of DS-1. Similarly, “other” tags, or those which are either unknown by the researcher, nonsensical or the author’s intent is unclear, only comprise 5.3% of DS-2 and are entirely absent from DS-1.

IMPLICATIONS AND CONCLUSIONS

This study is a much-needed first step in the realm of popular music folksonomy analysis. Though many tag-analysis studies have been performed on other folksonomy-based websites, such as Del.icio.us and Flickr, popular music folksonomies have received very little attention from the scholarly community. More tag analysis in this domain will help us to gain a better understanding of how users describe their music collections, and therefore aid with music information retrieval and discovery. The development and implementation of more consistent methodologies for tag analysis in the genre of popular music, such as the one presented in this study, will also help achieve this goal. This study takes on that challenge.

Although the sample size was limited due to practical research constraints, the results provide insight into popular music tagging behaviors and trends. The findings in this study illustrate that hierarchical vocabulary structures are clearly evident within the Last.fm folksonomy. Additionally, the results show there are notable similarities between popular music folksonomy structures and those for images. Language and vocabulary choice play key roles in any collection of non-textual items, and the propensity towards uncontrolled vocabularies in both image and musical object tagging systems is paramount. Furthermore, the results show obvious differences as well, such as the element of mood, which plays a much stronger role in musical description than for

images. Most importantly, the findings also show that tagging data is more reliable in representing musical genre and subject than previously speculated, indicating that with proper analysis and coding, social tag data could be harvested to provide genre-level metadata for popular music titles. A contribution of this work is that it has mapped out a methodology for further study of this topic, specific to music folksonomies and vocabularies, which may also be useful for other disciplines.

Future research in this area could include more analysis of data from AllMusic and/or Last.fm. Additional networking sites for popular music, such as Pandora.com or MySpace.com, would also be worthwhile to study. Analysis with the facet system developed in this study using data from other websites and sources would be beneficial for refining or expanding the current facet system. Description of popular music is a difficult task due to the fluid nature of the field, combined with a lack of rigorous and codified structures for analysis of genre and style. More research into how people describe and organize their own music collections can help librarians and researchers to better understand the classification and organization of popular music for information retrieval.

BIBLIOGRAPHY

- Aucoutrier, J. & Pachet F. (2003) Representing musical genre: A state of the art. *Journal of New Music Research* 32(1), 83-93.
- Al-Khalifa, H. & Davis, H. (2007) Towards a better understanding of folksonomic patterns, in *Eighteenth ACM Conference on Hypertext and Hypermedia* (HT 07, 10-12 September, Manchester, UK), ACM Press, New York, NY, 163-166.
- AllMedia Guide (2008). *About us*. Retrieved April 14, 2008, from http://www.allmusic.com/cg/amg.dll?p=amg&sql=32:amg/info_pages/a_about.html
- Audioscrobbler Ltd. (2008). *AudioScrobbler – The music technology playground from Last.fm*. Retrieved April 14, 2008, from <http://www.audioscrobbler.net/>
- Beaudoin, J. (2007) Flickr image tagging: patterns made visible. *Bulletin of the American Society for Information Science and Technology* 34 (1), 26-29.
- No. 1 on the charts. (2008, March 8). *Billboard*, 120(10), 3.
- Cronquist, M. (2004) *Subject access to music: An evaluation of the Music Thesaurus Project*. Unpublished master's paper, University of North Carolina at Chapel Hill.
- Donkar, A. (2007) I'm in ur bookmarks, stealin' ur tags: Closed communities and their influence on consistent vocabularies in user developed folksonomies. Unpublished master's paper, University of North Carolina at Chapel Hill.
- Dye, J. (2006). Folksonomy: A game of high-tech (and high-stakes) tag. *E-Content Magazine*, 29(3), 38-43.
- Geleijnse, G., Schedl, M. & Knees, P. (2007) The quest for ground truth in musical artist tagging in the social web era, in *Proceedings of the 8th International Conference on Music Information Retrieval* (ISMIR 2007, September 23 - 27, Vienna, Austria), Austrian Computer Society, 525-530.
- Gruber, T. (2007) Ontology of folksonomy: A mash-up of apples and oranges. *International Journal on Semantic Web & Information Systems* 3(2). Retrieved April 14, 2008, from <http://tomgruber.org/writing/ontology-of-folksonomy.htm>

- Guy, M., & Tonkin, E. (2006) Folksonomies: Tidying up tags? *D-Lib Magazine*, Vol. 12(1). Retrieved April 14, 2008, from <http://www.dlib.org/dlib/january06/guy/01guy.html>
- Gruber, T. (2007) Ontology of folksonomy: A mash-up of apples and oranges. *International Journal on Semantic Web & Information Systems* 3 (2). Accessed April 14, 2008, from <http://tomgruber.org/writing/ontology-of-folksonomy.htm>
- Hu, X. & Downie, J. (2007) Exploring mood metadata: Relationships with genre, artist and usage metadata, in *Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007, September 23 - 27, Vienna, Austria)*, Austrian Computer Society, 67-72.
- Kome, S. (2005) *Hierarchical subject relationships in folksonomies*. Unpublished master's paper, University of North Carolina at Chapel Hill.
- Last.fm Ltd. (2008) *Frequently asked questions – Last.fm*. Retrieved April 14, 2008, from <http://www.last.fm/help/faq/>
- Levy, M. & Sandler, M. (2007) A semantic space for music derived from social tags, in *Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007, September 23 - 27, Vienna, Austria)*, Austrian Computer Society, 411-416.
- Marlow, C., Naaman, M., Boyd, D. & Davis, M. (2006) HT06, tagging paper, taxonomy, Flickr, academic article, to read, in *HYPERTEXT '06: Proceedings of the seventeenth conference on Hypertext and hypermedia* (Odense, Denmark, Aug 22-25, 2006), ACM Press, New York, NY, 31-40.
- Mathes, A. (2005) *Folksonomies - Cooperative Classification and Communication Through Shared Metadata*: Retrieved April 14, 2008, from <http://www.adammathes.com/academic/computer-mediated-communication/folksonomies.html>
- McKay, C. & Fujinaga, I. (2006) Musical genre classification: Is it worth pursuing and how can it be improved? In *Proceedings of the Seventh International Conference on Music Information Retrieval (ISMIR '06, October 2006, Victoria, Canada)*, 101-106.
- Ménard, E. (2007) Image indexing: How can I find a nice pair of italian shoes? *Bulletin of the American Society for Information Science and Technology* 34 (1), 21-25.
- McKnight, M., Griscom, R. & Bradford Young, J. (1989) Improving access to music: A report of the MLA Music Thesaurus Project Working Group. *Notes* 45 (4), 714-721.

- Miller, P. (2005). Web 2.0: building the new library. *Ariadne*, 45 (October, 2005). Retrieved April 14, 2008, from <http://www.ariadne.ac.uk/issue45/miller/>
- Morrison, P.J. (2007) Why Are They Tagging, and Why Do We Want Them To? *Bulletin of the American Society for Information Science and Technology* 34 (1), 12-15.
- Neal, D. (2007) Folksonomies and Image Tagging: Seeing the Future? *Bulletin of the American Society for Information Science and Technology* 34 (1), 7-11.
- Nielson Business Media, Inc. (2008) *Billboard.com – About us*. Retrieved April 14, 2008, from http://www.billboard.com/bbcom/about_us/index.jsp
- Peterson, E. (2006) Beneath the metadata: Some philosophical problems with folksonomy. *D-Lib Magazine* 12(11). Retrieved April 14, 2008, from <http://www.dlib.org/dlib/november06/peterson/11peterson.html>
- Shirky, C. (2005) *Ontology is Overrated: Categories, Links, and Tags*. Retrieved April 14, 2008, from http://shirky.com/writings/ontology_overrated.html
- Spilker, J. (2005) Toward an international music thesaurus. *Fontes Artis Musicae* 52 (1), 29-44.
- Vander Wal, T. (2007) *Folksonomy*. Retrieved April 14, 2008, from <http://vanderwal.net/folksonomy.html>