Towards a Music Recommendation Infrastructure

Daniel McEnnis
University of Waikato
Hamilton, New Zealand
dm75@cs.waikato.ac.nz

ABSTRACT
Contrary to opinions widely voiced in the popular media, music is not a universal language. Music exists in all cultures, but what it means and how it is interpreted differs on many levels. This paper seeks to break away from data-driven analysis that makes implicit assumptions about musical properties common among all cultures and replace it with a framework built on social sciences research. The Relational Analysis Toolkit (RAT) provides structure and support for building musical recommendation systems that are not ethno-centric and integrates differences in musical taste and understanding across every level of analysis: intra-personal, inter-personal, geographical and cultural levels.

Categories and Subject Descriptors
E.1 [**Data Structures**]: Graphs and networks; J.5 [**Arts and Humanities**]: Music

General Terms
music recommendation

Keywords
music recommendation, framework, graph

1. INTRODUCTION

> "Though we recognize the importance of sound musical analysis as an essential step in our research, we have widened our horizon to the point where we consider music not merely as an esthetic object, but as a product and symbol of human behavior, inextricably associated and interrelated with the other elements of culture."
> W. Rhodes, 1963 p179 [35].

Traditionally, music recommendation researchers have avoided in-depth psychological or sociological analysis. Typically, the waveforms of the music are analyzed and higher-level features such as genre are predicted exclusively from the audio content. Other methods such as using web searches provide a 'context-based' (i.e. not in a musical signal) approach, augmenting existing music recommendation systems. This attempt to analyze musical taste and prepare music recommendation systems suffers from assumptions about culturally-invariant properties of musical taste. These assumptions make the algorithms developed silently dependent on only a given subsection of all music or are only accurate for a given subculture. In order to provide good recommendations across as much music as possible for as many people as possible, a new approach to music recommendation is needed that includes all relevant social factors without making assumptions about universal properties of musical taste.

The social sciences provide a radical, culturally-grounded approach. While computer science and music information retrieval (MIR) has largely restricted itself to the data easily collected, researchers in the social sciences have been working from a completely different set of assumptions that directly contradict the assumptions prevalent in MIR and computer science literature. However, in order to successfully use this information, a new framework for musical analysis is needed that provides concrete algorithmic functionality for analyzing this social data. The end product of these tools are concrete musical recommendations produced in a way that can be empirically validated. To this end, a toolkit is proposed that provides the tools needed to construct music recommendation systems that take advantage of the findings related to music analysis conducted by researchers in the social sciences.

This paper provides an overview of the findings and assumptions of the social science literature related to musical taste in Section 2. In Section 3, the existing music recommendation systems are described. In Section 4, the design requirements and existing toolkits are described. In Section 5 the Relational Analysis Toolkit (RAT) is described. Section 6 describes initial recommendations produced using this toolkit using LiveJournal friend-of-a-friend data. Section 7 provides an overview of work completed and Section 8 provides the direction of future work.

2. MOTIVATION FROM THE SOCIAL SCIENCES

While much of research in computer science has focused on getting working systems, there exists research in the social sciences that is focused on why we listen to the music that we do and is dependent on factors from many levels. The different levels of analysis as described by Doise [11] are: intra-personal, inter-personal, group factors, and cultural factors. Each level of analysis involves a different way of studying people and their reactions. These factors are given a concrete representation in social network analysis by Scott [38]. He describes these levels as studying the properties of an individual (actor), the properties of a group of individuals (actors), the properties of a sub-graph of individuals (actors), and studying the properties of an entire dataset as a whole [38].

Because of this overlap between a precise graph-based representation of society and the research findings that map to precise graph-structures, We have adopted this paradigm for music recommendation. Each level is described along with the research that analyzes aspects of musical taste at this level within the psychological literature and the sociological literature.

2.1 Intra-personal Factors

Intra-personal factors are those factors that occur within an individual. Each person is considered as an individual with no influences from factors of interactions with other people. The graph equivalent of this level assumes each actor is independent and studies the distribution of properties across each actor.

Individuals have differing preferences for music depending on their moods and their current tasks. The most comprehensive study is by North and Hargreaves [25] who found that the differences of preferences by mood and task vary so greatly that there is little correlation between the musical tastes of an individual in different mood/task combinations. Their paper follows initial research by Konecni that found that listening while performing other tasks causes a preference for less complex music [17]. The comprehensive study was extended who found that people have specific style preferences for certain tasks [29].

In the specific task of exercise music, North and Hargreaves showed that, at least in the part of the United Kingdom where the study was conducted, listeners show a clear difference when the task at hand is relaxation (yoga), exercise, or relaxing after exercise [24, 26]. Those participants assigned to yoga showed no clear preference, while exercisers chose up-tempo music while exercising and slow tempo music when relaxing after exercise.

On a more general scale, North and Hargreaves [28] showed a correlation between lifestyle choices and musical preference, lending itself to predicting musical taste from a person's lifestyle information.

Each of these studies provide evidence of the kind of information that can be extracted from studying the ways an individual's musical taste is influenced by different factor such as mood and task. This corresponds to analyzing and categorizing the properties of a given actor.

2.2 Inter-personal Factors

Inter-personal factors are the influences from one's local peer group that one interacts with on a regular basis. Inter-personal factors are especially well studied in terms of adolescents, where this effect is especially important. Interpersonal factors exclude those individuals not in contact with the group members on a regular basis. This level involves both studies of one's immediate circle of friends (immediate peer group) as well as all individuals one is in contact with (peer group).

For example, Aebischer[1] demonstrated that opinions of peers matter, even if the opinion is not from the immediate peer group. Ironically, opinions from non-immediate peers showed the greatest influence over extended time periods, but no change occurred immediately regardless of the source. This finding is consistent with descriptions of communication in social network analysis [38].

Tarrant et al. [42] demonstrated that adolescents use music as an identity marker. These findings are corroborated by North et al. [30] and Lamont et al. [18]. This effect is so pronounced that Tarrant et al. [41] found that musical tastes are so closely linked to group membership that adolescents musical preferences are correlated with overall well-being. This research indicates that factors that correlate with identity will be useful in predicting musical tastes.

Each of these studies provides insights into the kind of information that can be extracted with this level of analysis. This corresponds to analyzing the local neighborhood of an individual for the flow of musical tastes in peer groups.

2.3 Group Factors

Group factors are those influences that are larger than one's peer group but are not studied as an entire culture. Individuals need not be in regular contact with one another, but should identify with a similar larger-than-local subculture. Subcultures such as the Goth community [14] as well as less extreme examples such as gender differences [10] and participants in dance culture [43] provide examples of group factors that transcend one's immediate peer group but are not describing factors consistent across an entire culture. Empirical analysis at this level is lacking, but examples of the kind of information that can be extracted are demonstrated by the following case study of Hip Hop.

Hip Hop shows changes in its meaning and, hence, the social factors that are correlated with it, when a sub-group in a non-US culture reinterprets the meaning of this style of music in terms of their own cultural background. For Newcastle, England hip hop is an expression of working class values and working class problems as a reaction against classism. In Frankfurt, Germany, hip hop is an expression of nationalism for racial minorities, countering expressions that its participants are not really German. Both of these interpretations of hip hop music is far removed from original meaning of exploring white-black racial conflict within the United States.

> "Moreover, such appropriation have in each case involved a reworking of hip hop in ways that engage with local circumstances. In every respect then, hip hop is both a global and a local form." [4]

This quote demonstrates that the musical meaning of cultural subgroups require analysis as distinct subgroups. Even if the musical content of an apparently global, context-free music is the same, the music can not be analyzed without reference to a specific sub-cultural community. This music must be reinterpreted in terms of the sub-culture appropriating the music, since meanings and the associations it has with other musics and lifestyle choices will differ.

2.4 Cultural Factors

Cultural factors are trends and opinions that are generally consistent across an entire country or region. One example is Tarrant's study [40] which showed differences in the reasons that United Kingdom and United States youth listened to music. As with group factors, empirical analysis of this phenomenon is lacking.

A good example of cultural differences is how the Polish composer Chopin is seen by various different cultures.

> "Many foreign authors claim that his [Chopin's] works were really cosmopolitan and international, inspired by personal, romantic emotions rather than by something in the outside world. For Poles, to cite Norman Davies, 'Chopin's works were built on his experiences in the formative years in Warsaw, distilled from the Polish melodies, harmonies and rhythms that he heard in his youth, and inspired by a bitter sweet nostalgia for the land of his birth; they represent the quintessence of 'Polishness'.' " [22]

Different musical traits are chosen for each analysis, producing different groupings of artists. This phenomenon has been discovered in MIR [2] but the factors creating differences have not as yet been explored. By discovering algorithms for defining musical preferences that are culturally consistent but globally inconsistent, one can create recommendation engines that are culturally aware. This awareness allows one to tailor recommendations based on culturally specific interpretations of musical similarity and its underlying meanings and values.

2.5 Disseminating Musical Information

Beyond the influences of word-of-mouth communication, the media also has an important role in defining musical tastes. We are constantly bombarded by advertisements for music, often in the form of radio broadcasts, as well as less intrusive influences such as music reviews. This factor is often neglected when creating music recommendation systems that are, in effect, a new form of experiencing music and the culture this music helps define that itself influences opinions about music. This process results in the dangerous phenomenon that recommendations works in a lab, but fails in practice as the system alters the community it serves, causing destructive feedback loops. Understanding how media constructs or destroys musical preferences at different levels of analysis is an essential part of a robust system that can survive real world deployment.

2.5.1 Construction of Musical Taste

Media has traditionally been critical in constructing musical taste. People often hear music as a result of broadcast media, both voluntarily and involuntarily [9]. The influence of repeatedly forcing people to hear a song increases its popularity in many cases [12, 36], implicitly influencing musical taste.

The effect goes beyond reinforcing or modifying tastes to actually constructing and defining subcultures and their musical tastes within a culture.

> "Just to be clear, the point is not to assess the accuracy of such accounts, but merely to emphasize that by what they [media] choose to include and exclude, they played an important gatekeeper role in reinforcing and developing the value system of the subculture [shared musical tastes]." [14] p183.

The same effect occurs in more mainstream sub-cultural groups as well. In United Kingdom, Thorton studied dance music in the late 90's:

> "Rather, media and other culture industries are there and effective right from the start. They are central to the process of subcultural formation [a group with similar musical tastes], integral to the way we 'create groups with words'." [43] p117.

This effect of media labelling is so powerful that simply giving a label from a media source influences how the music is interpreted. North and Hargreaves [27] demonstrated that labelling a song about suicide as either suicide-inducing or suicide-preventive significantly altered whether the song was viewed positively or negatively. Thus, media is influencing and even creating musical communities—an effect that can not be ignored.

2.5.2 Music Recommendation as a Media Source

The effects of music recommendation systems as a media source is little understood. Providing good recommendations for everyone is the goal, but the act of creating recommendations broadcasts musical preferences of its listeners. This strips control of the dissemination of culture from its participants. This loss of can be destructive to the very community these systems are attempting to create.

Particularly for sub-cultures, the loss of control over who gets to participate in their community is a serious problem:

> "... its [the club scene in London's] main antagonist is not the police (who arrest and imprison) but the *media* who continually threaten to release its cultural knowledge to other social groups." italics in original [43] p90.

However, providing musical information does not necessarily mean its dissemination:

> "The following respondent ... found it hard, initially, to get anyone to respond to his messages because he had not learned what sort of topics and modes of behavior people were interested in." [14] p180.

Last.FM[1]—the largest on-line music recommendation system—routinely fails to deploy new algorithms for music recommendation because the effects of the algorithm change would create a feedback loop as the user base changes in response.[2] All of these factors demonstrate that recommendation systems should consider the possible consequences of feedback loops when constructing models of musical taste.

3. EXISTING MUSIC RECOMMENDATION SYSTEMS

There have been a number of music recommendation systems described in the literature. Generally, these have used either content-based analysis or collaborative filtering analysis for generating playlists. In addition, a few systems have used user-supplied metadata and web-based data to augment content-based approaches. No system provides analysis at a level higher than the group.

Logan [21, 20] uses a purely content-based approach, producing playlists either directly from similarity, or similarity to a set of songs. Pauws [39] utilizes a more formal definition of constraints based on content-based features to generate playlists. Pampalk et al. [33] also uses content-based analysis but the playlists are altered by users skipping songs. Pampalk and Gasser [31] extend their system by replacing skipping behavior with explicit user ratings. Pandora[3] is an example of a commercial content-based recommendation system based on hand-crafted musical descriptions in its database. Tiemann et al. [44] produced a hybrid system which uses iterative classifications to tag music for recommendation. Two of Chen and Chen's [7] recommendation algorithms utilize purely content-based approaches.

One of Chen and Chen's [7] algorithms utilizes a pure collaborative filtering approach. (Collaborative filtering implicitly performs group-level analysis assuming the entire system is a unique subgroup.) This filtering finds users with similar tastes and suggests the music one listens to but the other does not. Crossen [8] utilizes user recommendations to determine the music to play in a shared space, filtering over hand-picked genre classifications of music. Yoshii et al. [46] provide a hybrid approach where content-based analysis and collaborative filtering are calculated separately and then later integrated. Yoshii et al. extended this in [47]. Anglade et al. [3] produced a peer-to-peer recommender that clustered users with similar tastes and streamed from their shared musical collection. Last.FM uses collaborative filtering based on watching user's listening habits directly with modifications made through explicit negative and positive ratings.

Other approaches have supplemented the traditional pure content-based approaches with web-derived context. Pauws and Eggen [34] produced a systems that generates playlists from a single song using clustering songs similar in audio content but defines these playlists by learning from what songs the user removed. Celma et al. [6] constructs networks of artists derived from targeted web searches to provide recommendations for users based on their FOAF (Friend Of A Friend) profiles. Sandvold et al. [37] extended this with tagging and content-based analysis. Pampalk and Goto [32] extended Sandvold et al.'s work by combining three similarity metrics, one of which is audio similarity.

None of these systems provide analysis at either the interpersonal or cultural level.

4. GRAPH ANALYSIS FOR MUSIC RECOMMENDATION

Each level of analysis requires information not utilized at other levels and generalizes information from lower levels. Furthermore, these different levels of analysis need to interact with one another in order to produce meaningful recommendations. This is not present in any existing music recommendation system.

A graph representation that supports executing a number of different algorithms derived from social network analysis and social psychology at each level of analysis is needed. Tools are needed to differentiate nodes and links by type and apply arbitrary numbers of multi-value properties. It should provide functionality for creating and nesting subgraphs. The tools must support algorithms from social network analysis as well as general graph algorithms. This toolkit must also support an execution framework that is scriptable. Furthermore, the toolkit must keep detailed book keeping [15] of the inputs and outputs of all the algorithms executed that is both documented in human-readable and machine-readable formats.

4.1 Existing Graph Toolkits

While no toolkit exists specifically for music recommendation, there exists a number of existing toolkits for graph structures. Table 1 presents a brief summary of prominent toolkits. The Hypergraph, Piccolo, Giny , Guess, Jung2, Prefuse, and JGraph toolkits are all geared towards visualization of graphs without any analysis capabilities. Jung2 provides a set of analysis tools, but does not provide support for multiple modes or relations. Proximity provides support for modes and relations, but does not include an algorithm base and has limited support for paths and properties.

5. RELATIONAL ANALYSIS TOOLKIT (RAT)

A toolkit is under development to meet these requirements. This toolkit provides basic graph functionality with properties for vertices (actors) and graphs. Web crawlers and parsers are provided that parse various online data sources. A collection of algorithms are also provided to perform experiments against this data. This toolkit is available online at `http://www.sf.net/projects/graph-rat.`

RAT provides actors (a.k.a. vertices or nodes), links (a.k.a. edges or arcs) and graphs. Graphs can be nested arbitrarily deep where each subgraph contains a strict subset of the links and actors of the parent. All actors, links, and graphs

[1]http://www.last.fm
[2]conversations with employees of LastFM
[3]http://www.pandora.com

Table 1: Existing graph toolkits

Name	URL	Description
Hypergraph	hypergraph.sourceforge.net	Hyperbolic tree applet for visualizing graph structure in an XML file
Piccolo	www.cs.umd.edu/hcil/piccolo	Toolkit for 2D graphics, especially zooomable interfaces
prefuse	www.prefuse.org	Toolkit for data visualization including graphs
Giny	csbi.sourceforge.net	framework for interactive information visualization
Guess	www.graphexploration.cond.org	analysis and visualization tool for graphs; uses piccolo, prefuse, and Jung. Includes Jython interpreter
Jung2	jung.sourceforge.net	framework for modeling, analyzing and visualizing graph-based data
JGraph	www.jgraph.com/jgraph.html	graph visualization component
Proximity	kdl.cs.umass.edu/software	graph analysis package

are named with an ID. For graphs, this ID is globally unique. For links and actors, this id is unique within its type (a.k.a. mode or relation). Graphs also contain PathSets which are cached sets of paths through the graph.

The RAT framework provides a structure around the low level features to make it easier to construct experiments. Creating a graph is conducted by Data Acquisition objects that parse various different types of data sources (see Section 5.1). The structure of the system is a variant of the Blackboard pattern in Patterns of Software Architecture [5]. This toolkit deviates from the pattern by only providing deterministic schedulers.

Each Data Acquisition object has an output descriptor which describes what graph output it creates. Data Acquisition objects also have parameters that can be modified to control their execution. Graphs can also be loaded from XML or read from a database.

After the graph is created, algorithms are run against the data set. Each algorithm has specified input and output descriptors, providing book keeping as given in [15]. These descriptors are present both in written documentation and in machine readable form for automated validation. Likewise, algorithms have parameters that can alter how the algorithm is created. Each algorithm instance also has a unique name and a unique order in the execution framework.

The graph type, the data acquisition modules to execute, and the algorithms to execute are controlled by a scheduler object. Schedulers execute scripts in the XML format, eliminating the need to hard-wire experiments. These schedulers can be deterministic or non-deterministic depending on which scheduler is specified in the XML file.

5.1 Parsing Data Sources for Music Recommendation

RAT provides a number of parsing tools for acquiring data from a number of sources. Each of these sources provides information that is useful for constructing music recommendations.

LiveJournal.com provides access to a network of online users in Friend-of-a-friend (FOAF) format. Each profile contains information about the person such as geographical location and lists of interests as well as a list of people this person knows.

Epinions.com provides music reviews of albums. These reviews are given with the user-id of the person who wrote the review. Also, the site provides a 'web-of-trust' where each person lists which reviewers give opinions trusted by a person, creating a separate graph of relationships between users on the site.

Last.FM provides a number of REST based web services (providing XML files over HTTP instead of following the SOAP or CORBA syntax) providing access to information such as artist-to-artist recommendations, the list of the top 100 tags for a particular artist and statistics on the most popular artists on the site.

MusicBrainz.org provides an online database of nearly every compact disc recorded. In addition to providing information about each CD, the site provides lists of all artists that are in the database and artist-to-artist similarity, along with album and song information.

Yahoo.com provides REST access to their search engine which allows for web-based searching that allows the use of web-mining techniques for generating metadata about artists for music recommendation. Currently RAT uses Yahoo! for determining if a string represents an artist name or not.

5.2 Algorithms

The goal of this toolkit is to incorporate as many algorithms as possible to give a music-recommendation researcher as many ways to analyze the data as possible. The algorithms currently implemented are roughly grouped into four categories: calculating prestige, producing recommendations, producing sub-graphs via clustering, and miscellaneous. In all following descriptions, N is the number of actors in the graph and V is the number of links in the graph.

5.2.1 Prestige and Centrality

The prestige and centrality algorithms implemented in RAT are as follows:

1. **Degree**
   The in-degree (number of incoming links) is degree prestige, normalized by the maximum possible in-degree (all actors linking to this actor) so that the value is

between 0 and 1. Similarly out-degree (the number of outgoing links) is degree centrality. This is an O(n) algorithm in both space and time [13].

2. **Closeness**
The closeness prestige is a measure of the distance from any other actor to this actor which is then normalized to be between 0 (Infinite distance to nodes) and 1 (all nodes are distance 1 away). Similarly closeness centrality is a measure of the distance from this actor to every other actor. In the case of a disconnected graph, closeness is only calculated on the component it is a member of. This algorithm is O(n^2) in time and O(n^2) in space if paths are calculated separately, but O(n) in space if the algorithm mixes shortest-path calculations while calculating closeness [13].
$$Closeness_i = \frac{N}{\sum_{j=1}^{N} \min distance(i,j)} \quad (1)$$
where N is the total number of nodes. While this should be ∞ for graphs that do not have paths between some nodes, this is set to be closeness within its component (largest connected subgraph) instead.

3. **Betweenness**
The number of times this actor appears on the shortest path between any two actors, normalized so that the value is between 0 (not on any paths) and 1 (on every shortest path.) Similar to closeness, it is O(n^2)in time and O(n^2) in space if paths are calculated separately, but O(n) in space if the algorithm mixes shortest-path calculations while calculating betweenness [13].
$$Betweenness_i = \frac{\sum_{j=1}^{N} \sum_{k=1}^{N} i \subseteq Path(j,k)}{N(N-1)} \quad (2)$$
where N is the total number of nodes and $Path(j,k)$ is the shortest path between nodes j and k.

4. **Page Rank**
This algorithm calculates prestige for an actor based on the prestige of the actors linking to it. The prestige value for each actor is the value of the first eigenvector of the normalized link matrix with an additional phantom node. The normalization is that the sum of outgoing links from every actor is 1. For this phantom actor, every actor is connected to by every other node with a weight of 15% (100% for those without outgoing links) of all links from the actor and to all nodes equally. Directly calculating the eigen-matrix, this is an O(n^2) algorithm in space and time, but an O(n) in time and O(v) in space using the power method [19]. The O(n^2) algorithm is useful because it also populates the eigen-matrix for use in clustering algorithms.

5. **HITS**
This algorithms generates Hubs (points to a prestigious page) and Authorities (is a prestigious page). The hubs are the first eigenvector of the LL^T where L is the link matrix. The authorities are the values of the first eigenvector of $L^T L$. Calculating the eigenvectors directly, this algorithm is O(n^2) in both time and space. Using a variation of the power method, it is O(n) in time and O(v) in space [16]. Like with PageRank, The O(n^2) algorithm is useful for clustering.

5.2.2 Clustering
Clustering algorithms are algorithms that subdivide a graph into a set of subgraphs.

1. **Maximal Cliques**
A clique is a fully connected sub-graph and a maximal clique is any clique such that no clique exists that contains the maximal clique as a strict subset. This algorithm finds all maximal cliques and runs in O($nlog(n)k!$) where k is the maximum clique size in the graph.

2. **Bipartite Clustering**
This identifies all subgraphs that have the property that all elements are connected by at least 2 fully independent paths (i.e. the paths share no common actor beyond the beginning and ending node).

3. **Strongly Connected Components**
This identifies all strongly connected sub-graphs within a graph. A subgraph is strongly connected if $\forall a, b \in V$, a->b exists and b->a exists.

4. **Weakly Connected Components**
This identifies all weakly connected components. A subgraph is weakly connected if, should all links on the sub-graph be made bidirectional, the sub-graph is strongly connected.

5. **Traditional Edge Betweenness**
This algorithm performs hierarchal clustering on a graph. Edges are removed in order of their betweenness score until all actors are isolated.

6. **Norman-Girvan Edge Betweenness**
This algorithm is identical to traditional edge betweenness clustering except the betweenness of the edges are re-calculated every time a link is removed[23].

5.2.3 Local Recommendation
The recommendation algorithms implemented in RAT are as follows:

1. **Interest Comparisons**
This algorithm compares the 'interest' property pulled from LiveJournal FOAF pages between two actors that know one another (via FOAF declaration). Similarity is compared using an exponential metric that gives a numeric measure of the similarity between the two actor's list of interests which can be both positive or negative.

2. **Music References**
This algorithm determines which interests listed in actor's interest lists are the names of music artists, creating a new actor (of type `Artist`) and creates a link between the actor `User` and the actor `Artist`. This algorithm uses the MusicBrainz data set to identify artist names in the interest lists.

3. **Music Comparisons**
This algorithm creates a numeric measure of similarity between the music links of two actors that know each other. This similarity measure is calculated identically to the Interest Comparison method.

4. **Music Recommendation**
This algorithm recommends new music to an actor using a linear combination of the musical tastes of the actors that know him or her in proportion to the numeric similarity of their interests and music. This generates a numeric recommendation whose length is the number of unique elements across all friends.

5. **Evaluation**
This algorithm compares the music recommended by the Music Recommendation algorithm and compares it against the given music likes. The results are the mean and standard deviation of precision, recall, and ranking for both positive and negative recommendations.

6. **WekaEvaluate**
This algorithm propositionalizes a graph and then learns with the chosen Weka classifier [45], and then generates precision and recall. The instances are constructed as follows: For all friends of an actor, the instance is the music and interest numeric value followed by a binary vector of all artists in the graph this friend listens to. For every artist in the graph, if the artist is listened to by the user, all friends graphs for this artist classifier are true, otherwise it is false. This generates one classifier for every artist in the graph.

6. EVALUATION
An initial data set of 1000 FOAF articles were crawled from the LiveJournal web site using nine seed users that expressed an interest in 'music'. This data set was obtained using the 'snowball' sampling method where all friends from the seeds are traversed in a breadth first fashion until the required sample size is obtained. This biases the sample set by providing full connections for the initial few users, but also guarantees a fully connected graph for analysis. This data set was then loaded into the RAT system where every FOAF user was represented by an actor. Each 'Knows' reference in the FOAF description of each user became a link, the 'interest' property of the FOAF document became a property of the actor, and every property of the actor that was a musical artists generated a new artist actor and a link between the FOAF user and the artist user.

The Music Recommendation algorithm recreated the original music tastes of its users with 4% precision and 16% recall. The J48 classifier from weka performed at 1% precision and 62% recall.

7. CONCLUSIONS
In conclusion, performing music recommendation integrating differing levels of analysis requires a toolkit that can create graph structures that hold the different areas of musical information that is available online for use by music recommenders. This information is analyzed using a collection of algorithms within a scriptable framework. Initial results have been obtained, but much work remains to complete.

8. FUTURE WORK
More clustering algorithms are needed. The propositionalization of Weka needs more results from more classifiers. More methods of integrating the results of one algorithm into another are also needed. Additional exploration should include representing actor graphs as pipes with sources and sinks of musical taste and other forms of iterative recommendation. Additional thought on the consequences of feedback loops and their implications on subgroup formation is also needed.

Beyond additional algorithms, a Facebook application is also planned. FaceBook provides a rich social network environment where it is easy to create a network of friends from a variety of sources such as email accounts. In addition, the system provides a mechanism that permits uploads of front ends to collect data from users that choose to opt-in. One attractive possibility is the possibility of allowing users to link accounts on Facebook with accounts on other sites such as LastFM, Epinions, and LiveJournal.

9. ACKNOWLEDGMENTS
The author would like to thank the University of Waikato for its generous support through the Waikato Doctoral Scholarship and Andrea Schweer for assistance in finding graph packages.

10. REFERENCES
[1] Aebischer, Verena, and M. Hewstone. Minority influence and musical preference: Innovation by conversion, not coercion. *European Journal of Social Psychology*, 14:23–33, 1984.
[2] H. Allan, D. Mullensiefen, and G. Wiggins. Methodological considerations in studies of musical similarity. *International Conference on Music Information Retrieval*, pages 473–478, 2007.
[3] A. Anglade, M. Tiemann, and F. Vignoli. Virtual communities for creating shared music channels. *International Conference on Music Information Retrieval*, pages 95–100, 2007.
[4] A. Bennet. *Popular Music and Youth Culture*. MacMillan Press Ltd, London, 2000.
[5] F. Buschmann, R. Meunier, H. Rohnert, P. Sommerlad, and M. Stal. *Pattern-Oriented Software Architecture: A System of Patterns*. Wiley, Sussex, 1996.
[6] O. Celma, M. Ramìrez, and P. Herrera. Getting music recommendations and filtering newsfeeds from FOAF descriptions. *International Conference on Music Information Retrieval*, 2005.
[7] H.-C. Chen and A. L. P. Chen. A music recommendation system based on music data grouping and user interests. *CIKM*, pages 231–238, 2001.
[8] A. Crossen, J. Budzik, and K. J. Hammond. Flytrap: Intelligent group music recommendation. *IUI*, 2002.
[9] S. J. Cunningham, D. Bainbridge, and D. McKay. Finding new music: a diary study of everyday encounters with novel songs. *International Conference on Music Information Retrieval*, pages 83–88, 2007.
[10] N. Dibben. *Musical Identities*, chapter Gender Identity and Music. Oxford University Press, Oxford, 2002.
[11] W. Doise. *Levels of Explanation in Social Psychology*. European Monographs in Social Psychology. Cambridge University Press, Cambridge, 1986.
[12] M. Erdèlyi. The relation between "radio plugs" and sheet sales of popular music. *Journal of Applied Psychology*, 24:696–702, 1940.

[13] L. C. Freeman. Centrality in social networks. *Social Networks*, 1:215–239, 1979.

[14] P. Hodkinson. *Goth: Identity, Style and Subculture*. Berg, Oxford, 2002.

[15] D. Jensen. Unique challenges of managing inductive knowledge. In *Proceedings of the 1997 AAAI Spring Symposium on Artificial Intelligence in Knowledge Management*, Stanford, March 1997. American Association of Artificial Intelligence.

[16] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46(5):604–632, 1999.

[17] V. J. Konecni and D. Sargent-Pollack. Choice between melodies differing in complexity under divided-attention conditions. *Journal of Experimental Psychology: Human Perception and Performance*, 3(2):347–56, 1976.

[18] A. Lamont, D. J. Hargreaves, N. A. Marshall, and M. Tarrant. Young people's music in and out of school. *British Journal of Music Education*, 20(3):229–341, 2003.

[19] A. N. Langville and C. D. Meyer. Deeper inside PageRank. *Internet Mathematics*, 1(3):335–380, 2003.

[20] B. Logan. Content-based playlist generation: Exploratory experiments. *International Conference on Music Information Retrieval*, 2002.

[21] B. Logan. Music recommendation from song sets. *International Conference on Music Information Retireval*, 2004.

[22] Z. Mach. *Ethnicity, Identity, and Music: The Musical Construction of Place*, chapter National anthems: the case of Chopin as a national composer. Berg, Oxford, 1994.

[23] M. E. J. Newmann and M. Girvan. Finding and evaluating community structure in networks. *Physical Review E*, 69, 2004.

[24] A. C. North and D. J. Hargreaves. Responses to music in aerobic exercise and yogic relaxtion classes. *British Journal of Psychology*, 87:535–47, 1996.

[25] A. C. North and D. J. Hargreaves. Situational influences on reported musical preferences. *Psychomusicology*, 15:30–45, 1996.

[26] A. C. North and D. J. Hargreaves. Musical preferences during and after relaxation and exercise. *American Journal of Psychology*, 113(1):43–67, 2000.

[27] A. C. North and D. J. Hargreaves. Brief report: Labelling effects on the percieved deleterious consequences of pop music listening. *Journal of Adolescence*, 28:433–40, 2005.

[28] A. C. North and D. J. Hargreaves. Lifestyle correlates of musical preference: 1. relationships, living relationships, beliefs and crime. *Society for Research in Psychology of Music and Music Education*, 35(1):58–87, 2007.

[29] A. C. North, D. J. Hargreaves, and J. J. Hargreaves. Uses of music in everyday life. *Music Perception*, 22(1):41–77, 2004.

[30] A. C. North, D. J. Hargreaves, and S. A. O'Neill. The importance of music to adolescents. *British Journal of Educational Psychology*, 70:255–72, 2000.

[31] E. Pampalk and M. Gasser. An implementation of a simple playlist generator based on audio simlarity measures and user feedback. *International Conference on Music Information Retrieval*, 2006.

[32] E. Pampalk and M. Goto. Musicsun: A new approach to artist recommendation. *International Conference on Music Information Retrieval*, pages 101–104, 2007.

[33] E. Pampalk, T. Pohle, and G. Widmer. Dynamic playlist generation based on skipping behavior. *International Conference on Music Information Retrieval*, 2005.

[34] S. Pauws and B. Eggen. Pats: Realization and user evaluation of an automatic playlist generator. *International Conference on Music Information Retrieval*, 2002.

[35] W. Rhodes. A decade of progress. *Ethnomusicology*, pages 178–181, 1963.

[36] P. A. Russel. Effects of repetition on the familiarity and likeability of popular music recordings. *Psychology of Music*, 15:187–197, 1987.

[37] V. Sandvold, T. Aussenac, O. Celma, and P. Herrera. Good vibrations: Music discovery through personal musical concepts. *International Conference on Music Information Retrieval*, 2006.

[38] J. Scott. *Social Network Analysis: A Handbook*. Sage Publications, London, 2nd edition, 1991.

[39] M. V. Steffen Pauws, Wim Verhaegh. Fast generation of optimal music playlist using local search. *International Conference on Music Information Retrieval*, 2006.

[40] M. Tarrant. English and american adolescents' reasons for listening to music. *Society for Research in Psychology of Music and Music Education*, 28(2):166–72, 2000.

[41] M. Tarrant, L. MacKenzie, and L. A. Hewitt. Friendship group identification, multidimensional self-concept, and experience of developmental tasks in adolescence. *Journal of Adolescence*, 29:627–40, 2006.

[42] M. Tarrant, A. C. North, and D. J. Hargreaves. *Musical Identities*, chapter Youth Identity and Music. Oxford University Press, Oxford, 2002.

[43] S. Thorton. *Club Cultures: Music, Media and Subcultural Capital*. Weslyan University Press, Middletown, 1996.

[44] M. Tiemann, S. Pauws, and F. Vignoli. Ensemble learning for hybrid music recommendation. *International Conference on Music Information Retrieval*, pages 179–181, 2007.

[45] I. H. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, San Francisco, 2nd edition, 2005.

[46] K. Yoshii, M. Goto, K. Komatani, T. Ogata, and H. G. Okuno. Hybrid collaborative and content-based music recommendation using probabilistic model with latent user preferences. *International Conference on Music Information Retrieval*, 2006.

[47] K. Yoshii, M. Goto, K. Komatani, T. Ogata, and H. G. Okuno. Improving effeciency and scalability of model-based music recommender system based on incremental training. *International Conference on Music Information Retrieval*, pages 89–94, 2007.