

Analysis and Exploitation of Musician Social Networks for Recommendation and Discovery

Ben Fields, *Student Member, IEEE*, Kurt Jacobson, *Student Member, IEEE*, Christophe Rhodes, Mark d'Inverno, Mark Sandler, *Senior Member, IEEE*, and Michael Casey, *Member, IEEE*

Abstract—This paper presents an extensive analysis of a sample of a social network of musicians. The network sample is first analyzed using standard complex network techniques to verify that it has similar properties to other web-derived complex networks. Content-based pairwise dissimilarity values between the musical data associated with the network sample are computed, and the relationship between those content-based distances and distances from network theory explored. Following this exploration, hybrid graphs and distance measures are constructed, and used to examine the community structure of the artist network. Finally, results of these investigations are shown to be mostly orthogonal between these distance spaces. These results are considered with a focus recommendation and discovery applications employing these hybrid measures as their basis.

Index Terms—Content-based retrieval, graph theory, music information retrieval, shortest path problem, social network services: MySpace.

I. INTRODUCTION

AS more freely-available audio content continues to become accessible, listeners require more sophisticated tools to aid them in the discovery and organization of new music that they will find enjoyable. This need, along with the advent of Web-based social networks and the increasing accuracy of signal-based music information retrieval, have created an opportunity to exploit both social relationships and acoustic similarity in recommender and discovery systems. However, current systems have tended to use one of these techniques in isolation. In our view, combining these techniques provides

Manuscript received September 15, 2010; revised January 07, 2011; accepted January 13, 2011. Date of publication February 04, 2011; date of current version July 20, 2011. This work was supported in part by the Engineering and Physical Sciences Research Council via the Online Music Recognition And Searching II (OMRAS2) project, reference numbers EP/E017614/1 and EP/E02274X/1, and as part of the Networked Environments for Music Analysis (NEMA) project, funded by The Andrew W. Mellon Foundation. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Xian-Sheng Hua.

B. Fields, C. Rhodes, and M. d'Inverno are with the Intelligent Sound and Music Systems group, Department of Computing, Goldsmiths, University of London, New Cross, London SE14 6NW, U. K. (e-mail: b.fields@gold.ac.uk).

K. Jacobson and M. Sandler are with the Centre for Digital Music, School of Electronic Engineering and Computer Science, Queen Mary University of London, Mile End Road, London E1 4NS, U. K.

M. Casey is with the Intelligent Sound and Music Systems group, Department of Computing, Goldsmiths, University of London, New Cross, London SE14 6NW, U. K., and also with the Bregman Music Audio Research Studio, Music Department, Dartmouth College, Hanover, NH 03755 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMM.2011.2111365

a means to improving the understanding of the complex relationship between song objects that ultimately will lead to improved song recommendation. The most obvious way to do that is to base recommendations on more information than is provided by a single distance measure between songs. This would allow the production of systems capable of mediating content-based recommendations with given social connections and the construction of socially structured playlists.

Motivated by this, we examine the Myspace artist network. Though there are a number of music-oriented social-networking websites (e.g. Soundcloud,¹ Jamendo,² etc.), Myspace³ is the *de facto* standard for web-based music artist promotion. For the purpose of this paper, *artist* and *artist page* are used interchangeably to refer to the collection of media and social relationships found at a specific Myspace page residing in Myspace's artist subnetwork. Although exact figures are not made public, recent estimates suggest there are over 8 million artist pages⁴ on Myspace.

The Myspace social network, like most social networks, is based upon undirected relational links between *friends* designating some kind of association. A link is created when a user makes a request and another accepts the request to become friends; both users are then friends and an undirected link is established. Within each Myspace user's friends, there is a subset of between 8 and 40 *top friends*. While generic friends are mutually confirmed, individual users unilaterally elevate friends to become top friends from the generic friends set. It is these top friends that are displayed in a user's profile page—other friends require one or more *click-throughs* to access them. In addition, any user can declare themselves as an *artist* which requires them to provide audio or video content. In our work, we concern ourselves only with these artist users to limit the scope of our investigation to only those nodes on the graph that have audio content.

Social networks present a way for nearly anyone to distribute their own media. As a result, there is an ever-larger amount of available music from an ever-increasing array of artists.

- 1) Given that this music is published within a relational space, how can we best use all of the available information to discover new music?
- 2) Can both social metadata and content-based comparisons be exploited to improve discovery of new material?

¹<http://www.soundcloud.com>.

²<http://www.jamendo.com>.

³<http://www.myspace.com>.

⁴<http://techradar1.wordpress.com/2008/01/11/facebookmyspace-statistics>.

- 3) Can this crowd-sourced tangle of social networking ties provide insights into the dynamics of popular music?
- 4) Does the structure of a network of artists have any relevance to music-related studies such as music recommendation or musicology?

To work towards answering the questions posed above, we explore a subset of the artist network and consider only their top friend connections. We analyze this network and measures of acoustic distance (according to techniques using content-based analysis, which we will describe in the next section) between these artists. Furthermore, we identify communities of artists based on the Myspace network topology and attempt to relate these community structures to musical genre. Finally, we present a prototype system of music playlist generation, with particular attention paid to the means for its evaluation.

Immediately following this section is a review of relevant literature from complex network theory and signal-based music analysis in order that the reader can understand the contributions of our work. Next, a detailed discussion of our data acquisition methods to build a sampled data set is presented, followed by a broad analysis of this data set in Section III. The initial experiments into the relationship between the social connectivity and the acoustic feature space are described and their results presented and discussed in Section IV. Finally, the implications of this work are explored in Section V followed by directions for future work in Section VI.

II. BACKGROUND

To provide a base of understanding for the work that follows, we begin the background with a discussion of existing tools for the analysis and manipulation of networks in Section II-A. This subsection covers complex network analysis, network flow analysis, particular issues pertaining to networks of musicians and community structure. In Section II-B, we examine highlights of past work in audio content-based music similarity. We assume the reader to be familiar with the particulars of comparing non-normal distributions, particularly mutual information. More detail in this area as it relates to this work can be found in [1].

A. Existing Tools for Networks

1) *Complex Networks*: Complex network theory deals with the structure of relationships in complex systems. Using the tools of graph theory and statistical mechanics, physicists have developed models and metrics for describing a diverse set of real-world networks—including social networks, academic citation networks, biological protein networks, and the World-Wide Web. It has been shown that these diverse networks often exhibit several unifying characteristics such as small worldness, scale-free degree distributions, and community structure [2].

A given network G is described by a set of *nodes* N connected by a set of *edges* E . Each edge is defined by the pair of nodes (i, j) it connects. This pair of nodes are *neighbors* via edge $E(i, j)$. If the edges imply directionality, i.e., $(i, j) \neq (j, i)$, the network is a *directed network*. Otherwise, it is an *undirected network*. Since we are dealing primarily with the *top friends* sub-network of Myspace artists, in this paper, all edges are directed




Fig. 1. Simple flow network with directed weighted edges. Edge width is representative of node capacity, which is also labelled on each edge. Treating node A as the source and node F as the sink, the Maximum Flow is 4.

unless otherwise stated. In some graphs, each edge $E(i, j)$ will have an associated label $w(i, j)$ called the *weight*. This weight is sometimes thought of as the cost of traversing an edge, or an edge's resistance. The number of edges incident to a node i is the *degree* k_i . In a directed network, there will be an *indegree* k_i^{in} and an *outdegree* k_i^{out} corresponding to the number of edges pointing into the node and away from the node, respectively. The *geodesic* d_{ij} is the shortest path distance from i to j in number of edges traversed.

We will discuss some of the characteristics of the Myspace artist network in Section III-B. For a more in-depth discussion of complex network-analysis techniques, the reader is referred to [2] and [3].

2) *Network Flow Analysis*: The basic premise in network flow analysis is to examine a network's nodes as sources and sinks of some kind of *traffic* [4]. Typically, though not exclusively, flow networks are directed, weighted graphs. Many useful measures for determining the density of edge connectivity between sources and sinks can be found in this space [5]. One of the most common among them is the Maximum Flow, which is a means of measuring the maximum capacity for fluid to flow between a source node to a sink node or, equivalently, the smallest sum of edge weights that must be *cut* from the network to create exactly two subgraphs, one containing the source node and one containing the sink node. This equivalence is the Maximum Flow/Minimum Cut Theorem [6]. If the edges in a graph are unweighted, this value is also equivalent to the number of paths from the source to the sink which share no common edges. Mature algorithms, incorporating a number of optimization strategies, are available for computing the maximum flow between nodes [4], [7].

An example of Maximum Flow can be seen on the network in Fig. 1. The narrowest flow capacity from node A to node F are the edges $E(a, b)$ and $E(a, c)$, where $E(a, b) + E(a, c) = 4$. The Maximum Flow can simply be found by taking the sum of the magnitude of each edge in the minimum cut set.

The few examples of network flow analysis being applied in music informatics deal primarily with constructing playlists using segments of a complete solution to the Traveling Salesman Problem [8]. Others use exhaustive and explicit textual metadata without comparisons to content-based metrics [9].

3) *Musician Networks*: Networks of musicians have been studied in the context of complex network theory—typically viewing the artists as nodes in the network and using either collaboration, influence, or similarity to define network edges. These networks of musicians exhibit many of the properties expected in social networks [10]–[12]. However, these studies all

examine networks created by experts (e.g., All Music Guide⁵) or via algorithmic means (e.g., Last.fm⁶) as opposed to the artists themselves, as is seen in Myspace and other similar networks. Networks of music listeners and of listeners connected to artists have also been studied [13], [14].

4) *Community Structure*: Recently, as more data-heavy complex networks have been created across many domains, there has been a significant amount of interest in algorithms for detecting community structures in these networks. These algorithms are meant to find dense subgraphs (communities) in a larger sparse graph. More formally, the goal is to find a partition $\mathcal{P} = \{C_1, \dots, C_c\}$ of the nodes in graph G such that the proportion of edges inside C_k is high compared to the proportion of edges between C_k and other partitions.

Because our network sample is moderately large, we restrict our analysis to use more scalable community detection algorithms. We make use of the greedy modularity optimization algorithm [15] and the walktrap algorithm [16]. These algorithms are described in detail in Section III-C.

B. Content-Based Music Analysis

Many methods have been explored for content-based music analysis, attempting to characterize a music signal by its timbre, harmony, rhythm, or structure. One of the most widely used methods is the application of Mel-frequency cepstral coefficients (MFCC) to the modeling of timbre [17]. While a number of other spectral features have been used with success [18], when used in combination with various statistical techniques, MFCCs have been successfully applied to music similarity and genre classification tasks [19]–[22].

A simple and prevalent means to move from the high dimensional space of MFCCs to single similarity measure is to calculate the mean and covariance of each coefficient across an entire song and take the Euclidean distance between these mean and covariance sets [20]. In the Music Information Retrieval Evaluation eXchange (MIREX) [23], [24] competitions of both 2007⁷ and 2009,⁸ this method, as employed by the Marsyas software suite, was shown to do a reasonable job of approximating human judgements of content-based similarity. A slightly more complex approach for computing timbre-based similarity between two songs or collections of songs creates Gaussian mixture models (GMM) describing the MFCCs and comparing the GMMs using a statistical distance measure. Often the Earth Mover's Distance (EMD), a technique first used in computer vision [25], is the distance measure used for this purpose [22], [26]. The EMD algorithm finds the minimum work required to transform one distribution into another. While the EMD-GMM approach models distance better than a simple Euclidean distance between averages of feature values, the simpler method may be sufficient and is considerably less computationally complex.

⁵<http://www.allmusic.com>.

⁶<http://www.lastfm.com>.

⁷See http://www.music-ir.org/mirex/2007/index.php/Audio_Music_Similarity_and_Retrieval_Results entry by Tzanetakakis.

⁸See http://www.music-ir.org/mirex/2009/index.php/Audio_Music_Similarity_and_Retrieval_Results entry by Tzanetakakis.

Web graphs combined with content-based features have recently been shown to improve the image classification problem [27]. A 12% performance improvement (as measured by the area under the ROC curve) for an adult-content-recognition task was obtained using a combination of text and image features in conjunction with a connected sub-graph, of the entire Web, that contained a number of images labelled as offensive or not. Text features were obtained using a latent semantic indexing (LSI) model and image features used deep belief networks (DBN) and principle components analysis (PCA) to yield 500 dimensional features. The Web graph contained 83 k web pages with 211 k attached images for a total of 295 k nodes. While our problem concerns a different content-domain (music) and application (recommendation and discovery), this related work is highly encouraging.

III. DATA SET ACQUISITION AND ANALYSIS

Now that we have a foundational understanding of complex networks, we need to gather our data set. For reasons we discuss shortly, it is not feasible to capture the entire Myspace artist network. We therefore take a sample which we show to be representative. In this section, we report on our sampling of the Myspace network, describing our method in Section III-A and properties of this sample in Section III-B. In order to examine the topography of our sample and the distribution of connectivity within the sample, we describe our methods for detecting community structure in Section III-C.

A. Sampling Myspace

The Myspace social network presents a variety of challenges. Firstly, its size prohibits analyzing the graph in its entirety, even when considering only the artist pages: therefore we sample a small yet sufficient portion of the network. Secondly, the Myspace social network is filled with noisy data—plagued by spammers and orphaned accounts: we limit the scope of our sampling in a way that minimizes this noise. Finally, there currently is no published interface for easily collecting the network data from Myspace. Our data are collected using web crawling and HTML document scraping techniques.⁹

1) *Artist Pages*: It is important to note we are only concerned with a subset of the Myspace social network—the Myspace *artist* network. Myspace artist pages are different from standard Myspace pages in that they include a distinct audio player application containing material uploaded by that user. Standard practice (and a requirement of the End User License Agreement) is that this material has been generated by this user. We therefore use the presence or absence of this player to determine whether or not a given page is an artist page where, as stated in Section I, *artist page* is used to refer to the collection of social links and audio material assumed to be generated by the same person or group of people.

A Myspace page will include a top friends list. This is a hyperlinked list of other Myspace accounts explicitly specified by the user and, unlike generic friends, need not be a reciprocal

⁹Myspace scraping is done using tools from the MyPySpace project available at <http://mypyspace.sourceforge.net>.

relationship. The top friends list is limited in length with a maximum length of 40 friends (the default length is 16 friends). In constructing our sampled artist network, we use the top friends list to create a set of directed edges between artists. Only top friends who also have artist pages are added to the sampled network; standard Myspace pages are ignored. We also ignore the remainder of the friends list (i.e., friends that are not specified by the user as top friends), assuming these relationships are not as relevant. Our sampling method is based on the assumption that artists specified as top friends have some meaningful musical connection for the user—whether through collaboration, stylistic similarity, friendship, or artistic influence. This artificially limits the outdegree of each node in such a way as to only track social connections that have been selected by the artist to stand out, beyond the self-promoting noise of their complete friend list. Further, it is also a practical reduction as top friends can be scraped from the same single HTML document as all the other artist metadata. Fifty friends are displayed per page, so gathering a full friend list would require $N/50$ pages to be scraped,¹⁰ significantly increasing the number of page requests required to sample the same number of artists.

In addition to these social connections, we also gather metadata about each artist. This metadata includes the name of the artist, the number of page views, and genre labels associated with the artists. Each artist selects from 0 to 3 genres from a list of 119 given by Myspace. The audio files associated with each artist page in the sampled network are also collected for feature extraction. Note that genre tags collected are at the level of artists, rather than audio files; therefore all audio files associated with that artist will have the same genre labels applied (see Section IV-C).

2) *Snowball Sampling*: There are several network sampling methods; however, for the networks like the Myspace artist network, snowball sampling is the most appropriate method [28], [29]. In this method, the sample begins with a seed node (artist page), then the seed node’s neighbors (top friends), then the neighbors’ neighbors, are added to the sample. This breadth-first sampling is continued until the fraction of nodes in the sample reaches the target or *sampling ratio*. Here, we randomly select a seed artist¹¹ and collect all artist nodes within 6 edges to collect 15 478 nodes. If the size of the Myspace artist network is around 7 million, then this is close to the 0.25% sampling ratio suggested for accurate degree distribution estimation in sampled networks. Note that the sampling ratio is not sufficient for estimating other topological metrics such as the clustering coefficient and assortativity [30]; such global measures are not required for this paper.

With snowball sampling, there is a tendency to over-sample hubs because they have many links and are typically picked up early in the breadth-first sampling. This effect reduces the degree distribution exponent by introducing a higher proportion of nodes with high connectivity than are seen in the complete network, producing a heavier tail but preserving the overall power-law nature of the network [29].

¹⁰Where N is the number of friends, typically 10^3 but in some cases of the order 10^7 .

¹¹The artist is *Karna Zoo*, Myspace <http://www.myspace.com/index.cfm?fuseaction=user.viewProfile&friendID=134901208>.

TABLE I

NETWORK STATISTICS FOR THE MYSPACE ARTIST NETWORK SAMPLE WHERE n IS THE NUMBER OF NODES, m IS THE NUMBER OF EDGES, $\langle k \rangle$ IS THE AVERAGE DEGREE, l IS THE MEAN GEODESIC DISTANCE, AND d_{max} IS THE DIAMETER, AS DEFINED IN SECTION II-A1

	n	m	$\langle k \rangle$	l	d_{max}
undirected	15478	91326	11.801	4.479	9
directed	15478	120487	15.569	6.426	16

B. Network Analysis of the Myspace Artist Network Sample

The Myspace artist network sample exhibits many of the network characteristics common to social networks and other real-world networks. Some of the network’s statistics are summarized in Table I.

We see that the Myspace artist network is like many other social networks in its “small world” characteristics—having a small diameter and geodesic distance. Additionally, in previous work, it has been shown that the Myspace artist network is assortative with respect to genre labels—that is, artists preferentially form connections with other artists that have the same genre labels [31].

Although the network is constructed as a directed network, for some of our experiments, we convert to an undirected network to simplify analysis. This conversion is done to reduce complexity for analysis and to better examine the reflexive properties that are present in the broader mutual friend connections of the whole Myspace network. Each edge is considered bi-directional, that is $(i, j) = (j, i)$, and if a reflexive pair of edges existed in the directed graph, only one bi-directional edge exists in the undirected graph.

The degree distribution for this undirected reduction network is plotted in Fig. 2 on a log-log scale. As mentioned earlier, it is common to find a power-law degree distribution in social networks [2]. However, exponential degree distributions have been reported previously in some types of music recommendation networks [10]. This is especially true for networks with imposed degree limits. For moderate degree values ($35 < k < 200$), our sample shows a power-law distribution. For lower degree values, the distribution is closer to exponential. This may be related to the fact that our network has an out degree limit imposed by Myspace restricting the maximum number of top friends ($k_{out} \leq 40$). The power-law fit also breaks down for high values of k —most likely due to the limited scope of our sample. Similar “broad-scale” degree distributions have been reported for citation networks and movie actor networks [32]. A more detailed analysis of this Myspace artist network can be found in [31].

C. Community Structure

We apply two community detection algorithms to our network sample—the greedy optimization of modularity [15] and the walktrap algorithm [16]. Both of these algorithms are reasonably efficient for networks of our size, and both algorithms can be easily adapted to incorporate audio-based similarity measures (see [33] and Section III-C).




Fig. 2. Cumulative degree distributions for the Myspace artist network sample. For moderate values of k , the distribution follows a power-law (indicated by the dotted line), but for low and high values the decay is exponential.

1) *Greedy Modularity Optimization*: *Modularity* is a network property that measures the appropriateness of a network division with respect to network structure. Modularity can be defined in several different ways [3]. In general, the modularity Q captures the relationship between the number of edges within communities and the expected number of such edges. Let A_{ij} be an element of the network's adjacency matrix and suppose the nodes are divided into communities such that node i belongs to community C_i . We choose the definition of modularity Q as the fraction of edges within communities minus the expected value of the same quantity for a random network of the same size and degree distribution. Then Q can be calculated as follows:

$$Q = \frac{1}{2m} \sum_{ij} \left[A_{ij} - \frac{d_i d_j}{2m} \right] \delta_{C_i C_j} \quad (1)$$

where the $\delta_{C_i C_j}$ function is 1 if $C_i = C_j$ and 0 otherwise, m is the number of edges in the graph, and d_i is the *degree* of node i —that is, the number of edges incident on node i . The sum of the term $d_i d_j / 2m$ over all node pairs in a community represents the expected fraction of edges within that community in an equivalent random network where node degree values are preserved.

If we consider Q to be a benefit function we wish to maximize, we can then use an agglomerative approach to detect communities—starting with a community for each node such that the number of partitions $|\mathcal{P}| = n$ and building communities by amalgamation. The algorithm is greedy, finding the changes in Q that would result from the merge of each pair of communities, choosing the merge that results in the largest increase of Q , and then performing the corresponding community merge. It can be proven that if no community merge will increase Q , the algorithm can be stopped because no further modularity optimization is possible [15]. Using efficient data structures based on sparse matrices, this algorithm can be performed in time $\mathcal{O}(m \log n)$.

2) *Random Walk: Walktrap*: The walktrap algorithm uses random walks on G to identify communities. Because communities are more densely connected, a random walk will tend to be “trapped” inside a community—hence the name “walktrap”.

At each time step in the random walk, the walker is at a node and moves to another node chosen randomly and uniformly from its neighbors. The sequence of visited nodes is a *Markov chain* where the states are the nodes of G . At each step, the transition probability from node i to node j is $P_{ij} = A_{ij}/d_i$ which is an element of the transition matrix P for the random walk. We can also write $P = D^{-1}A$ where D is the diagonal matrix of the degrees ($\forall i, D_{ii} = d_i$ and $D_{ij} = 0$ where $i \neq j$).

The random walk process is driven by powers of P : the probability of going from i to j in a random walk of length t is $(P^t)_{ij}$ which we will denote simply as P_{ij}^t . All of the transition probabilities related to node i are contained in the i th row of P^t denoted as $P_{i\bullet}^t$. We then define an inter-node distance measure for a given value of t :

$$r_{ij} = \sqrt{\sum_{k=1}^n \frac{(P_{ik}^t - P_{jk}^t)^2}{d_k}} = \left\| D^{-\frac{1}{2}} P_{i\bullet}^t - D^{-\frac{1}{2}} P_{j\bullet}^t \right\| \quad (2)$$

where $\|\cdot\|$ is the Euclidean norm of \mathbb{R}^n . This distance can also be generalized by averaging to a distance between communities: $r_{C_i C_j}$ or to a distance between a community and a node: $r_{C_i j}$.

We then use this distance measure in our algorithm. Again, the algorithm uses an agglomerative approach, beginning with one partition for each node ($|\mathcal{P}| = n$). We first compute the distances for all adjacent communities (or nodes in the first step). At each step k , two communities are chosen based on the minimization of the mean σ_k of the squared distances between each node and its community:

$$\sigma_k = \frac{1}{n} \sum_{C_i \in \mathcal{P}_k} \sum_{i \in C_i} r_{i C_i}^2. \quad (3)$$

Direct calculation of this quantity is known to be NP-hard [16], so instead we calculate the variations $\Delta\sigma_k$. Because the algorithm uses a Euclidean distance, we can efficiently approximate these variations as

$$\Delta\sigma(C_1, C_2) = \frac{1}{n} \frac{|C_1||C_2|}{|C_1| + |C_2|} r_{C_1 C_2}^2. \quad (4)$$

The community merge that results in the lowest $\Delta\sigma$ is performed. We then update our transition probability matrix

$$P_{(C_1 \cup C_2)\bullet}^t = \frac{|C_1| P_{C_1\bullet}^t + |C_2| P_{C_2\bullet}^t}{|C_1| + |C_2|} \quad (5)$$

and repeat the process updating the values of r and $\Delta\sigma$ then performing the next merge. After $n - 1$ steps, we get one partition that includes all the nodes of the network $\mathcal{P}_n = \{N\}$. The algorithm creates a sequence of partitions $(\mathcal{P}_k)_{1 \leq k \leq n}$. Finally, we use modularity to select the best partition of the network, calculating $Q_{\mathcal{P}_k}$ for each partition and selecting the partition that maximizes modularity.

Because the value of t is generally low (we use $t = 4$, selected empirically), this community detection algorithm is quite scalable. For most real-world networks, where the graph is sparse, this algorithm runs in time $\mathcal{O}(n^2 \log n)$ [16]. Note though, the optimized greedy modularity algorithm scales significantly better for sparse graphs than the walktrap algorithm— $\mathcal{O}(m \log n)$ versus $\mathcal{O}(n^2 \log n)$ —and in our implementation is faster by an order of magnitude on our sample graph.

D. Summary

In an effort to create an experimental dataset, the Myspace social network’s artist network was sampled. The sample was taken via random entry and a breadth-first walk. Basic analysis of this sample set shows it to conform to norms of other studied social networks. Further an explanation of community structural analysis techniques were laid out, from which to perform multimodal analysis and measurement of the sample. With this understanding of the basic properties of our data set, we can now go forward with experimentation using hybrid distance techniques.

IV. HYBRID METHODS OF DISTANCE ANALYSIS

To move towards well-formed uses of both social and acoustic notions of distance, a better understanding of the relationship between these two spaces is required. We therefore conduct a series of experiments to analyze the effect of combining social and content-based distance. Our first two experiments are concerned with distance between pairs of nodes (both artists and songs) in our graph; the third experiment looks into the effect that acoustic distance measurements have in the detection of community structure. These experiments are presented as follows.

- 1) The geodesic distance between all pairs of artists within the sample are compared to the acoustic similarity of songs associated with each artist.
- 2) Maximum flow analysis is used to analyze the artist social space.
 - a) This measure is compared to the same artist-based acoustic similarity used in item 1.
 - b) An additional song-to-song acoustic metric generated by the Marsyas software suite is also used.
- 3) Community segmentation and structural analysis are explored as a further means of understanding the interaction between these two spaces.

Some of this work requires a network of songs rather than artists (as we sampled in Section III-A). An unweighted graph between songs can be constructed by simply applying the artist connections to their associated songs; weights can be assigned to these song-to-song edges individually, for example based on acoustic dissimilarity between pairs of songs computed with the methods described in Section IV. These node relationships are illustrated in Fig. 3.

MFCCs are extracted from each audio signal using a Hamming window on 8192 sample FFT windows with 4096 sample




Fig. 3. Comparison of sampled and song expanded means of representing the relationship between artists. (a) Sampled artist to artist relationship and (b) expanded artist relationship, with songs as nodes. Note that the connections of song k and song l have been omitted for clarity.

overlap. All MFCCs are created with the `fftExtract` tool.¹² For each artist node, a GMM is built from the concatenation of MFCC frames for all songs found on each artist’s Myspace page. Generally artists have between 1 and 4 songs, although some artists have many more. The mean number of songs is slightly more than 3.5 per artist. An $n \times n$ matrix is populated with the Earth Mover’s Distance λ_{ij} between the GMMs corresponding to each pair of nodes in the sample. As a second acoustic dissimilarity measure, the software suite Marsyas¹³ is used in the exact configuration that was used in the MIREX 2009 Audio Similarity and Retrieval¹⁴ task to generate MFCC-based average value vectors per song and then to generate an $n \times n$ Euclidean distance matrix of these songs. These distance matrices are used to draw λ values to compare against the song expanded graph as detailed above.

¹²Source code at <http://omras2.doc.gold.ac.uk/software/fftextract>.

¹³<http://marsyas.info/>.

¹⁴<http://music-ir.org/mirex/2009/results/abs/GTfinal.pdf>.




Fig. 4. Box and whisker plot showing the spread of pair-wise artist dissimilarity grouped by geodesic distance as found on the artist graph. The whiskers cover the second and seventh octiles beyond the inner quartiles covered in each box.

A. Geodesic Paths


The relation between audio signal dissimilarity and the geodesic path length is first examined using a box and whisker plot. The plot is shown in Fig. 4. These dissimilarities are grouped according to the geodesic distance in the undirected network between the artist nodes i and j , d_{ij} . There appears to be no clear correlation between these λ values and geodesic distance. The Pearson product-moment correlation coefficient confirms this giving a ρ of -0.0016 . This should be viewed in the context of the number of pairwise relationships used, implying it is stable, at least for the community of artists found via this sample of the network. Further, it should be noted that our approach to audio-based dissimilarity results in measures which are mostly orthogonal to network structure [34].

B. Maximum Flow

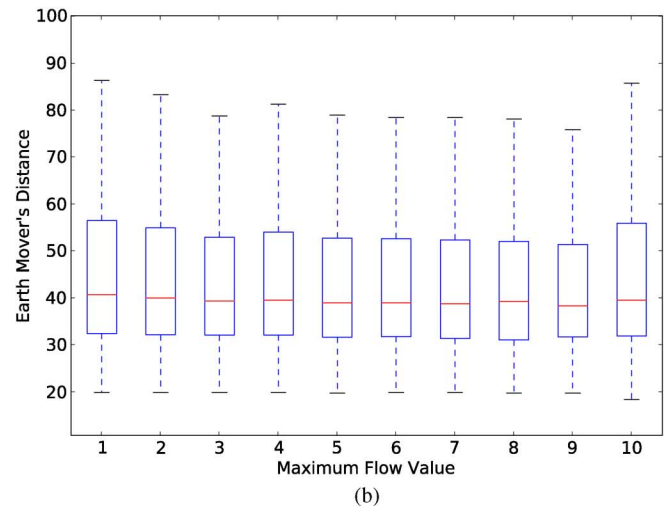
In our Myspace top friends graph, the maximum flow is measured on the directed and undirected reduction of the unweighted graph from the source artist node to the sink artist node. This extends the work of [35] by applying an additional acoustic distance measure (that of the Marsyas entries into MIREX) and examining all the results via means of mutual information.

1) *Experiment*: The Maximum Flow value is calculated, using the snowball sample entry point as the fixed source against every other node in turn as a sink, yielding the number of edges connecting each sink node to the entry point node at the narrowest point of connection. The acoustic distances are then be compared to these Maximum Flow values.

In order to better understand a result from analysis of our Myspace sample, a baseline for comparison must be used. To that end, we examine random permutations of the node locations. In order to preserve the overall topology present in the network, we perform this randomization by shuffling the artist label and associated music attached to a given node on the network. This is done ten fold, creating a solid baseline to test



(a)



(b)

Fig. 5. Box and whisker plots showing the distribution of EMD grouped by maximum-flow value between artists on the Myspace social graph and the randomized permutations of the graph. (a) EMD distribution on the sampled graph. (b) EMD distribution on the random permutations of the graph, maintaining the original edge structure.

the null hypothesis that the underlying community structure is not responsible for any correlation between Maximum Flow values and λ_{ij} from either of the two acoustic dissimilarity measures.

2) *Results*: Figs. 5 and 6 summarize the distributions of the acoustic distance measures we are considering between the snowball sample entry point and the other nodes in the sample, given the Maximum Flow value between those nodes, both for the sampled graph and for the random permutation. Although the variations in the sampled graph for both distance measures [Figs. 5(a) and 6(a)] appear to the eye to be larger than in the random permutation [baseline, Figs. 5(b) and 6(b)] cases, the magnitude of the variations of the medians (summarized in Table II and Fig. 7) are not large, and performing a Kruskal-Wallis test for the differences of medians in the distributions of acoustic distance measures given the Maximum Flow (Table III) reveals that the observed variation of medians is not strongly out of line with chance ($p \approx 0.2$, compared with the expected baseline value of $p \approx 0.5$).

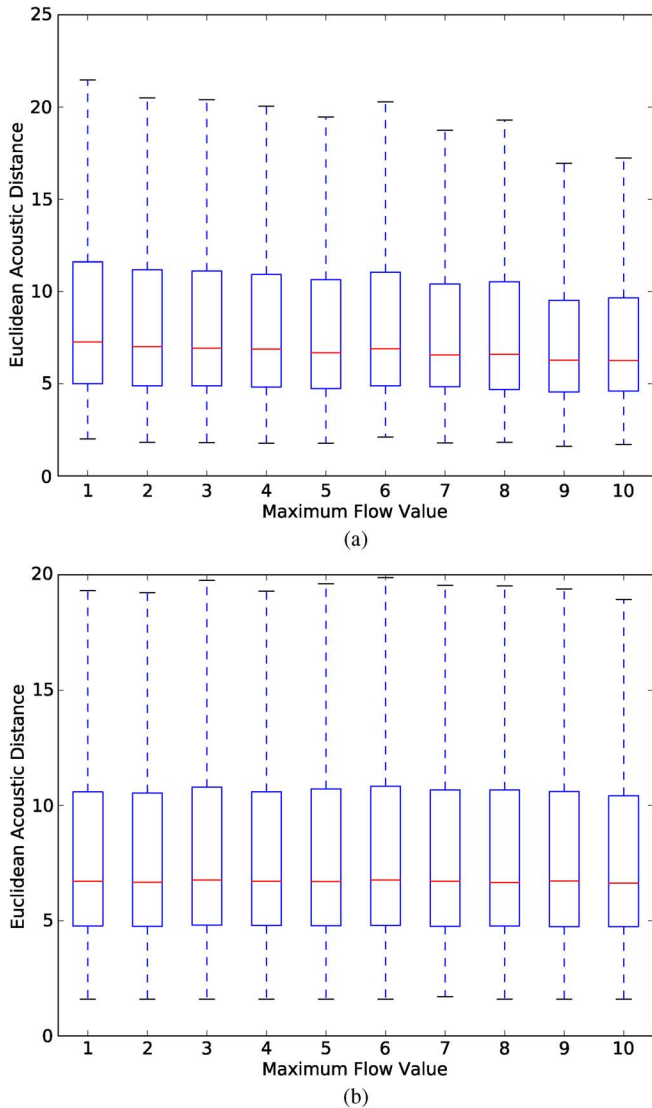


Fig. 6. Box and whisker plots showing the distribution of Euclidean distance grouped by Maximum Flow value between artists on the Myspace social graph and the randomized permutations of the graph. (a) Euclidean distance distribution on the sampled graph. (b) Euclidean distance distribution on the random permutations of the graph, maintaining the original edge structure.

This negative result (that we do not have evidence that the acoustic distances vary substantially with the social distance characterized by Maximum Flow) leads to the question of whether the two distance measures are in fact independent—that is, does knowing one (e.g., an acoustic distance) give any information at all about the other (the social relationship)? In order to answer this question, we compute the marginal and conditional entropies¹⁵ of the various distance distributions (Table IV). Here the Maximum Flow value distributions are the same, having an entropy of 3.1 bits. Further, the audio distance distributions have similar entropies (about 8.29 bits for the Euclidean distance and 8.65 bits for the GMM/EMD distance), and also that the conditional entropies

¹⁵All mutual information and related entropy calculations in this work are calculated using `pyentropy`, available at <http://code.google.com/p/pyentropy/>, a python library for performing information theoretic analysis on data distributions [36].




Fig. 7. Deltas from the global median for each Maximum Flow value group of acoustic distance values, from the sampled graph and the randomized graph.

for the Maximum Flow value given the acoustic measure are only very slightly lower—knowing the Euclidean distance reduces the uncertainty of the Maximum Flow by 0.1 bits (and knowing the GMM/EMD reduces this uncertainty by 0.4 bits). This small amount of mutual information can be interpreted as “almost” independence—for practical purposes, knowing how similar songs sound gives almost no information about the social relationships between them (or vice versa).

C. Using Audio in Community Detection

Both community detection algorithms described in Section III-C are based on the adjacency matrix A of the graph. This allows us to easily extend these algorithms to include audio-based similarity measures. We simply insert an inter-node similarity value for each non-zero entry in A . We calculate these similarity values using both the Earth-Mover’s Distance and Marsyas’ audio-based analysis methods described in Section IV. Dissimilarity values from these methods must be converted to similarity values to be applied to the community detection algorithms. We do this by taking the reciprocal of each dissimilarity:

$$A_{ij} = \begin{cases} \lambda_{ij}^{-1}, & \text{if nodes } i \text{ and } j \text{ are connected} \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

1) *Genre Entropy*: Now that we have several methods for detecting community structures in our network, we need a means of evaluating the relevance of these structures in the context of music. Traditionally, music and music artists are classified in terms of *genre*. If the structure of the Myspace artist network is relevant to music, we would expect the communities identified within the network to be correlated with musical genres. That is, communities should contain nodes with a more homogeneous set of genre associations than the network as a whole.

In our sampling of the Myspace network (described in Section III-A, we collected genre tags that are associated with each artist. In order to measure the diversity of each community

TABLE II
 NODE PAIRS OF MEDIAN ACOUSTIC DISTANCE VALUES GROUPED BY ACTUAL MINIMUM CUT VALUES AND RANDOMIZED MINIMUM CUT VALUES, SHOWN WITH DEVIATIONS FROM THE GLOBAL MEDIANS OF 39.53 FOR EMD AND 6.6848 FOR EUCLIDEAN DISTANCE. EMD WEIGHTS ARE ON THE LEFT AND EUCLIDEAN DISTANCES AS GENERATED BY MARSYAS ARE ON THE RIGHT

Max Flow	Earth Movers Distance				Marsyas generated Euclidean Distance			
	median	deviation	randomized	deviation	median	deviation	randomized	deviation
1	40.80	1.26	39.10	-0.43	7.256	0.571	6.710	0.025
2	45.30	5.76	38.34	-1.19	7.016	0.331	6.668	-0.016
3	38.18	-1.35	38.87	-0.66	6.932	0.247	6.764	0.079
4	38.21	-1.32	38.64	-0.89	6.872	0.187	6.707	0.022
5	40.00	0.47	39.11	-0.42	6.673	-0.011	6.695	0.010
6	41.77	2.25	39.02	-0.51	6.896	0.211	6.761	0.076
7	39.94	0.41	39.24	-0.29	6.568	-0.116	6.714	0.029
8	39.38	-0.15	38.76	-0.77	6.597	-0.087	6.660	-0.023
9	38.50	-1.03	38.87	-0.66	6.270	-0.414	6.717	0.032
10	39.07	-0.46	40.85	1.32	6.253	-0.431	6.623	-0.061

TABLE III
 KRUSKAL-WALLIS ONE-WAY ANOVA TEST RESULTS OF EMD AGAINST MAXIMUM FLOW FOR BOTH THE SAMPLED GRAPH AND ITS RANDOM PERMUTATIONS. THE H-VALUES ARE DRAWN FROM A CHI-SQUARE DISTRIBUTION WITH 10 DEGREES OF FREEDOM

	H-value	P-value
From sample	12.46	0.19
Random permutations	9.11	0.43

TABLE IV
 ENTROPY VALUES FOR THE ACOUSTIC DISTANCES AND MAXIMUM FLOW VALUES. X IS THE SET OF MAXIMUM FLOW VALUES, Y IS THE SET OF AUDIO DISTANCE MEASUREMENTS

audio distance type	$H(X)$	$H(X Y)$	$H(Y)$	$I(X;Y)$
Euclidean distance	3.10	3.00	8.29	0.100
GMM/EMD	3.10	2.72	8.65	0.375

with respect to genre, we use a variant of Shannon entropy we call *genre entropy* S . This approach is similar to that of Lambiotte [14]. For a given community C_k , we calculate genre entropy as

$$S_{C_k} = - \sum_{\gamma \in C_k} P_{\gamma|C_k} \log P_{\gamma|C_k} \quad (7)$$

where $P_{\gamma|C_k}$ is the probability of finding genre tag γ in community C_k . As the diversity of genre tags in a community C_k increases, the genre entropy S_{C_k} increases. As the genre tags become more homogeneous, the value of S_{C_k} decreases. If community C_k is described entirely by one genre tag, then $S_{C_k} = 0$. We can calculate an overall genre entropy S_G by including the entire network sample. In this way, we can evaluate each community identified by comparing S_{C_k} to S_G . If the community structures in the network are related to musical genre, we would expect the communities to contain more homogeneous mixtures of genre tags. That is, usually, we would expect $S_{C_k} \leq S_G$. However, as community size decreases the genre entropy will tend to decrease because fewer tags are available. To account for this, we create a random partitioning of the graph that results in the same number of communities with the same number

of nodes in each community and calculate the corresponding genre entropies S_{rand} to provide a baseline.

If an artist specified no genre tags, this node is ignored and makes no contribution to the genre entropy score. In our data set, 2.6% of artists specified no genre tags.

2) *Results*: The results of the various community detection algorithms are summarized in Fig. 8 and Table V. When the genre entropies are averaged across all the detected communities, we see that for every community detection method, the average genre entropy is lower than S_G as well as lower than the average genre entropy for a random partition of the graph into an equal number of communities. This is strong evidence that the community structure of the network is related to musical genre.

It should be noted that even a very simple examination of the genre distributions for the entire network sample suggests a network structure that is closely related to musical genre. Of all the genre associations collected for our data set, 50.3% of the tags were either ‘‘Hip-Hop’’ or ‘‘Rap’’ while 11.4% of tags were ‘‘R&B’’. Smaller informal network samples, independent of our main data set, were also dominated by a handful of similar genre tags (i.e., ‘‘Alternative’’, ‘‘Indie’’, ‘‘Punk’’). In context, this suggests our sample was essentially ‘‘stuck’’ in a community of Myspace artists associated with these particular genre inclinations. However, it is possible that these genre distributions are indicative of the entire Myspace artist network. Regardless, given that the genre entropy of our entire set is so low to begin with, it is an encouraging result that we could efficiently identify communities of artists with even lower genre entropies.

Without audio-based similarity weighting, the greedy modularity algorithm (gm) and the walktrap algorithm (wt) result in genre entropy distributions with no statistically significant differences. However, the walktrap algorithm results in almost five times as many communities, which we would expect to result in a lower genre entropies because of smaller community size. Also note that as discussed in Section III-C, the optimized greedy modularity algorithm is considerably faster than the walktrap algorithm.

With audio-based similarity weighting, we see mixed results. Applying audio weights to the greedy modularity algorithm (gm+a) actually increased genre entropies but the differences between gm and gm+a genre entropy distributions are not statistically significant. Audio-based weighting applied to the

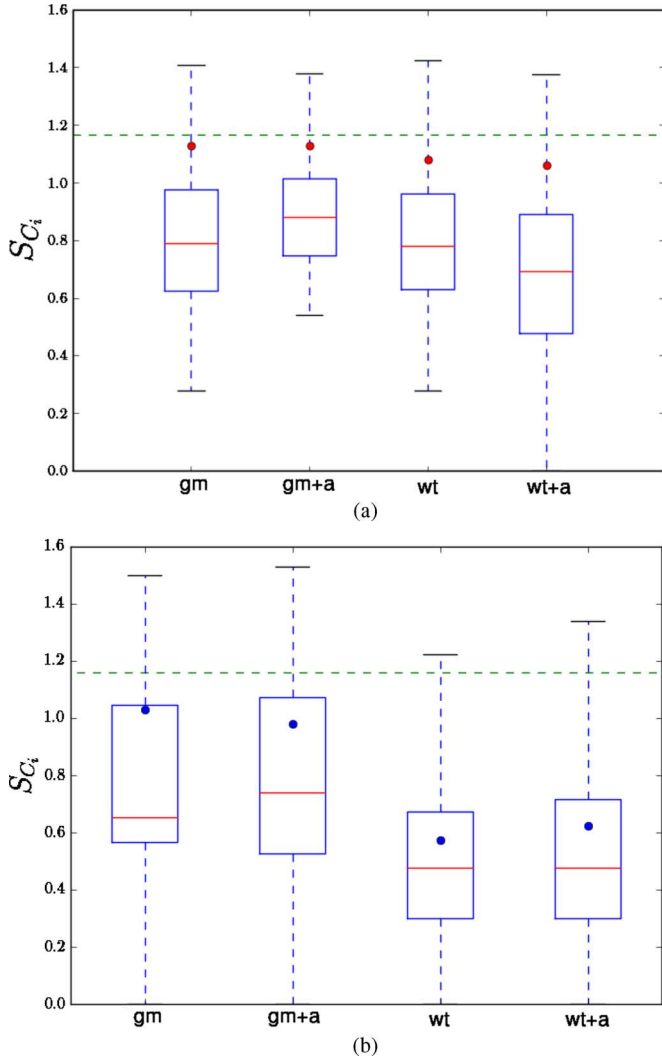


Fig. 8. Box and whisker plots showing the spread of community genre entropies for each graph partition method where gm is greedy modularity, gm+a is greedy modularity with audio weights, wt is walktrap, and wt+a is walktrap with audio weights. The horizontal line represents the genre entropy of the entire sample. The circles represent the average value of genre entropy for a random partition of the network into an equivalent number of communities. (a) uses the Earth Mover’s Distance for audio weight. (b) uses Euclidean distance from Marsyas.

TABLE V

RESULTS OF THE COMMUNITY DETECTION ALGORITHMS WHERE c IS THE NUMBER OF COMMUNITIES DETECTED, $\langle S_C \rangle$ IS THE AVERAGE GENRE ENTROPY FOR ALL COMMUNITIES, $\langle S_{rand} \rangle$ IS THE AVERAGE GENRE ENTROPY FOR A RANDOM PARTITION OF THE NETWORK INTO AN EQUAL NUMBER OF COMMUNITIES, AND Q IS THE MODULARITY FOR THE GIVEN PARTITION AS DEFINED IN (1)

algorithm	c	$\langle S_C \rangle$	$\langle S_{rand} \rangle$	Q
none	1	1.16	-	-
gm	42	0.81	1.13	0.61
gm+a	33	0.90	1.13	0.64
wt	195	0.80	1.08	0.61
wt+a	271	0.70	1.06	0.62

walktrap algorithm (wt+a) results in a statistically significant decrease in genre entropies compared to the un-weighted walktrap algorithm ($p = 4.2 \times 10^{-4}$).

D. Summary

In an effort to better understand and leverage both social connectivity and content-based dissimilarity, we conducted a number of experiments, in two categories: pairwise distance and community segmentation. Our pairwise distance work showed little in terms of a linear correlation, and when examining the mutual information of the two distance distributions, it becomes clear that the two encode largely independent spaces, with a very small information content overlap. When looking into community segmentation, we used genre entropy to see if using acoustic distance would improve the quality of segmentation. This addition of the content-based distance made only a slight difference to the segmentation; however, it is clear that the social structure tightly corresponds to the self-applied genre labels.

V. CONCLUSIONS

We have presented an analysis of the community structures found in a sample of the Myspace artist network. We have applied two efficient algorithms to the task of partitioning the Myspace artist network sample into communities and we have shown how to include audio-based similarity measures in the community detection process. We have evaluated our results in terms of genre entropy—a measure of genre tag distributions—and shown the community structures in the Myspace artist network are related to musical genre. The communities detected have lower entropy over genre labels than a graph with randomly permuted labels.

We compared social space of the Myspace sample with content-based acoustic space in two ways in Section IV. First the geodesic distances of pairs of artists were compared to the acoustic distance between these pairs of artists. Then Maximum Flow between pairs of artists was compared to both the acoustic distance between the artists and amongst the artists’ songs. While not perfectly orthogonal, the artist social graph and the acoustic dissimilarity matrix clearly encode different relational aspects between artists. This can be clearly seen in the small amount of mutual information shared between the sets of distances. The implication is that using both of these spaces in applications driven by similarity measures will result in much higher entropy in the data available to such an application. This suggests that a recommendation or discovery system that can use both domains well has the potential to perform much better than a similar system that relies on only one domain in isolation.

To understand more completely contributions, we revisit the questions posed in the introduction: what have we learned?

1) *Given That This Music is Published Within a Relational Space, How Can We Best Use All of the Available Information to Discover New Music?:* Broadly speaking, our work presented two potential ways to combine the disparate domains of social and content-based space. By weighting the social graph with a measure of acoustic distance, various techniques from complex analysis can be applied. Here we focused on pathfinding and community segmentation. Given the breadth of available techniques from complex networks, we cannot yet say if these works are *best*; however, pathfinding is a natural fit to the construction

of playlists and previous work has shown playlists to be excellent vehicles for music discovery [37]–[39].

2) *Can Both Social Metadata and Content-Based Comparisons Be Exploited to Improve Discovery of New Material?:* While this is similar to the previous question, when looked at this way, we can be considerably more definitive. When looking at the entirety of our experimentation, especially the mutual information across distributions seen in Section IV-B2, the answer to this question is a clear yes. While a complete end-user oriented system remains to be developed, this work shows that such a system would be better served drawn from social and acoustic driven notions of distance and similarity. This may well hold with a variety of different applications and data-sets as similar conclusions are reached using text-ranking combined with an audio-similarity measure to improve text-based music search results [40].

3) *Can This Crowd-Sourced Tangle of Social-Networking Ties Provide Insights Into the Dynamics of Popular Music?:* On this point, a clear conclusion from this work is that the expected linear correlation between social and acoustic distance is not present. So does an artist sound like their friends? While perhaps not what one would first guess, it appears the answer is that while an artist may sound like (i.e., similar to) his/her friends, he/she does not sound significantly dissimilar to artists that are not (i.e., artists which have a high social distance are only slightly further away in acoustic terms than those with a low social distance).

4) *Does the Structure of a Network of Artists Have Any Relevance to Music-Related Studies Such as Music Recommendation or Musicology?:* This work lays out the parts with which an engaging recommender system could be built or musicological study conducted. This compels further study. As the Myspace artist network is of interest to other researchers, we have converted our graph data to a more structured format. We have created a Web service¹⁶ that describes any Myspace page in a machine-readable Semantic Web format. Using FOAF¹⁷ and the Music Ontology¹⁸ [41], the service models a Myspace page in RDF and serializes it as XML RDF. This will allow future applications to easily make use of Myspace network data (e.g., for music recommendation).

While it is unclear how to best use all the available information from the wide range of artists and musicians, what this work makes clear is that there are advantages to complex multi-domain notions of similarity in music. By using both acoustic and social data, recommender systems have more avenues available to find new material to suggest to users in a transparent way. Whether either of these spaces can provide insight into the other remains an open question, though our work tends to show the likely predictability of one space from the other is low. In spite or perhaps because of this separation, and given the sheer quantity of data available on the web, it seems inevitable that these domains will be used in tandem in future music recommendation and musicological study.

¹⁶Available at <http://dbtune.org/myspace>.

¹⁷<http://www.foaf-project.org/>.

¹⁸<http://musicontology.com/>.

VI. FUTURE WORK

We explore two distinct yet related efforts to extend this work: extending and improving sampling to better understand network structures and end-user focused applications based on the construction of playlist which walk the captured network.

A. Understanding Network Ecology

In future work, we plan to examine community detection methods that operate locally, without knowledge of the entire network. We also plan to address further directed artist graph analysis, bipartite networks of artists and listeners, different audio analysis methods, and the application of these methods to music recommendation.

Many of these tasks require the expansion of our sample network. The goal of any effort to expand the sample size of a network such as Myspace is best focused on ways to make the sample set more indicative of the whole. While it is impossible to assess this without capturing the entire graphs, some assumptions can be made. Snowball sampling has a tendency to over-sample hubs. Given this, a better expanded network is likely to result through the selections of new starting seed artist (most likely at random) and proceeding via a breadth-first crawl until that crawl results in overlap with the known network. It is reasonable to assume that this method, when used over multiple hubs, will produce a lower proportion of high centrality hubs than simply continuing further with the existing breadth-first crawl. With a lower proportion of these over-sampled hubs, the social structure of the sample would better match that of the whole.

B. Engineering Playlist-Based Applications

In an effort to create domain-specific recommender and discovery systems, we outline some possible ways to extend this work to end-listener applications. The playlist is an ideal means for this (e.g., [38], [39]) and such applications could then be evaluated using recommender system standard practice [42].

1) *Max Flow Playlist:* To build playlists using both acoustic and social-network data, the Earth Mover's Distance is used between each pair of neighbors as weights on the Myspace sample network. Two artists are then selected, a starting artist as the source node and a final artist as the sink node. One or more paths are then found through the graph via the Maximum Flow value, generating the list and order of artists for the playlist. The song used for each artist is the most popular at the time of the page scrape. In this way, playlists are constructed that are both influenced by timbre similarity and bound by social context, regardless of any relationship found between these two spaces found via the work discussed in Section IV. Playlists generated using this technique were informally auditioned, and were found to be reasonable on that basis.

There is clearly potential in the idea of the Maximum Flow playlist. When using either audio similarity measure as a weight, the results appear to be quite good, at least from a qualitative perspective. The imposed constraint of the social network alleviates to some extent shortcomings of a playlist built purely through the analysis of acoustic similarity by moving more toward the balance between uniformly acoustically-similar works and completely random movement.

2) *Steerable Optimized Self-Organizing Radio*: Using the song-centric graph, the following system is in development as a means of deployment and testing. This system is designed to play a continuous stream of songs via an Internet radio stream. The playback system begins with an initial seed song and destination song and then constructs a playlist. While this playlist is being broadcast, anyone tuning into the broadcast is able to vote via a web-based application on the next song to serve as the destination. In order to produce a usable output, the vote system presents a list of *nominees*, each selected as a representative track from various communities as segregated via means discussed in Section IV-C.

Once the current destination song begins to broadcast, the voting for the next cycle ceases. This destination song is then considered the seed song for the next cycle and the song with the most votes becomes the new destination, then the next playlist will be calculated and its members broadcast. This process will continue for the duration of the broadcast. Once this automatic playlist creation system is allowed to run for a sufficient amount of time, a great deal of user data will be recorded. This would include direct preference feedback, voting behavior, average length of time continuously listened, and whether listeners (or at least IP addresses) return. This provides a built-in means of human listener evaluation for these playlists.

It is hoped that this system, or one like it, will provide an application driven means to evaluate the usability of the measures explored in this work in task of music discovery and recommendation.


ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers: their comments have helped strengthen this work significantly.

REFERENCES


- [1] R. Steuer, J. Kurths, C. Daub, J. Weise, and J. Selbig, "The mutual information: Detecting and evaluating dependencies between variables," *Bioinformatics* vol. 18, no. Suppl.2, pp. S231–S240, 2002. [Online]. Available: <http://direct.bl.uk/bld/PlaceOrder.do?UIN=132115618&ETOC=RN&from=search.engine>.
- [2] M. E. J. Newman, "The structure and function of complex networks," *SIAM Rev.* vol. 45, p. 167, 2003. [Online]. Available: <http://www.citebase.org/abstract?id=oai:arXiv.org:cond-mat/0303516>.
- [3] L. F. Costa, F. A. Rodrigues, G. Travieso, and P. R. V. Boas, "Characterization of complex networks: A survey of measurements," *Adv. Phys.*, vol. 56, p. 167, 2007, 10.1080/00018730601170527.
- [4] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Upper Saddle River, NJ: Prentice-Hall, 1993.
- [5] H. Nagamochi and T. Ibaraki, "Computing edge-connectivity in multi-graphs and capacitated graphs," *SIAM J. Discret. Math.*, vol. 5, no. 1, pp. 54–66, 1992.
- [6] P. Elias, A. Feinstein, and C. Shannon, "A note on the maximum flow through a network," *IEEE Trans. Inf. Theory*, vol. 2, no. 4, pp. 117–119, Dec. 1956.
- [7] A. V. Goldberg and R. E. Tarjan, "A new approach to the maximum-flow problem," *J. ACM*, vol. 35, no. 4, pp. 921–940, 1988.
- [8] P. Knees, T. Pohle, M. Schedl, and G. Widmer, "Combining audio-based similarity with web-based data to accelerate automatic music playlist generation," in *Proc. 8th ACM Int. Workshop Multimedia Information Retrieval*, 2006, pp. 147–154.
- [9] M. Alghoniemy and A. Tewfik, "A network flow model for playlist generation," in *Proc. IEEE Int. Conf. Multimedia and Expo*, 2001. [Online]. Available: citeseer.ist.psu.edu/alghoniemy01network.html.
- [10] P. Cano, O. Celma, M. Koppenberger, J. M., and Buldu, "The topology of music recommendation networks," *Chaos: Interdiscipl. J. Nonlin. Sci.* 2006. [Online]. Available: <http://www.citebase.org/abstract?id=oai:arXiv.org:physics/0512266>; [Online]. Available: <http://arxiv.org/abs/physics/0512266v1>.
- [11] P. Gleiser and L. Danon, "Community structure in jazz," *Adv. Complex Syst.*, vol. 6, no. 4, pp. 565–573, 2003.
- [12] J. Park, O. Celma, M. Koppenberger, P. Cano, and J. M. Buldu, "The social network of contemporary popular musicians," *Int. J. Bifurcat. Chaos* vol. 17, pp. 2281–2288, 2007. [Online]. Available: <http://www.citebase.org/abstract?id=oai:arXiv.org:physics/0609229>.
- [13] A. Anglade, M. Tiemann, and F. Vignoli, "Virtual communities for creating shared music channels," in *Proc. Int. Symp. Music Information Retrieval*, 2007.
- [14] R. Lambiotte and M. Ausloos, "On the genre-fication of music: A percolation approach (long version)," *Eur. Phys. J. B*, vol. 50, no. 183, 2006, 10.1140/epjb/e2006-00115-0.
- [15] A. Clauset, M. E. J. Newman, and C. Moore, "Finding community structure in very large networks," *Phys. Rev. E*, vol. 70, no. 6, p. 066111, Dec. 2004.
- [16] P. Pons and M. Latapy, "Computing Communities in Large Networks Using Random Walks (Long Version)," arXiv:physics/0512106v1, 2005. [Online]. Available: <http://www.citebase.org/abstract?id=oai:arXiv.org:physics/0512106>.
- [17] B. Logan, "Mel frequency cepstral coefficients for music modeling," in *Proc. Int. Symp. Music Information Retrieval*, 2000.
- [18] M. Casey, R. Veltkamp, M. Goto, M. Leman, C. Rhodes, and M. Slaney, "Content-based music information retrieval: Current directions and future challenge," *Proc. IEEE*, vol. 96, no. 4, pp. 668–696, Apr. 2008.
- [19] A. Berenzweig, B. Logan, D. P. W. Ellis, and B. P. W. Whitman, "A large-scale evaluation of acoustic and subjective music-similarity measures," *Comput. Music J.*, vol. 28, no. 2, pp. 63–76, 2004.
- [20] G. Tzanetakis, "Marsyas: A case study in implementing music information retrieval systems," *Inf. Sci. Ref.*, 2007. [Online]. Available: http://marsyas.sness.net/pdfs/0000/0007/imis_bookchapter.pdf.
- [21] B. Logan and A. Salomon, "A music similarity function based on signal analysis," in *Proc. IEEE Int. Conf. Multimedia and Expo (ICME 2001)*, Aug. 22–25, 2001, pp. 745–748.
- [22] E. Pampalk, "Computational models of music similarity and their application in music information retrieval," Ph.D. dissertation, Tech. Univ. Wien, Wien, Germany, May 2006.
- [23] J. S. Downie, "The music information retrieval evaluation exchange (mirex)," *D-Lib Mag.* Dec. 2006. [Online]. Available: <http://dlib.org/dlib/december06/downie/12downie.html>.
- [24] J. S. Downie, "The music information retrieval evaluation exchange (2005–2007): A window into music information retrieval research," *Acoust. Sci. Technol.* vol. 29, no. 4, pp. 247–255, 2008. [Online]. Available: http://www.jstage.jst.go.jp/article/ast/29/4/29_247_article.
- [25] Y. Rubner, C. Tomasi, and L. J. Guibas, "The earth mover's distance as a metric for image retrieval," *Int. J. Comput. Vis.*, vol. 40, no. 2, pp. 99–121, 2000.
- [26] J.-J. Aucouturier and F. Pachet, "Improving timbre similarity: How high's the sky?," *J. Negative Results Speech Audio Sci.*, 2004.
- [27] D. K. Mahajan and M. Slaney, "Image classification using the web graph," in *Proc ACM Multimedia*, Florence, Italy, 2010, ACM.
- [28] Y.-Y. Ahn, S. Han, H. Kwak, S. Moon, and H. Jeong, "Analysis of topological characteristics of huge online social networking services," in *Proc. 16th Int. Conf. World Wide Web*, Banff, AB, Canada, May 2007, IW3C2.
- [29] S. H. Lee, P.-J. Kim, and H. Jeong, "Statistical properties of sampled networks," *Phys. Rev. E*, vol. 73, pp. 102–109, Jan. 2006.
- [30] H. Kwak, S. Han, Y.-Y. Ahn, S. Moon, and H. Jeong, "Impact of Snowball Sampling Ratios on Network Characteristics Estimation: A Case Study of Cyworld, KAIST, Tech. Rep. CS/TR-2006-262, Nov. 2006.
- [31] K. Jacobson and M. Sandler, "Musically meaningful or just noise, an analysis of on-line artist networks," in *Proc. CMMR*, 2008, pp. 306–314.
- [32] L. A. N. Amaral, A. Scala, M. Barthélemy, and H. E. Stanley, "Classes of small-world networks," in *Proc. Nat. Acad. Sci.*, 2000.
- [33] K. Jacobson, B. Fields, and M. Sandler, "Using audio analysis and network structure to identify communities in on-line social networks of artists," in *Proc. Int. Symp. Music Information Retrieval*, 2008.
- [34] B. Fields, K. Jacobson, M. Casey, and M. Sandler, "Do you sound like your friends? exploring artist similarity via artist social network relationships and audio signal processing," in *Proc. Int. Computer Music Conf.*, Aug. 2008.

- [35] B. Fields, K. Jacobson, C. Rhodes, and M. Casey, "Social playlists and bottleneck measurements: Exploiting musician social graphs using content-based dissimilarity and pairwise maximum flow values," in *Proc. Int. Symp. Music Information Retrieval*, Sep. 2008.
- [36] R. A. Ince, R. S. Petersen, D. C. Swan, and S. Panzeri, "Python for information theoretic analysis of neural data," *Front Neuroinform.* vol. 3, no. 4, 2009. [Online]. Available: <http://www.hubmed.org/display.cgi?uids=19242557>.
- [37] S. J. Cunningham, D. Bainbridge, and A. Falconer, "More of an art than a science: Supporting the creation of playlists and mixes," in *Proc. Int. Conf. Music Information Retrieval*, 2006.
- [38] A. Flexer, D. Schnitzer, M. Gasser, and G. Widmer, "Playlist generation using start and end songs," in *Proc. Int. Symp. Music Information Retrieval*, Oct. 2008.
- [39] P. Lamere and D. Eck, "Using 3d visualizations to explore and discover music," in *Proc. Int. Conf. Music Information Retrieval*, 2007.
- [40] P. Knees, T. Pohle, M. Schedl, D. Schnitzer, K. Seyerlehner, and G. Widmer, "Augmenting text-based music retrieval with audio similarity: Advantages and limitations," in *Proc. Int. Conf. Music Information Retrieval*, 2009.
- [41] Y. Raimond, S. Abdallah, M. Sandler, and F. Gaisson, "The music ontology," in *Proc. Int. Conf. Music Information Retrieval*, 2007.
- [42] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Trans. Knowl. Data Eng.*, vol. 17, no. 6, pp. 734–749, Jun. 2005.



Ben Fields (S'03) received the B.S. degree in computer engineering from the University of California, Santa Cruz, in 2004 and the M.S. degree in music engineering technology from the Frost School of Music, University of Miami, Miami, FL. He is currently pursuing the Ph.D. degree in computer science at Goldsmiths, University of London, London, U.K., where he expects to complete the degree in the spring of 2011.

His research interests include sequentially-ordered music recommenders, large-scale music informatics, the boundary between computational analysis and linked data, and modeling notions of similarity.



Kurt Jacobson (S'08) received the Ph.D. degree in electrical engineering from the Centre for Digital Music at Queen Mary University of London, London, U.K., in 2011.

His research interests include the application of network analysis methods and semantic web technologies to music informatics tasks. He is currently working as a knowledge engineer for The Echo Nest corporation in Somerville, MA.



Christophe Rhodes received the B.A. degree in physics and the Ph.D. degree in applied mathematics from the University of Cambridge, Cambridge, U.K., in 2000 and 2004, respectively.

His Ph.D. research formed part of the Planck satellite science program. He lectures in Computer Science at Goldsmiths, University of London, London, U.K., where his current research interests span musical information retrieval and representation, user interfaces, and language design and implementation. He is co-founder and Chief Scientist of Teclio Net-

works AG, a company delivering unique optimization and acceleration products for mobile broadband networks.



Mark d'Inverno received the B.A. degree in mathematics and the M.Sc. degree in computation from Oxford University, Oxford, U.K., in 1986 and 1988, respectively, and the Ph.D. degree from University College London, London, U.K., in 1998.

He became a Professor of computer science in 2001. His background is in the application of formal methods in providing frameworks for new and emerging computer science collaborations but joined Goldsmiths College in 2006 as a Professor of computing to pursue an increasing interest in interdisciplinary research including the crossover with art, music, design, biology, and energy conservation. In 2007, he became head of the Computing Department which has built a reputation as one of the leading European institutions in interdisciplinary computer science research. He has co-authored and edited five books, including *Computers and Creativity* (New York: Springer, to be published), and published around 100 papers.

Dr. d'Inverno is co-founder and first chair of the European Workshop on Multi-Agent Systems.



Mark Sandler (SM'98) was born in 1955. He received the B.Sc. and Ph.D. degrees from the University of Essex, Essex, U.K., in 1978 and 1984, respectively.

He is a Professor of signal processing at Queen Mary, University of London, London, U.K., and Head of the School of Electronic Engineering and Computer Science. He has published over 300 papers in journals and conferences.

Dr. Sandler is a Fellow of the Institute of Electronic Engineers (IEE) and a Fellow of the Audio Engineering Society. He is a two-time recipient of the IEE A. H. Reeves Premium Prize.



Michael Casey (M'99) received the Ph.D. degree from the Media Laboratory, Massachusetts Institute of Technology, Cambridge, in 1998.

He is Visiting Professor of computer science at Goldsmiths College, University of London, London, U.K. He was the Principal Investigator at Goldsmiths for two MIR-related grants: Online Music Recognition and Searching II (OMRAS2) and Hierarchical Segmentation and Markup of Musical Audio (SeMMA). He was a Contributing Editor of the MPEG-7 international standard for content-based multimedia description ISO15938-4 (audio) 2002 while a Research Scientist at Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, before taking a position at Goldsmiths, in 2004. In 2008, he joined the faculty of the Music Department at Dartmouth College, Hanover, NH.