

Exploring Music Artists via Descriptive Terms and Multimedia Content

Markus Schedl and Tim Pohle
music@jku.at

Department of Computational Perception
Johannes Kepler University
Linz, Austria
<http://www.cp.jku.at>

Abstract. This paper presents an approach to browse collections of web pages about music artists by means of descriptive terms and multimedia content. To this end, a user interface called *Three-Dimensional Co-Occurrence Browser (3D-COB)* is introduced. 3D-COB automatically extracts and weights terms from artist-related web pages. This textual information is complemented with information on the multimedia content found on the web pages. For the user interface of 3D-COB, we elaborated a three-dimensional extension of the Sunburst visualization technique. The hierarchical data to be visualized is obtained by analyzing the web pages for combinations of co-occurring terms that are highly ranked by a term weighting function.

As for evaluation, we investigated different term weighting strategies in a first user study. A second user study was carried out to assess ergonomic aspects of 3D-COB, especially its usefulness for gaining a quick overview of a set of web pages and for efficiently browsing within this set.

1 Introduction

Automatically finding descriptive terms for a given music artist is an important task in music information retrieval (MIR). Such terms may describe, for example, the genre or style of the music performed by the artist under consideration and enable a wide variety of applications, e.g., enriching music players [19], recommending unknown artists based on the user's favorite artists [24], or enhancing user interfaces for browsing music collections [13, 17, 10, 15, 22].

One possibility for assigning musically relevant terms to a given artist is manual annotation by music experts or communities, as it is usually employed by music information systems like *allmusic* [5] and *last.fm* [4] or interfaces for music search like *musiclens* [3]. However, this is a very labor-intensive task and barely feasible for huge music collections. An alternative way, which we follow here, is to exploit today's largest information source, the World Wide Web. Automatically deriving information about music artists from the web is advantageous since it incorporates the opinions of a large number of different people, and thus embodies a kind of cultural knowledge.

The *Three-Dimensional Co-Occurrence Browser (3D-COB)* presented here automatically indexes a set of web pages about music artists according to a dictionary of musically relevant terms and organizes these web pages by creating a number of subsets, each of which is described by a set of terms. The terms that describe a particular subset are determined by a term weighting function. The subsets are then visualized using a variant of the Sunburst technique [8, 21].

The purpose of 3D-COB is threefold. First, it facilitates getting an *overview* of the web pages related to a music artist by structuring them according to co-occurring terms. Second, since the descriptive terms that most often occur on web pages related to a music artist X constitute an individual profile of X , 3D-COB is also suited to reveal various *meta-information* about the artist, e.g., musical style, related artists, or instrumentation. Third, by visualizing the amount of *multimedia content* provided at the indexed web pages, the user is offered a means of exploring the audio, image, and video content of the respective set of web pages.

2 Related Work

This paper is mainly related to the two research fields of *web-based music information retrieval* and *information visualization of hierarchical data*, which will be covered in the following.

Determining terms related to a music artist via web-based MIR has first been addressed in [23], where Whitman and Lawrence extract different term sets (e.g., noun phrases and adjectives) from artist-related web pages. Based on term occurrences, individual term profiles are created for each artist. The authors then use the overlap between the term profiles of two artists as an estimate for their similarity. A quite similar approach is presented in [12]. Knees et al. however do not use specific term sets, but create a term list directly from the retrieved web pages. To this list, a term selection technique is applied to filter out less important terms. Hereafter, the TF-IDF measure is used to weight the remaining words and subsequently create a weighted term profile for each artist. Knees et al. propose their approach for artist-to-genre classification and similarity measurement. Both approaches presented so far barely address the potential of utilizing the large amount of terms which, though not directly related to the artist, occur on many artist-related web pages. In contrast, Pampalk et al. in [16] use a dictionary of about 1,400 musically relevant terms to index artist-related web pages. Different term weighting techniques are applied to describe each artist with some terms. Furthermore, the artists are hierarchically structured using a special version of the Self-Organizing Map. The authors show that considering only the terms in the dictionary for term weighting and clustering outperforms using all terms found on the extracted web pages. An approach to assign descriptive terms to a given artist is presented in [19]. Schedl et al. use co-occurrences derived from artist-related web pages to estimate the conditional probability for the artist name under consideration to be found on a web page containing a specific descriptive term and vice versa. To this end, a set of predefined genres

and other attributes, like preferred tempo or mood of the artist’s performance, is used. The aforementioned probabilities are then calculated, and the most probable value of the attribute under consideration is assigned to the artist. Independent of Schedl et al., Geleijnse and Korst present in [9] an approach that differs from [19] only in regard to the normalization used.

The 3D-COB proposed here uses a dictionary like [16] to extract artist-related information from web pages. However, the clustering is performed in a very different way and on a different level (for individual web pages instead of artists). Related work on visualizing hierarchical data primarily focuses on the *Sunburst* approach, as the 3D-COB extends the Sunburst in various aspects. The Sunburst as proposed in [8, 21] is a circular, space-filling visualization technique. The center of the visualization represents the highest element in the hierarchy, whereas elements on deeper levels are illustrated by arcs further away from the center. Child elements are drawn within the angular borders of their parent, but at a more distant position from the center. In almost all publications related to the Sunburst, its usual application scenario is browsing the hierarchical tree structure of a file system. In this scenario, directories and files are represented by arcs whose sizes are proportional to the sizes of the respective directories/files. In the case of the 3D-COB, however, some *constraints for the size of the visualization* are necessary. Furthermore, the arc sizes are determined by the *term weighting function*, which is applied to select the most important terms (for the clustering). Moreover, the 3D-COB allows for *encoding an additional data dimension* in the height of each arc. This dimension is used to visualize the amount of multimedia content provided by the analyzed web pages. As three different types of multimedia content are taken into account (audio, image, and video), the Sunburst stack of the 3D-COB consists of three individual Sunburst visualizations.

Other space-filling visualization techniques for hierarchically structured data include the *Treemap* [11]. In contrast to the Sunburst, the Treemap uses a rectangular layout and displays elements further down in the hierarchy embedded in the rectangle of their parent element. The Sunburst, however, displays all elements that reside on the same hierarchy level on the same torus, which facilitates getting a quick overview of the hierarchy.

3 The Three-Dimensional Co-Occurrence Browser (3D-COB)

To get a first impression of the appearance of the 3D-COB user interface, the reader is invited to take a look at Figure 1. This figure shows a stack of three three-dimensional Sunburst visualizations created from 161 web pages of the band “Iron Maiden”. Details on the information gathering process, the creation of the visualization, and the user interaction possibilities are provided in the following subsections.

3D-COB has been implemented using the *processing* environment [6] and the *CoMIRVA* framework [18]. CoMIRVA already contained an implementation of the two-dimensional version of the Sunburst. We heavily extended this version by

heaving it to the third dimension and incorporating various multimedia content in the visualization. To this end, we particularly had to elaborate new data representation and user interaction models.

3.1 Retrieval and Indexing

Given the name of an artist, we first query *Google* with the scheme “*artist name*” +*music* +*review* to obtain the URLs of up to 1,000 web pages related to the artist, whose content we then retrieve. Subsequently, a term analysis step is performed. To this end, we use a dictionary of musically relevant terms, which are searched in all web pages of every artist, yielding an inverted file index. For the conducted experiments, a manually compiled dictionary that resembles the one used in [16] was utilized. It was assembled using various sources such as [7, 5, 1] and contains music genres and styles, instruments, moods, and other terms which are somehow related to music. Altogether, the dictionary contains 1,506 terms.

As for indexing the multimedia content of the web pages, we first extract a list of common file extensions for audio, image, and video files from [2]. We then search the HTML code of each web page for links to files whose file extension occur in one of the extracted lists. Finally, we store the URLs of the found multimedia files and the inverted file index gained by the term analysis in an XML data structure.

3.2 Creation of the Visualization

Using the inverted index of the web pages of an artist X , we can easily extract subsets $S_{X,\{t_1, \dots, t_r\}}$ of the web page collection of X which have in common the occurrence of all terms t_1, \dots, t_r .

Starting with the entire set of web pages $S_{X,\{\}}$ of X , we use a term weighting function (e.g., document frequency, term frequency, TF·IDF) to select a maximum number N of terms with highest weight, which are used to create N subsets $S_{X,\{t_1\}}, \dots, S_{X,\{t_N\}}$ of the collection. These subsets are visualized as arcs $A_{X,\{t_1\}}, \dots, A_{X,\{t_N\}}$ around a centered cylinder which represents the root arc $A_{X,\{\}}$, and thus the entire set of web pages. The angular extent of each arc is proportional to the weight of the associated term t_i , e.g., to the number of web pages containing t_i when using document frequencies for term weighting. To avoid very small, thus hardly perceivable, arcs, we omit arcs whose angular extent is smaller than a fixed threshold E , measured in degrees. Furthermore, each arc is colored with respect to the relative weight of its corresponding term t_i (relative to the maximum weight among all terms). The term selection and the corresponding visualization step are recursively performed for all arcs, with a maximum R for the recursion depth. This eventually yields a complete Sunburst visualization, where each arc at a specific recursion depth r represents a set of web pages $S_{X,\{t_1, \dots, t_r\}}$ in which all terms t_1, \dots, t_r co-occur.

As for representing the multimedia content found on the web pages, in each layer of the Sunburst stack, the amount of a specific category of multimedia files is

depicted. To this end, we encode the relative number of audio, image, and video files in the height of the arcs (relative to the total number represented by the root node of the respective layer). For example, denoting the audio layer as L_A , the image layer as L_I , and the video layer as L_V and focusing on a fixed arc A , the height of A in L_I shows the relative number of image files contained in the web pages that are represented by arc A , the height of A in L_V illustrates the relative number of video files, and the height of A in L_A the relative number of audio files. The corresponding multimedia files can easily be accessed via the user interface of 3D-COB.

3.3 User Interface and User Interaction

Figure 1 depicts a screenshot of 3D-COB’s user interface for 161 web pages retrieved for “Iron Maiden”. The constraints were set to the following values: $N = 6$, $R = 8$, $E = 5.0$ (cf. Subsection 3.2). Document frequencies were used for term weighting. Each arc $A_{X,\{t_1, \dots, t_r\}}$ is labeled with the term t_r that subdivides the web pages represented by the arc’s parent node $A_{X,\{t_1, \dots, t_{r-1}\}}$ into those containing t_r and those not containing t_r . Additionally, the weight of the term t_r is added in parentheses to the label of each arc $A_{X,\{t_1, \dots, t_r\}}$. The topmost layer illustrates the amount of video files found on the web pages, the middle one the amount of image files, and the lower one the amount of audio files. In the screenshot shown in Figure 1, the arc representing the web pages on which all of the terms “Iron Maiden”, “guitar”, and “metal” co-occur is selected. Since document frequencies were used for this screenshot, determining the exact number of web pages represented by a particular arc is easy: 74 out of the complete set of 161 web pages contain the mentioned terms.

User interaction is provided in several ways. First, the mouse can be used to *rotate* the Sunburst stack around the Y-axis, i.e., the vertical axis going through the root nodes of all Sunbursts in the stack. *Zooming* in/out (within predefined boundaries) is provided as well as *changing the inclination of the stack*, which is limited to angles between a front view and a bird’s eye view. To select a particular arc, e.g., to access the multimedia content of the corresponding web pages, the arrow keys can be used to navigate in the hierarchy. The currently selected arc is highlighted by means of drawing a white border around it and coloring its label in white. So are all previously selected arcs at higher hierarchy levels. This facilitates tracing the selection back to the root arc and quickly recognizing all co-occurring terms on the web pages represented by the selected arc.

In addition to the basic interaction capabilities described so far, the following functionalities are provided.

- Creating a new visualization based on the subset of web pages given by the selected arc.
- Restoring the original visualization that incorporates all web pages in its root node.
- Showing a list of web page URLs which are represented by the selected arc.
- Displaying and browsing a list of audio, image, or video files, which are found on the web pages of the currently selected arc.




Fig. 1. The user interface of 3D-COB for a collection of web pages about the band “Iron Maiden”.

- Opening the web pages or the available multimedia files represented by the selected arc.

4 Evaluation of the Term Weighting Functions

We experimented with three different term weighting functions (document frequency, term frequency, TF-IDF) for term selection in the Sunburst creation step, cf. Subsection 3.2. Given a set of web pages S of an artist, the *document frequency* DF_t of a term t is defined as the absolute number of pages in S on which t appears at least once. The *term frequency* TF_t of a term t is defined as the sum of all occurrences of t in S . The *term frequency inverse document frequency* measure $TF \cdot IDF_t$ of t is calculated as $TF_t \cdot \ln \frac{|S|}{DF_t}$. To assess the influence of the term weighting function on the quality of the hierarchical clustering, the hierarchical layout, and thus on the visualization of 3D-COB, we conducted a user study as detailed in the following.

4.1 Setup

For the user study, we chose a collection of 112 well-known artists (14 genres, 8 artists each). Indexing was performed as described in Subsection 3.1. To create the evaluation data, for each artist, we calculated on the complete set of his/her retrieved and indexed web pages, the 10 most important terms using each of the three term weighting functions. To avoid biasing of the results, we combined, for each artist, the 10 terms obtained by applying every weighting function. Hence, every participant was presented a list of 112 artist names and, for each of the name, a set of associated terms (as a mixture of the terms obtained by the three weighting functions). Since the authors had no a priori knowledge of which artists were known by which participant, the participants were told to evaluate only those artists they were familiar with. Their task was then to rate the associated terms with respect to their appropriateness for describing the artist or his/her music. To this end, they had to associate every term to one of the three classes $+$ (*good description*), $-$ (*bad description*), and \sim (*indifferent or not wrong, but not a description specific for the artist*).

Due to time constraints, we had to limit the number of participants in the user study to five. Three of them are computer science students, the other two researchers in computer science. All of them are male and all stated to listen to music often.

4.2 Results and Discussion

We received a total of 172 assessments for sets of terms assigned to a specific artist. 92 out of the 112 artists were covered. To analyze the results, we calculated, for each artist and weighting function, the sum of all points obtained by the assessments. As for the mapping of classes to points, each term in class $+$ contributes 1 point, each term in class $-$ gives -1 point, and each term in class \sim yields 0 points.

The complete results for each artist assessed cannot be depicted here due to space limitations.¹ However, Table 1 shows the sum of the points (over all artists) for each term weighting function as well as the average score obtained for an arbitrarily chosen artist. Investigating the averaged points for different artists (which are not shown in the table) reveals that the quality of the terms vary strongly between the artists. Nevertheless, we can state that, for most artists, the number of descriptive terms by far exceeds that of the non-descriptive ones. Due to the performed mapping from classes to points, the averaged score values can be regarded as the average excess of the number of good terms over the number of bad terms. Hence, overall, we assume that the document frequency measure performed best, the term frequency second best, and the TF-IDF worst. To test for the significance of the results, we performed Friedman’s non-parametric two-way analysis of variance [20]. This test is similar to the two-way ANOVA, but does not assume a normal distribution of the data. The test yielded a p of

¹ Detailed results per artist can be found at [omitted due to blind review].

0.000024. Therefore, it is highly probable that the variance differences in the results are significant. Moreover, pairwise comparisons between the results given by the three term weighting functions showed that TF-IDF performed significantly worse than both TF and DF, whereas no significant difference could be made out between the results obtained using DF and those obtained using TF. The laborious task of combining and analyzing the different assessments of the participants in the user study allowed us to take a qualitative look at the terms. Overall, the majority of the terms was judged descriptive. However, we discovered an interesting source of error: the erroneous assignment of a term to an artist if the term is part of artist, album, or song name. Examples for this problem are “infinite” for the artist “Smashing Pumpkins” as well as “human” and “punk” for the artist “Daft Punk”.

	TF	DF	TF-IDF
sum	386.00	413.00	271.00
avg	2.22	2.44	1.53

Table 1. For each term weighting function, the summed up user ratings and the average ratings per artist.

5 Evaluating the User Interface

To investigate the usefulness of 3D-COB for gaining a quick overview of a set of artist-related web pages and efficiently browsing within this set, we conducted a second user study that primarily focuses on ergonomic aspects of 3D-COB.

5.1 Setup

We formulated the following tasks, which we believe are important for the mentioned purposes, and evaluated them in a quantitative manner:

1. Which are the five top-ranked terms that occur on the web pages mentioning “Iron Maiden”?
2. Indicate the number of web pages containing all of the terms “Iron Maiden”, “metal”, and “guitar”.
3. Show a list of web pages that contain the terms “Iron Maiden” and “british”.
4. Considering the complete set of web pages, which are the three terms that co-occur on the highest number of web pages?
5. How many web pages contain the terms “Iron Maiden” and “metal”, but not the term “guitar”?
6. Display a list of audio files available at web pages containing the term “Iron Maiden”.

7. Which terms co-occur on the set of web pages that contains the highest number of image files in hierarchy level three?
8. Indicate the URL of one particular web page that contains image files but no video files.
9. How many web pages does the complete collection contain?
10. Find one of the deepest elements in the hierarchy and select it.
11. Generate a new visualization using only the web pages on which the terms “bass” and “heavy metal” co-occur.

The tasks 1–8 are general ones that are likely to arise when analyzing and browsing collections of web pages. In particular, tasks 1–5 address the co-occurring terms, whereas tasks 6–8 deal with the multimedia content extracted from the web pages. In contrast, the tasks 9–11 relate to the structure of the Sunburst tree.

After having explained the interaction functionalities provided by 3D-COB to our participants, they had five minutes to explore the user interface themselves with a visualization gained for “Britney Spears”. During this warm-up, the participants were allowed to ask questions. After the exploration phase, we presented them the visualization obtained when using the web page collection of “Iron Maiden”, cf. Figure 1. We consecutively asked them each of the questions and measured the time they needed to finish each task. Each participant had a maximum time of three minutes to complete each task. The constraints were set as follows: $N = 8$, $R = 8$, and $E = 3.0$ (cf. Subsection 3.2).

Due to time limitations, we had to restrict the number of participants in the user study to six (five males, one female). All of them were computer science or business students at [omitted due to blind review] and all stated to have a moderate or good knowledge of user interfaces and to be very interested in music. All participants performed the user study individually, one after another.

5.2 Results and Discussion

As for the results of the study, Table 2 shows the time, in seconds, needed by the participants (A–F) to finish each task. In general, the tasks related to structural questions were answered in a shorter time than those related to browsing the collection. Among the structural questions, solely task 11 required a quite high average time. This is explained by the fact that the term “bass” was not easy to find on all layers. The same holds for the term “british” requested in task 3.

For the questions related to browsing in the hierarchy, it was observed that tasks requiring intensive rotation of the Sunburst stack (1, 3, 4, 5, 7) yielded worse results than those for which this was not the case (2, 6). In general, users spent a lot of time rotating the Sunburst stack to a position at which the label of the selected arc was readable. This process should be automatized in future versions of 3D-COB.

The relatively high average time required to perform the first task may be attributed to the fact that most participants needed some seconds to get used to the new visualization of “Iron Maiden” after having explored the web pages

of “Britney Spears” in the exploration phase. In spite of the fact that task 3 was solved in only 37 seconds on average, we realized that some participants had problems locating the arc “british” since it was hardly perceivable due to its position behind a much higher arc. As both task 4 and 2 required finding the same arc, it was quite interesting that the averaged times differed considerably. As for task 5, two participants were not sure which number to subtract from which other. Except for one participant, who chose a correct but time-consuming solution, task 6 was generally solved quickly. Solving task 7 took the second highest average time since it required finding and navigating to the Sunburst that illustrates the amount of image files and comparing the heights of all arcs in hierarchy level three of this Sunburst. Task 8 yielded the worst results as no arc on the video layer had a height of zero, which confused most of the participants. It was obviously not clear that a positive height of an arc on the video layer does not necessarily mean that each web page represented by this arc offers video content.

To conclude, the user study assessing ergonomic aspects showed that 3D-COB can be efficiently used for tasks related to browsing sets of web pages. Although barely comparable to the user study on similar tree visualization systems conducted in [14], due to a different application scenario, a very rough comparison of the average total performance times for the tasks shows that this time is much shorter for 3D-COB (45 sec) than for the best performing system of [14] (101 sec). Therefore, our results are promising. However, it can be stated that the user interaction functionalities provided by 3D-COB need some improvements. This will be addressed in future versions of the implementation.

Task	1	2	3	4	5	6	7	8	9	10	11
A	28	13	45	47	36	61	172	180	2	12	25
B	69	23	46	52	14	15	68	76	6	12	62
C	15	3	39	27	22	3	34	68	1	9	31
D	132	1	57	30	117	14	43	180	5	12	40
E	110	9	16	8	163	7	12	148	2	38	74
F	36	14	21	46	44	12	79	180	3	5	61
Mean	65	11	37	35	47	19	68	97	3	15	54

Table 2. For each participant, the time (in seconds) needed to finish each task of the user study on ergonomic aspects. Inverse numbers indicate that the given answer was wrong. The mean was calculated excluding the wrong answers.

6 Conclusions and Future Work

In this paper, we presented the *Three-Dimensional Co-Occurrence Browser (3D-COB)*, a user interface for browsing collections of music artist-related web pages in a novel way. 3D-COB automatically extracts musically relevant

terms from web pages about artists, applies a term weighting function, organizes the web pages according to co-occurring terms, and finally employs a variant of the Sunburst visualization technique to illustrate not only the extracted terms, but also the amount of multimedia files, grouped in different categories (audio, image, video).

Moreover, we conducted two user studies: one to evaluate the performance of different term weighting strategies for finding descriptive artist terms, the second to assess ergonomic aspects of 3D-COB's user interface. From the first study, we learned that using TF·IDF yielded significantly worse results than the simple TF and DF measures with respect to the appropriateness of the terms to describe the music artists used in our experiments. In contrast, comparing the measures TF and DF, no significant difference in their performance was detected. The second user study showed that 3D-COB offers valuable additional information about web pages that cannot be discovered by the standard list-based representation of search results, which is commonly used by web search engines.

As for future work, we will elaborate alternative ways to navigate in the visualization, e.g., using alternative input devices. We are also developing a focused web crawler in combination of which 3D-COB may be used to browse a set of web pages related to a certain topic (not necessarily related to music) without relying on existing search engines. The 3D-COB could also be employed to domains like movies, literature, painting, or news. Moreover, we will improve the user feedback provided by 3D-COB, e.g., by showing all terms that co-occur in the selected set of web pages, independently from the labels of the arcs. Finally, smoothly embedding the multimedia content directly in the user interface instead of opening it in external applications would be a desirable feature for future versions of 3D-COB.

References

1. <http://dir.yahoo.com/Entertainment/Music/Genres> (access: 03-2008).
2. http://en.wikipedia.org/wiki/List_of_file_formats (access: 12-2007).
3. <http://finetunes.musiclens.de> (access: 03-2008).
4. <http://last.fm> (access: 12-2007).
5. <http://www.allmusic.com> (access: 11-2007).
6. <http://www.processing.org> (access: 03-2008).
7. <http://www.wikipedia.org> (access: 12-2007).
8. K. Andrews and H. Heidegger. Information Slices: Visualising and Exploring Large Hierarchies using Cascading, Semi-Circular Discs. In *Proc. of IEEE Information Visualization 1998*, Research Triangle Park, NC, USA, 1998.
9. G. Geleijnse and J. Korst. Web-based Artist Categorization. In *Proc. of the 7th Int'l Conf. on Music Information Retrieval*, Victoria, Canada, 2006.
10. M. Goto and T. Goto. Musicream: New Music Playback Interface for Streaming, Sticking, Sorting, and Recalling Musical Pieces. In *Proc. of the 6th Int'l Conf. on Music Information Retrieval*, London, UK, 2005.
11. B. Johnson and B. Shneiderman. Tree-Maps: A Space-Filling Approach to the Visualization of Hierarchical Information Structures. In *Proc. of the 2nd IEEE Conf. on Visualization*, San Diego, CA, USA, 1991.

12. P. Knees, E. Pampalk, and G. Widmer. Artist Classification with Web-based Data. In *Proc. of the 5th Int'l Symposium on Music Information Retrieval*, Barcelona, Spain, 2004.
13. P. Knees, M. Schedl, T. Pohle, and G. Widmer. An Innovative Three-Dimensional User Interface for Exploring Music Collections Enriched with Meta-Information from the Web. In *Proc. of the 14th ACM Conf. on Multimedia 2006*, Santa Barbara, CA, USA, 2006.
14. A. Kobsa. User Experiments with Tree Visualization Systems. In *Proc. of the 10th IEEE Symposium on Information Visualization 2004*, Austin, Texas, USA, 2004.
15. F. Mörchen, A. Ultsch, M. Nöcker, and C. Stamm. Databionic Visualization of Music Collections According to Perceptual Distance. In *Proc. of the 6th Int'l Conf. on Music Information Retrieval*, London, UK, 2005.
16. E. Pampalk, A. Flexer, and G. Widmer. Hierarchical Organization and Description of Music Collections at the Artist Level. In *Proc. of the 9th European Conf. on Research and Advanced Technology for Digital Libraries*, Vienna, Austria, 2005.
17. E. Pampalk and M. Goto. MusicSun: A New Approach to Artist Recommendation. In *Proc. of the 8th Int'l Conf. on Music Information Retrieval*, Vienna, Austria, 2007.
18. M. Schedl, P. Knees, K. Seyerlehner, and T. Pohle. The CoMIRVA Toolkit for Visualizing Music-Related Data. In *Proc. of the 9th Eurographics/IEEE VGTC Symposium on Visualization*, Norrköping, Sweden, 2007.
19. M. Schedl, T. Pohle, P. Knees, and G. Widmer. Assigning and Visualizing Music Genres by Web-based Co-Occurrence Analysis. In *Proc. of the 7th Int'l Conf. on Music Information Retrieval*, Victoria, Canada, 2006.
20. D. J. Sheskin. *Handbook of Parametric and Nonparametric Statistical Procedures*. Chapman & Hall/CRC, Boca Raton, 3rd edition, 2004.
21. J. Stasko and E. Zhang. Focus+Context Display and Navigation Techniques for Enhancing Radial, Space-Filling Hierarchy Visualizations. In *Proc. of IEEE Information Visualization 2000*, Salt Lake City, UT, USA, 2000.
22. F. Vignoli, R. van Gulik, and H. van de Wetering. Mapping Music in the Palm of Your Hand, Explore and Discover Your Collection. In *Proc. of the 5th Int'l Symposium on Music Information Retrieval*, Barcelona, Spain, 2004.
23. B. Whitman and S. Lawrence. Inferring Descriptions and Similarity for Music from Community Metadata. In *Proc. of the 2002 Int'l Computer Music Conference*, Göteborg, Sweden, 2002.
24. M. Zadel and I. Fujinaga. Web Services for Music Information Retrieval. In *Proc. of the 5th Int'l Symposium on Music Information Retrieval*, Barcelona, Spain, 2004.