

How Last.fm Illustrates the Musical World: User Behavior and Relevant User-Generated Content

Ya-Xi Chen
Media Informatics
University of Munich
yaxi.chen@ifi.lmu.de

Sebastian Boring
Media Informatics
University of Munich
sebastian.boring@ifi.lmu.de

Andreas Butz
Media Informatics
University of Munich
andreas.butz@ifi.lmu.de

ABSTRACT

Over the last few years, online multimedia exchange platforms have experienced a rapid growth. They allow users to share their own content and access other's in turn and hence form very large public collections of User-Generated Content. While research is mostly looking at photo sharing platforms, such as Flickr, much less is known about online music communities. In this paper we present the results of an observational user study followed by a large-scale online survey, which investigated the behavior and the relevant content generated by the users of Last.fm, one of the most popular music communities. Based on the analysis of the results, we present implications for the usage of User-Generated Content in online music communities. Then we developed a first prototype based on the implications for improving semantic understanding of collaborative tags. We believe our study gives insights for developing information visualization and recommender systems for online music communities.

Author Keywords

Online music community, User-Generated Content, user behavior, Last.fm.

ACM Classification Keywords

H5.2. Information interfaces and presentation (e.g., HCI): User Interfaces.

INTRODUCTION

Most of the current research on public multimedia exchange platforms is focusing on the behavior around photos in online communities, such as searching, tagging and sharing. Much less is known about how people define their musical taste and how User-Generated Content (UGC) helps online music communities to make more sense of music. We be-

lieve, that an investigation of online music communities could lead to a better understanding of people's behavior surrounding music in general and bring valuable insights on how to successfully harness the metadata contributed by the users of these music communities.

There are several online music communities. Similar to artist map proposed by Gulik and Vignolo in [7], Musicoverry¹ is an interactive radio station, for which the user can define the current mood, time range, desired tempo and genre. Live365² is a radio network, in which the user can generate a personalized radio station. The recommendations are organized and characterized by genre. Similar radio functionalities are also provided in Jamendo³. Imeem⁴ is a social media community offering a variety of media types, such as music, video, photos and blogs.

Last.fm⁵ is one of the largest and most popular online music communities with a large user group and abundant services. According to Wikipedia⁶, Last.fm has over 30 million active users spreading over 200 countries. As Last.fm claims, they focus on playing the right songs to the right people. Its functionality can be extended based on a released API and a series of applications have already been proposed. However, there is little research focusing on the user behavior and relevant UGC in those music communities. To obtain implications for better use of UGC, such as providing personalized recommendations and facilitating discovery of new music, we chose Last.fm as our experimental platform and conducted a user study based on it.

RELATED WORK

There are studies about users' behavior with music, for example, searching, sharing and tagging. Some research also focuses on music recommendations. All these studies reveal

¹ Musicoverry, <http://www.musicoverry.com/>

² Live365, <http://www.live365.com>

³ Jamendo, <http://www.jamendo.com>

⁴ Imeem, <http://www.imeem.com/>

⁵ Last.fm, <http://www.last.fm>

⁶ Wikipedia, <http://en.wikipedia.org/wiki/Last.fm>

the nature of our experience with music and help to understand the users' desires regarding music-related technologies.

Searching

People often do not explicitly search in media collections. They are rather looking for something that satisfies certain (possibly vague) criteria, instead of one specific item. Others follow a different strategy by first picking up some candidates and then making a final decision among these pre-selections. Vignoli [24] claimed that non-expert users have strong difficulties to express their musical preferences in a formal way, and that they often change their minds during the search process.

Kim et al. [14] investigate people's perception of music and observe that both in the description and in searching, users tend to combine music with events and emotions. Similar implications were derived in [5] based on the analysis of respective requests posted to a music-related newsgroup.

Collaborative tagging

With the rapid growth of the next-generation Web, many websites allow the users to make contributions by tagging digital items. This collaborative tagging has become a fashion on many websites. The user-contributed tags are not only an effective way to facilitate personal organization, but also provide a possibility for the users to search for information or discover new things.

A TagCloud (see figure 3) is a visual presentation of the most popular tags, in which tags are usually displayed in alphabetical order and text attributes, such as font size, weight or color are used to represent features (e.g., font size for prevalence and color brightness for recentness). As a result of collaborative tagging, TagClouds have a more accurate meaning than those assigned by a single person, and reflect the general interests among a broad demography [9, 23]. Due to their easy understandability and aesthetical presentation, TagClouds have become a fashion on many websites. However, they still have some intrinsic disadvantages and many researchers have been dedicated to improve their aesthetical presentation [1, 13, 20] or semantic understanding [8, 15].

Sharing

One important activity around music is sharing, which facilitates social communication and information exchange, but also helps to maintain personal images in front of others. One of the few detailed investigations [2] compared music sharing behavior with offline and online sharing systems such as Napster, and then explored in detail a system named Music Buddy for browsing other people's music collections. The study showed that music sharing is tightly bonded with social activities, and it suggested that music should be shared in a more collaborative and community-related environment. Volda et al. [25] explored practices surrounding the iTunes music sharing functionality and made several improvement suggestions.

Transparency of recommender systems

Many online music communities, such as Pandora.com, iTunes Genius and Amazon, offer music recommendations, and the mechanisms behind them vary from content analysis to the users' listening or purchasing patterns.

Transparency is a crucial issue in recommender systems. Herlocker et al. [10] suggested that the explanations of recommendations can make the system more understandable and involve the user more in it, and thus improve the user's satisfaction. In contrast to previous research focusing on statistical accuracy of the algorithm, Swearingen and Sinha [22] emphasized interface issues from the user perspective. They claimed that users like and feel more confident about recommendations with transparency, especially for new items. SIMAC [11] is one of the few existing systems, which addressed the issue of transparency. In SIMAC six semantic descriptors were designed in order to solve the semantic gap. The weights of all descriptors were visualized in a radial graph in which the radial distance presents the value of weight. The user can change the weight by moving the descriptor manually.

USER-GENERATED CONTENT IN LAST.FM

In Last.fm, each user has a personal profile integrated with library and playlists, charts of listened music, social networks such as friends and groups. Users can listen to music online, receive recommendations from the system and from other users, and they are also allowed to tag all music items. Based on the music-surrounding behavior, there is abundant data generated by users, such as personal listening history, tags and social network, which work as the fundament of the Last.fm services for personal charts, system recommendations and tag-based search.

Listening history

The listening history is automatically recorded when the user listens to Last.fm music. It serves as the statistical basis of Last.fm's main functionalities of charts and system recommendations.

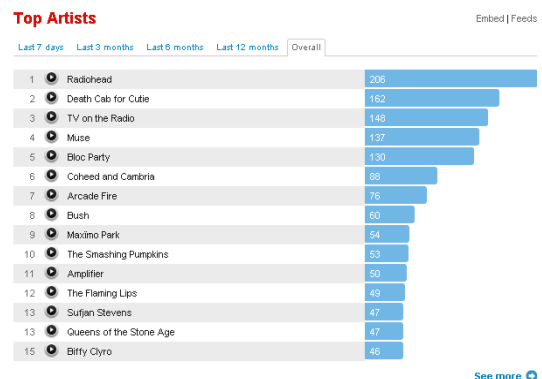


Figure 1. Personal chart for top artists.

Charts are statistical presentations of the listening history. Personal charts are displayed as a list of recently played music, ordered by play count. Figure 1 is an example chart

for top artists. Similarly, there are public charts calculated based on all users' listening histories.

Based on the aggregation of all users' listening histories, the system provides recommendations of similar artists for each artist and neighbors who share a similar musical taste with the user. If the user further browses each neighbor's profile, the similarity of musical taste between these two users is represented as a bar slider called musical compatibility (see Figure 2).




Figure 2. System recommendation of neighbors.

Tags

Last.fm allows users to tag each track, album and artist with free form texts, which can then be used for tag-based visualizations and search.

Last.fm offers TagCloud visualization of the top tags generated by users. As shown in figure 3, most of the popular tags are genre-related.




Figure 3. TagCloud for top tags in Last.fm.

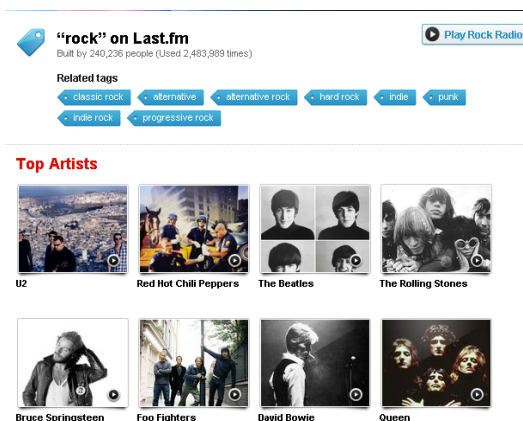


Figure 4. Retrieval results of the tag "rock".

Based on these user generated tags, the user can conduct tag-based searching and Last.fm will return a page for the respective tag, in which related tags and the top artists for this tag will be displayed. Figure 4 is the retrieval results of the tag "rock".

Social network

The user can add other users as friends, and join groups, in which people with common interests gather. Similar to the personal profile, Last.fm generates a profile for each group. A group radio is created based on the overall listening history of the whole group.

Besides system recommendations, the user can also recommend music to other users by sending internal textual message, which is called "sharing" in Last.fm.

INTERVIEW

As already discussed, UGC forms the fundamental basis for Last.fm. In order to gain more insights on the effective use of metadata contributed by the users, the following essential issues need to be explored: the performance of system recommendations based on the users' listening histories, other useful information which can be extracted from the listening history, the features and benefits of music-related tags, and the user's social network activities.

In order to answer these questions, we first conducted interviews with Last.fm users.

Participants

We recruited 13 participants in the Last.fm online forum, 3 female and 10 male. Their age ranged from 18 to 26 with an average age of 23 years. Most of the participants were students and all of them have common knowledge about computers and the Internet. Participants are all music amateurs and rated themselves to be experienced Last.fm users with an average score of 4.2 (5 for very experienced).

Settings and procedure

During the interview, the participants were equipped with a PC, keyboard and mouse. They could freely browse the Last.fm website and relevant applications, such as the desktop radio. One visualization tool for listening histories was installed beforehand.

First, the participants were asked to fill out a pre-questionnaire about their personal information and general experience with music. Then they joined an interview about their personal experience with Last.fm, which mainly covered the issues of system recommendation, personal profile, tagging and searching behavior, and social network. Participants could freely browse their personal profiles and other services of Last.fm. On average the user study lasted about 1 hour per participant. It was conducted in English and recorded on video. The Think-Aloud protocol was applied.

The questions were grouped into four categories. To learn about the participant's general experience with Last.fm, we

asked about the services that were considered as most useful, the main source for discovering new music and the quality of the system recommendations. Example questions are: “How often do you visit the Last.fm website?”, “Do you also use other desktop or portable applications?”, “Which functionalities do you think are most useful?”, “How do you discover new music?” and “What do you think of the system recommendation of artists and neighbors?”.

In the next step, participants answered questions related to their personal profiles, which helped to understand their musical tastes. Example questions were: “How would you describe your musical taste?”, “Do you think it is hard to express musical taste verbally?”, “How well does your Last.fm library present your musical taste?” and “Do you mind your personal profile being public in Last.fm?”.

Another explored key issue was the tagging and searching behavior and relevant user-generated tags. Example questions for searching were: “How often do you search for music in Last.fm?”, “How often do you use tags for searching?” and “What do you think about TagClouds of Last.fm?”. About the tagging behavior, some example questions were: “How often do you tag music in Last.fm?”, “Which kind of tags do you use for tagging?” and “Do you think tagging music is difficult?”.

Since Last.fm offers functionalities for social networking, such as friends and groups, we also discussed those with the participants. Some example questions were: “How many of your Last.fm friends are also friends in your daily life?”, “How do you find new friends and groups?”, “How often do you receive music recommendations from other users?” and “How often do you recommend music to other users?”

Results

Based on the analysis of the questionnaire and the recorded video, the following results were discovered:

Personal music experience

All participants own portable music devices with normally more than 500 songs. When asked about the general sources for discovering new music, all of them chose Last.fm as the main online source, other sources being music services such as napster, amazon, iTunes and youTube. 9 out of 13 receive recommendations from friends and only 4 mentioned conventional means, such as CD stores, TV programs or newspapers.

Regarding devices for listening to music, the PC seems to be the dominant device. Most of the participants listen through the PC much longer (4.9 hours/day) than through portable devices (1.8 hours/day), such as an MP3 player or mobile phone. Regarding the listening situations, the four equally mentioned main situations are background music for working, during the commute, social events such as parties, and pure enjoyment.

General experience with Last.fm

Besides frequently visiting the website, participants also use other Last.fm applications. 8 of them are regular user of AudioScrobbler, a plugin for desktop music players, which automatically transfers statistics of the user’s listening history to the personal charts in Last.fm. The two participants who own an iPhone or iPod Touch also use the Last.fm mobile applications. Regarding useful functionalities in Last.fm, the top three are AudioScrobbler, personal charts and the system recommendation for similar artists and neighbors.

Since the system recommendations and the discovery of new music are remarkably important for the participants, we discussed these two issues in more detail. All of the participants mainly discover new music from the system recommendation of similar artists. The other means are recommendations by social contacts, such as friends or groups, and by browsing neighbors’ profiles. Only one participant uses the searching functionality to find music of a certain genre. Generally all the participants appreciated the system recommendations and scored higher for recommendation of similar artists ($M=4.33$, $SD=0.65$) than neighbors ($M=3.66$, $SD=0.49$). There were two main reasons for the lower score of neighbor recommendation: besides a list of neighbors with the relevant shared artists, the participants would have liked an additional detailed description of the neighbors’ musical preferences; the current recommendation is based on the latest weekly listening history. The user might get different neighbors if the weekly interests change. Although this reflects the continuously changing nature of musical taste, some participants still expressed the wish to get neighbors with overall similar taste.

User 4: the biggest part of my music is funk, others are electronic and classical. However, I only get funk neighbors.

User 13: My girlfriend and I intentionally listen to similar music but our weekly musical compatibility is unstable, maybe because of the different listening sequences.

Personal profile in Last.fm

When asked to describe the personal musical taste with free text, all participants came up with short descriptions and most of them were genre-related. Most of the participants have a relatively stable preference. When asked how hard it was to express musical taste verbally, 8 out of 13 scored higher than 3 (5 for very difficult).

Although the participants did not concern about the profile being public, some of them still applied different strategies to maintain their personal images. For example, one participant has two players, one for free personal usage with his whole collection, the other one with representative music with plugged scrobbler which automatically transfers the listening history of these songs to his Last.fm personal charts.

Since the personal listening history is essential for both the user and the system, some applications are developed for the visualization of personal listening histories. Most of them use a flow metaphor to represent how the personal musical taste changes over time. Extra Stats⁷ is an application, which visualizes the top artists as colored waves on a timeline (see figure 5). Each wave presents one artist and the width represents the play count of this artist in each time period. Other similar visualizations can be found in LastGraph⁸ and Last.fm Spiral⁹. During the interview, the participants were asked to observe the visualization results of their own listening history and one of another participant's. A consistent pattern appeared in all the visualization results: there were always bursts when the user found new artists and listened to them very often in a short time period. After a while, these discoveries fell into the normal flows.

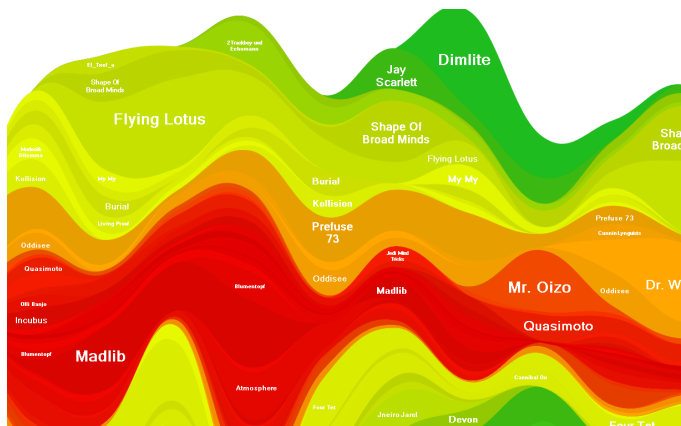


Figure 5. Extra Stats: flow visualization of the personal listening history.

All participants thought the visualization was useful and they also learnt additional information from the visualization. For example, they noticed the break period during their usage of Last.fm, and also received new insights with their own listening behavior and other's musical taste:

- Recall of relevant social activities:

User 1: (point at one peak) I just returned from vacation and I met a girl there. I listened a lot to the music she liked.

- Re-discovery of forgotten music:

User 3: there was a band I once liked very much but they never came again. Maybe I should listen to them again.

- Understanding of personal listening behavior:

User 8: Drops down in august, maybe I was not so often at home in summer.

⁷ Extra Stats, <http://build.last.fm/item/34>

⁸ LastGraph, lastgraph3.aeracode.org

⁹ Last.fm Spiral, <http://build.last.fm/item/377>

- Understanding of other's musical taste:

User 5: He likes rock and pop music. I don't think he sticks to any specific artists.

In these comments, we can see that Last.fm helps to discover new music and that the listening history contains rich information. It also works as a self-reflection and helps to understand other's musical taste.

Searching and Tagging

Most of the participants use the search functionality frequently, with the exception of one, who finds music by browsing the charts for popular artists. Besides the standard keywords such as the name of artist, album and song, tags are less used for searching and the scores for the usage frequency were rather low ($M=2.18$, $SD=1.08$, on a 5-point Likert-scale where 1 stands for "never"). The top three types of tags used for searching are genre, mood and artist biography. The aspects of tags are diverse, but currently in Last.fm the user cannot combine multiple tags for specific searching.

User 3: It is a pity that I cannot use more than 1 tag as keywords, for example, to find a tiny part between punk and indie electronic.

All the participants felt that the too general tags might make the user getting lost among abundant results and thus find nothing specific.

User 5: Tags are too subjective and heavily depend on the personal musical taste. For example, for your favorite song, others might think it is awful. It is not suitable to describe the essence of music.

User 12: "seen live" doesn't help me at all. It's like asking for the way to the Eiffel Tower and someone tells you "in Europe".

When asked to give comments of the top tags shown in Figure 3, one prominent comment was the redundancy, for example "favorite" and "favourite". Since music is difficult to express verbally, and there is no standard category for genre, people have different definitions of genres and even have different understanding of the same genre, which leads to remarkable redundancy and even errors with genre-related tags.

User 4: I noticed that some people think IDM (Intelligent Dance Music) and electronic are the same so they always appear in a pair. But actually they are different.

The participants do not tag so often and the average tagging frequency is 1.09 ($SD=0.83$). Similar to the description of personal musical taste and tags used for searching, most of their generated tags were also related to genre, mood and artist biography. Some other participants also use personalized tags for quick relocating, such as "listen again" and "Sunday morning". The majority of participants thought that tagging music is hard.

User 1: Talking about music is just like dancing with a poem. It is hard to describe music with words.

Social network in Last.fm

Besides music, Last.fm also offers functionalities for social networking, such as friends and groups. Most of the participants use Last.fm only for music, since they already have other social networks. Adding users as friends either actively or passively is determined by the social contacts with them. For the users who have no daily contacts, most of them will be added on their requests. The participants' friend lists showed that most of them are real friends.

Compared with friends, group-related activity is less popular. Generally the themes of the groups are related to a location (affiliation, city, country) or genre. Which group to join and how to find a suitable group is determined by the personal music experience or influenced by friends, geographic and cultural factors.

User 5: Groups are very useful because my musical taste is special and in daily life I don't know too many people sharing the same taste.

Although last.fm offers functionality for recommending music by sending a message, it is seldom used and participants rarely recommend music explicitly. Only 2 participants once received recommendations from others and only 2 occasionally send recommendations.

ONLINE SURVEY

In order to verify the results of the interview, we conducted an online survey in English which lasted for two months. The questions asked in the survey were consistent with the interview, mainly covered the demographic information, general experience with Last.fm, system recommendations, searching and tagging behavior, and social network.

In total we received 228 complete questionnaires, 93 female and 133 male (two gender identifiers were left blank). Their age ranged from 16 to 36 with an average age of 22 years. Most of the participants were students and employees from North America and Europe. Participants rated themselves to be experienced Last.fm users with an average score of 3.8 (5 for very).

Results

In general, the results of the online survey are consistent with those derived during the interview.

Personal music experience

About the general sources for discovering new music, the online source was very popular ($M=4.47$, $SD=0.93$, on a 5-point Likert-scale where 1 stands for "daily") and the most often mentioned websites were Last.fm, iTunes and YouTube. The other two main sources were recommendations from others ($M=3.69$, $SD=1.08$), and traditional sources ($M=2.65$, $SD=1.18$).

The most often used devices for playing music were PC ($M=4.75$, $SD=0.55$), portable digital player ($M=3.96$, $SD=1.33$) and mobile phone ($M=2.23$, $SD=1.44$). The main listening situations were consistent with the answers in the interviews.

General experience with Last.fm

Besides Last.fm website, other frequently used applications were AudioScrobbler ($M=4.18$, $SD=1.45$), desktop radio station ($M=1.99$, $SD=1.25$) and MobileScrobbler ($M=1.65$, $SD=1.31$).

The main means of discovering new music were system recommendations ($M=3.69$, $SD=1.30$), browsing friends' profiles ($M=3.67$, $SD=1.28$), recommendations from friends ($M=3.20$, $SD=1.46$), browsing neighbors' profile ($M=2.96$, $SD=1.49$) and recommendations from group ($M=2.53$, $SD=1.43$). The system recommendations were appreciated and received higher for recommendation of similar artists ($M=4.11$, $SD=1.07$) than neighbors ($M=3.28$, $SD=1.18$).

Personal profile in Last.fm

Participants believed that their libraries well represented their tastes ($M=4.25$, $SD=0.75$). For the description of personal taste, 173 out of 228 participants proposed genre-related texts. The general attitude toward public nature of the personal profile was rather neutral ($M=2.95$, $SD=1.32$).

Concerning the listening behavior, they always play music from own library ($M=3.6$, $SD=1.30$) and a repetitive listening pattern was revealed: They tend to repeatedly listen to certain artists, albums and songs.

The visualization of personal listening history in Extra Stats was commented as useful in supporting understanding taste changes over time, artist re-discovery and reflection of listening patterns.

Searching and Tagging

Participants look for music in Last.fm very frequently ($M=3.99$, $SD=1.15$, on a 5-point Likert-scale where 1 stands for "daily"), but they more likely browse with no clear goal rather than specific search. Different from participants in the interview, keyword based search was less conducted ($M=1.80$, $SD=1.10$) and participants mostly search music-related information such as artist, album and song ($M=4.04$, $SD=1.31$), and less about social aspects such as group, user or event.

The Last.fm TagClouds was commented as useful to gain an overall impression of the most popular items but similar linguistic problems were also noticed. The majority of participants seldom tag. They mainly tag music in their own libraries and most of their generated tags were genre-related. Different from participants in the interview, they consider tagging as rather easy ($M=2.22$, $SD=0.09$, 5 for very difficult). The top motivations for tagging were facilitating browsing and searching, facilitating personal organization, and helping others to understand music.

Social network in Last.fm

Last.fm was considered more of a music website ($M=4.59$, $SD=0.69$, 5 for highly agree) than a social network ($M=3.44$, $SD=1.14$) and the most popular social networks among the participants were facebook, myspace and twitter.

The number of friends varied from 0 to 322 with average number of 32 ($SD=41.40$). Different from participants in the interview, the Last.fm friends also known in daily life were much less ($M=6$, $SD=9.05$). Most of the friends were added on their requests. The number of group also varied a lot from 0 to 60 ($M=28$, $SD=66.50$). Compared with friends, the group-relevant activities were less popular. And the popular group themes were genre, artist, geo-location and events. The functionality of recommending music to others was less used.

IMPLICATIONS

Based on the results of the interview and online survey, some implications about the user's behavior surrounding online music and relevant UGC were revealed:

General experience with music

The PC dominates as the main music device and portable devices show a noticeable potential when people are "on the way" and thus relevant applications should receive more attention. A smart music recommendation system should recognize the context, choose and switch songs smoothly, for example as Cunningham et al. mentioned in [4], shuffle by genre, which might be more appealing than existing random shuffle mode.

System recommendation

Current system recommendations of similar artists is generally appealing and it could be further improved, for example, by taking the recency factor into account.

Last.fm recommendations of neighbors are based on the latest weekly charts. When the user has an unstable musical taste, especially when discovering new bursts and sticking to them for a while, the neighbors keep changing. Although the system offers a list of neighbors with a high musical compatibility score, more detailed explanation is expected, which also helps to build self-reflection and to understand others' musical taste. When the user wants a neighbor recommendation based on his or her overall musical taste, the system should offer a more flexible and smart recommendation scheme, in which the user's requirements could be dynamically integrated. The system could, for example, let the user choose a time period or select some of the neighbors as examples, which help to discover new matching neighbors.

Listening history

Personal listening history is the key issue of Last.fm which helps to formulate the charts and system recommendations. As the title of [4], music is more of an art than science, which illustrates that musical taste is hard to express efficiently by purely statistical methods. Compared with statistical charts, the graphical visualization for the listening his-

tory offers better understanding about how the musical taste changed over time. Users can get abundant information from the visualization which helps to discover personal listening behavior, re-discover forgotten music and understand others' musical tastes. Since some users might have a long history, the visualization should offer a better overview while helping to construct a complete mental model conveniently. Although existing visualization tools receive positive feedback, more interactions should be introduced to enhance the understandability. Most of the current tools only target single users and it might be appealing to offer users an intuitive way to browse and compare multiple users' listening histories, which in turn could improve the system transparency.

Tags and relevant tagging behavior

People do not tag music so often and they tag for different reasons. Some people take music very seriously and want others to know more about their favorite music through tags. Some users annotate music with special tags for personal use. Others simply make a contribution or offer knowledge by tagging.

In Last.fm, most of the top tags are related to genre, mood or artist biography. There is less chance for users to be 'educated' since the personal understanding of genre and emotion is subjective and according to different musical experiences, the users might come up with different tags for the same music. Therefore, searching by tags is not common in Last.fm because freely generated tags are normally too general to help users narrowing down the results. More neat and organized tags with less redundancy would be more useful and the option of combining multiple tags in the searching process might help the user to harness the searching direction.

Social network

Most of the participants use Last.fm only for music and the social-related activities are mainly passive, such as receiving recommendations from others, adding friends or joining groups. Active music recommendation is not popular in last.fm, even though the system offers a sharing functionality. Although the personal profile being public is not a big issue, some users still want to maintain personal images, for example, by keeping the Last.fm library or charts in a representative and neat way.

EXPERIMENT BASED ON IMPLICATIONS

Based on the implications derived from our user study, applications for information visualization and recommender systems can be built: for example, illustrating the worldwide musical trends, improving semantic understanding of tags, and facilitating discovery of new music and people sharing similar tastes.

As the results of the user study showed, TagClouds contains redundancies and errors with freely generated tags and can not support semantic understanding of the relationships

among tags. Therefore, we developed an aggregation of TagClouds named TagClusters (see Figure 6).

The hierarchical structure and positions of tags are achieved based on a semantic analysis. Text analysis is first applied to produce a semantic clustering of similar tags: After removal of separators such as “_” and “&”, the Porter algorithm [19] is applied to detect the stem of each tag. Tags with the same stem words are clustered in the same group. For example, metal related tags such as “heavy metal”, “gothic metal” and “melodic death metal” are grouped into one metal cluster. After semantic grouping of similar tags into genre-clusters, the hierarchical structure in each cluster is determined based on the tag length because of the characteristic feature of genre-related tags: the tag in lower semantic level always contains the tag in the higher level and the length of tag is proportional with its semantic level, for example, “death metal” and “brutal death metal”.

The location of each tag is determined by the semantic similarity (see Equation 1). It equals to the ratio between the number of resources in which a pair of tags A and B co-occur and the number of resources in which any of these two tags appears.

$$Sim(A, B) = |A \cap B| / |A \cup B| \quad (1)$$

After this semantic analysis, semantically similar tags are clustered into groups and their visual distance represents their semantic similarity, thus the visualization offers a better hierarchical understanding of collaborative tags.




Figure 6. TagCluster: Aggregation of TagClouds [3].

A comparative evaluation was conducted with TagClouds and TagClusters based on the same Last.fm tag collection. 12 participants were recruited and were required to conduct 6 tasks (each task is consisted of two similar sub-tasks): locating one single item, sorting tags by popularity, grouping similar tags, driving group structure, finding relation between tags and judging their similarity. The complete time and the answer precision were measured. After completed each task, the participants were asked to score the easiness of each task and the usefulness of both systems. After completing all the tasks, the participants filled out a post-questionnaire which concerns the overall impression of

both systems. The analysis of both quantitative and qualitative data indicated that TagClusters performed overall better and have advantages in supporting semantic understanding, impression formation and matching. In our future work, we will explore using TagClusters to support tag recommendation and multiple-tags-based searching.

CONCLUSION AND FUTURE WORK

In this paper we conducted a preliminary user study with Last.fm, an online music community. We investigated key issues about User-Generated Content, such as listening history, tags and social network, based on which Last.fm offers services of charts, system recommendations of similar artists and neighbors. Based on an analysis of relevant user behavior and relevant generated data, implications for usage of UGC were derived. We developed our first prototype for improving semantic understanding of tags. We believe our user study could bring insights for better usage of UGC and help users to get better understanding of the Last.fm musical world. In our future work, we plan to develop prototypes based on the derived implications, mainly in the realm of information visualization and recommender systems. Based on the accumulated experience with the prototype development we expect to obtain general design guidelines with UGC in online music communities.

ACKNOWLEDGMENTS

This research was funded by the China Scholarship Council (CSC) and by the German state of Bavaria. We would like to thank the participants of our study.

REFERENCES

1. Ahlberg C., Shneiderman B. Visual information seeking: tight coupling of dynamic query filters with starfield displays. In *Proc. CHI 1994*, ACM Press (1994), 313-317.
2. Brown B., Geelhoed E., Sellen A. Music sharing as a computer supported collaborative application. In *Proc. ECSCW 2001*, ACM Press (2001), 179-198.
3. Chen Y.-X., Santamaria R., Butz A., Theron R. TagClusters: Semantic Aggregation of Collaborative Tags beyond TagClouds. In *Proc. SG 2009*, Springer Press (2009), 56-67.
4. Cunningham S., Bainbridge D., Falconer A. ‘More of an art than a science’: supporting the creation of playlists and mixes. In *Proc. ISMIR 2006*.
5. Downie J. S., Cunningham S. Toward a theory of music information retrieval system design implications. In *Proc. ISMIR 2002*.
6. Graham, A., Garcia-Molina, H., Paepcke, A., Winograd, T.: Time as essence for photo browsing through personal digital libraries. In *Proc. JCDL 2002*, ACM Press (2002), 326-335.
7. Gulik, R van, Vignoli, F. Visual Playlist Generation on the Artist Map. In *Proc. ISMIR 2005*.

8. Hassan M. Y., Herrero S. V. Improving tag-clouds as visual Music Retrieval interfaces. In *Proc. INSCIT 2006*.
9. Hearst M. A. What's up with Tag Clouds? http://perceptualedge.com/articles/guests/whats_up_with_tag_clouds.pdf. Accessed December 30, 2008.
10. Herlocker, L., Konstan J., Riedl J. Explaining Collaborative Filtering Recommendations. In *Proc. CSCW 2000*, ACM Press (2000), 241-250.
11. Herrera P., Bello J., Widmer G., Sandler M., Celma O., Vignoli F., Pampalk E., Cano P., Pauws S., Serra X. SIMAC: semantic interaction with music audio contents. *Journal of Intelligent Information Systems*, 2005.
12. Huynh, D., Drucker, S., Baudisch, P., Wong, C.: Time quilt: Scaling up zomable photo browsers for large, unstructured photo collections. *Ext. Abstracts CHI 2005*, ACM Press (2005), 1937-1940.
13. Kaser O., Lemire D. Tag-cloud drawing: algorithms for cloud visualization. In *Proc. WWW 2007*, ACM Press (2007).
14. Kim J-Y., Belkin J. Categories of Music Description and Search Terms and Phrases Used by Non-Music Experts. In *Proc. ISMIR 2002*.
15. Li R., Bao S., Yu Y., Fei B., Su Z. Towards effective browsing of large scale social annotations. In *Proc. WWW 2007*, ACM Press (2007), 943-952.
16. Pampalk E., Goto M. MusicRainbow: a new user interface to discover artists using audio-based similarity and web-based labeling. In *Proc. ISMIR 2006*.
17. Pampalk E., Goto M. MusicSun: a new approach to artist recommendation. In *Proc. ISMIR 2007*.
18. Platt, J., Czerwinski, M., Field, B.: Phototoc: automatic clustering for browsing personal photographs. *Microsoft Research Technical Report MSR-TR-2002-17*, 2002.
19. Porter, M.F.: An algorithm for suffix stripping. *Program* 14(3), 130-137 (1980).
20. Seifert C., Kump B. Kienreich W. On the beauty and usability of tag clouds. In *Proc. IV 2008*, IEEE Press (2008), 17-25.
21. Shneiderman, B., Kang, H. Direct annotation: a drag-and-drop strategy for labeling photos. In *Proc. IV 2000*, IEEE Press (2000), 88-98.
22. Swearingen K., Sinha R. Beyond algorithms: An HCI perspective on recommender systems. In *ACM SIGIR 2001 Workshop on Recommender Systems*, ACM Press (2001), 1-11.
23. Viégas F. B., Wattenberg M. Tag Clouds and the Case for Vernacular Visualization. *Interactions*. 15(4), 49-52, 2008.
24. Vignoli F. Digital Music Interaction concepts: a user study. In *Proc. ISMIR 2004*.
25. Volda, A. Grinter, R. E. Ducheneaut, N. Edwards, W. K. Newman, M. W. Listening In: Practices Surrounding iTunes Music Sharing. In *Proc. CHI 2005*, ACM Press (2005), 191-200.