SmartMusic: An Online Music Recommendation System Based on Semantic Web Technology *

Wei Hu[1,2], Kaidi Yan[2], Cunxin Jia[2], and Junwei Wu[2]

[1] State Key Laboratory for Novel Software Technology, Nanjing University, China
[2] Department of Computer Science and Technology, Nanjing University, China
whu@nju.edu.cn, {yankaidi,jiacunxin,wujunwei}@smail.nju.edu.cn

Abstract. SmartMusic is a Chinese online music recommendation system for college students in Nanjing University based on Semantic Web technology. Music metadata are firstly collected with a Web crawler from diverse websites, and then processed and transformed into the RDF representation. By computing semantic distances between the descriptions of singers, songs and albums, content relevance is revealed. Additionally, by analyzing user behaviors, the system is able to speculate users' preferences on music. With these two features combined, SmartMusic offers better recommendation results than some existing music recommendation systems.

1 Introduction

With the rapid development of Semantic Web, we have witnessed many exciting applications that use Semantic Web technology (see *Semantic Web Challenges*), including some well-designed semantic recommendation systems. *Music recommendation*, being a branch of recommendation systems, also has some successful practices [1,5,8]. These systems depend on music ontologies and data in English. However, in the areas of data collection, processing and further recommendation for *Chinese* music, we have not seen such good applications.

In this paper, we introduce an online music recommendation system named *SmartMusic* (`http://ws.nju.edu.cn/smartmusic/`), which is based upon Semantic Web technology and has been deployed in Nanjing University for college students without paying any internet charge.

The remainder of this paper is organized as follows. In the following section, we begin with the collection and refinement of raw music metadata. In Section 3, we provide two ways of recommendation and the combination strategy, in order to offer personalized recommendation. Section 4 shows the system interface and reports experimental results on the efficiency and accuracy of the recommendation. Related studies and applications are discussed in Section 5. Finally, Section 6 summarizes this paper.

* This work is supported in part by the NSFC under Grant No. 61003018 and in part by the Nanjing University Undergraduate Innovation Program 2010.

2 Music Metadata Extraction and RDFization

2.1 Extraction Using a Web Crawler

At present, no free Chinese music datasets are offered. In order to collect necessary information on singers, songs and albums, we developed a *Web crawler* in Java. This crawler targets on specific, structured data sources for the purpose of getting accurate raw music data as many as possible. We surveyed a number of famous Chinese music websites, and finally decided to crawl on Google Music (`http://www.google.cn/music/`), Baidu MP3 (`http://mp3.baidu.com/`) and QQ Music (`http://music.qq.com/`)[3], as they provide abundant, well-structured music metadata and have a large amount of Chinese users.

Through comparison between the three information sources, we found that Baidu MP3 has the most raw data but not well structured; Google Music does not have sufficient raw data; while QQ Music provides the second largest raw data with high accuracy and the best structure. Higher quality of raw data will greatly release our burden on data pre-processing and improve recommendation accuracy. Therefore, we decided to use QQ Music as the main data source while using the rest two ones as its complement. In total, we retrieved 4,679 records on singers, 137,882 records on songs and 12,945 on albums.

Furthermore, we crawled data from *Wikipedia Infobox* so as to enrich the descriptions on singers and songs. For instance, through visiting and analyzing Andy Lau's Chinese wiki page,[4] we can supply the descriptions about his debut date, most famous songs, etc.

2.2 Refinement

The collected raw data have to be refined due to a few abnormal structures, redundancy, and encoding problem. To solve the redundancy problem, we proposed a similarity comparison algorithm based on virtual documents [6] and Vector Space Model (VSM) [7] to determine whether two raw metadata are actually very similar. Formally, we consider an attribute A_i for a song s as a vector in the virtual space, and we use a weight α_i ($\alpha_i \in [0, 1]$) to indicate the importance of A_i to describe songs. With a higher weight, the attribute shows higher importance. Therefore, the description of s can be defined as: $\text{Desc}(s) = \sum_{i=1}^{n} \alpha_i * A_i$.

If two records of songs are actually the same, the necessary condition is that the song names are similar, and so are their singers, album names and other attribute values. Since the importance of each attribute varies, we assign specific values to different α_i according to our empirical tests.

We created the description vectors based on ICTCLAS, an integrated Chinese lexical analysis system (`http://ictclas.org/`), and the similarity between songs can be determined by their cosine value. We assign 0.95 to the similarity threshold θ. If the cosine value is larger than 0.95, the duplicate song records will be deleted. Same procedure is applied to singer and album records.

[3] Baidu is the largest search engine in China, and QQ is the largest social networking and instant messaging platform in China.

[4] Chinese names are translated to English throughout the paper for English audience.

```
<foaf:Person rdf:about="http://ws.nju.edu.cn/smartmusic/artist/1830">
  <rdf:type rdf:resource="http://musicontology.com/MusicArtist" />
  <foaf:name>Jacky Cheung</foaf:name>
  <foaf:age>49</foaf:age>
  <mo:composed_in>
    <mo:Composition>
      <mo:produced_work>
        <mo:MusicWork>
          <mo:performed_in>
            <mo:Performance>
              <mo:recorded_as>
                <mo:Signal>
                  <mo:published_as>
                    <mo:Record rdf:about="http://ws.nju.edu.cn/smartmusic/record/7534" />
                    ...
```

Fig. 1. An RDF fragment about Jacky Cheung

2.3 Conversion to RDF

We used Jena (`http://jena.sourceforge.net/`) to convert the data into RDF format, and MusicOntology (MO, `http://musicontology.com/`) to describe the data. MO is designed for sharing and searching all the information related to music, so our system can reach good compatibility with existing standards. For instance, Fig. 1 shows an RDF fragment about a Chinese singer Jacky Cheung.

3 Music Recommendation

3.1 Basic Strategies

Our goal is to provide accurate personalized recommendation. Our system predicts users' preferences based on the analysis of user actions, calculates the semantic distances between songs and provides recommendations using specific algorithms. Users are also allowed to react to these recommendations. Over and over again, the system can form a better idea of a user's preference.

We applied two recommendation strategies together. One is recommending possible beloved songs of a user, and the other one is randomly picking up songs in the database, recommending them if the system believes that the user would not dislike the songs. Such combination could provide recommendation of fresh songs plus beloved old songs, thus attract users.

3.2 Analysis on User Actions and Preferences

Information of user actions on the system helps us to predict users' preferences on music, so it is the foundation of our recommendation. We recorded and analyzed user actions. The information contained in the actions determines how the system should react (see Table 1): an active expression means that a user clearly likes a specific singer or song; a neutral expression indicates that a user shows vague interest on a specific singer or song; and a passive expression implies that it is possible that a user dislikes a singer or song.

Table 1. Information delivered from user actions

Information delivered	Typical user actions
Active expression	Add to favorites; dig a singer; search a specific song/singer
Neutral expression	Sequentially listening
Passive expression	delete/skip a song; bury a singer

Moreover, the system provides a small psychology test designed according to the work in [4], in order to form a general view about the user music preference. At present the test can only speculate the preference very generally, for example whether the user likes pop music, or rock music. Because the test somehow lacks credibility, we only consider the results as neutral expression.

3.3 Semantic Distance Computation

We need to calculate the semantic distances between songs to know how close their relationships are. A song is a very special information carrier. It contains several key attributes, such as singer names, genres, styles, which are enough to speculate whether user like this song. We reused the method in Section 2.2 to calculate the semantic distances using the three above attributes.

Based on our tests, when the similarity is greater than 0.8, our recommendation results are satisfactory but limited. For instance, if a user is clearly fond of a singer M, then the system will find some singers closely related to M and recommend only the songs by those singers. This leads to a situation that all the other singers are ignored. So, we decreased this threshold, and saved more corresponding results. The connections between these songs are weak, but perhaps meaningful. In fact, the threshold is not fixed in the system, but fluctuates according to the user fondness of songs. If a user clearly states that he likes a song, then the threshold is slightly lower than the case that the user's preference is vague. Such assignment can provide abundant and reasonable results.

4 System Interface and Evaluation

4.1 System Interface

Our system interface is mainly divided into four parts, as shown in Fig. 2.

A *Music Player*. We built our system on an open source software named CMP Player (`http://bbs.cenfun.com/`). Almost every action a user performs on the player (e.g., switch or skip a song) is listened and then recorded.
B *Information Board*. Click "Open Infoboard" so as to check the favorites and search result of the logged user. Moreover, a music preference test is offered.
C *Interaction Zone*. User can provide feedbacks to the recommended songs and corresponding singers.
D *Information Bar*. Tips and feedback information are displayed.

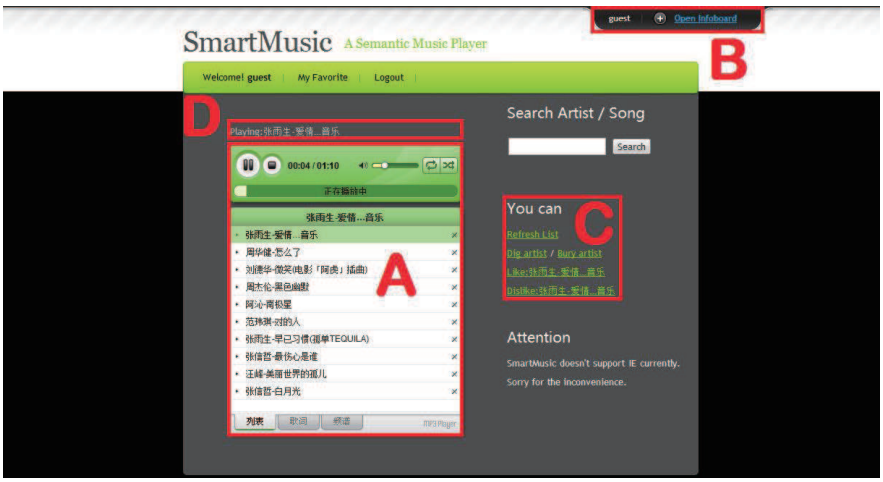

Fig. 2. System interface

4.2 Evaluation

To analyze and evaluate the recommendation offered by SmartMusic, we adopted the evaluation method described in [3]. We invited 9 volunteers (3 female and 6 male college students) to assess the results. First, we randomly divided the nine students into three groups, each with three people. Each user in every group has to provide information about 5–10 their beloved/disgusted songs. After the system gets the information, it recommends 10 songs to each student. The students have to score every recommended song (score 1 represents hatred and score 5 represents fondness). Same test was repeated for three times. Moreover, we asked the volunteers to do exactly the same job on Douban FM (`http://douban.fm/`) and Baidu FM (`http://fm.mp3.baidu.com/`), in order to make comparison.

We found that recommendation improves with more information provided. The average satisfaction degree of our system is 66.4%. It is not an extraordinary result in the area of music recommendation. Also, the recommendation offered by SmartMusic is slightly worse than Douban FM, but better than Baidu FM. Two explanations are offered. Firstly, our recommendation algorithm needs to be further tested and improved, since our system lacks enough preference data; Secondly, as the music databases of those commercial music websites such as Douban FM are huge and accurate, their recommendations might have gone through long time of selection and thus fit the public appetite. Although the test cannot reach a completely justifiable conclusion, it did reflect the trend that recommendation will be improved if user preference information is enriched.

5 Related Work

We investigated several music recommendation systems in China. Online music players such as Yige (`http://www.1g1g.com/`), Douban FM and Baidu FM can

automatically play music without any input. These websites record user actions and recommend based upon action analysis. 8box (`http://8box.com/`) is the first website in China offering music recommendation, which could import contact information from a user's mailbox, and speculate the user's preference by the preferences of her friends. Additionally, it provides a music appetite test. Collaborative filtering [2] is a commonly used strategy for all these websites.

Outside China many mature systems for music recommendation exist. Music sharing website Last.fm is established in 2002 and it has 40 million users now. It adopts a recommendation system called Audioscrobbler, which creates a very detailed record for every user. Audioscrobbler records the songs that a user has listened to in Last.fm, and gives recommendation by the analysis of these songs as well as its huge social network. dbrec [5] is a singer recommendation system based on DBpedia, which calculates the semantic distances between singers by an algorithm called LDSD. Its most attractive feature is that it can give users a clear explanation on the recommending results.

6 Conclusion

In this paper, we introduced SmartMusic, an online music recommendation system based upon Semantic Web technology. Metadata collection, refinement and RDFization, semantic distance calculation, and user preference analysis are covered. Through half a year's use in our university, we believe that SmartMusic can provide relatively satisfactory music recommendation for our college students.

References

1. Celma, Ò.: Foafing the Music: Bridging the Semantic Gap in Music Recommendation. In: Cruz, I., et al. (eds.) ISWC 2006. LNCS, vol. 4273, pp. 927–934. Springer, Heidelberg (2006)
2. Goldberg, D., Nichols, D., Oki, B.M., Terry, D.: Using Collaborative Filtering to Weave an Information Tapestry. Communications of the ACM 35(12), 61–70 (1992)
3. Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating Collaborative Filtering Recommender Systems. ACM Transactions on Information Systems 22(1), 5–53 (2004)
4. North, A.C., Hargreaves, D.J.: Lifestyle Correlates of Musical Preference: 2. Media, Leisure Time and Music. Psychology of Music 35(2), 179–200 (2007)
5. Passant, A.: dbrec — Music Recommendations using DBpedia. In: Patel-Schneider, P.F., et al. (eds.) ISWC 2010, Part II. LNCS, vol. 6497, pp. 209–224. Springer, Heidelberg (2010)
6. Qu, Y., Hu, W., Cheng, G.: Constructing Virtual Documents for Ontology Matching. In: WWW 2006. pp. 23–31. ACM Press, New York (2006)
7. Raghavan, V.V., Wong, S.K.M.: A Critical Analysis of Vector Space Model for Information Retrieval. Journal of the American Society for Information Science 37(5), 278–287 (1986)
8. Stegers, R., Fekkes, P., Stuckenschmidt, H.: MusiDB: A Personalized Search Engine for Music. Journal of Web Semantics 4(4), 267–275 (2006)

A Compliance with the Minimal Requirements

1. *The application has to be an end-user application.*
   - ✓ SmartMusic is a Chinese online music recommendation system for college students in Nanjing University. End-users do not need to be aware of the behind semantic technology, and can use it as a simple music player.
2. *The information sources used*
   (a) *should be under diverse ownership or control.*
      - ✓ SmartMusic crawled music metadata on Google Music, Baidu MP3, QQ Music as well as Wikipedia. None of the resources are under the control of SmartMusic itself.

   (b) *should be should be heterogeneous.*
      - ✓ The music metadata from Google Music, Baidu MP3 and QQ Music have different webpage structures with varying quality. Their used attributes are heterogeneous as well.

   (c) *should contain substantial quantities of real world data.*
      - ✓ SmartMusic collects thousands of singer records, hundreds of thousands of song records. All the data are real-world data retrieved from popular music websites. In addition, the corresponding physical contents are also crawled and stored locally.
3. *The meaning of data has to play a central role.*
   (a) *Meaning must be represented using Semantic Web technologies.*
      - ✓ SmartMusic uses MusicOntology to provide formal semantics to music metadata. Ontology matching techniques are proposed to calculate the semantic distances between singers, songs and albums.

   (b) *Data must be manipulated/processed in interesting ways to derive useful information.*
      - ✓ Wikipedia is used to enrich the semantic descriptions of music metadata. Music recommendation is offered based on semantic distance calculation and user preference speculation.

   (c) *This semantic information processing has to play a central role in achieving things that alternative technologies cannot do as well, or at all.*
      - ✓ Without combining semantic techniques, SmartMusic cannot formalize metadata and compute semantic distances for recommendation. Semantic information processing offers added-value to our system.

B Compliance with the Additional Desirable Features

1. *The application provides an attractive and functional Web interface.*
   - ✓ SmartMusic provides a simple system interface, where users can interact with the system explicitly through feedbacks or implicitly through clicks, search, etc.
2. *The application should be scalable.*
   - ✓ Thousand of singers and hundreds of thousands of songs from diverse data sources are added in SmartMusic. Furthermore, we continue crawling new data for update.

3. *Rigorous evaluations have taken place that demonstrate the benefits of semantic technologies, or validate the results obtained.*
   – We conducted empirical comparison with two existing online music services. Additionally, SmartMusic has been deployed in our university and used for half a year.
4. *Novelty, in applying semantic technology to a domain or task that have not been considered before.*
   × Music recommendation is not a new topic, and applying semantic technology to music recommendation has also been studied in some literature before, e.g., [1,5,8].
5. *Functionality is different from or goes beyond pure information retrieval.*
   ✓ By using the semantics that is encoded in ontologies, SmartMusic offers recommendation different from pure information retrieval.
6. *The application has clear commercial potential and/or large existing user base.*
   ✓ We have thousands of students in our university, which constitute a large existing user base for SmartMusic. Moreover, the students do not need to spend any internet charge for this on-campus music service.
7. *Contextual information is used for ratings or rankings.*
   ✓ Semantic distances between songs and user actions/preferences are used for ranking recommended results.
8. *Multimedia documents are used in some way.*
   ✓ Music is a kind of multimedia resources, so SmartMusic uses multimedia documents in some way.
9. *There is a use of dynamic data, perhaps in combination with static information.*
   – Although existing music websites change their data frequently, SmartMusic only performs a periodical update strategy.
10. *The results should be as accurate as possible.*
    ✓ SmartMusic offers recommendation based on semantic distance calculation, user action analysis and appetite test. Empirical evaluation shows that the recommended results are satisfactory.
11. *There is support for multiple languages and accessibility on a range of devices.*
    ✓ We currently aim at Chinese college students but we also implemented an English user interface for demonstration. Please note that most music in SmartMusic are sung in Mandarin or Cantonese.