Culture-aware Music Recommendation

CITY UNIVERSITY LONDON

Daniel Wolff, Tillman Weyde and Andrew MacFarlane
School of Informatics, Northampton Square, EC1V 0HB London
daniel.wolff.1@soi.city.ac.uk

In the proposed PhD project, a parametrised music similarity model based on multiple feature types is being developed. Adding to the models and algorithms currently available, our similarity model will adapt to a user's cultural context. Motivation comes from cultural models of music perception and interpretation: Instead of assuming an identical understanding for all users, regardless of their cultural, especially musical context, various information sources are combined according to present knowledge about the user‘s cultural context.

New model for music similarity

The main contribution of this work will be an extension to the way computational music similarity is specified: Culture-specific information is integrated into the similarity estimation process.

- **Song similarity will be parametrised**
- **New parameter: cultural indicators**
- **Adapt to users` perception**

Central assumption: within different **cultural contexts**, different musical aspects are perceived. A comparison of songs is thus influenced by culture, with an impact on the musical impression for a listener. Using culturally adapted similarity models may lead to recommendation results being more **satisfactory to the user**. The scope of such contexts may vary and reasonable choices have to be determined yet [UIT2002], [MCK2006].

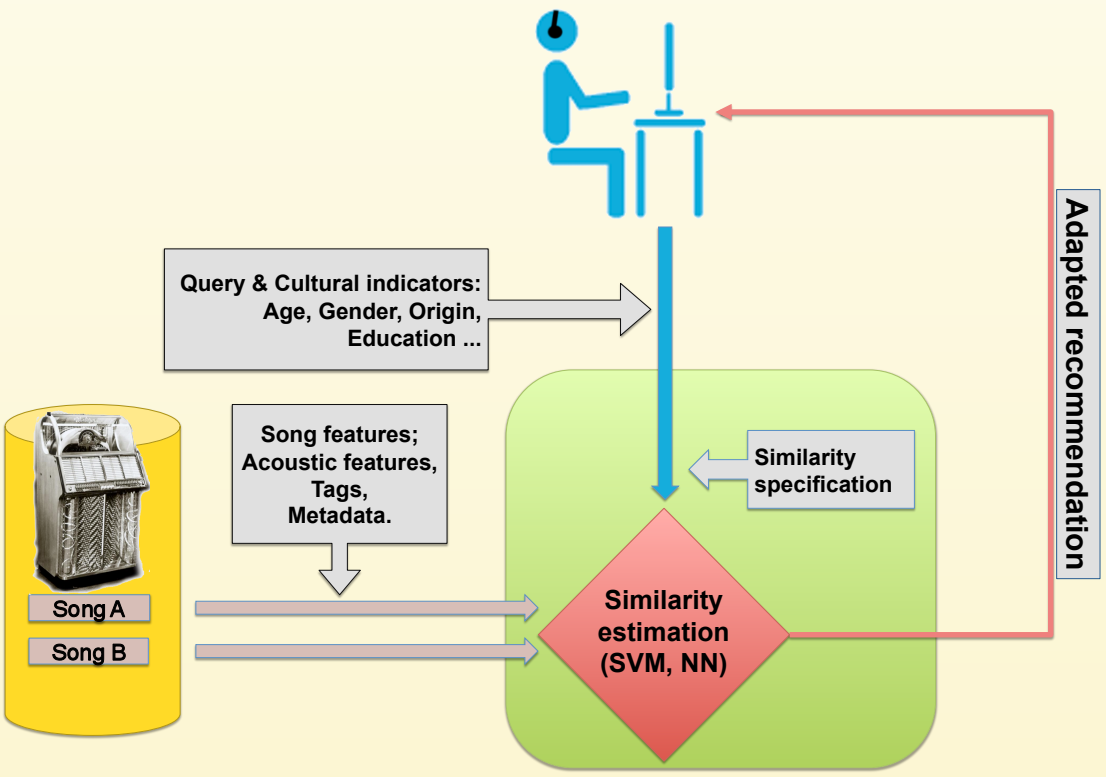

Extending current systems

Most of concurrent music recommendation engines rely on **collaborative filtering** data. Other popular recomendation sites use **folksonomies**, user-driven weighted tag graphs, to access more diverse properties of music. Especially for the recommendation of apriori unknown songs, **acoustic features** have been investigated to provide suitable similarity measurements for audio recommendation.

Combinations of these features have been used for recommendation tasks [PIC2001], [DIN2010], and recently research has started in contextualizing music recommendation engines [MIL2010]. Still, such systems assume an **identical similarity perception** for different users. Customisation of results is only performed on the data end, by selecting different sample songs or genres.

Model architecture

The envisaged similarity model should **adapt its interpretation** of the input song features according to the cultural indicators describing the current user. These are used to associate the user to a specific similarity measure within a similarity space. For modelling the parametrised measurements, we will experiment with the following trainable classificator models:

- **Kernel Based Support Vector Machines**
- **Bayes Classifier**
- **Neural Networks**

The SVM kernel structure has been shown to facilitate post-training analysis of the learned similarity measure [BAR2009]. The kernel space itself may serve for describing the resulting similarity space, whereas the other models provide more straightforward means to include the additional cultural user information.

Collection of data

For training the above model, the following data needs to be collected:

Indicators for cultural contexts per user and associated similarity ratings: In order to collect this data, we plan to use **human computation** games. Two collaborative games will be deployed, engaging the users to either **rate the similarity of songs** or to determine the **relevance of tags** produced by the game.

Existing music databases with human annotation data available, like the MagnaTagATune [LAW2009] database, will be used as a data souce for the games. Such rich databases enable a later analysis of the similarity results **under observation of additional parameters as genre** or specific tags.

Song-based music features: Content-based acoustic features will be extracted directly from the acoustic material. We will consider a full range starting from low-level acoustic descriptors to culturally shaped features considering melody, harmony and rhythm.

References

BAR2008 Barrington, L., Yazdani, M., Turnbull, D. and Lanckriet, G. (2008). Combining Feature Kernels for Semantic Music Retrieval. In Proc. International Symposium on Music Information Retrieval 2008, 614—619.

BAR2009 Turnbull, Douglas R., Barrington, Luke, Lanckriet, Gert and Yazdani, Mehrdad (2009). Combining audio content and social context for semantic music discovery. In Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval, 387—394. ACM.

DIN2010 Dingding Wang and Tao Li (2010). Are Tags better than audio features? The effect of joint use of tags and audio content features. In Proc. International Symposium on Music Information Retrieval 2010.

LAW2009 Edith Law and Luis Von Ahn (2009). Input-agreement: A New Mechanism for Collecting Data Using Human Computation Games. Proc. of CHI . ACM Press.

MCK2006 Cory McKay and Ichiro Fujinaga (2006). Musical genre classification: Is it worth pursuing and how can it be improved?. In Proc. International Symposium on Music Information Retrieval. 2006, 101—106.

MIL2010 Scott Miller, Paul Reimer, Stephen Ness and George Tzanetakis (2010). Location-aware, content-based music browsing using self-organizing tag clouds. In Proc. International Symposium on Music Information Retrieval 2010.

PIC2001 Jeremy Pickens (2001). A Survey of Feature Selection Techniques for Music Information Retrieval. In (Eds.), Proceedings of the 2nd International Symposium on Music Information Retrieval 2001.

UIT2002 Alexandra L. Uitdenbogerd and Ron G. van Schyndel (2002). A Review of Factors Affecting Music Recommender Success. In (Eds.), Proc. International Symposium on Music Information Retrieval 2002.