Adaptive Personalization: A Multi-Dimensional Approach to Boosting a Large Scale Mobile Music Portal

Erich Gstrein [a], Florian Kleedorfer [a], Robert Mayer [d], Christoph Schmotzer [d], Gerhard Widmer [b], Oliver Holle [a], Silvia Miksch [a c]

[a] Smart Agents Technologies, Research Studios Austria, ARC Seibersdorf research GmbH, Vienna, Austria.
{erich.gstrein, florian.kleedorfer, silvia.miksch, oliver.holle} @researchstudios.at

[b] Department of Computational Perception, Johannes Kepler University Linz
and
Austrian Research Institute for Artificail Intelligence, Vienna, Austria
Gerhard.widmer@jku.at

[c] Institute of Software Technology and Interactive Systems of the Faculty of Informatics,
Vienna University of Technology, Favoritenstrasse 9 - 11, A-1040 Vienna

[d] three united mobile solutions ag, Vienna, Austria
{r.mayer, c.schmotzer} @3united.com

ABSTRACT
Although mobile commerce is a rapidly growing business, personalization has yet not the importance it should have. Bad usability of mobile devices and the enormous mass of data induces that the user gets frustrated and refuses the service. This is especially true in the context of mobile music distribution where additional, domain dependent problems arise. The approach we present in this paper tries to tackle the personalization problem as a whole with the aim to optimize the performance from the users point of view. We demonstrate how the combination of the results of several areas of research like *music information retrieval*, *interaction design*, etc. will lead to a better system.

Keywords
Personalization, recommender systems, collaborative filtering, content filtering, music information retrieval

1. INTRODUCTION
The electronic distribution of music has become a very hot topic in e-commerce and especially m-commerce, where it has the potential to become a 'killing app'. Following this trend, many portals offering music (e.g. amazon.com) try to improve their service by providing personalization for better supporting the user in finding – and buying – the appropriate music items.

However, the content domain 'music' bears some specific traps these systems must overcome:

a. *Genres*: An often disgraced concept, but indispensable to a music portal. The most serious problem with genre is that they are not standardized and that they tend to be a source of dispute. For example, when it comes to music styles the *AllMusicGuide* offers 531, *Amazon* 719 und *MP3.COM* about 430 different genres [8].

b. *Content volume*: Music portals often use huge music archives with rapidly increasing content (e.g., www.napster.com promises 1.000.000 titles).

c. *Content life-cycle*: The music industry produces more and more nine days' wonders, like the annual 'summer hits'. From a recommender's point of view, the complexity of the 'new item problem' [7] rises. [1]

d. *Cultural dependency*: The cultural background of the music interested persons plays an important role [8], because it influences many dimensions of the selecting, profiling, and recommending components.

Many personalization solutions are based on *collaborative* or *content filtering* concepts [13, 7]. While the former will have problems with (b), (c) and (d) the latter will mainly suffer from (a).

In this paper we will present a personalization concept, called *adaptive personalization*, which tries to overcome these problems by

- using a hybrid recommender system, based on a multi-dimensional profile model
- incorporating a 'audio similarity measure' based on Music Information Retrieval (MIR) methods [3,4,5] for classifying items

---

[1] The ‚new item' problem is common in recommender systems based on item ratings. The problem is: How to recommend an item which was never rated/bought before?

- incorporating 'cultural similarity' based on analyses of web pages for considering the *cultural dependency*.[2,9,10]

The concepts of this paper were implemented in large parts in the personalization system of *Ericsson's Media Suit – Music* (working title), a large international mobile music portal, which will go online at the end of summer 2005.

2. RELATED WORK

Many personalization solutions are based on *collaborative* filtering [13, 7] or *content-based filtering* concepts [18]. The concept behind the content-based approach is to suggest items to the user which are similar to ones they liked in the past. To compute this similarity a description of the items must be provided in form of profiles or feature vectors. One drawback of this approach is the 'over specialization', based on that kind of recommendations. A broadening of the user's horizon is not supported.

In contrast, collaborative filtering (CF) systems take the similarity of users as a basis for generating recommendations. The CF systems asks the user to rate a presented item - so the knowledge 'who likes what' is gathered. When asked for recommendations a list of items, which where high rated by similar users in the past, is generated by the CF system. The similarity between users is calculated based on the rating behaviour of common items, mostly using the 'P*earson correlation coefficient*' [15]. A good survey about CF systems is presented in [7].

The main drawback of the CF approach is known as the 'cold start problem'. New users have no or only a very poor 'behaviour profile' (e.g. rated or bought items) thus the definition of similar users is not possible. This problem applies e.g. to the *Pearson correlation coefficient*, the most popular measure value used in CF systems, which can only be applied to ratings of items, which all concerned users have rated! Further problem areas of CF are the 'new item problem' – how to present an item which was not rated before? - and the 'sparsity' problem, which occurs when the number of rated (and therefore useable) items is very small.

To overcome the drawbacks of these two approaches hybrid systems, combining the advantages of both, got popular in the past [19, 20] and led to very successful applications. Amazon.com is one of the most popular hybrid systems. Beside the improvements of algorithms the interaction design for recommender systems as well as trust concepts got very popular in the research community [1, 17, 12] for a better consideration of the 'human factor'.

A very interesting contest of music recommender engines called RECO.ENGINE.04, organized by www.musicline.de, proved that producing good recommendations is by far not sufficient in order to achieve high user acceptance. For more information see http://www.musicline.de/de/recoengine04.

3. ADAPTIVE PERSONALIZATION

Our personalization approach is based on a hybrid, self-adapting recommendation system combining the advantages of *collaborative* and *item-based filtering* systems. The major building blocks of our concept are

- a highly sophisticated profile system combining model-based and behaviour based approaches
- a well-defined set of recommendation strategies based on Data Mining and other Artificial Intelligence algorithms focusing on the demands and needs of the users
- the self-adapting, or learning, behaviour based on instance-based learning and Data Mining algorithms
- the integration of meta-information provided by classifier systems.

In contrast to many other systems the collaborative filtering part is implemented by using a *model-based* as well as a *behaviour-based* approach both mapped to the profile structure.

In knowledge-based approaches profiles are created statically by collecting/defining the 'relevant attributes' describing best the model for the given problem domain. The definition of these attributes is a critical task, because the model must be as accurate and abstract as possible. Some modelling techniques are presented in [21].

Behaviour-based approaches usually construct the user model using machine-learning techniques to discover useful patterns in the behaviour. Behavioural logging is employed to obtain the data necessary from which to extract patterns.

The coarse system architecture is presented in figure 1.

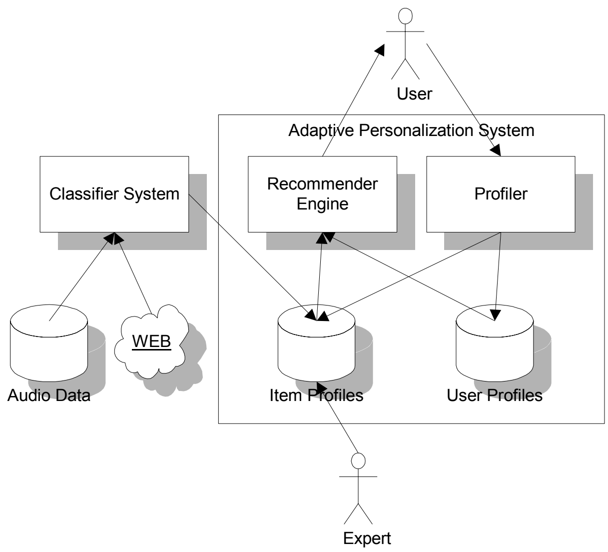

Figure 1: Adaptive Personalization

The *Recommender Engine* generates recommendations based on the *Item–* and *User Profiles*. The *Profiler* 'observes' the user's behaviour and refines the appropriate profiles. The *Expert* initializes the *Item Profiles* either manually or based on third party information like catalogues. The *Classifier System* is used – by the Expert – to generate music meta information (stored in the item profiles) based on audio data – e.g. like music archives - and WEB data.

3.1 Multi-Dimensional Profile System

3.1.1 Modelling the User

Among the user's needs, a distinction can be made between well defined needs where the user is able to characterize an appropriate means of satisfaction, and ill defined needs where the user does not know how to satisfy or even how to define them. Furthermore, users are normally not isolated during the use of personalization systems; therefore an effect is created within the relevant community. These effects can be very multifarious, ranging from deliberate interactions, like the placement of ratings or recommendations, to being a (passive) example for other users or Data Mining algorithms.

These considerations led to a multi-layer model where each user is modelled by a profile comprised of three different views:

- self assessment
- system observation
- community assessment (assessment of others)

This structure is presented in figure 2.

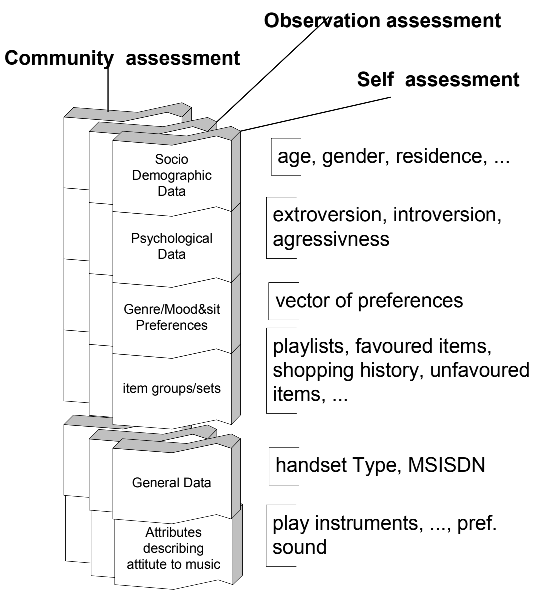

Figure 2: Structure of User Profile

The *self assessment* view, or sub profile, is used to model the 'self-portrait' of the user concerning information like preferences, socio-demographic data like age, gender, etc. mainly used to serve the 'well-defined needs' but also psychological attributes, taken from models like *big-five-factors* [14] or from the *Myers-Briggs Type Indicator* [15], in order to get a broader impression/image of the user. These psychological dimensions together with some Data Mining strategies are used to overcome the problem of 'ill defined' needs by a way to fill them.

The *system observation* is a view, where the behaviour-based profile of a user is stored. The system observes the user while using the system and so a 'dynamic' profile is created. In contrast to self assessment, where users give an image of themselves (which is subjective, of course), the system observation profile represents what they actually do.

The *community assessment*, not always as important as the other two sub profiles, represents how a user is seen by others and can be used as a feedback or rating on the *self assessment* view.

This complex model, combining knowledge– and behaviour-based approaches, forms the basis on which a wide range of needs can be served. The information of the self-portrait can be used to satisfy the 'obvious' needs - even (and especially important) when this description is somewhat 'idealized'. The sub profile created and automatically refined through *system observation* view is used to identify and satisfy behaviour-based needs.

The user profile structure implemented in the *Ericsson's Media Suit – Music* supports the following groups of attributes:

1. Socio-demographic data, like age, gender and place of residence
2. Music genre preferences, currently referring to the genres used in the STOPMP model [11]
3. Music connotation preferences based on a collection of moods and situations (e.g. car driving)
4. Compilations like favoured/unfavoured artists, tracks or playlists
5. Important aspects, like preferred sound, importance of lyrics, preferred instruments
6. History data describing bought or viewed items

Where to place an attribute mainly depends on how this information can be elicitated:

- socio-demographic data (1) and important aspects (5) can only be defined by asking the user, therefore they are part of the *self assessment* view.
- The different kinds of preferences (2, 3, 4) are part of the *self assessment* and the *system observation* view.
- The history data (6) is collected only by the system and is therefore only part of the *system observation* view.

In the case of recommendations the three views can be used separately or by combining them to a 'weighted-sum' profile. The latter approach is implemented in *Ericsson's Media Suit – Music*. The community view is not yet used as well as the psychological data but will be introduced in the context of community features in future releases.

3.1.2 Modelling the Items

A similar structure is also provided for the items to be recommended, where the affiliation of items to certain clusters (like genres, see below) is modelled with the help of three different views:

- the assessment of a domain expert
- the assessment of the (user) community
- affiliations calculated by a classifier systems as described later (see chapter 'The Role of MIR)

The structure is presented in figure 3.

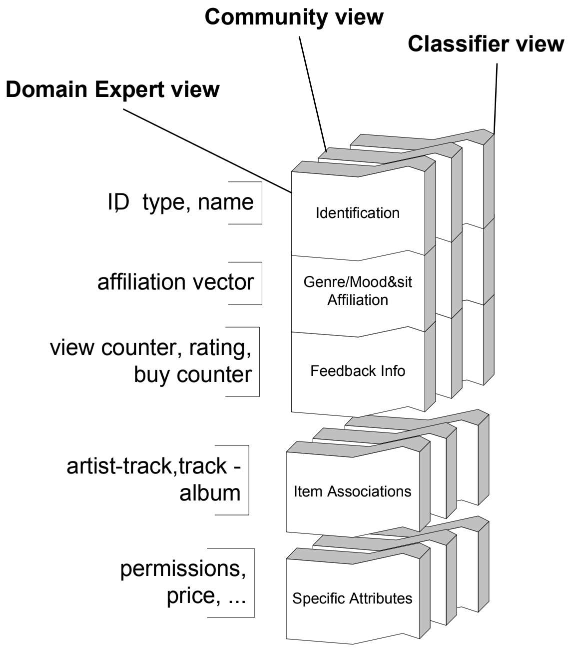

Figure 3: Structure of Item Profile

The *domain expert* view often represents the opinion or assessment of the content owner while the *community view* reflects how the content is seen by the consumers/users. By providing both views the content owner gets an important feedback and a better explanation model can be provided for the user concerning the recommended items.

Especially in the context of music the affiliation of artists or tracks to some given genres is a very controversial topic, leading to arbitrary classifications and thus to hardly acceptable recommendations.

The *classifier view* can be seen as an extension of the *domain expert* view, where 'third party' information is used to refine the item profile. This information can be provided simply by a catalogue or even an appropriate classifier system (see chapter 'The Role of Music Information Retrieval'). Within the *Ericsson's Media Suit – Music* an audio classifier system, based on Music Information Retrieval methods, is used.

In the case of recommendations the three views can be used separately or a 'weighted-sum' profile can be created. The latter approach is implemented in *Ericsson's Media Suit – Music*.

3.1.3 Clustering the Itemspace

The basic idea behind the structuring of the item space is to support a flexible and individual classification of the items to be recommended.

A cluster can be seen as a named container for items sharing commonalities in some respect. Between clusters associations can exist expressing some special relations, like 'is-sub-cluster-of', 'is similar-to', etc. The number and kinds of associations is mainly guided by the problem domain. The current implementation supports two sets of clusters addressing music genres following the STOMP model as presented in [11] and some client specific 'mood and situations' categories. 'Mood and situations' is a popular classification scheme to categorize music along some consumer relevant topics like: 'Car driving', 'Candle light Dinner', 'Feeling Down', etc.

Clusters can be created manually by the administrator or can be created by some cluster analyzing programs operating on the item set. The assignment 'Item – Cluster' can be made by hand, with the help of some classifier systems or simply by using existing information about domain specific clustering on item level (e.g., catalogues).

Furthermore this structure is also used during profile refinement and supports the learning or adaptive behaviour of the system.

3.2 Profile Initialization and Refinement

The *self assessment* view of the user profile is initialized during an optional registration process and is refined by the ongoing re-definition of favoured items (e.g. favoured artists, songs) by the users themselves.

The *system observation* view of the user profile is refined by observing the users behaviour. When the user 'rates' an item, his/her cluster preferences are adapted according to the cluster affiliation vector of the rated item and the kind of rating. Several user actions, like pre-listening or buying tracks, viewing artist/track pages (where the track/artist is presented) are seen as *ratings* with different weights.

The item profile will be initialized by a domain expert in most cases, either by manually classifying the items – affecting the *domain expert* view – or by using third party information (e.g. provided by catalogues or classifier systems) – affecting the *classifier* view.

The *community assessment* view is initialized and/or refined by the feedback of the users concerning the classifications of item. In *Ericsson's Media Suit – Music* the user can give feedback about the cluster affiliation of items (tracks, artists, etc.)

3.3 Recommendation Strategies

Another important factor is to support the right set of recommendation strategies the user is expecting from an intelligent music recommender system. We use the following strategies which are partly followed from the proposals of Swearingen and Sinha [1]:

1. Reminder recommendations.
2. 'More like this' recommendations
3. Recommend new/'hot' items
4. 'Broaden my horizon' recommendations
5. 'Similar' users like (e.g. view, buy,)
6. Shopping Cart recommendations

Reminder recommendations should help the user not to forget or oversee some important items he/she was willing to use or buy in the past. Theses recommendations are based on a list which is maintained by the user (for example, think of a 'black board' feature).

More-like-this recommendations – probably the most common one – should help the user to find similar items. Hereby similarity is defined by

- another track of the same artist
- another track/artist out of the same cluster (e.g., genre)
- another track 'sounding similar' (to a given one)
- a similar artist as defined by a classifier system (see chapter 'The Role of Music Information Retrieval)

starting from a given item.

This recommendation is generated by using explicitly defined item relations (artists-track), the similarity of the cluster affiliations of the items and similarity relations.

Hot-Item-recommendation should support users to be up-to-date within the range of their preferences. These recommendations help to satisfy community needs, where a user wants to be best informed within his/her social environment (e.g., 'more accurate than friends'). In *Ericsson's Media Suit – Music*, new tracks of favoured artists or new tracks of the preferred genres are recommended.

This recommendation is generated by using the artist – track relation (new tracks of favoured artists) and by matching the cluster preferences of the user with the cluster affiliation of the items.

Broaden my Horizon is a very important recommendation strategy because it supports the user to explore his/her taste and it helps to sharpen or verify the users profile, avoiding to get stick in a 'local maximum' of the profile/item space. Starting from the well defined needs (favoured artists, preferred genres, etc.) the user can explore his/her taste by allowing more and more offensive recommendations. The direction of this 'broadening process' is mainly defined by the structure of the item space.

Also this recommendation is generated by using the artist – track relation and by matching the user's cluster preferences with the cluster affiliation of the items.

The *Similar Users like* recommendations address the social aspects, where users want to know what others do. The similarity between users is defined by the similarity of their profiles or a subset of these attributes. Several different similarity relations can be defined either based on the user's behaviour – for example, buying/rating history – and/or on the model dimensions, like socio-demographic data. Therefore the collaborative aspect can be implemented very flexible and robust.

This recommendation is generated by finding the 'k nearest neighbours' based on the user's profile [16]. Having found these 'similar users', their favoured songs/artists are taken.

The *Shopping cart recommendations* – a strategy used by many personalization systems – try to find associations between items which are not modelled within the profiles. This recommendation is generated on a basis of association rules using the latest buys of the users (compare [6]).

3.4 Self Adapting Capability

The self adapting or learning behaviour of the personalization system is realized by using Data Mining and instance-based learning algorithms. This 'learning behaviour' is realized on three different levels:

1. The individual level, where profiles of users are permanently refined based on implicit (e.g., navigation observation) and explicit feedback (e.g., ratings, buying behaviour, etc.). For frequent users, the quality of the profile will increase over time.
2. The collaborative level, where the community ratings of items will improve the recommendation quality as well as the refinement of profiles will lead to an improved 'similarity' relation (e.g., 'items, similar users like'). Furthermore, an evolutionary aspect is introduced, also considering the refinement path of a profile and not only the current characteristics.
3. The statistic level, where data mining algorithms are applied, like association rules [6] to generate new recommendations. Because these algorithms operate on data based on huge amount of user-behaviour (e.g., shopping history, compilations of favoured items) the quality will improve over time.

4. THE ROLE OF MUSIC INFORMATION RETRIEVAL (MIR)

As described above, a basic knowledge concerning the item space - for example, the number and relations of clusters - is an important aspect of our approach. In many problem domains basic classification schemes exist. But what about music, concerning the problems mentioned before? How to classify items without having a sound set of classes? How to deal with the fact that some content providers do not provide classification information? And what about cultural diversity – how to address this topic?

As a designer of a personalization system you should take any information you can get, but you must face the worst case. In our context, the minimal sources of information are:

- a set of audio files, each having a title and the name of the artist
- the web with an unforeseeable number of pages, related to some music topics, like fan-pages, artist's home pages, etc.

But how to extract the needed meta information to feed the personalization system? The field of Music Information Retrieval (MIR) has recently started to investigate questions like these. In particular, the definition of similarity measures between music items is a hot topic in the MIR community, currently focusing on two different approaches:

1. definition of similarity based on audio data [3, 4, 5]
2. definition of similarity based on cultural aspects [2, 9, 10]

The audio-based similarity is generated by first extracting feature vectors out of the audio files and then computing the distances between these vectors. The approach used in *Ericsson's Media Suit – Music* is based on 'timbre similarity' [4] which leads to a 'sounds similar' relation of audio tracks.

The similarity based on cultural aspects is generated by analyzing web pages. Based on the results of a Internet search (e.g., using a search engine with the name of an artist as the search key) text retrieval methods are applied to get the most important (discriminatory) words forming the 'result profile'. In a next step the distances between these 'result profiles' are generated.

These similarities are used manifold in the adaptive personalization concept:

- based on the distances between the items, automatic clustering algorithms are applied. These clusters can be used to generate 'similar items' recommendations (e.g., 'sounds similar'). This will lead to more intuitive classifications than genres, because the key factor is the timbre similarity.
- Even the common genre classification can be provided by introducing a prototype-based genre definition. A specific class – e.g., genre – is defined by a set of prototypes, determined by a domain expert. The classification of an item to a class is then defined by the distance to the prototypes of each class.

The audio-based similarity definition is used to define the cluster affiliation of the items (tracks, artists) and mainly affects the classifier view. The web-based similarity relation reflects the opinion of the community and therefore refines the community view of the items.

Furthermore the web-based approach can be used to implement cultural dependent recommendations and to detect trends early. The incorporation of the web approach is scheduled for the next releases of *Ericsson's Media Suit – Music*.

5. EXAMPLE USE CASE

In this section we will present a short example use case as implemented in a prototype of *Ericsson's Media Suit – Music*.

Figure 4: Demonstrator

The screen of the mobile device (figure 4) on the left side shows the personalized welcome page for a registered user with a specific nickname. Below the welcome bar, a top recommendation ("Shiny happy people") is presented to the user together with the cover poster of the track/album.

Below further recommendation strategies are provided to the user:

- "Hot Music For You" offers more top recommendations
- "Explore Your Taste" implements the broaden your horizon strategy
- "Similar Users Like" provides the best rated tracks of users who are similar in respect to the socio-demographic data and genre preferences.

Selecting the "Explore Your Taste" link will lead to the screen presented on the mobile device(figure 4) at the right hand side.

Below the (explanation) text on the top, a list of tracks is presented calculated on the basis of the user's preferences. The user now can shift his/her musical perspective by using a slider concept, here implemented with two buttons <Minus> and <Plus>, for broaden his/her horizon. With this slider the user can define how unconventionally/open minded this recommendation should be, starting from the focal point of the explicitly defined preferences. (e.g. favored artists/songs).

6. CONCLUSIONS, AND FUTURE STEPS

The concepts behind the *adaptive personalization* aim to tackle the problem of personalization as a whole, ranging from generation of meta information of the content to user modeling aspects and interaction design for recommender systems. The focus of our research is to optimize the overall performance of a personalization system from the users point of view, especially for the mobile world.

Initial user tests concerning the usability and the acceptance of the provided recommendation strategies gave promising results. We will verify them with the huge amount of data acquired from our real world working system, which is a very challenging task.

Another interesting aspect which can only be tested with the real world application is the acceptance of the classifications based on MIR concepts. To apply these techniques in the large some problems mainly concerning the performance and the handling of prototype based classifications must be improved.

Future research will also address trust (e.g., web-of-trust) concepts as presented in [12] for improving the recommendation quality and to increase the robustness of the personalization system against attacks.

7. ACKNOWLEDGMENTS

This work was supported by the Austrian Federal Ministry of Economics and Labor.

8. REFERENCES

[1] K. Swearingen, R. Sinha. Interaction Design for Recommender Systems. University of California, Berkley, USA 2001

[2] Knees, P., Pampalk, E., and Widmer, G. (2004). Artist Classification with Web-based Data. In Proceedings of the 5th International Conference on Music Information Retrieval(ISMIR'04), Barcelona, Spain, October 10-14, 2004.

[3] Pampalk, E, Dixon, S., and Widmer, G. (2003). On the Evaluation of Perceptual Similarity Measures for Music. In Proceedings of the 6th International Conference on Digital Audio Effects (DAFx-03), London.

[4] Aucouturier, J.J and Pachet, F. (2002). Music Similarity Measures: What's the Use?In Proceedings of the International Conference on Music Information Retrieval (ISMIR'02), Paris, France.

[5] Aucouturier, J.-J. and Pachet, F. (2004). Improving Timbre Similarity: How High isthe Sky? Journal of Negative Results in Speech and Audio Sciences 1(1).

[6] Agrawal, R., and R. Srikant. 1994. Fast algorithms for mining associations rules in large databases. In: Bocca, J., M. Jarke, and C. Zaniolo (Publ.), Proc. International Conference on Very Large Data Bases, Santiago, Chile. San Francisco: Morgan Kaufmann, S.478-499

[7] Herlocker, J., L., and Konstan, J., A., Terveen, L., G., and Riedl, J., T., (2004). Evaluating Collaborative Filtering Recommender Systems. ACM Transactions on Information systems, Vol. 22 No. 1, January 2004, Pages 5 – 53

[8] A. Uitdenbogerd, R.v. Schyndl. A Review of Factors Affecting Music Recommender Systems. Dep. of Computer Science, Melbourne, Australia.

[9] Whitman, B. and Lawrence, S. (2002). Inferring Descriptions and Similarity for Musicfrom Community Metadata. In Proceedings of the 2002 International Computer Music Conference, pp 591-598. 16-21 September 2002, Göteborg, Sweden.

[10] Baumann, S. and Hummel, O. (2003). Using Cultural Metadata for Artist Recommendation. In Proceedings of the International Conference on Web Delivery of Music (WedelMusic), Leeds, UK.

[11] P.J. Rentfrow, S. Gosling. The Do Re Mi's of Everyday Life: The Structure and Personality Correlates of Music Preferences; Journal of Personality and Social Psychology, 2003, Vol. 84, No. 6, 1236–1256

[12] P.Massa, B. Bhattacharjee. Using Trust In Recommender Systems: an Experimental Analysis. 2004 iTrust conference

[13] J. Herlocker, J.A.Konstan, J. Riedl. Explaining Collaborative Filtering Recommendations. ACM 2000; Conference on Computer Supported Collaborative Work.

[14] Tupes, E.C., and Christal, R.E. (1961). Recurrent personality factors based on trait ratings (Tech. Rep. ASD-TR-61-97). Lackland Air Force Base. TX: US Air Force.

[15] L.Q. Naomi. Essentials of Myers-Briggs Type Indicator Assesment. John Wiley & Sons, 1999.

[16] E. Gstrein. The BPnD-Tree. An efficient search structure for high dimensional profile based recommendation systems. Technical Report. Researchstudios Austria 2003.

[17] D. Cosley, et al. Is Seeing Believing? How Recommender Interfaces Affect Users' Opinions. ACM 2003. CHI 2003.

[18] Meteren, R. v.; Someren, M. v.: Using Content-Based Filtering for Recommendation. In: Proceedings of ECML 2000 Workshop: Machine Learning in New Information Age. Barcelona, Spanien 2000. S. 47-56.

[19] Burke, R. Hybrid recommender systems: Survey and experiments. In User Modeling and User-Adapted Interaction, 2002.

[20] Tran, T. and Cohen, R.: 2000, 'Hybrid Recommender Systems for Electronic Commerce'. In Knowledge-Based Electronic Markets, Papers from the AAAI Workshop, AAAI Technical Report WS-00-04. pp. 78-83. Menlo Park, CA: AAAI Press.

[21] KOBSA,A. 1993. User modeling: Recent work, prospects and Hazards. In Adaptive User Interfaces: Principles and Practice, M. Schneider-Hufschmidt, T. K¨uhme, and U. Malinowski, Eds. North-Holland, Amsterdam, The Netherlands.