InCarMusic: Context-Aware Music Recommendations in a Car

Linas Baltrunas[1], Marius Kaminskas[1], Bernd Ludwig[1], Omar Moling[1], Francesco Ricci[1], Aykan Aydin[2], Karl-Heinz Lüke[2], and Roland Schwaiger[2]

[1] Faculty of Computer Science
Free University of Bozen-Bolzano
Piazza Domenicani 3, 39100 Bolzano, Italy
fricci@unibz.it

[2] Deutsche Telekom AG, Laboratories
Innovation Development
Ernst-Reuter-Platz 7, D-10587 Berlin, Germany
Karl-Heinz.Lueke@telekom.de

Abstract. Context aware recommender systems (CARS) adapt to the specific situation in which the recommended item will be consumed. So, for instance, music recommendations while the user is traveling by car should take into account the current traffic condition or the driver's mood. This requires the acquisition of ratings for items in several alternative contextual situations, to extract from this data the true dependency of the ratings on the contextual situation. In this paper, in order to simplify the in-context rating acquisition process, we consider the individual perceptions of the users about the influence of context on their decisions. We have elaborated a system design methodology where we assume that users can be requested to judge: a) if a contextual factor (e.g., the traffic state) is relevant for their decision making task, and b) how they would rate an item assuming that a certain contextual condition (e.g., traffic is chaotic) holds. Using these evaluations we show that it is possible to build an effective context-aware mobile recommender system.

1 Introduction

Recommender Systems (RSs) are software tools and techniques providing suggestions for items to be of use to a user [11]. In this paper we focus on a particular approach for RSs, Collaborative Filtering (CF). In CF, explicit ratings for items, given by a population of users, are exploited to predict the ratings for items not yet evaluated by the users [6]. Often, system generated recommendations can be more compelling and useful if the contextual situation of the user is known. For instance, in a music recommender, the traffic condition or the mood of the driver may be important contextual conditions to consider before suggesting a music track to be played in her car. Context-aware recommender systems (CARSs) are gaining ever more attention and various techniques have been introduced to improve their performance [1].

C. Huemer and T. Setzer (Eds.): EC-Web 2011, LNBIP 85, pp. 89–100, 2011.

To adapt recommendations to the user's context, the dependency of the user preferences (i.e., the ratings in CF) on the contextual situations must be modeled. Hence, a major initial issue for the correct design of CARSs is the assessment of the contextual factors that are worth considering when generating recommendations. This is not an easy problem: it requires informed conjectures to be formulated regarding the influence of some data, before collecting the data. Moreover, after a meaningful set of contextual factors is identified, a model, which predicts how the ratings will change depending on the contextual factors, must be built. For a set of items, this step requires the collection of explicit ratings from a population of users under several distinct contextual situations.

The main contribution of this paper is the description of a methodology for supporting the development cycle of a Context-Aware Collaborative Filtering system, as sketched above. This methodology has been previously applied to a completely different application scenario, for recommending places of interest [2], and it is adapted here to the problem of recommending music tracks to a group of users in a car. The methodology comprises four steps: context factors relevance assessment; in-context acquisition of ratings; context-aware rating prediction; and context-aware recommendation generation and visualization for a user. Each of these steps is supported by a specific system and technique. First, in order to quantitatively estimate the dependency of the user preferences on a candidate set of contextual factors, we developed a tool for acquiring context relevance subjective judgments. Second, we developed a user interface that actively asks the users to rate items under certain contextual conditions. Next, a predictive model was built, which has the goal of predicting the user's ratings for items under target contextual situations where these ratings are not known. We show that this model, which extends classical matrix factorization, can generate accurate recommendations, i.e., can better predict the true ratings, compared with a system that does not take into account the contextual information. Finally, a mobile recommender system (InCarMusic) was built to present the recommendations to the user. The recommendations have the highest predicted rating for the user's contextual situation with the joint preferences of all the passengers in the car considered, i.e., providing group recommendations [5,3].

The rest of this paper is organized as follows: in Section 2 we discuss some of the related work. In Section 3, we introduce our context-aware recommender system prototype (InCarMusic) to give immediately the motivations of our technological development. In Section 4, we explain our approach for acquiring the data describing the relationships between user preferences and contextual situations. In Section 5, we present our algorithm for context-aware recommendations and we illustrate the results of the evaluation of the proposed model. We finally draw our conclusions and list some open issues that call for future work.

2 Related Work

Context-awareness in recommender systems as a research topic has been receiving considerable attention in the last years [1]. To the best of our knowledge, the specific problem of in-car context-aware music recommendation has not been addressed until now. There is a body of work, however, on the related problem of context-aware music recommendation, which typically addresses different recommendation scenarios. For instance, [8] has improved a music recommender service with context awareness using case-based reasoning. The used context factors include the season, month, weekday, weather and temperature information. Listening cases have been obtained by aligning users' listening history data with weather bureau information. In [10] a context-aware music recommender for urban environments is presented. The context factors include the location of the user (in terms of a ZIP code), time of day, weekday, noise/traffic level, temperature and weather data. The system was bootstrapped by manually annotating the tracks in the user's library with the values of the selected contextual factors.

A common feature of these systems is the usage of a generic context model, mostly consisting of time- and weather-related information. We note that these research works do not formally address the issues of context factor selection and system bootstrapping as we do in the presented work. The choice of the most informative context factors has not been informed by any data mining experiment, and the impact of individual context factors on music perception has not been investigated.

Another area of context-aware music recommendation is dedicated to adapting music content to other types of multimedia, e.g., web pages [4] or images [9]. These systems typically use machine learning methods for learning relations between music and the context information (i.e., text or images).

3 InCarMusic Mobile Application

InCarMusic is a mobile application (Android) offering music recommendations to the passengers of a car after they have entered ratings for some items using a web application that will be illustrated in the next section. If the user did not previously enter any ratings, then the recommendations are adapted solely to the contextual situation and not to the user long term preferences described by her ratings.

First, the *Channels* tab allows the user to specify and edit channels (see Fig. 1(a)). A channel is meant to provide a certain kind of music to the user. In the channel specification the user can detail, for instance, that the channel "Happy-Classic" is appropriate when she is "happy" and would like to listen mostly to classical music and a bit of jazz. Creating such a channel enables the user to quickly switch to this type of music whenever she likes. A default channel is also provided for recommending music without asking the user to create one. Second, the *Passengers* tab allows the user to identify the passengers that are present in the car (see Fig. 1(b)). We note that the user, is always included in the passengers list. Passengers can be imported from the local contacts and should have previously provided some ratings, as it is requested to the user (see Fig. 1(c)). This means that they should have registered to the Web portal that provides the music content to our system.

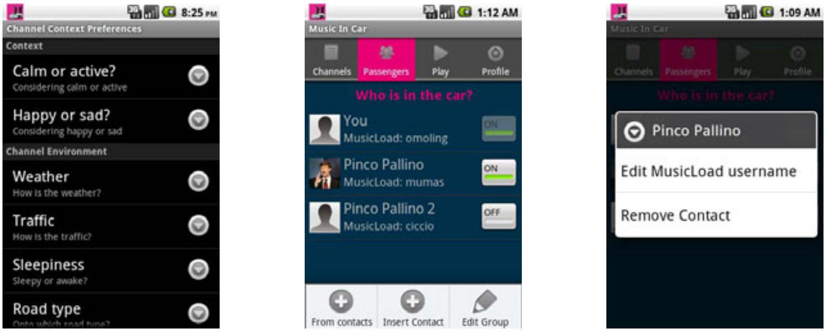

(a) Editing a Channel (b) Setting the Group (c) Editing the Passengers

Fig. 1. InCarMusic user interface

The *Play* tab allows the user to retrieve music recommendations (tracks) adapted to the selected channel and the passengers (see Fig. 2(a)). Due to lack of space, in this paper we will not explain how the recommendations are adapted to the group of passengers. For that purpose, we exploit recommendation aggregation techniques illustrated in [3]. Hence, for the rest of this paper we will consider only the scenario where a single user is present in the car. While the user is listening to a music track, she can rate it (see Fig. 2(b)). These ratings are "in-context", i.e., as we explained in the introduction, the system collects the ratings together with the description of the current contextual situation of the user. We note that these ratings are immediately uploaded to the recommender server component and can be exploited for the computation of the next recommendations.

Finally, the *Profile* tab allows the user to modify her profile and define some application settings (see Fig. 2(c)). In particular, the user can set her current contextual situation and current music genre preferences (see Fig. 2(d)). These settings are used in the default channel, if the user has not selected a particular channel. We note that this last interface is pretty similar to that used for channel configuration (see Fig. 1(a)), as the operation is the same: here the user is just configuring a particular channel, the default one.

4 Rating Acquisition

In order to offer the service described in the previous section we collected the users' assessment of the effect of context on their music preferences using two web applications that are described here . In fact, there was no ready-to-use application for collecting ratings from car drivers and other passengers while in the car. As any effort to record these conditions during a trip in a car was considered not easily solvable, we developed two web-based tools, which were

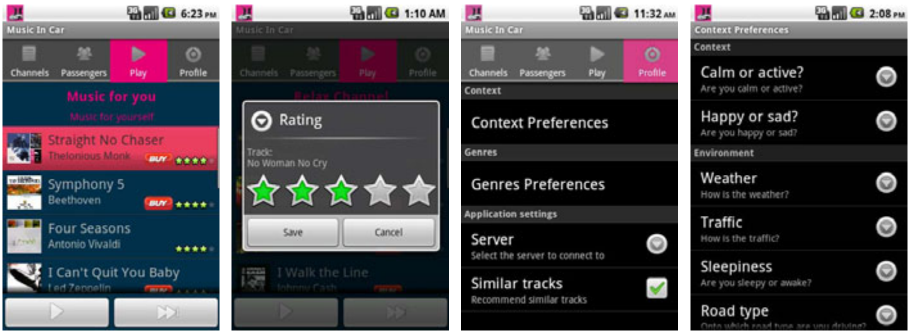

(a) Tracks Proposed to Play (b) Rating a Track (c) Editing the User Profile (d) Configuring the Recommender

Fig. 2. InCarMusic user interface (cont)

used in two consecutive phases, for simulating situations occurring in a car. In the first phase, the users were asked to evaluate the effect of certain contextual conditions on the propensity to listen to music of a particular genre, while in the second phase the users entered ratings for tracks assuming that certain contextual conditions hold (see below for more details).

4.1 Context Model and Music Track Corpus

In order to understand the influence of context on the music preferences of car passengers, context was modeled as a set of independent contextual factors. The factors are assumed to be independent in order to get a tractable mathematical model. This assumption, even if it is clearly false, as in other probabilistic models such as the naive Bayes classifier, still does not prevent the generation of reliable rating predictions. We identified the following factors and their possible values, *contextual conditions*, as potentially relevant for in car music recommendations:

Contextual Factor	**Contextual Conditions**
driving style (DS)	relaxed driving, sport driving
road type(RT)	city, highway, serpentine
landscape (L)	coast line, country side, mountains/hills, urban
sleepiness (S)	awake, sleepy
traffic conditions (TC)	free road, many cars, traffic jam
mood (M)	active, happy, lazy, sad
weather (W)	cloudy, snowing, sunny, rainy
natural phenomena (NP)	day time, morning, night, afternoon

Music tracks were of ten different genres. We observe that there is no unified music genre taxonomy, and we have chosen to use the genres defined in [12]: classical, country, disco, hip hop, jazz, rock, blues, reggae, pop and metal. For phase one, i.e., the relevance assessment of different contextual factors, five representative tracks per genre were manually selected. This resulted in

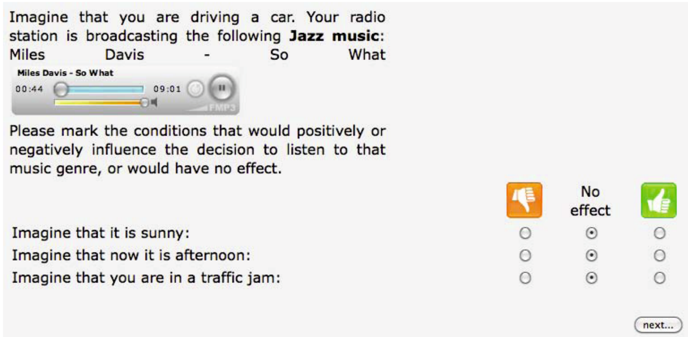

(a) Interface for Acquiring Context Relevance Judgments

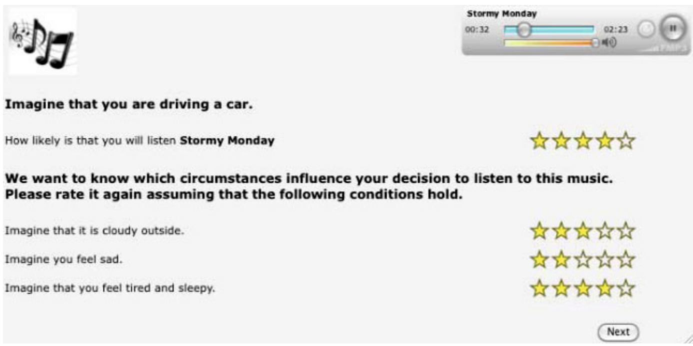

(b) Interface for Collecting Ratings with And without Context

Fig. 3. Tools for collecting user data in phase one and two

a dataset of 50 music tracks. For phase two, i.e., the assessment of the impact of contextual conditions for particular tracks, 89 additional tracks (belonging to pop, disco and hip hop genres) were added to the dataset from the MusicLoad (`http://www.musicload.de/`) download site.

4.2 Relevance of the Contextual Factors

In order to estimate the relevance of the selected contextual factors, we developed a tool for acquiring subjective judgments about the impact of these factors on the users' listening preferences. For this purpose, the users were requested to evaluate if a particular contextual condition, e.g., "today is sunny", has a positive or negative influence in her propensity to listen to music of a particular genre (see Figure 3(a)). In phase one, we acquired 2436 evaluations from 59 users with the help of our web based interview tool.

Then, for estimating the relevance of the considered contextual factors, we computed the probability distribution $P(I|F,G)$, where I is the random (response) variable of the user's answer (one out of $+1$ "increase", -1 "decrease", or 0 "no effect"), F is a contextual factor (the value of this random variable may

Table 1. Relevance of contextual factors rel(I, F, G) for different music genres

Blues music	Classical music	Country music	Disco music	Hip Hop music
DS 0.324193188	DS 0.77439747	S 0.469360938	M 0.177643232	TC 0.192705142
RT 0.216609802	S 0.209061123	DS 0.363527911	W 0.17086365	M 0.151120854
S 0.144555483	W 0.090901095	W 0.185619311	S 0.147782999	S 0.105843345
TC 0.118108963	NP 0.090509983	M 0.126974621	TC 0.129319405	NP 0.105765981
NP 0.112002402	M 0.088905397	L 0.112531867	DS 0.098158779	W 0.066024976
L 0.107824176	L 0.055675749	RT 0.109261318	RT 0.057335072	L 0.049526929
W 0.085346042	RT 0.020526969	TC 0.098999258	NP 0.049819373	DS 0.047180469
M 0.063156392	TC 0.015991764	NP 0.037183774	L 0.048588262	RT 0.01483038

Jazz music	Metal music	Pop music	Reggae music	Rock music
S 0.168519565	DS 0.462220717	S 0.418648658	S 0.549730059	TC 0.238140493
RT 0.127974728	W 0.264904662	DS 0.344360938	DS 0.382254696	S 0.224814184
W 0.106333439	S 0.196577939	RT 0.268688459	TC 0.321430505	DS 0.132856064
DS 0.100983424	L 0.122791055	TC 0.233933032	M 0.167722198	L 0.111553065
NP 0.08421736	TC 0.096436983	M 0.137086672	L 0.145512313	RT 0.096436983
L 0.053389487	M 0.06953522	NP 0.098963857	W 0.131936343	M 0.087731308
TC 0.04519683	RT 0.05580976	W 0.072377398	NP 0.105242236	W 0.083079089
M 0.035043738	NP 0.046507175	L 0.051131981	RT 0.07481265	NP 0.078288105

be any of the contextual conditions assigned to this dimension – see previous section), and G is the genre of the item. The effect of F can be measured by comparing $P(I|F, G)$ with $P(I|G)$ that does not take any context into account. For this purpose, we computed the normalized mutual information rel(I, F, G) of the random variables I and F for each music genre G:

$$\mathrm{rel}(I, F, G) = \frac{H(I|G) - H(I|F, G))}{H(I|G)}$$

where $H(X)$ is the entropy of the discrete random variable X taking values from $\{1, \ldots, n\}$: $H(X) = -\sum_{i=1}^{n} P(X = i) \log(P(X = i))$. rel$(I, F, G)$ gives a measure of the relevance of the contextual factor F: the bigger this value, the greater the relevance. In Table 1, we rank the contextual factors, for each genre, according to their influence on I, as measured by rel(I, F, G). These figures indicate the contextual factors that are likely to influence a recommendation either positively of negatively. In particular, the factors F with higher rel(I, F, G) (for each genre G) are those providing more information to the knowledge of the influence variable I (representing the change of the propensity to listen to that music). But these values do not say what conditions, i.e., values of the factors, are likely to produce positive or negative influences I. To find out these conditions we searched for the values that maximize the probability to have a positive (negative) influence, i.e., the contextual conditions c_p and c_n such that: $c_p = \mathrm{argmax}_c P(I = +1|F = c)$ and $c_n = \mathrm{argmax}_c P(I = -1|F = c)$. Due to space constraints, we present, for each genre, only the two most influential contextual conditions (see Table 2). In fact, these results could be immediately used in a context-aware recommender system: given a particular contextual condition one can look in Table 2 and find the music genres, which are preferred or not (high or low probability) by the user in that condition.

Table 2. Influence of context on the driver's decision to select a certain Genre

genre	F	c_n	$P(-1\|c_n)$	c_p	$P(+1\|c_p)$
Blues	DS	sport driving	0.89	relaxed driving	0.6
	RT	serpentine	0.44	highway	0.6
Classics	DS	sport driving	0.9	relaxed driving	0.4
	S	sleepy	0.6	awake	0.33
Country music	S	sleepy	0.67	sleepy	0.11
	DS	sport driving	0.6	relaxed driving	0.67
Disco music	M	sad	0.5	happy	0.9
	W	cloudy, rainy	0.33	sunny	0.8
Hip Hop music	TC	many cars, traffic jam	0.22	free road	0.6
	M	sad	0.56	happy	0.78
Jazz music	S	sleepy	0.7	awake, sleepy	0.2
	RT	city, highway	0.4	highway	0.4
Metal music	DS	relaxed driving	0.56	sport driving	0.7
	W	snowing	0.56	cloudy	0.78
Pop music	S	sleepy	0.8	awake	0.44
	DS	relaxed driving	0.5	sport driving	0.67
Reggae music	S	sleepy	0.5	awake	0.44
	DS	sport driving	0.5	relaxed driving	0.89
Rock music	TC	traffic jam	0.8	free road, many cars	0.44
	S	sleepy	0.44	awake	0.44

4.3 The Impact of Contextual Conditions on Ratings

The aim of phase one was to find out the contextual factors that are more influential in changing the propensity of the user to listen to music of different genres. Conversely, in the second phase of our study, we were interested in individual tracks and their ratings, and we wanted to measure if there were any differences in these ratings in the two following cases: without considering any contextual condition, and under the assumption that a certain contextual condition holds. Therefore, we implemented a second web tool, where we asked the users to rate a track without assuming any particular context and also imagining three different contextual conditions (see Fig. 3(b)). The users rated the played tracks on a scale from 1 (*I do not like the track at all*) to 5 (*I like the track very much*). The contextual factors occurred in the questionnaires randomly but proportionally to their relevance as assessed in phase one.

In this phase, 66 different users rated music tracks; overall, 955 interviews (see the screenshot in Fig. 3(b)) were conducted. As in each interview three ratings in context were collected, the data consists of 955 ratings without context and 2865 ratings with context. In Table 3, we present the analysis of the collected data. We compare the average rating for all the items: rated under the assumption that the given context factor holds (*Mean with context* – MCY) and rated without assuming any contextual condition (*Mean without context* – MCN). We conducted t-tests in order to find out the contextual conditions that produce significant differences between MCN and MCY. The table illustrates that for many contextual conditions there are statistically significant differences. This illustrates that in this application context-awareness is relevant, as the user rating behavior is dependent on context. This hypothesis will be further validated in the next section.

Table 3. Influence of contextual conditions on the average rating of music tracks

Condition	ratings	p-value	MCN	MCY	Influence	Significance
- Driving style						
relaxed driving	167	0.3891	2.382876	2.275449	↓	
sport driving	165	0.3287	2.466782	2.345455	↓	
- Landscape						
coast line	119	0.6573	2.420207	2.487395	↑	
country side	118	0.02989	2.318707	2.033898	↓	*
mountains/hills	132	0.1954	2.530208	2.348485	↓	
urban	113	0.02177	2.456345	2.141593	↓	*
- Mood						
active	97	0.01333	2.552778	2.154639	↓	*
happy	96	0.5874	2.478322	2.385417	↓	
lazy	97	0.07	2.472376	2.185567	↓	.
sad	97	0.01193	2.552632	2.134021	↓	*
- Natural phenomena						
afternoon	92	0.9699	2.407186	2.413043	↑	
day time	98	0.09005	2.381215	2.132653	↓	.
morning	98	0.6298	2.559441	2.479592	↓	
night	90	0.1405	2.516224	2.777778	↑	
- Road type						
city	123	0.551	2.479029	2.398374	↓	
highway	131	0.2674	2.457348	2.618321	↓	
serpentine	127	0.07402	2.542066	2.291339	↓	.
- Sleepiness						
awake	69	0.3748	2.561437	2.739130	↑	
sleepy	80	0.0009526	2.60371	2.01250	↓	* * *
- Traffic conditions						
free road	117	0.7628	2.491131	2.538462	↑	
many cars	132	0.3846	2.530444	2.409091	↓	
traffic jam	127	1.070e-06	2.478214	1.850394	↓	* * *
- Weather						
cloudy	103	0.07966	2.647727	2.378641	↓	.
rainy	77	0.6488	2.433453	2.519481	↑	
snowing	103	0.02056	2.601759	2.252427	↓	*
sunny	97	0.6425	2.570236	2.649485	↑	

Significance: * * *: $p < 0.001$; * *: $0.001 \le p < 0.01$; *: $0.01 \le p < 0.05$; .: $0.05 \le p < 0.1$

It is also notable that in the majority of the cases, context has a negative influence on the users' ratings. This may be a consequence of the low overall rating for the music tracks that we observed in the study: for the average user who did not like the tracks, there was no context that could change this attitude. We observe however, that for single users who provided many ratings and had a more positive attitude towards the tracks we could find several contextual factors that had a positive influence on the ratings.

5 Prediction Model

The rating prediction component computes a rating prediction for all the items, while assuming that the current user context holds. The current context is partially specified by the user, using the system GUI (as we illustrated in Section 3). Then the items with the highest predicted ratings are recommended. In this section, we present this algorithm, which extends Matrix Factorization (MF), and incorporates contextual information to adapt the recommendation to the user's contextual situation.

In [6] the authors present a Matrix Factorization approach to CF that uses "baseline" parameters, i.e., additional model parameters for each user and item. They indicate the general deviation of the rating of a user for an item from the global average. So for instance, a user baseline will be positive if it refers to a user that tends to rate higher than the average users' population. Baseline parameters can also be used to take into account the impact of context. This has been already shown by [7], where the authors introduced baseline parameters to model the time dependency of the ratings.

We have extended and adapted this approach to the music domain by incorporating the selected contextual factors into the MF model. We have introduced one model parameter for each contextual condition (value for a factor) and music track genre pair. This provides an opportunity to learn how a contextual condition affects the ratings and how they deviate from the standard personalized prediction. This deviation is the *baseline* for that contextual condition and genre combination. In principle, we could introduce parameters for each contextual condition and music track, however, this would require much more data to train the model.

More formally, in the collected context-aware rating data base a rating $r_{uic_1\ldots c_k}$ indicates the evaluation of the user u for the item i made in the context $c_1, \ldots, c_k$, where $c_j \in \{0, 1, \ldots, z_j\}$ is the set of possible (index) values of the contextual factor j, and 0 means that the contextual factor j is unknown. The tuples $(u, i, c_1, \ldots, c_k)$ for which $r_{uic_1\ldots c_k}$ is known are stored in the set $R = \{(u, i, c_1, \ldots, c_k) | r_{uic_1\ldots c_k} \text{ is known}\}$. Note that in our collected data set, there are ratings where only one contextual condition is known and all others are unknown. We recall that MF aims at factorizing the ratings matrix into two $m \times d$ and $n \times d$ dimensional matrices V and Q respectively. A user is then represented with a vector $\boldsymbol{v}_u$ and an item i with a vector $\boldsymbol{q}_i$. We propose the following model for the computation of a personalized context-dependent rating estimation.

$$\hat{r}_{uic_1\ldots c_k} = \boldsymbol{v}_u \cdot \boldsymbol{q}_i + \bar{\imath} + b_u + \sum_{j=1}^{k} b_{g_i j c_j} \tag{1}$$

where $\boldsymbol{v}_u$ and $\boldsymbol{q}_i$ are d dimensional real valued vectors representing the user u and the item i. $\bar{\imath}$ is the average of the item i ratings in the data set R, $b_{g_i j c_j}$ is the baseline of the contextual condition c_j and genre g_i of item i. If a contextual factor is unknown, i.e., $c_j = 0$, then the corresponding baseline $b_{g_i j c_j}$ is set to 0. In this way, one can learn the influence only of the known contextual conditions.

Model Training. In order to generate rating predictions, the model parameters should be learned using the training data. We defined the learning procedure as an optimization problem:

$$\min_{v_*, q_*, b_*} \sum_{r \in R} [(r_{uic_1\ldots c_k} - \boldsymbol{v}_u \cdot \boldsymbol{q}_i - \bar{\imath} - \sum_{j=1}^{k} b_{g_i j c_j})^2 + \lambda(\|\boldsymbol{v}_u\|^2 + \|\boldsymbol{q}_i\|^2 + \sum_{j=1}^{k} b_{g_i j c_j}^2)]$$

where $r = (u, i, c_1, \ldots, c_k)$. For better generalization performance, a regularization term, λ, is added, as it is usual in this type of models. As λ grows the model

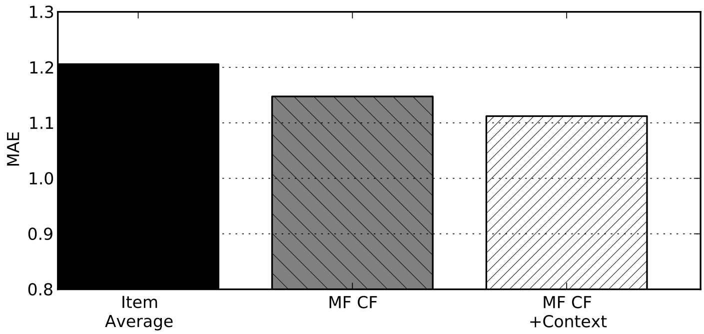

Fig. 4. Mean Absolute Error of different prediction models

becomes more "rigid", and fits less the variability in the training data. Model parameters were learned using stochastic gradient descent, which has been already proven to be efficient [6].

The Mean Absolute Error (MAE) of the considered models is shown in Figure 4. The largest improvement with respect to the non-personalized model based on the item average is achieved, as expected, by personalizing the recommendations ("MF CF" in the figure). This gives an improvement of 5%. However, the personalized model can be further improved by contextualization ("MF CF + Context") producing an improvement of 7% with respect to the item average prediction, and a 3% improvement over the personalized model. We conclude that the modeling approach and the rating acquisition process described in the previous sections can substantially improve the rating prediction accuracy when taking into account the contextual information.

6 Conclusions

In this paper we have illustrated a methodology for acquiring subjective evaluations about the relevance and the impact of certain contextual conditions on the ratings for music tracks. We have shown that using this approach a useful and effective set of ratings can be collected and a context-aware recommender system can be bootstrapped. The off-line evaluation of the predictive model, which extends Matrix Factorization (MF), has shown that it can substantially improve a non-personalized prediction, but also a classical personalized prediction based on MF, hence showing the practical advantage of the proposed approach. The mobile application that we have developed can offer context-aware and personalized music recommendations to users in a car scenario.

In the future we plan to perform a field study to validate the usability of the prototype and to incorporate a technique for extrapolating the item ratings from user actions on the items; e.g., listening to a track for a certain time in a contextual situation may be interpreted as a graded sign that this context is suited for the track. The challenge here is to filter noisy signs and build a reliable predictive model of the rating by using the user actions as predictive features.

References

1. Adomavicius, G., Tuzhilin, A.: Context-aware recommender systems. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P. (eds.) Recommender Systems Handbook, pp. 217–256. Springer, Heidelberg (2011)
2. Baltrunas, L., Ludwig, B., Peer, S., Ricci, F.: Context relevance assessment and exploitation in mobile recommender systems. Personal and Ubiquitous Computing (to appear, 2011)
3. Baltrunas, L., Makcinskas, T., Ricci, F.: Group recommendations with rank aggregation and collaborative filtering. In: RecSys 2010: Proceedings of the 2010 ACM Conference on Recommender Systems, pp. 119–126 (2010)
4. Cai, R., Zhang, C., Wang, C., Zhang, L., Ma, W.Y.: Musicsense: contextual music recommendation using emotional allocation modeling. In: MULTIMEDIA 2007: Proceedings of the 15th International Conference on Multimedia, pp. 553–556. ACM, New York (2007)
5. Jameson, A., Smyth, B.: Recommendation to groups. In: Brusilovsky, P., Kobsa, A., Nejdl, W. (eds.) Adaptive Web 2007. LNCS, vol. 4321, pp. 596–627. Springer, Heidelberg (2007)
6. Koren, Y.: Factorization meets the neighborhood: a multifaceted collaborative filtering model. In: KDD 2008: Proceeding of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 426–434. ACM, New York (2008)
7. Koren, Y.: Collaborative filtering with temporal dynamics. In: KDD 2009: Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 447–456. ACM, New York (2009)
8. Lee, J.S., Lee, J.C.: Context awareness by case-based reasoning in a music recommendation system. In: Ichikawa, H., Cho, W.-D., Chen, Y., Youn, H.Y. (eds.) UCS 2007. LNCS, vol. 4836, pp. 45–58. Springer, Heidelberg (2007)
9. Li, C.T., Shan, M.K.: Emotion-based impressionism slideshow with automatic music accompaniment. In: MULTIMEDIA 2007: Proceedings of the 15th International Conference on Multimedia, pp. 839–842. ACM Press, New York (2007)
10. Reddy, S., Mascia, J.: Lifetrak: music in tune with your life. In: HCM 2006: Proceedings of the 1st ACM International Workshop on Human-Centered Multimedia, New York, NY, USA, pp. 25–34 (2006)
11. Ricci, F., Rokach, L., Shapira, B.: Introduction to recommender systems handbook. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P. (eds.) Recommender Systems Handbook, pp. 1–35. Springer, Heidelberg (2011)
12. Tzanetakis, G., Cook, P.: Musical genre classification of audio signals. IEEE Transactions on Speech and Audio Processing 10(5), 293–302 (2002)