ELEKTROTEHNIŠKI VESTNIK 78(4): 177–180, 2011
ENGLISH EDITION

Emotional properties of latent factors in an image recommender system

Marko Tkalčič, Andrej Košir, Štefan Dobravec, Jurij Tasič

Univerza v Ljubljani, Fakulteta za elektrotehniko, Tržaška 25, 1000 Ljubljana, Slovenija
E-pošta: marko.tkalcic@fe.uni-lj.si

Abstract. In this paper we analyze the relations between the latent factors with high variance description and affective parameters in an image recommender system. Using the matrix factorization approach we identify the main two factors in the user-item rating database. We exploit the affective metadata related to each item to identify relations between the main factors and the affective metadata. Results show that the first latent factor is strongly related with the valence and dominance while the arousal does not appear to be related. The second factor, however, shows no relation with the affective parameters.

Keywords: recommender systems, affective computing, matrix factorization, latent factors

1 INTRODUCTION

Recommender systems for multimedia content (e.g. films, images, music, books, etc.) are systems that exploit the knowledge about an observed user's preferences (the user profile) and the knowledge about the multimedia item properties (the item profile) to filter a limited set of multimedia content suited to the user's tastes. Recommender systems have outgrown the laboratory environment and found their place in commercial applications. Amazon, for example, has been using their solution to recommend shopping items on their online store [5]. Research work has been following mainly two paths: (i) the algorithmic path, where the existing algorithms have been improved, or new algorithms have been proposed, and (ii) the descriptor path, where new features have been sought in order to increase the variance.

The first algorithms were very simple content-based recommenders (CBR) [6] and collaborative filtering recommenders (CF) [1]. However, during the Netflix competition (http://www.netflixprize.com/), the matrix factorization approach turned out to outperform other approaches in recommender systems [2], [3]. So far, the matrix factorization approach is regarded as the best approach when the training dataset is big enough.

On the other hand, researchers were looking out for new features to include in the user and item profiles that would carry more information for predicting relevant content to recommend. The first features used were generic metadata like the genre, actors, director, etc. [1], [7]. As the algorithms were improving, they were able to get the maximum information from these metadata and researchers started exploring new kinds of features for improving the performance of recommender systems. In our previous work we identified emotional metadata to describe considerable variance in user's data thus improving the performance of the CBR system [10].

Received August 26, 2011
Accepted September 25, 2011

The introduction of the matrix factorization approach in recommender system replaced the former human understandable features (e.g. genre, etc.) with latent features, which are not necessarily human understandable. The first two latent factors in the Netflix dataset turned out to be explained as intellectual-shalllow and masculine-feminine, respectively [3], [2]. This line of reasoning, the interpretation of the properties of the main latent factors, is interesting for further research because it opens different perspectives on users' preferences and their modeling.

1.1 Problem statement

The goal of this paper is to explore the properties of the main latent factors of a recommender system dataset in terms of the emotions that the content items induce in end users. The problem adressed in this paper arises from two presumptions: (i) the matrix factorization algorithm identifies the main latent factors that describe the variance in the user-item rating matrix and (ii) the items' affective parameters (the parameters that describe the emotion that an item induces in the user) vary along the main factors' axes. Thus we aim at using the explanatory factor analysis to identify the affective properties that characterize the items at the ends of the main latent factors, as we depicted in Fig. 1. We denote the groups of items as the groups $G_{1.1}, G_{1.2}, G_{2.1}$ and $G_{2.2}$. The result is the visualization and interpretaion of these properties.

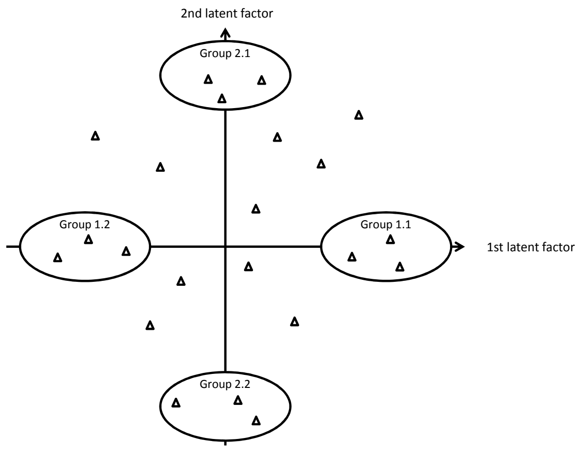

Figure 1. The items (denoted by the triangles) are scattered in the space of the latent factors. The groups of items $G_{1.1}, G_{1.2}, G_{2.1}$ and $G_{2.2}$ represent the items that are at the extremes of the first two latent factors.

2 EXPERIMENT

The flow of the experiment is shown in Fig. 2. In the dataset acquisition phase with real users we built the user-item rating matrix and we performed the matrix factorization algorithm providing the main latent factors of the dataset. In the next step we computed the basic statistics of the items lieing at the ends of the first two factors.

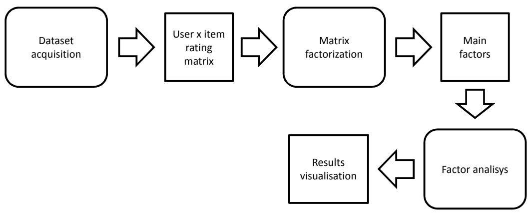

Figure 2. Flow of the experiment

2.1 Dataset acquisition

We built a full items-users rating matrix (depicted in Tab. 1) in the dataset acquisition session. In this session we had 52 users, denoted with $u \in \{1 \dots 52\}$, who rated 70 images, denoted with $i \in \{1 \dots 70\}$. We denoted the rating given by the user u to the item i as $r(u,i) \in \{1 \dots 5\}$. The images were selected from a larger set of images, the IAPS dataset [4]. Each image was annotated with metadata that described the induced emotion in end users. The parameters used to describe a single emotive response were the valence, arousal and dominance parameters, which are frequently used [8]. The valence describes whether an emotion is positive (e.g. happy) or negative (e.g. sad). The arousal describes the strength of the observed emotion. The dominance describes whether the user is in control of the emotion (e.g. calm) or not (e.g. strong anger). Thus, each image i was annotated with the average of the valence, arousal and dominance over a set of users that consumed the images, which we denoted as the triple $md(i) = \{\bar{v}, \bar{a}, \bar{d}\}$. Please note that the set of users used to compute the averages were different than the set of users that assigned ratings to the images. Further information on the dataset used is given in [11].

u \ i	1	2	3	...	70
1	$r(1,1)$	$r(1,2)$	$r(1,3)$		$r(1,70)$
2	$r(2,1)$	$r(2,2)$	$r(2,3)$		$r(2,70)$
3	$r(3,1)$	$r(3,2)$	$r(3,3)$		$r(3,70)$
...					
52	$r(52,1)$	$r(52,2)$	$r(52,3)$		$r(52,70)$

Table 1. The user-item rating matrix: contains the ratings $r(u,i) \in \{1 \dots 5\}$ given by the user $u \in \{1 \dots 52\}$ to the item $i \in \{1 \dots 70\}$.

2.2 Matrix factorization

The matrix factorization approach aims at identifying a space of latent factors that can be used to describe the input data in a more compact way. In our case the input data was the user-item rating matrix (see Tab. 1). Usually, the preferred method is singular value decomposition (SVD), but in recommender systems there are two reasons why it is not usable: (i) the sparsity of the data and (ii) the size of the input space (lots of users and lots of items). Thus we calculated the main factors using the stochastic gradient descent algorithm, which is frequently used [2], [3], [9]. When using matrix factorization, each user is described with a user vector of latent factors p_u and each item with an item vector of latent factors q_i.

The user profile, which describes the user's preferences, is now in the form of a vector of latent features p_u. The values in this vector reflect the degree of how much the user likes the respective latent property. For example, in the Netflix dataset, the first latent factor appears to be intellectual-shallow. A high value in the first factor of the user profile would thus mean that the observed user likes intellectual films.

Following the same reasoning, the values in the item profile q_i reflect the degree of how much the observed item contains the specific latent property. For example, again in the Netflix dataset, a high value in the first factor of the item profile q_i means that the observed item has the quality of being intellectual.

The vectors have the same length F, which is the number of the factors used. The goal of a recommender system is to predict the scalar rating that an observed user u would give to a specific item i. The rating

group	Number of items
$G_{1.1}$	31
$G_{1.2}$	13
$G_{2.1}$	15
$G_{2.2}$	24

Table 2. Number of items in each of the four observed groups

predicition $\hat{r}(u,i)$ is calculated as a scalar product of both vectors $\hat{r}(u,i) = p_u \cdot q_i^T$. The stochastic gradient algorithm learns the feature values in the vectors p ang q by minimizing the root mean sqare error (RMSE) between the real rating $r(u,i)$ and the predicted rating $\hat{r}(u,i)$:

$$argmin \sum_{u,i} (r(u,i) - p_u \cdot q_i^T) + \lambda(\|q_i\|^2 + \|p_u\|^2) \quad (1)$$

where λ is the regularization factor for limiting over-fitting. We performed the stochastic gradient descent algorithm using several learning iterations.

We observed the performance of the predicted ratings in terms of the RMSE at various numbers of features $F \in \{1 \dots 300\}$ and various numbers of learning epochs $e \in \{1 \dots 500\}$.

For selecting the first two factors we used two factors $F = 2$ and 2000 learning iterations. Thus each user and item were characterized with a vector of two features from which we calculated the predicted ratings

$$\hat{r}(u,i) = (q_{i,1}, q_{i,2}) \cdot (p_{u,1}, p_{u,2}) \quad (2)$$

where the features $q_{i,1}, q_{i,2}, p_{u,1}$ and $p_{u,2}$ represent the values of the first two latent factors of the user u and item i.

2.3 *Group members selection*

For each group of items, we chose those that had the respective feature above 70% of the maximum factor value for the groups 11 and 21, and less than 30% of the maximum factor value for the groups 12 and 22. This approach yielded different amounts of items in the groups, which we show in Tab. 2.

3 Results

First, as an interesting side result, we observed the performance of the matrix factorization algorithm as a function of the number of features F (see Fig. 3) and the number of learning iterations e (see Fig. 4). Ten-fold cross validation was used.

Figs. 5 to 7 show the distributions of the emotive parameters valence, arousal and dominance, respectively, in the four groups in the form of boxplots. On each boxplot, the central mark is the median, the edges of the box are the 25th and 75th percentiles, the whiskers extend to the most extreme data points not considered outliers, and outliers are plotted individually.

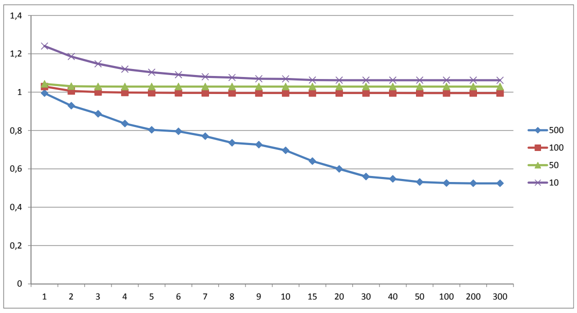

Figure 3. RMSE (y axis) as a function of the number of features F (x axis) at different numbers of learning epochs $e \in \{10, 50, 100, 500\}$

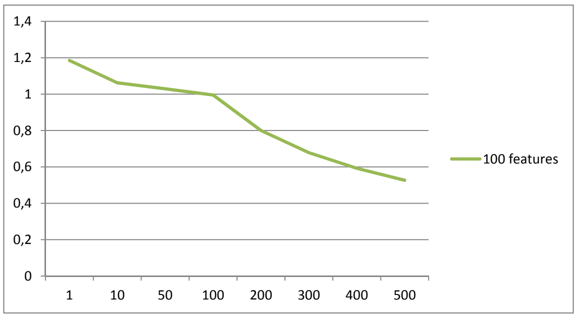

Figure 4. RMSE (y axis) as a function of the number of epochs e (x axis) at a fixed number of features ($F = 100$)

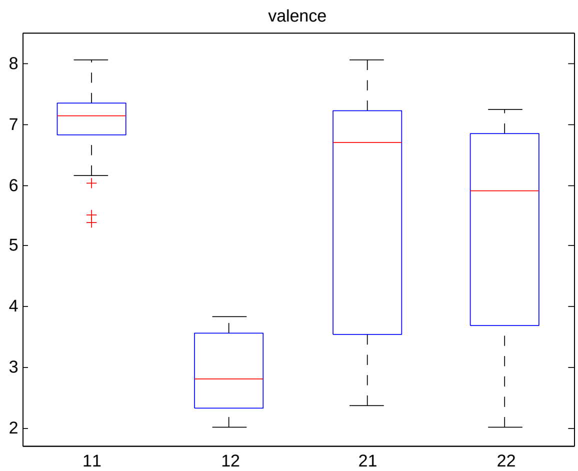

Figure 5. Distributions of the parameter valence in the four groups.

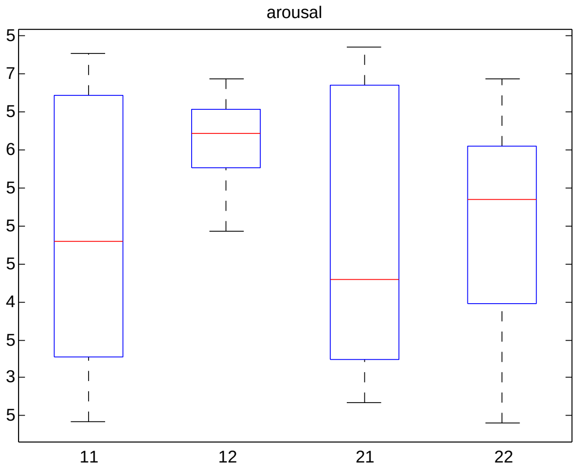

Figure 6. Distributions of the parameter arousal in the four groups.

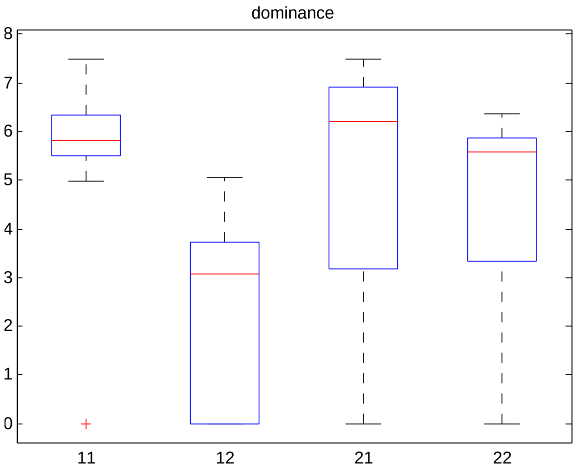

Figure 7. Distributions of the parameter dominance in the four groups.

4 Discussion and conclusion

The experimental results indicate that the valence and dominance are related to the first latent factor that accounts for the majority of the variance in users' ratings. Furthermore, the arousal is very high in the group $G_{1.2}$, but equally distributed in the group $G_{1.1}$. On the other hand, no visible relations were found between the second latent factor (groups $G_{2.1}$ and $G_{2.2}$) and the affective parameters.

In this paper we identified the affective properties of the items lieing at the extremes of the first two latent factors that account for the majority of the variance in the users' ratings.

Acknowledgements

The work has been supported by the Slovenian Research Agency http://arrs.gov.si under the contract P2-0246. The authors would like to thank the students and the staff of the Gimnazija Poljane for their involvement in the dataset acquisition and then the colleagues from the LDOS group for helping in the experimental part.

References

[1] Gediminas Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749, June 2005.

[2] Simon Funk. Netflix Update: Try This at Home, 2006.

[3] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix Factorization Techniques for Recommender Systems. *Computer*, 42(8):30–37, August 2009.

[4] Peter J Lang, M M Bradley, and B N Cuthbert. International affective picture system (IAPS): Affective ratings of pictures and instruction manual. Technical Report A-8. Technical report, University of Florida, 2005.

[5] G. Linden, B. Smith, and J. York. Amazon.com recommendations: item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1):76–80, January 2003.

[6] Matevz Pogacnik, Jurij Tasic, Marko Meza, and Andrej Kosir. Personal Content Recommender Based on a Hierarchical User Model for the Selection of TV Programmes. *User Modeling and User-Adapted Interaction: The Journal of Personalization Research*, 15(5):425–457, 2005.

[7] Matevz Pogacnik, Jurij Tasič, and Andrej Košir. Optimization of Multi-attribute User Modeling Approach. *AEU - International Journal of Electronics and Communications*, 58(6):402–412, 2004.

[8] Jonathan Posner, James a Russell, and Bradley S Peterson. The circumplex model of affect: an integrative approach to affective neuroscience, cognitive development, and psychopathology. *Development and psychopathology*, 17(3):715–34, January 2005.

[9] Yue Shi, Martha Larson, and Alan Hanjalic. Mining mood-specific movie similarity with matrix factorization for context-aware recommendation. *Proceedings of the Workshop on Context-Aware Movie Recommendation*, pages 34–40, 2010.

[10] Marko Tkalčič, Urban Burnik, and Andrej Košir. Using affective parameters in a content-based recommender system for images. *User Modeling and User-Adapted Interaction: The Journal of Personalization Research*, pages 1–33–33, September 2010.

[11] Marko Tkalčič, J Tasič, and Andrej Košir. The LDOS-PerAff-1 Corpus of Face Video Clips with Affective and Personality Metadata. *Proceedings of Multimodal Corpora: Advances in Capturing, Coding and Analyzing Multimodality (Malta, 2010), LREC*, page 111, 2009.

Marko Tkalčič is a research scientist at the Faculty of electrical engineering of the University of Ljubljana . His research areas include affective computing, user modeling, recommender systems and human computer interaction.

Andrej Košir is a professor at the Faculty of electrical engineering of the University of Ljubljana.

Štefan Dobravec is a research scientist at the Faculty of electrical engineering of the University of Ljubljana.

Jurij Tasič is a professor at the Faculty of electrical engineering of the University of Ljubljana.