Towards a Generic Architecture for Recommenders Benchmarking

Mohamed Ramzi Haddad[1], Hajer Baazaoui[1], Djemel Ziou[2] and Henda Ben Ghezala[1]
[1]Riadi-Gdl Laboratory, École Nationale des Sciences de l'Informatique, Université de la Manouba, Manouba 2010, Tunisia
[2]Centre de Recherche MoIVRe, Département d'Informatique, Faculté des Sciences, Université de Sherbrooke
2500 Boul. Université, Sherbrooke, J1K 2R1, Canada

Keywords: Hybrid Recommender Systems, User Modeling, Consumption Behaviors Prediction, Benchmarking.

Abstract: With current growth of internet sales and content consumption, more research efforts are focusing on developing recommendation and personalization algorithms as a solution for the choice overload problem. In this paper, we first enumerate several state-of-the-art recommendation algorithms in order to highlight their main ideas and methodologies. Then, we propose a generic architecture for recommender systems benchmarking. Using the proposed architecture, we implement and evaluate several variants of existing recommendation algorithms and compare their results to our unified recommendation model. The experiments are conducted on a real world dataset in order to assess the genericity of our recommendation model and its quality. At the end, we conclude with some ideas for further development and research.

1 INTRODUCTION

As e-commerce and content delivery businesses are increasingly popular on the web, it became necessary to address the information overload problem and provide an intelligent and personalized access to the available goods and content. Several social and consumption psychology studies such as (Schwartz, 2005), (Jacoby et al., 1974) and (Iyengar and Lepper, 2000) confirmed the existence of an overchoice problem (also known as choice overload) and stressed the need to reduce customers' choices to facilitate the decision making.

Scientific and industrial communities proposed several tools to address this challenge such as personalization and recommendation systems. The goal of a recommender system is to propose a set of interesting items for a user based on heuristics or on acquired knowledge. Suggestions of books on Amazon and movies on Netflix are real world examples of recommender systems. Each recommendation algorithm is based on a set of assumptions about consumption behaviors in order to predict individuals' interests and future purchases. Therefore, system designers may rely on hybridization of several recommendation algorithms in order to take into consideration all the facts that drive consumers' purchase decisions and interests. In this context, we believe that it may be more relevant to propose a new class of multi-faceted recommendation algorithms, capable of adapting their recommendations based on the domain-driven consumption behaviors.

In this paper, we present a new generic contextual recommendation model encompassing the main ideas and the hypothesis of state of the art algorithms. Such model would be able to describe and predict more precisely consumers interests and purchases by using all the available data in the targeted field. Besides, we develop a generic recommendation architecture that is used to implement, study and assess the quality state-of-the-art algorithms by integrating all the necessary components for content recommendation. The proposed recommendation model is also implemented within the architecture in order to evaluate its performances.

This paper is organized as follows. Next, in Sect. 2, we present the major state of the art of recommendation approaches and describe briefely our previous works on statistical context-aware recommendation models. Afterwards, in Sect. 3 we present a generic architecture for recommendation algorithms benchmarking. Finally, in Sect. 4, the experimentation context and the obtained results are presented and discussed. The paper is concluded by summarizing our proposition and presenting some future research perspectives.

Ramzi Haddad M., Baazaoui H., Ziou D. and Ben Ghezala H..
Towards a Generic Architecture for Recommenders Benchmarking.
DOI: 10.5220/0005216904350442
In *Proceedings of the International Conference on Agents and Artificial Intelligence* (ICAART-2015), pages 435-442
ISBN: 978-989-758-074-1

2 RECOMMENDATION ALGORITHMS OVERVIEW AND MOTIVATIONS

Recommender systems are a class of personalization systems whose main objective is to predict users' interests towards the available informational content in the application domain. To achieve this goal, several approaches and methodologies were proposed in the literature.

Collaborative filtering (CF) is widely adopted in e-commerce and is based on the assumption that users who, in the past, had the same attitudes towards the same items would eventually agree in the future (Goldberg et al., 1992). In user-based CF, as detailed in (Herlocker et al., 1999), current user's missing ratings are predicted by aggregating the ratings of a neighborhood of similar users. Due to a lack of scalability of this approach, item-to-item CF was proposed. Hereby, instead of matching similar users, the algorithm matches a user's rated items to similar items. Item-based CF approaches, as proposed in (Linden et al., 2003) and (Deshpande and Karypis, 2004), are based on the assumption that users prefer items that are rated similarly, or correlated to the items they already know and like. In practice, item-based CF leads to faster systems and delivers a better recommendation quality (cf. (Sarwar et al., 2001)).

Several researches consented to modeling item's features in the recommendation data model which led to content-based filtering (CBF) approaches. CBF focuses on eliciting users' preferences towards items' features in order to use them to score unknown items. Thus, in (Mooney and Roy, 2000), (Balabanović and Shoham, 1997) and (Pazzani, 1999), the recommended items are the ones that have the most interesting features regarding the user past interactions. However, users' preferences can be augmented by their likeminded neighbors' preferences in a collaborative filtering manner (Berkovsky et al., 2008).

Demographic filtering (DF) techniques adopt a generalized user and item representations. In fact, users are described by a set of demographic attributes for a better handling of users similarity. Moreover, items may be described by their features as in CBF approaches. DF generalizes CF and CBF and may then reuse their recommendation generation methodologies (Krulwich, 1997; Pazzani, 1999).

Several researches on recommender systems pointed out the importance of the context in individuals' choices and perceived relevancy. This led to a new class of context-aware recommendation approaches that are able to identify cases where the context (mainly time and location) implies some common consumption behaviors and may then respond with more accurate recommendations (Woerndl and Groh, 2007; Boutemedjet and Ziou, 2008).

Finally, several recent researches focus on hybrid recommendation approaches. Recommender systems hybridization relies on the assumption that the aggregation of several recommendation techniques improves their efficiency and helps overcoming their shortcomings. In (Burke, 2007), the author argues hybrid recommender systems and defines several hybridization strategies.

In previous works, we proposed a new unified contextual recommendation approach based on a statistical model of consumers behaviors (Haddad et al., 2012). The model predicts users purchases and interests based on a set of factors issued from consumers psychology researches without being restricted to any of the underlying assumptions of existing approaches. To achieve this goal using the proposed model, we compute the probability of observing a user u assigning a rating e to an item x in the context q and then recommend the top N ones with the highest expected rating (Haddad et al., 2012). The proposed model unifies the main ideas of existing recommendation approaches and is able to dynamically select the most appropriate inference technique based on probabilities learned from existing data.

Based on the study of consumer psychology researches and existing recommendation approaches, we believe that system designers would benefit from a generic architecture integrating all the necessary component for content recommendation. Our work is motivated by the fact that such architecture would make it possible to benchmark existing recommendation algorithms and develop more generic ones that encompass all the relevant ideas, methodologies and practices of the application domain. In this context, the originality of this work consists on the development of a generic architecture for content recommendation and on the benchmarking of existing algorithms in order to validate our unified contextual recommendation model.

3 A GENERIC ARCHITECTURE FOR RECOMMENDERS BENCHMARKING

In this work, we propose a generic architecture for recommender systems benchmarking in order to evaluate existing algorithms and compare their results to our unified recommendation model. In this section, first the main components of our architecture are de-

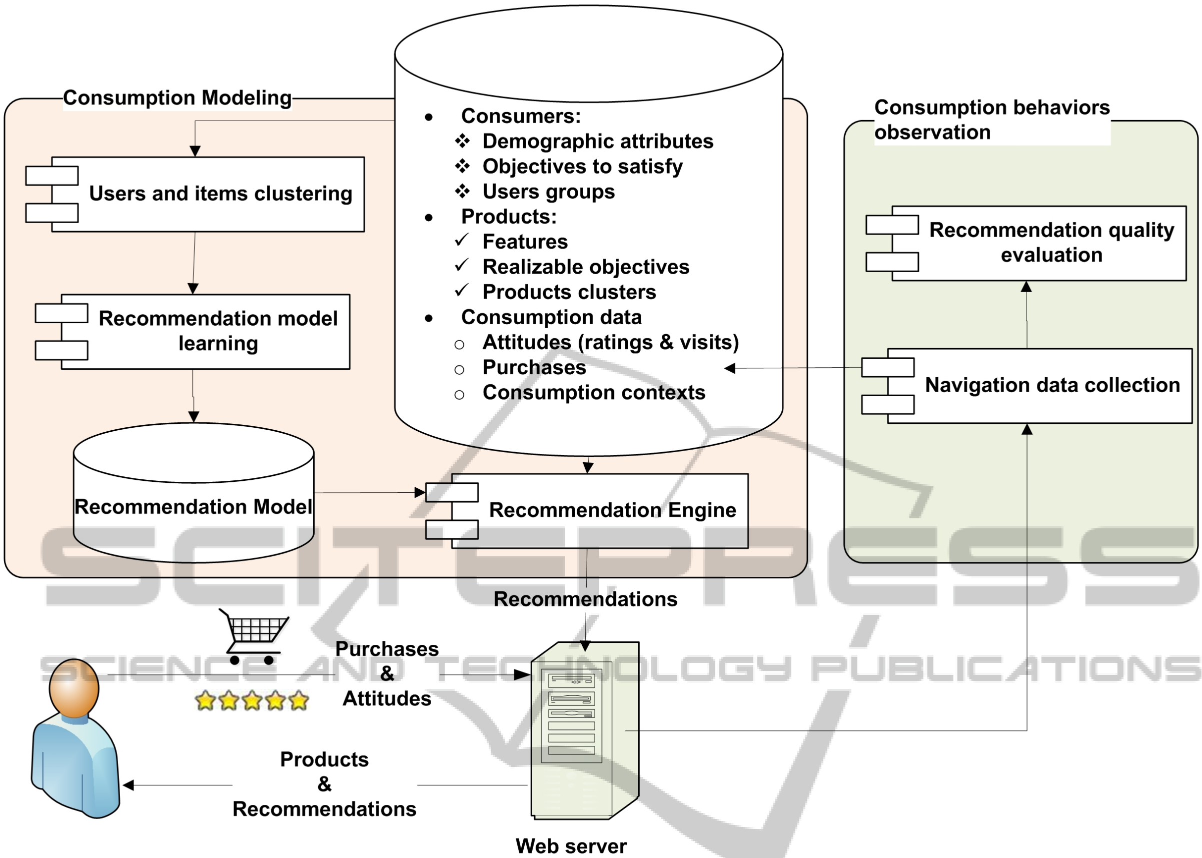

Figure 1: Generic recommendation architecture.

tailed (c.f. figure 1). Then, the implemented algorithms and their different variants are enumerated.

3.1 Architecture Components

Figure 1, represents the components of the proposed architecture which are the following :

1. Data repository : regroups the raw data on which the recommendation process is based (e.g. users' demographics, products features, ratings, etc...).
2. Clustering algorithms: this component regroups a set of clustering algorithms such as k-means (Kaufman and Rousseeuw, 1990) and c-means (Bezdek, 1981) for users and items clustering.
3. Recommendation models: this component acts as a repository for recommendation models that are generated by model-based approaches and used to infer recommendations.
4. Recommendation engines: are responsible for recommendations generation either from raw data or from previously learned model. In our case, recommendation engines are used to infer users' ratings in order to recommend the items that are the most likely to interest the active user.
5. Data collection processes: is responsible for collecting users' navigation data which may indicate their interests (e.g. ratings, purchases, consultations, etc...). Such data needs to be captured and stored in order to be used as an input for the recommendation algorithms.
6. Quality measures: several quality measures may be used in order to evaluate and compare the recommendation algorithms. In this work, the Mean Absolute Error (MAE) and the Root Mean Square Error (RMSE) measures were implemented to evaluate the quality of the predicted ratings.

3.2 Implemented Algorithms

In this work, four recommendation approaches were implemented and compared to our model. Hereby, the objective is to predict the rating r_{ac} that an active user u_a would assign to a candidate item x_c. The implemented algorithms are the following:

3.2.1 User-centered Collaborative Filtering (UC-CF)

UC-CF as detailed in (Herlocker et al., 1999) and (Resnick et al., 1994) is based the assumption that users with similar preferences will rate products similarly. Thus, the implemented variants of UC-CF undergo the following stages:

- Compute $USim(u_a, u_i) = USim(\vec{r_a}$ and $\vec{r_i})$, the similarities between the active user u_a and all the other users u_i based on their common ratings $\vec{r_a}$ and $\vec{r_i}$ assigned to the same items.
- Select S_{u_a}, the set of the k most similar users to the active user u_a.
- Estimate $\bar{r}_{ac}$ using a rating estimator such as the mean or the weighted mean (c.f. equations 2 and 3) on the set of ratings assigned by S_{u_a} to x_c.

3.2.2 Item-centered Collaborative Filtering (IC-CF)

IC-CF approaches such as proposed in (Sarwar et al., 2001) and (Deshpande and Karypis, 2004) are based on the assumption that users prefer items that are correlated to the items they already know and like. Thus, all the implemented IC-CF algorithms persue the following generic pattern:

- Compute $ISim(x_c, x_i) = ISim(\vec{r_c}, \vec{r_i})$, the similarities between the candidate item x_c and all the other items x_i based on the ratings $\vec{r_c}$ and $\vec{r_i}$ they were assigned by the same users.
- Select S_{x_c}, the set of the k most similar items to the candidate item x_c. S_{x_c} is also referred to as the neighborhood of x_c.
- Estimate $\bar{r}_{ac}$ by aggregating the ratings assigned by u_a to the items in S_{x_c} using a rating estimator such as the mean and the weighted mean (c.f. equations 2 and 3).

3.2.3 Content Based Filtering (CBF)

In CBF, as presented in (Mooney and Roy, 2000), (Balabanović and Shoham, 1997) and (Pazzani, 1999), the recommended items are those having similar features to the ones that the user have already liked or purchased. Consequently, the implemented CBF variants persue the following generic algorithm:

- Compute features similarities $ISim(x_c, x_i)$ between the candidate item x_c and all the other items x_i based on their sets of features $\vec{f}_i$.
- Select S_{x_c}, the neighborhood of the k most similar items to the candidate item x_c.
- Estimate $\hat{r}_{ac}$ based on the ratings assigned by u_a to items in S_{x_c} (c.f. equations 2 and 3).

3.2.4 Demographic Filtering (DF)

DF relies on predicting ratings based on both users' demographic and items' features similarities. Several variants of DF are detailed in (Krulwich, 1997) and (Pazzani, 1999). The general pattern of the evaluated DF variants is as follows:

- Compute features similarities $ISim(x_c, x_i)$ between x_c and all the other items x_i as in CBF.
- Select S_{x_c}, the set of the k most similar items to the candidate item x_c.
- Compute demographic similarities $USim(u_a, u_i)$ between the active user u_a and all the other users u_i based on their demographic attributes $\vec{d}_i$.
- Select S_{u_a}, the set of the L most demographically similar users to u_a.
- Estimate $\bar{r}_{ac}$ by aggregating the ratings assigned by users in S_{u_a} to the items in S_{x_c}.

3.2.5 The Proposed Frequentist Model (FM)

Our recommendation model predicts consumers' interests and purchases using a statistical methodology based on a set of variables that can be easily collected in e-commerce platforms (Haddad et al., 2012). First, users and items are clustered into groups and categories. Then, the probability $p(e_k|u,x,q)$ of observing an evaluation (i.e. rating) $e_k \in \{e_1, e_2 \ldots, e_{N_e}\}$ being assigned by a user u belonging to the group g_u to an item x of the category c_x in the context q is calculated for each possible value of the rating variable e. Finally, the ratings probabilities are used as an input to estimate the rating $\hat{r}_{ux}$ that the user u would assign to an unknown item x. In our model, we assume that the users' demographics and behaviors are induced by their respective groups and that items' attributes depend only on their categories. Besides, we assume that users' interests modeled by the rating variable e are induced by their groups, the products' categories and the consumption context. Those hypothesis define the conditional dependencies (and independencies) between our model's variables and enable us to develop and simplify the rating probability as follows:

$$p(e_k|u,x,q) = p(e_k|g_u, c_x, q)p(g_u|u)p(c_x|x) \quad (1)$$

$p(c_i|x)$ (resp. $p(g_j|u)$) represent the membership degree of an item x (resp. a user u) to the category

c_i (resp. to the group g_j). $p(c_i|x)$ and $p(g_j|u)$ are related to users and items clustering. However, our first experiments using fuzzy clustering have shown that such probabilistic clustering approaches decrease the model's prediction quality (Haddad et al., 2012). Besides, the term $p(e_k|g_u,c_x,q)$ is closely related to collaborative, content based and demographic filtering approaches and unifies their main ideas while adding the contextual aspect.

3.3 Recommendation Engines

Several ratings estimators were implemented for the recommendation algorithms in order to estimate the rating r_{ac} that the active user u_a would assign to a candidate item x_c.

- Mean: let N be the number of existing ratings $r_{i,j}$ such as $u_i \in S_{u_a} and x_i \in S_{x_c}$. The mean estimator is as follows:
$$\hat{r}_{ac} = \frac{1}{N} \sum_{u_i \in S_{u_a}} \sum_{x_i \in S_{x_c}} r_{i,j} \quad (2)$$
- Weighted mean: ratings may be weighted using $USim(u_a,u_i)$, the similarity of a user u_i to the active user u_a and/or $ISim(x_c,x_i)$, the similarity of an item x_i to the candidate item x_c. In DF, both similarities are used (c.f. equation 3).

$$\hat{r}_{ac} = \frac{\sum_{u_i \in S_{u_a}} \sum_{x_j \in S_{x_c}} USim(u_a,u_i) \cdot ISim(x_c,x_j) \cdot r_{i,j}}{\sum_{u_i \in S_{u_a}} \sum_{x_j \in S_{x_c}} USim(u_a,u_i) \cdot ISim(x_c,x_j)} \quad (3)$$

In UC-CF, items similarities are not taken into account which leads to considering $S_{x_c} = \{x_c\}$ and $ISim(x_c,x_c) = 1$ in equation 3. Analogically, in IC-CF and CBF, we consider that $S_{u_a} = \{u_a\}$ and $USim(u_a,u_a) = 1$.

In the proposed frequentist model (FM), once users and items are clustered, we used the calculated rating probabilities $p(e_k|u,x,q)$ as an input to the rating estimators. Hereby, the estimator calculates the expected value of the rating variable such as

$$\hat{r}_{ux} = E[p(e_k|u,x,q)] = \sum e_k \cdot p(e_k|g,c,q) \quad (4)$$

3.4 Clustering Methodologies

The clustering step of our recommendation approach was carried using k-means and Expectation maximization techniques in order to build items categories and users groups (Kaufman and Rousseeuw, 1990). In order to cluster items and users based on the available data, several methodologies were used each feeding different set of variables to the clustering algorithm. Items were clustered using one of the following methodologies:

1. Features clustering : clusteing items based on their features leads to homogeneous clusters. This favors the recommendation of items that are similar to the ones the user already appreciated.
2. Ratings clustering : using only items' ratings for clustering leads to clusters containing diverse but correlated items as in IC-CF. This favors the recommendation and the discovery of new relevant items with different features.
3. Mixed clustering : using features and ratings for clustering leads to recommending items that are similar or correlated to the ones that the user likes.

Similarly to items categories, users groups were generated using three clustering methodologies:

1. Demographic clustering: using only users' demographic attributes for clustering favors the recommendation of items that are popular within a given demographic class of users.
2. Ratings clustering: using only users' ratings for clustering helps regrouping like-minded users with similar interests without necessarily being similar demographically.
3. Mixed clustering: using both users' demographics and assigned ratings for clustering helps regrouping users that are demographically similar and/or having similar rating patterns. This favors recommending not only demographically interesting items, but also the ones that are correlated to the user's interests.

Several similarity and distance measures were implemented for clustering such as Euclidian distance, cosine, adjusted cosine, mean absolute error, root mean square error and Manhattan distance.

4 EXPERIMENTATION

In order to evaluate the performances of our recommendation model and assess the genericity of our recommendation architecture, we implemented several state of the art recommendation approaches and conducted a series of experimentations on a real world dataset. In the following, the experimentations details and results are presented and discussed.

4.1 Dataset

In this work, we adopted the MovieLens dataset since it is the closest to our requirements. In fact, this dataset includes ratings assigned by a set of users to the movies they have watched. The data also includes users' demographics in addition to items' features. The dataset includes the following data:

- 1700 users : each user is described by his identifier, age, gender and occupation (cf. table 1).
- 950 films : each movie is descibed by its identifier, title, release date and the genres it belongs to among the 19 predefined ones (e.g. Action, Animation, etc...). Table 2 represents a data sample describing a set of movies. The i^{th} binary value of the last column in table 2 is 1 if the corresponding movie belongs to the i^{th} predefined genre.
- 100000 ratings assigned by users to the available movies with their respective timestamps. Hereby, each user has at least 20 ratings.

Table 1: Users data sample.

ID	Age	Gender	Occupation
1	24	M	technician
3	23	M	writer

Table 2: Items data sample.

ID	Title	Genres
22	Braveheart	0100000010000000010
23	Taxi Driver	0000000010000000100
29	Batman Forever	0110011000000000000

4.2 Experimental Results and Benchmarking

Several configurations were evaluated for each approach using different combinations of similarity measures and ratings estimators. Figure 2 aggregates the obtained mean average error values for each approach in a box-and-whisker plot in order to display the variability of its recommendation quality. The bottom and top of each box indicate the upper and the lower quartiles (respectively the 0.25 and 0.75 quantiles) of the obtained MAE values for a given approach, whereas, the band inside the box indicates the median (the 0.5 quantile). Outlier values outside the upper and lower whiskers are plotted with dots.

Table 3 enumerates the configurations of the proposed recommendation model giving the best recommendations with regard to the MAE and the RMSE measures. Meanwhile, table 4 regroups the best results obtained using the candidate algorithms.

Our experiments show that UC-CF is the least stable approach since it depends heavily on its underlying parameters, similarity measure and rating estimator. However, the results obtained from demographic filtering are similar and show that it is less sensitive to its configuration. Meanwhile, item-centered collaborative filtering, content-based filtering and our frequentist model present similar prediction quality ranges for the experimented configurations.

Item-centered collaborative filtering and content based filtering present the results that are the closest to ours. Although our recommendation model is better with regard to MAE measure, IC-CF and CBF have better results when considering RMSE as a quality measure. IC-CF's results depend mainly on the neighborhood size k determining the number of items considered as similar to a candidate one. In fact, when considering a small neighborhood of similar items, IC-CF is able to recommend highly targeted but not varied items which reflects a higher precision and a lower recall. Besides, even if a larger neighborhood of similar items leads to better balance between the quality of the recommended items and their novelty, such configuration would decrease the algorithms responsiveness due to the number of ratings to be aggregated in order to make a prediction.

Content based filtering makes predictions based on items similarities. The similarity measure and the rating estimator have less influence on the algorithm's prediction quality. Similarly to IC-CF, our experiments show that CBF responsiveness relies heavily on the chosen neighborhood size. However, this parameter do not significantly influence the algorithm's prediction quality.

Demographic filtering shows similar results independently of the employed similarity measure, rating estimator. However, predictions quality depends mainly on items and users neighborhood sizes which need to be fine tuned in a way that maximizes performances and predictions quality.

The presented experimentations show that the proposed recommender model captures consumers' purchase behaviors and is able to predict them. In fact, by unifying the main ideas and hypothesis of the experimented state of the art algorithms, our model is able to predict and quantify consumers' interests whether if they are driven by demographics, items' features or by the context.

5 CONCLUSION AND FUTURE WORKS

In this work, we proposed a generic architecture for recommender systems development and benchmarking including common recommendation algorithms, clustering techniques and similarity measures. Using the proposed architecture, we compared our recommendation model with several existing algorithms and showed that it encompasses their main ideas and may

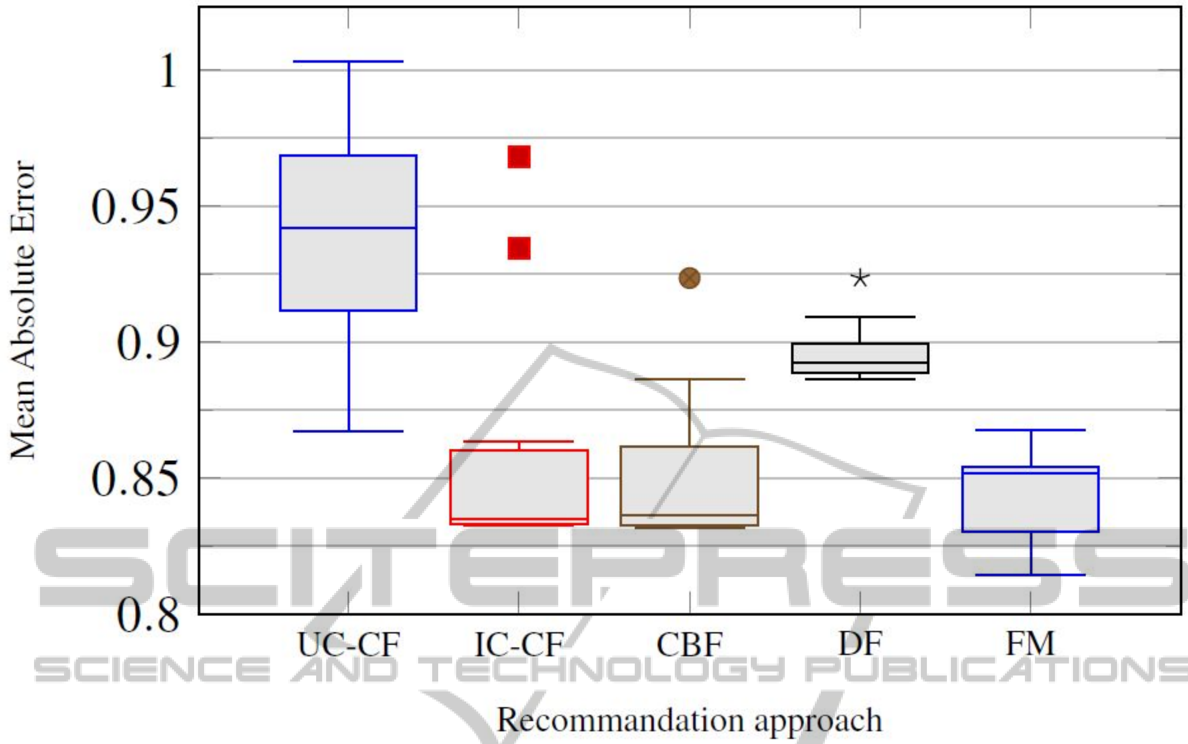

Figure 2: Recommendation algorithms benchmarking results.

Table 3: Quality of the proposed recommendation model.

N_G	N_C	Users clustering	Items clustering	Similarity measure	MAE	RMSE
15	30	Demographics and ratings	Ratings	Euclidean distance	**0.8145**	1.0856
15	35	Demographics and ratings	Ratings	Euclidean distance	0.8452	**1.0539**
15	48	Demographics and ratings	Features and ratings	Euclidean distance	0.8195	1.0908
30	25	Ratings	Ratings	Euclidean distance	0.8171	1.0930

Table 4: Quality of state of the art algorithms.

Algorithm	Neighborhood size	Similarity measure	Ratings estimator	MAE	RMSE
UC-CF	150	Adjusted Cosine	Mean	0.9445	1.1842
UC-CF	150	Cosine	Mean	0.9264	1.1712
UC-CF	150	Adjusted Cosine	Weighted Mean	0.8784	1.1075
UC-CF	150	Cosine	Weighted Mean	**0.8674**	**1.0981**
IC-CF	300	Euclidean	Mean	**0.8325**	**1.0432**
IC-CF	300	Adjusted cosine	Mean	0.8330	1.0434
CBF	300	Euclidean	Mean	0.8325	**1.0432**
CBF	300	Euclidean	Weighted Mean	**0.8318**	1.0446
DF	users: 150 , items: 10	Cosine	Weighted Mean	**0.8862**	1.1594
DF	users: 150 , items: 10	Adjusted Cosine	Weighted Mean	0.8939	1.1486
DF	users: 300 , items: 5	Manhattan	Weighted Mean	0.9235	**1.1225**

outperform their performances.
indent Similarly to the existing approaches, our model is not fitted for consumption behaviors prediction in static markets where purchases are driven by periodic needs instead of compulsive interests. In fact, a subset of consumers purchases are recurrent and often include the same items (i.e. groceries, food, tv series, phone plans, etc...). In this context, we proposed a recommendation model for predicting recurrent and periodic consumption behaviors whose goal

is to complement the model presented in this paper (Haddad et al., 2014). Future work will focus on aggregating the two propositions into a more unified recommendation model and on releasing the framework as an open source project in order to facilitate the benchmarking of other recommendation approaches.

REFERENCES

Balabanović, M. and Shoham, Y. (1997). Fab: content-based, collaborative recommendation. *Communications of the ACM*, 40(3):66–72.

Berkovsky, S., Kuflik, T., and Ricci, F. (2008). Mediation of user models for enhanced personalization in recommender systems. *User Model. User-Adapt. Interact*, 18(3):245–286.

Bezdek, J. C. (1981). *Pattern Recognition with Fuzzy Objective Function Algorithms*. Plenum Press, New York.

Boutemedjet, S. and Ziou, D. (2008). A graphical model for context-aware visual content recommendation. *IEEE Transactions on Multimedia*, 10(1):52–62.

Burke, R. D. (2007). Hybrid web recommender systems. In Brusilovsky, P., Kobsa, A., and Nejdl, W., editors, *The Adaptive Web, Methods and Strategies of Web Personalization*, volume 4321 of *Lecture Notes in Computer Science*, pages 377–408. Springer.

Deshpande, M. and Karypis, G. (2004). Item-based top-*N* recommendation algorithms. *ACM Transactions on Information Systems*, 22(1):143–177.

Goldberg, D., Nichols, D., Oki, B. M., and Terry, D. (1992). Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12):61–70.

Haddad, M. R., Baazaoui, H., Ziou, D., and Ben Ghezala, H. (2012). Towards a new model for context-aware recommendation. In *6th IEEE International Conference Intelligent Systems (IS), Sofia, Bulgaria*, pages 021 –027.

Haddad, M. R., Baazaoui, H., Ziou, D., and Ghezala, H. B. (2014). A predictive model for recurrent consumption behavior: An application on phone calls. *Knowledge-Based Systems*, 64(0):32 – 43.

Herlocker, L., J., Konstan, A., J., Borchers, A., and Riedl, J. (1999). An algorithmic framework for performing collaborative filtering. In *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Theoretical Models, pages 230–237.

Iyengar, S. S. and Lepper, M. R. (2000). When choice is demotivating: can one desire too much of a good thing? *Journal of Personality and Social Psychology*, 79(6):995–1006.

Jacoby, J., Speller, D. E., and Berning, C. A. K. (1974). Brand choice behavior as a function of information load: Replication and extension. *Journal of Consumer Research: An Interdisciplinary Quarterly*, 1(1):33–42.

Kaufman, L. and Rousseeuw, P. J. (1990). *Finding Groups in Data: an Introduction to Cluster Analysis*. Wiley.

Krulwich, B. (1997). Lifestyle finder: Intelligent user profiling using large-scale demographic data. *AI Magazine*, 18(2):37–45.

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

Mooney, R. J. and Roy, L. (2000). Content-based book recommending using learning for text categorization. In *Proceedings of the Fifth ACM Conference on Digital Libraries*, pages 195–204. ACM Press, New York.

Pazzani, M. J. (1999). A framework for collaborative, content-based and demographic filtering. *Artif. Intell. Rev*, 13(5-6):393–408.

Resnick, P., Iacovou, N., Suchak, M., Bergstorm, P., and Riedl, J. (1994). Grouplens: An open architecture for collaborative filtering of netnews. In *Proceedings of ACM 1994 Conference on Computer Supported Cooperative Work*, pages 175–186, Chapel Hill, North Carolina. ACM.

Sarwar, B. M., Karypis, G., Konstan, J. A., and Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. In *WWW*, pages 285–295.

Schwartz, B. (2005). *The paradox of choice: Why more is less*. Harper Perennial.

Woerndl, W. and Groh, G. (2007). Utilizing physical and social context to improve recommender systems. In *Web Intelligence/IAT Workshops*, pages 123–128. IEEE.