User and Context Aware Composite Item Recommendation

Krista Drushku
SAP
Paris, Levallois-Perret
krista.drushku@[sap.com,etu.univ-tours.fr]

Alexandre Chanson
University of Tours
Blois, France
alexandre.chanson@etu.univ-tours.fr

Ben Crulis
University of Tours
Blois, France
ben.crulis@etu.univ-tours.fr

Nicolas Labroche
University of Tours
Blois, France
nicolas.labroche@univ-tours.fr

Patrick Marcel
University of Tours
Blois, France
patrick.marcel@univ-tours.fr

ABSTRACT

We introduce a novel Composite Item Recommender algorithm named BFCM in a Business Intelligence application to provide users with customized recommendations to complete their reporting Tasks. To this extent, we propose a complete pipeline from the analysis of previous reports to the discovery of user intents to context-aware recommendations of Composite Items completing a report. Reported experiments show the importance of user profile in recommendation of composite items and the robustness of the proposed solution to the quality of the the user profile.

CCS CONCEPTS

• **Information systems → Recommender systems**; *Clustering*; *Business intelligence*;

KEYWORDS

Recommendation, Composite items, Bundles, Clustering, Business Intelligence

ACM Reference Format:
Krista Drushku, Alexandre Chanson, Ben Crulis, Nicolas Labroche, and Patrick Marcel. 2019. User and Context Aware Composite Item Recommendation. In *Proceedings of ComplexRec 2018 Second Workshop on Recommendation in Complex Scenarios.* ACM, New York, NY, USA, 4 pages.

1 INTRODUCTION

In today era of personal assistants [12] and with the growing need for intelligent data analysis tools that make data exploration less tedious, traditional search systems generally reach their limits for two main reasons.

The first reason relates directly to the format of the recommendation itself that is generally a ranked list of items: in this case the relevance of each item to the query is evaluated independently of the other items in the list and all items are treated equally. This kind of recommender may thus be inefficient in contexts where the list should be considered as a sequence and where the rank should reflect a relevance allowing to grasp a process, like in database exploration [1, 10], e-learning [9], or to recommend different types of items like in tourist itinerary planning [6] or finally to recommend complementary and diverse items at once like text, images and videos in web search engines [14]. In all the aforementioned cases, users expect more complex structures, that are called **composite items** or **bundles** [3, 4, 8]. These bundles group representative, cohesive, but still diverse and novel suggestions [16, 17], coming from different sources, possibly with several objectives in mind, and with distinct types of items, which makes the overall recommendation process more complex. Recent works in composite retrieval propose methods for building bundles of items around some central verticals as BOBO [4] or CPS [8] or as subset of clusters [2].

The second reason is that most of the traditional recommendation approaches do not introduce the context of the query or the user intent. Several works have been conducted to provide user intent model for recommendation, notably in the context of intelligent personal assistant [12], or for the recommendation of queries in the context of data exploration [11]. [2, 5, 15] are the first noticeable works to explicitly introduce a user intent term in a bundle recommendation process. In order to seamlessly compare items and users, the proposed method projects all items and users in a vector space of types [15] or topics [2], provided respectively by the item metadata or an LDA algorithm [7]. At the heart of the method is a constrained fuzzy c-means (FCM) algorithm that builds bundles around cluster centroids using a greedy function that aggregates several constraints like cohesiveness, personalization, diversity, etc. [3].

In this paper, we tackle the problem of completing Web Intelligence documents[1], each composed of several reports, with visualized queries, which asks to recommend items that are both conform to user interest and complementary to the current report. More precisely, the Web Intelligence platform of SAP aims at helping users constructing or completing their reports with queries already designed and shared by their colleagues. This way, they complete their reports faster, the existing reports become reusable and new information retrieved from the databases eventually become quickly visible. Our objective is then to group queries coming from different documents, that all together bring more information than a ranked list of independent queries.

To this aim, we propose an improved version of the work by [2] that introduces two new penalty terms related with i) the relevance to the user short-term interest and ii) the order of the queries in

ComplexRec 2018 Second Workshop on Recommendation in Complex Scenarios, October 7, 2018, Vancouver, Canada.

[1]https://help.sap.com/viewer/c95594c101a046159432081ca44d6b18/4.2.3/en-US/

the bundle. As we work with particular dynamic documents of the Web Intelligence platform, the user short term interest is its actual context, represented by the current report opened to edit. It completes the user long-term interest defined over their past actions. A user interest consists on a package of queries responding to the same user need. It is important to notice that because of the richness and the diversity of criteria used to build our bundles, more straightforward composite items recommendation algorithms like BOBO [4], that focuses on bundle cohesiveness, or CPS [8], that favors diversity among the items in the bundle, cannot reach the trade-off that we are looking for.

Similarly to [2], we define a vector space specifically tailored for this use case that allows to compute distances between items, between items and user and between items and report. Contrary to previous works, topics are discovered performing an efficient fuzzy k-medoid [13] over the past queries to learn global intents for all users. A proper metric is then learned on this vector space following the same methodology as presented in [11].

The paper is organized as follows: Section 2 introduces the definition of our user intents (or topics) space, how we define it and how to build a metric on this space. Section 3 formalizes our problem and describes the objective function of our modified fuzzy c-means algorithm. Finally Section 4 proposes some experiments on real data from the Web Intelligence platform at SAP, that show the importance of considering short and long term user profiles as well as an order on the queries in our context and Section 5 concludes and opens futures directions of research.

2 USER INTENT DISCOVERY

The Web Intelligence platform includes three elements that play an important role in the bundle construction: (i) the queries we can recommend, (ii) the user to whom we recommend and (iii) a report to complete, which represents a short-term interest. Each of these elements represents an *entity* of our platform which have to be projected in the same vector space to build the bundles, similarly to [2, 3, 5].

Definition 1. Let e be an *entity* and $I = \langle I_1, \ldots, I_N \rangle$ a set of N user intents or topics. An *entity profile* $I^e = \langle I^e_1, \ldots, I^e_N \rangle$ is a vector of weights $I^e_j \in [0, 1], \forall j \in [1, N]$, and such that $\sum_{j=1}^{N} I^e_j = 1$, defined over the set of intents. $I^e = \langle I^e_1, \ldots, I^e_N \rangle$ represents the relative importance of each intent $I_j, j \in [1, N]$ for a given entity e.

User intent space discovery as a clustering problem. Our objective is thus to define such user intent space that could represent the main information contained in the past queries. In [11], the authors identify user intents in the context of data exploration using a hard clustering algorithm on query representation. In our case, it can be observed in Definition 1 that the relation between the queries and the user intent space is more gradual. Indeed, I^e_j can be seen as the membership of an entity e to a user intent I_j. In this context, we use an efficient fuzzy k-medoid algorithm (FCMD) [13] to discover the user intents based on past queries. This algorithm needs a metric between queries to operate. As presented in Table 1 and following the methodology defined in [11] we define a set of features and their associated distance measures for each query: 3 features are built on the queries metadata (*same universe*, *same user*, *same folder*) and 2 others use topics discovered using LDA [7] over the query parts, defined in [11], or report and document titles, similarly to LDA topics in [5].

The overall distance $Dist(q_1, q_2)$ between queries q_1 and q_2 is defined as a linear combination of distance d_f for each specific feature f from the set of all features F as follows:

$$Dist(q_1, q_2) = \sum_{f \in F} \lambda_f \times d_f(q_1, q_2) \quad (1)$$

The learning of the appropriate metric corresponds to the learning of the weights λ_f. To this aim, we rely on the queries of labeled pairs of documents by SAP experts, who have judged for each pair if they represent the same intent or not. We train a Linear SVM classifier to learn the relative importance λ_f of each feature f.

Projection of queries, user and reports in user intent space. As per the definition of entity profile, a query, a report and a user profile can be represented by a weighted vector of importance of each interest I_j.

- Knowing the user intents as the clusters produced by the FCMD algorithm and the medoid of each cluster, it is possible to directly compute a membership vector for each new query based on the metric.
- A report r is basically a set of queries Q^r. It is thus possible to compute the coordinates of a report in the user intent space by averaging the coordinates of its queries as follows:

$$I^r = \frac{1}{|Q^r|} \sum_{q \in Q^r} I^q \quad (2)$$

- A user u can also be represented by the set of their previous queries Q^u. However, we take into account the frequency f_q of each query q in their past history as follows:

$$I^u = \frac{1}{\sum_{q \in Q^u} f_q} \sum_{q \in Q^u} I^q \times f_q \quad (3)$$

Table 1 details each feature used to compare two entities and the weights attributed by Linear SVM.

Feature f	Weight λ_f	Distance
same universe	0.24	MaxFrac.
same user	0.38	MaxFrac.
same folder	0.49	NormInt.
LDA query Parts	-0.56	Cosine
LDA titles	0.25	Cosine

Table 1: Query features description. Distance relates to metrics defined in [11]. Weights with the highest absolute score correspond to the most decisive features. 'same user' and 'same folder' favor the grouping of queries into the same user interest while 'LDA query parts' tend to discriminate among user intents.

3 COMPOSITE ITEM CONSTRUCTION

We can formulate the problem of building bundles as an optimization problem, that firstly aims at finding representative summaries of items and secondly selects the group of items respecting several constraints, to assure the relevance to the user profile, complementary and cohesion of this package. *Representativeness* assures that each bundle is built around a representative item of the dataset in order to cover the whole input data and it is ensured by applying a FCM over the items. Each cluster k is represented by a centroid c_k, which is projected in the same N-vector space, uniformly to the profile of all other entities. More precisely, given the set of items $|Q|$, the FCM algorithm returns K centroids and a partition matrix $M = \mu_{i,j} \in [0,1], i \in [1,|Q|], j \in [1,K]$ with $\sum_{k=1}^{K} \mu_{ik} = 1$ where each $\mu_{i,j}$ represents the degree to which a query i belongs to the cluster j.

We complete the objective function of [2] with new constraints to better fulfill user expectations. The only hard constraint we should respect for the creation of bundles is the number of queries they should contain: 5 in our use case of Web Intelligence reports, concluded by the experts as the adequate number of queries to recommend and easy to integrate in the interface of the existing BI platform.

The implementation of the penalty terms is different from the previous studies [2, 3, 5], as we work with different data. We define a distance function $dist()$, which measures the distance between entity profiles in the projected user intents space and a diversity function $div()$ that estimates the gain new items added in the bundle bring by presenting new visualization types.

Definition 2. Let $e1, e2$ be two different entities and I^{e_1} and I^{e_2} their corresponding profiles projected in the vector space, we define the function $dist$ of these entities as the squared Euclidean distance between their profiles:

$$dist(e_1, e_2) = \sum_{j \in [1,N]} (I_j^{e_1} - I_j^{e_2})^2 \tag{4}$$

Definition 3. Let r be the report to complete, B be the candidate bundle composed of Q^B queries and $viz(Q^B)$ the group of their visualizations. We define H_n as the normalized entropy function calculated over the visualization types of bundle and current report queries as follows:

$$div(Q^B \cup_B Q^r) = 1 - H_n(viz(Q^B \cup_B Q^r)) \tag{5}$$

Diversity was already explored by Amer-Yahia et al. in previous studies [2]. We differ in the implementation of diversity, which is computed using the type of query visualization (i.e. Pie chart, Graph, etc), assuring an orthogonal space different from other constraints.

Objective Function. Given a user u, an actual report r and the set of queries of all users Q, we aim to find (i) a set of K fuzzy clusters $C = \{C_1, ..., C_K\}$ of queries in Q, (ii) a membership function μ indicating the membership of each query to each cluster, and (iii) a set of K bundles $B = \{B_1 ... B_K\}$ with $B_k \in C_k$, $\#B = |B_k|$, which minimizes the following function:

$$\arg\min \ \frac{\alpha}{|Q|} \sum_{i=1}^{|Q|} \sum_{k=1}^{K} \mu_{ik}^m dist(q_i, c_k) \ +$$

$$\frac{1}{K} \sum_{k=1}^{K} \left(\frac{\beta}{\#B} \sum_{q \in B_k} dist(q, c_k) \ ① + \right.$$

$$\frac{\gamma}{\#B} \sum_{q \in B_k} dist(q, u) \ ② \ + \delta div(Q^{B_k} \cup_B Q^r) \ ③ +$$

$$\left. \frac{\rho}{\#B} \sum_{q \in B_k} dist(q, r) \ ④ \ + \omega \sum_{i=1}^{|Q^{B_k}|-1} dist(q_i, q_{i+1}) \ ⑤ \right)$$

where ① ensures the bundle $uniformity$ minimizing the distance of queries of the bundle to the center of cluster, ② guarantees that the suggested queries are *personalized* and correspond to the user long-term interests while ④ guarantees that they are *relevant* to the short-term interests, corresponding to the report they will be added in, ③ ensures the *diversity* of query visualizations and ⑤ enforces a logical *ranking* of items, ensuring the proximity between two consecutive queries in the bundle.

This definition of a minimization problem, using a distance measure can be changed to use a Similarity-based formulation by replacing the distance with a similarity measure, as Cosine Similarity for example, and $argmin$ by $argmax$.

This problem of constructing composite items reduces to the algorithm described in [2], following the standard FCM membership update and the modified centroids update rule and simply extending the greedy selection heuristics used in bundle composition to introduce our new penalty terms.

4 EXPERIMENTS

This section presents a set of experiments that illustrate the importance of considering the short and long terms interests in recommending for a final user and the significant weight of modeling a good user profile. We test their impact in recommending qualitative bundles. Due to space limitations, we limit the experimentation to a simple protocol with only a few settings for the objective function hyper parameters.

Data preparation. We use a selected set of 194 combinations of recent reports viewed by 46 users, containing more than 6 queries. They are separated in two parts: $future$ composed of the 5 last queries and *seed* containing the remaining queries of the beginning of the report. We run our algorithm over the *seed* and we try to recommend items close to the $future$, that follow the same logical ordering as well.

Experimental protocol. We evaluate our bundles in terms of precision and recall, comparing to the expected $future$. As it is unlikely that a query appears in several reports, we consider a similarity threshold, as defined in [1], above which two recommended queries are considered identical and the recommendation successful. We compare our algorithm to an adapted version of BOBO, based on our distance between entities, but that is agnostic of any user model or ordering of the items. We have used the same combination of hyper parameters for all our experiments: $\alpha = 0.1$, $\beta = 0.3$ and $\gamma = 0.3$, $\rho = 0.3$. Diversity δ and ranking ω are set to 0 unless otherwise stated.

Evaluating user constraint. We simulated an ideal user profile and degraded it with random noise, modifying the scores of membership to the learned intents. The ideal user profile is generated using the $future$ queries that should be discovered and the report profile is generated based the queries of the *seed*. We expect this setting to provide the highest precision score. To compare with a real context,

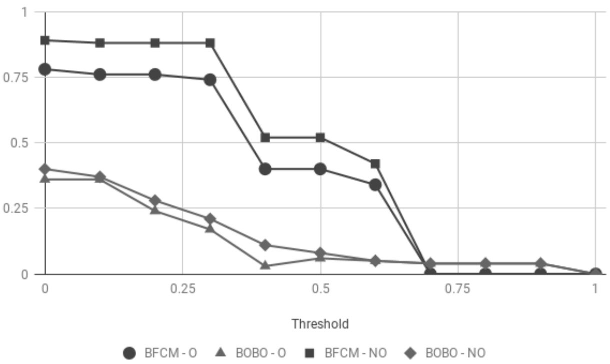

Figure 1: BOBO vs Constrained FCM

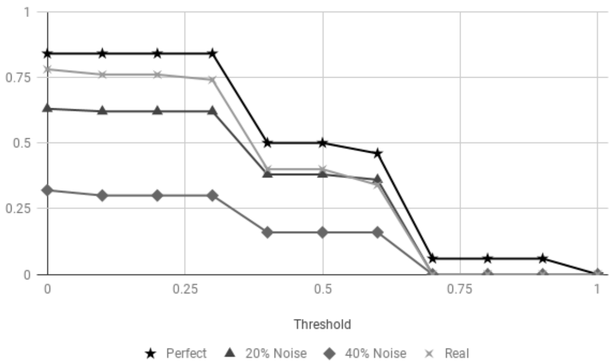

Figure 2: Precision with ordered measure

we learn the user profile as presented in Equation 3, based on the profiles of the queries she has consulted in the past.

In Figure 2 we compare the precision at different thresholds of prediction using respectively a perfect user profile, degraded with 20% noise, 40% noise and a real user profile extracted from the user's previous query usage.

Evaluating order constraint. As only our algorithm takes into account the potential order of items, we make two comparisons of the result to the $future$: (O) an ordered measure, where the n^{th} query of the recommendation is only compared to the n^{th} query of the $future$ to compute precision and (NO) an unordered measure, where we test each combination of pairs (predicted, expected) and keep the highest score.

5 RESULTS AND CONCLUSIONS

Importance of user profile. Results presented in Figure 1 show that BOBO, that is agnostic of user constraint, performs worst in this experiment compared to the real user profile as computed from user past history.

Robustness to user profile quality. Figure 2 shows that the best the quality of the user profile fed to our algorithm, the more relevant items are recommended. The perfect user profile allows very good recommendations for low similarity threshold, while the noise degrades the performances as expected. According to this test, it is possible to observe that our real user profile corresponds to approximately 20% of noise in the user profile. This is due to the lack of information in the query log used for this experiment.

Ordering items inside a bundle. As it can be seen in the Figure 1 for our algorithm BFCM, the order of items we recommend is close to the order of the expected queries as shown by the proximity of the plots for BFCM-O (ordered) and BFCM-NO (non-ordered).

Future work. We conducted several tests with different sets of hyper parameters, notably for ranking and diversity. However using the aforementioned precision measure we were unable to conclude on the contribution of the ranking as different ranking did not impact precision. This calls for a more subjective measure of quality of the bundle, that would be able to transcribe to which extent the bundle was beneficial for the user.

REFERENCES

[1] J. Aligon, E. Gallinucci, M. Golfarelli, P. Marcel, and S. Rizzi. 2015. A collaborative filtering approach for recommending OLAP sessions. *DSS* 69 (2015), 20–30.

[2] M. Alsaysneh, S. Amer-Yahia, E. Gaussier, V. Leroy, J. Pilourdault, R. M. Borromeo, M. Toyama, and J. Renders. 2017. Personalized and Diverse Task Composition in Crowdsourcing. In *Proceedings of the IEEE Transactions on Knowledge and Data Engineering, TKDE 2017.*

[3] S. Amer-Yahia, F. Bonchi, C. Castillo, E. Feuerstein, I. Méndez-Díaz, and P. Zabala. 2013. Complexity and algorithms for composite retrieval. In *22nd International World Wide Web Conference, WWW '13, Rio de Janeiro, Brazil, May 13-17, 2013, Companion Volume.* 79–80.

[4] S. Amer-Yahia, F. Bonchi, C. Castillo, E. Feuerstein, I. Méndez-Díaz, and P. Zabala. 2014. Composite Retrieval of Diverse and Complementary Bundles. *IEEE Trans. Knowl. Data Eng.* 26, 11 (2014), 2662–2675.

[5] S. Amer-Yahia, É. Gaussier, V. Leroy, J. Pilourdault, R. M. Borromeo, and M. Toyama. 2016. Task Composition in Crowdsourcing. In *2016 IEEE International Conference on Data Science and Advanced Analytics, DSAA 2016, Montreal, QC, Canada, October 17-19, 2016.* 194–203.

[6] I. Benouaret and D. Lenne. 2016. A Composite Recommendation System for Planning Tourist Visits. In *2016 IEEE/WIC/ACM International Conference on Web Intelligence, WI 2016, Omaha, NE, USA, October 13-16, 2016.* 626–631.

[7] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3 (March 2003), 993–1022.

[8] H. Bota, K. Zhou, J. M. Jose, and M. Lalmas. 2014. Composite retrieval of heterogeneous web search. In *23rd International World Wide Web Conference, WWW '14, Seoul, Republic of Korea, April 7-11, 2014.* 119–130.

[9] S. Changuel, N. Labroche, and B. Bouchon-Meunier. 2015. Resources Sequencing Using Automatic Prerequisite-Outcome Annotation. *ACM TIST* 6, 1 (2015), 6:1–6:30.

[10] M. Djedaini, N. Labroche, P. Marcel, and V. Peralta. 2017. Detecting User Focus in OLAP Analyses. In *Advances in Databases and Information Systems - 21st European Conference, ADBIS 2017, Nicosia, Cyprus, September 24-27, 2017, Proc.* 105–119.

[11] K. Drushku, J. Aligon, N. Labroche, P. Marcel, V. Peralta, and B. Dumant. 2017. User Interests Clustering in Business Intelligence Interactions. In *Advanced Information Systems Engineering - 29th International Conference, CAiSE 2017, Essen, Germany, June 12-16, 2017, Proceedings.* 144–158.

[12] R. Guha, V. Gupta, V. Raghunathan, and R. Srikant. 2015. User Modeling for a Personal Assistant. In *WSDM.* Shanghai, China., 275–284.

[13] R. Krishnapuram, A. Joshi, O. Nasraoui, and L. Yi. 2001. Low-complexity fuzzy relational clustering algorithms for Web mining. *IEEE Trans. on Fuzzy Systems* 9 (2001), 595–607.

[14] M. Lalmas. 2017. "Engage moi": From retrieval effectiveness, user satisfaction to user engagement. In *17ème Journées Francophones Extraction et Gestion des Connaissances, EGC 2017, 24-27 Janvier 2017, Grenoble, France.* 1–2.

[15] Vincent Leroy, Sihem Amer-Yahia, Éric Gaussier, and Seyed Hamid Mirisaee. 2015. Building Representative Composite Items. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management, CIKM 2015, Melbourne, VIC, Australia, October 19 - 23, 2015.* 1421–1430.

[16] Maximilien Servajean, Reza Akbarinia, Esther Pacitti, and Sihem Amer-Yahia. [n. d.]. Profile Diversity for Query Processing using User Recommendations. *Inf. Syst.* 48 ([n. d.]), 44–63.

[17] Cong Yu, Laks V. S. Lakshmanan, and Sihem Amer-Yahia. 2009. It takes variety to make a world: diversification in recommender systems. In *EDBT 2009, 12th International Conference on Extending Database Technology, Saint Petersburg, Russia, March 24-26, 2009, Proceedings.* 368–378.