

UNED @ CLEF-NEWSREEL 2014

A. Castellanos, A. García-Serrano, J. Cigarrán,

ETSI Informática, UNED. C/Juan del Rosal, 16 (Ciudad Universitaria) Madrid, Spain
{acastellanos, agarcia, juanci}@lsi.uned.es

Abstract. This paper summarizes our participation in the CLEF-NEWSREEL 2014 Challenge. The challenge focused on the recommendation of news articles. UNED's participation is in the "Recommend news articles in real-time" task. To address the recommendation tasks, a Formal Concept Analysis framework is proposed to first create the recommendation models and second to compute the recommendations. Our results prove that our FCA proposal outperforms the proposed baseline recommendation approaches. However, its performance is not still enough to be compared to other proposals for this task. In this sense some identified drawbacks, which prejudice the performance of our system, have been identified and possible solutions, to be addressed as future work, have been proposed.

Keywords: Formal Concept Analysis, Recommendation, News Recommender Systems.

1. Introduction

In recent times social networks have been the focus of most of the works related to the management of real-time streams in order to detect, review and characterize events or produce news reports. This area is especially challenging because even the best-performing algorithms in a theoretical environment can be useless or inefficient in a real environment. Experiment-based workshops organized different News Recommendation Tasks: the International News Recommender Systems Workshop and Challenge 2013 (NRS 2013), the ACM Conferences Series in Recommender Systems 2013 (RecSys 2013), together with contests, such as the plista Contest. The leitmotif of all of them has been the development of news recommender systems capable of working online in real environments. However, workshops and challenges not only offer the possibility to access to a real recommendation scenario, but also allow the evaluation of the algorithms according to metrics adapted to this context. Some novel works carried out within this context have been [1] or [2].

The CLEF-NEWSREEL offers a recommendation scenario for news articles. The organizers proposed two tasks: the first propose an offline scenario in which systems were provided with a collection to train the recommendation algorithms. The second task proposes a real recommendation scenario (the Online Recommendation Platform) in which the systems had to respond to the Open Recommendation Platform (ORP) request.

Our participation in the CLEF-NEWSREEL focuses on Task 2. To address the challenge proposed in the task, we propose a recommendation algorithm based on Formal Concept Analysis (FCA). FCA is a mathematical theory that allows content to be organized. We apply FCA to model items consumed by the users and then offer recommendations based on this modelling. FCA has been extensively used to model content and find out unknown knowledge [3]. This organization and discovering power has been applied in recommendation scenarios, although only in preliminary works. The aim of these works were to take advantage of the high performance of FCA to be able to offer better recommendations through better item modelling. [4, 5]. At this point, we propose the application of FCA for content-based recommendations in a real scenario, taking advantage of some of the conclusions of the aforementioned works.

The rest of the paper includes the description of the recommendation system, the experiments carried out and their results, and finally several conclusions are included.

2. System Description

In the following, the recommender system developed for our participation in the CLEF-NEWSREEL is detailed. The recommendation proposal follows a content-based approach: first the content of the items to be recommended (following a FCA approach) is modelled and secondly a recommendation step is carried out, by selecting “similar” content to that already consumed by the user.

2.1 FCA-Based Modelling

The textual information model is reached by means of Formal Concept Analysis (FCA). FCA is a mathematical theory of concept formation [6-9] derived from lattice and ordered set theories that provide a theoretical model to organize formal contexts. A formal context is defined as a set structure $\mathbb{K} := (G, M, I)$, where G is a set of (formal) objects, M a set of (formal) attributes and I a binary relationship between G and M , i.e. ($I \subseteq G \times M$), denoted by gIm , which is read as: the object g has the attribute m . From the point of view of a content-based recommender system, the FCA formal context can be seen as a recommendation context, in which the set of objects G is the set of items to be recommended, the set of attributes M is a set of textual features, representing the items, and the binary relationship I can be read as item g has the feature m . By textual features we refer to the most representative terminology associated to the items.

To set the “representativeness” of each term related to an item, we proposed the Kullback-Leibler Divergence [10] as a weighting measure, as described in [11]. Briefly explained, we compute the KLD value of the terms in each item by comparing them to the terms in the rest of the items; the terms with the higher KLD weight (those in the 1st tercile) are then selected as the representatives of the item. An example of this recommendation context is shown in Table 1.

Table 1. Example of Recommendation Context

	Feature 1	Feature 2	Feature 3	Feature 4	Feature 5
Item 1	X				X
Item 2		X	X		
Item 3	X			X	
Item 4					X

From the information reflected in the *recommendation context*, a set of *formal concepts* can be inferred. A *formal concept* is a pair (A, B) where $A \subseteq G$ is a set of objects (the *extent* of the formal concept) and $B \subseteq M$ is a set of attributes (the *intent* of the formal concept), which has the following properties:

- If an object a in A is tagged with an attribute b , then b must be included in B (i.e. $B = A'$ the intent of the formal concept includes all the attributes shared by the objects in the extent).
- Conversely, if an object a is tagged with all the attributes in B , then a must be included in A (i.e., $A = B'$: the extent of the formal concept includes all those objects filtered out by the intent).

To exemplify that, given the Formal Context in Table 2, the following Formal Concepts will be generated:

Table 2. Example of Formal Concepts

	Extent	Intent
C ₁	{O1, O2, O3, O4}	{∅}
C ₂	{O1, O2, O3}	{Attr1}
C ₃	{O1, O4}	{Attr5}
C ₄	{O2, O3}	{Attr1, Attr2}
C ₅	{O1}	{Attr1, Attr5}
C ₆	{O2}	{Attr1, Attr2, Attr3}
C ₇	{O4}	{Attr1, Attr2, Attr4}
C ₈	{∅}	{Attr1, Attr2, Attr3, Attr4, Attr5}

Formal concepts can be formally ordered in a *subconcept-superconcept- relationship* in accordance with their extents:

$$(A, B) \leq (A', B') : \Leftrightarrow (A, B) \subseteq (A', B') \Leftrightarrow A \subseteq A'$$

where (A', B') is called a *super-concept* of (A, B) and, conversely, (A, B) is a *sub-concept* of (A', B') (i.e., (A, B) is more specific than (A', B')). The order that results can be proven to be a *lattice*, which is called the *concept lattice*, denoted as $\mathcal{B}(G, M, I)$, associated to the *formal context*. Since *concept lattices* are ordered sets, they can be naturally displayed in terms of *Hasse diagrams* [7].

In a *Hassediagram*:

- There is exactly one node for each formal concept.
- If $C \subseteq C'$, then C' is placed above C (C is a *sub-concept* of C' or C' is a *super-concept* of C).
- If $C \subseteq C'$ but there is no other intermediary concept C'' such as $C \subseteq C'' \subseteq C'$, there is a line joining C and C' .




Fig.1. Example of Lattice

2.2 Recommendation Approach

The idea followed by the recommendation approach is along the lines of that proposed in [12]. The rationale is to take advantage of the relationships represented in the structure of the *concept lattice* to find suitable recommendations. Based on the FCA lattice, a content-based recommendation algorithm has been developed. The first step is the modelling of the items and second to recommend items similar to those already consumed by the users [13].

We have selected this approach because, in our view, a content-based approach appears to be the most suitable in the given context (news recommendations). In other contexts, such as movie recommendation, it is not clear whether the content of the items (e.g. the plot of the movies) is the main signal to get the user's interest (i.e. the user might like the movie because of the characters, director, special effects, etc., or the user interest could be related to some other issues such as the novelty of the movie or a direct recommendation of another user). However, even though other aspects might be also related, in news recommendation the content is the most important feature of an item on identifying the interest of the users. Several other aspects not directly related to the

content but to the recommendation context (e.g. novelty or popularity) are taken into account (see Section 3.2 where the experimental setup is explained).

The algorithm bases its operation on a navigation process across the lattice. Basically, given a target item, the algorithm looks for the most similar items to offer as recommendations; the algorithm takes the items already consumed by the user and offers similar items as recommendations. The navigation process starts at the *object concept* of the target item in the lattice. The *object concept* (γ_i) of i is the most specific concept (the smallest concept) including i in its extent. The *object concept* is selected as the starting point because it is the most specific *formal concept* (the one with the most attributes in the *intent*; that is, the one with the most information) in which the target item is included. The rationale is that the more information you have about an item the more accurate will be the recommendation based on similar items.

The navigation process is carried out by taking the *son concepts*: the sub-concepts (i.e. those linked immediately below in the lattice) of the target; and the *brother concepts*: the son concepts of the parent concepts (i.e. those linked immediately above in the lattice) of the target, except the target itself. The navigation is carried out as an iterative process, fixed by an N value that sets the number of levels that the algorithms should visit (up and down). All the items, still not consumed by the user, and belonging to the formal concepts included in the navigation path, are offered as recommendations.

3. Experiments and Results

Our participation in the CLEF-NEWSREEL is presented in the following. We finally participated only in Task 2, by applying the aforementioned FCA modelling and recommendation proposal, adapted to the specific requirements of the task.

3.1 Task Definition and the ORP

The task is based on the scenario of news recommendation in real-time. This scenario is especially challenging, beyond the specific requirements of the recommendation systems, it includes other issues such as: a high response speed is essential, scalability to be able to manage the real-time data stream, the ability to compute recommendation models in real-time to adapt them to the continuous information stream and to integrate them into the recommendation approach.

The scenario proposed for the task is the Open Recommendation Platform (ORP), operated by plista. ORP provides a framework in which the recommendation algorithms can be deployed. Subsequently, the servers will request recommendations from the systems. The ORP also provides an evaluation framework based on a real user study, based on the real interactions of the plista users with the recommendations offered. The evaluation will focus on click events: the absolute number of clicks and the relative number of clicks to recommendation requests.

3.2 Experimentation.

The experimentation focuses on testing the performance of the proposed FCA-based approach in a real recommendation scenario. Besides the aforementioned problems related to the real-time issue, there is another problem related to this scenario: the system has no previous information on the items or users at the beginning. The system should record all the information coming in to it (new users, new items, new item contents and new interactions) in order to compile background information to offer recommendations.

However, a problem arises in this situation: How much information should the system store? In principle, it might be thought that the more information stored, the better the performance of the system. However, this approach has some disadvantages. First of all, there is the problem of how to manage such a large amount of data. Given that recommendations should be made in real time, it is not possible to compute such an amount of data online: as can be seen in Fig. 2Fig.5, the systems should process between 12 and 35 thousand requests per day. Furthermore, the systems should also store and process the information related to these interactions. So, only offline computation is suitable in order to create a recommendation model. However, even offline computation is too expensive in terms of complexity. In this sense, instead of considering all the data, we propose two approaches: 1) consider the most novel 1,000 items, or 2) consider the top 1,000 scored items at the moment of computing the lattice.

Another disadvantage is related to the time dimension. Not all the data have the same importance: it is reasonable to think that the most novel (i.e. the latest data to arrive at the system) or the top scored (i.e. the data most consumed by the users) is more interesting for the users. Taking into account both the complexity problem and the time dimension, we decided take into account only the data belonging to the last previous 24 hours for the recommendation computation (i.e. every 24 hours the previous information is removed and a new FCA computation on the new data is carried out). But, considering that some information could be considered interesting for the users from day to day, the system keeps the most consumed items (the top 100) for the next day.

In the ORP scenario, the systems were requested to offer recommendations for a recommendation request. Both the ORP system and the challenge provide a common framework for the participants in the task in order to compare the system's performance. However to measure the suitability of our system, we also developed two baseline systems to show the improvement in the results due to the application of FCA. It is important to note that, both for the baselines and for the FCA based approach the input data is the same (the data of the last 24 hours plus the top scored items).

- **Baseline 1 - Most Novel Items:** Given a recommendation request, this system uses the set of most novel items; that is the last item that has arrived at the systems.
- **Baseline 2 - Top Scored Items:** Given a recommendation request, this system offers the set of top scored items. The score of an item is set by the number of times that it has been accessed or it has been clicked, when it has been offered as recommendation.

- **Approach 1: FCA-Based Recommendation – Most Novel:** To address the requests, the recommendation scenario proposed in Section 0 has been applied. That is, given a recommendation request, the item contained in the request is used for the system to look for similar items in the lattice structure. The information is updated daily based on the most novel items.
- **Approach 2: FCA-Based Recommendation–Top Scored:** To address the requests, the recommendation scenario proposed in Section 0 has been applied. That is, given a recommendation request, the item contained in the request is used for the system to look for similar items in the lattice structure. The information is updated daily based on the top scored items.

3.3 Results

The official results are shown in the following. Note that our systems were deployed for the whole month (May). However, the first weeks were only for development and debugging.

The results presented here correspond to the last weeks, when the systems were really in the production phase. Fig.2 and Fig.3 show the results of both baselines. The behavior of both is similar, however the results of the one based on most novel items outperform those of the approach based on the top scored. It points out that users prefer a novel although not-so-accurate recommendation.




Fig.2. Results for the Baseline 1 – Most Novel Items

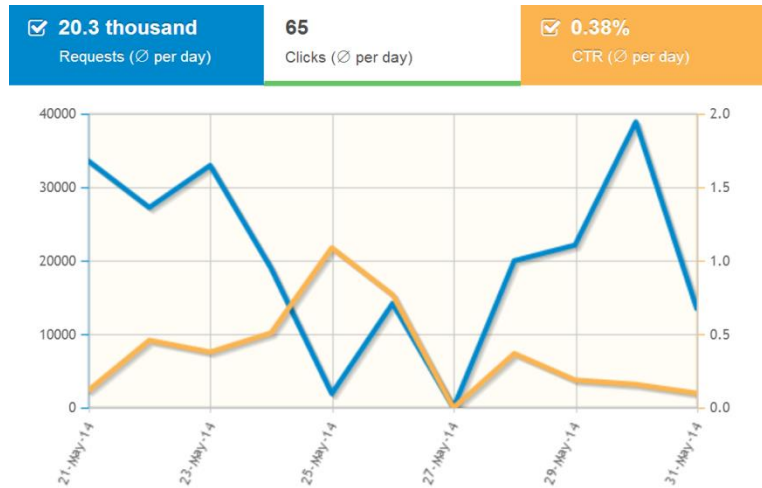


Fig.3. Results for the Baseline 2 - Top Scored Items

Once the baseline performance of the system has been outlined, the performance of the FCA-based approach is presented in Fig.4 and Fig.5. Note that Fig.4 has a time deviation with respect to other figures, due to problems with the computation of this approach. However the results can still be compared to the previous ones given that 1) the system, 2) the environment and 3) the amount of data processed is the same for all the approaches. The first point to highlight is that both approaches outperform their baselines (compare Fig.4 and Fig.2 for the most novel-based approaches and compare Fig.5 and Fig.3 for the top score-based approaches). As was highlighted in the previous results, the most novel-based approach again improves the results of the top scored.

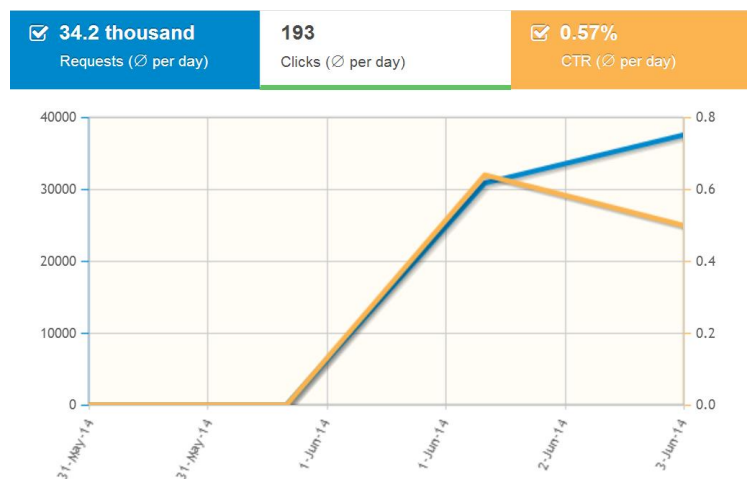


Fig.4. Result of the Approach 1: FCA Based Recommendation – Most Novel




Fig.5.Results of the Approach 2 : FCA-Based Recommendation – Top Scored

The previous analysis is only based on the comparison between our different approaches. Comparing our results with the rest of the participants in the task, it can be seen that our results (in bold) do not reach the general performance of the task (Table 3).

Table 3. Results for the last week

TEAM	Request	Clicks	CTR
abc	23,384	478	2.04
labor	386,526	7,203	1.86
plista GmbH	10,248	172	1.68
inbeat	236,115	3,803	1.61
riemannzeta	223,709	2,728	1.22
UNED	366,578	1,445	0.39
insight	84,391	317	0.38

As regards these results, firstly note that they combine the performance of all of our approaches and, consequently, the results are downgraded by the baselines. However, even taking the top-performing approach (FCA Based Recommendation – Top Scored), our results will not be among the best. During the experimentation we identified two aspects that could explain these results:

- **The need for a larger amount of information:** In the experimentation section we argued for the need to limit the amount of information to be computed in order to generate the recommendations. In this sense we only had to compute the 1,000 most suitable items. This number was set by our limitation in computing the lattices. A more extended item set to be taken into account to compute the lattices could lead to

a more informative representation and, consequently, a more accurate recommendation.

- **The time lag in the lattices:** The lattices were computed daily; that is, every day the most suitable items of the previous day were used in the computations. It meant that the recommendations had a one-day lag. If this lack of novelty were solved, the performance of the system might be better.

4. Conclusions

To address the recommendation of news articles in real-time task, based on the ORP Platform, a recommendation approach based on Formal Concept Analysis has been presented. The aim of this work was to take advantage of the high performance of the FCA in content modelling to apply it to a recommendation task.

To take the special requirements of this real-time environment into account and to be able to compute recommendations in the required time, our framework proposes a daily item modelling to create an FCA-based item modelling. This item modelling is used to look for similar items that related to the recommendation request. Some issues such as novelty and diversity have been also taken into account as explained.

For the experimentation, we computed two baselines covering two basic recommendation proposals: most-novel and top-scored items. Based on these two baselines we presented two FCA-based proposals, which are promising compared to our baselines. However if we take a look at the results of the other the participants, our FCA-based system does not reach the overall performance. In this sense, two aspects which prejudice our performance were identified: the need to use more data to create our FCA-based item representation and the need to solve the time lag. To sum up, we demonstrated the viability of FCA to be applied in such a challenging scenario, even if it is still far from the performance of the state-of-the-art systems in this field.

As regards the latter, some lines of future work come about. The main aspect to be addressed is related to solving the complexity problems of our approach. In this sense we are on the way to developing a FCA Big Data Framework, in order to take into account and compute the entire amount of data available in this kind of environment. In our approach, the FCA computation represents a bottleneck, even more so when the amount of data becomes as large as it does in this scenario. As a solution, we propose to select only some data (according to a criterion) to create the FCA models. It leads to a loss of knowledge related to this loss of data. A Big Data Framework would allow much more data to be computed, ideally leading to more informative models which could create more accurate recommendations.

Another issue to be addressed is the time lag problem. This aspect does not have such a direct solution; it could be addressed: 1) by taking a smaller time snapshot to compute the modelling, 2) by adapting the computation time to several variables (i.e. the amount of data available or period of the day), or 3) by applying a smoothing process in the removal of old data.

Acknowledgments. This work has been partially supported by MA2VIRMR (S2009/TIC-1542), and HOLOPEDIA (TIN 2010-21128-C02) Spanish projects.

References

1. Said, A. Bellogín, A. and de Vries, A.: News Recommendation in the Wild: CWI's Recommendation Algorithms in the NRS Challenge. In: Proceedings of the International News Recommender Systems Challenge (NRS 2013), at the 7th ACM Conference on Recommender Systems (RecSys 2013). (2013)
2. Garcin, F., Faltings, B.: Pen recsys: A personalized news recommender systems framework. In: Proceedings of the 7th ACM Conference on Recommender Systems. RecSys '13, New York, NY, USA, ACM (2013) 469–470
3. Poelmans, J., Elzinga, P., Viaene, S., Dedene, G.: Formal concept analysis in knowledge discovery: a survey. In: Proceedings of the 18th international conference on Conceptual structures: from information to intelligence. ICCS'10, Berlin, Heidelberg, Springer-Verlag (2010) 139–153
4. Ignatov, D., Kaminskaya, A., Bezzubtseva, A., Konstantinov, A., Poelmans, J.: FCA-Based Models and a Prototype Data Analysis System for Crowdsourcing Platforms. In Pfeiffer, H., Ignatov, D., Poelmans, J., Gadiraju, N., eds.: Conceptual Structures for STEM Research and Education. Volume 7735 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2013) 173–192
5. Li, X., Murata, T.: A knowledge-based recommendation model utilizing formal concept analysis and association. In: Computer and Automation Engineering (ICCAE), 2010 The 2nd International Conference on. Volume 4. (2010) 221–226
6. Belohlavek, R.: Introduction to formal concept analysis. Olomouc, UPOL, Faculty of Science, Department of Computer Science (2008)
7. Ganter, B., Wille, R., Franzke, C.: Formal concept analysis: mathematical foundations. Springer-Verlag New York, Inc. (1997)
8. Wille, R.: Concept lattices and conceptual knowledge systems. *Computers & mathematics with applications* **23**(6) (1992) 493–515
9. Wille, R.: Restructuring Lattice Theory: An Approach Based On Hierarchies Of Concepts. Volume 5548 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2009)
10. Kullback, S., Leibler, R.A.: On information and sufficiency. *The Annals of Mathematical Statistics* **22**(1) (1951) 79–86
11. Castellanos, A.: Recomendación de contenidos digitales basada en divergencias del lenguaje diseño experimentación y evaluación. Master's thesis, UNED (2013)
12. du Boucher-Ryan, P., Bridge, D.: Collaborative recommending using formal concept analysis. *Knowledge-Based Systems* **19**(5) (2006) 309–315
13. Ricci, F., Shapira, B.: Recommender systems handbook. Springer (2011)