

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/318789067>

Tailoring Recommendations for a Multi-Domain Environment

Conference Paper · September 2017

CITATIONS

0

READS

16

3 authors:



Emanuel Lacic

Graz University of Technology

22 PUBLICATIONS 82 CITATIONS

SEE PROFILE



Dominik Kowald

Graz University of Technology

45 PUBLICATIONS 213 CITATIONS

SEE PROFILE



Elisabeth Lex

Graz University of Technology

72 PUBLICATIONS 240 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Learning Layers [View project](#)



Web Information Quality Evaluation Initiative [View project](#)

All content following this page was uploaded by **Dominik Kowald** on 31 July 2017.

The user has requested enhancement of the downloaded file.

Tailoring Recommendations for a Multi-Domain Environment

Emanuel Lacic
Know-Center Graz
Graz, Austria
elacic@know-center.at

Dominik Kowald
Know-Center Graz
Graz, Austria
dkowald@know-center.at

Elisabeth Lex
Graz University of Technology
Graz, Austria
elisabeth.lex@tugraz.at

ABSTRACT

Recommender systems are acknowledged as an essential instrument to support users in finding relevant information. However, the adaptation of recommender systems to multiple domain-specific requirements and data models still remains an open challenge. In the present paper, we contribute to this sparse line of research with guidance on how to design a customizable recommender system that accounts for multiple domains with heterogeneous data. Using concrete showcase examples, we demonstrate how to setup a multi-domain system on the item and system level, and we report evaluation results for the domains of (i) LastFM, (ii) FourSquare, and (iii) MovieLens. We believe that our findings and guidelines can support developers and researchers of recommender systems to easily adapt and deploy a recommender system in distributed environments, as well as to develop and evaluate algorithms suited for multi-domain settings.

KEYWORDS

Recommender Systems; Multi-Domain Recommendation; Heterogeneous Data; Customizing Recommendation Approaches

1 INTRODUCTION

In the past decade, there has been a vast amount of research in the field of recommender systems. Most of these systems offer recommendations adapted for items belonging to a single domain (e.g., movies, music, news, etc.). However, supporting different domain-specific data models is still an open challenge, which is neglected by many recommender systems and that has just been recently taken up by the recommender systems' research community.

The work of [2] has actually been the first attempt to define the concept of a domain in the context of recommender systems. The authors distinguish between four different domain notations, namely (i) attribute level, where items are of the same type (e.g., movie genres), (ii) type level, where items are of similar types but share some attributes (e.g., movies and tv shows), (iii) item level, where items are not of the same type and differ in most or all attributes (e.g., movies and books) and, (iv) system level, where items and users belong to different systems (e.g., LastFM and MovieLens). Moreover, they distinguish between the concept of multi-domain recommendations and cross-domain recommendations. The goal of cross-domain recommendation is to utilize the knowledge derived

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

RecSysKTL Workshop @ ACM RecSys '17, August 27, 2017, Como, Italy

© 2017 Copyright held by the owner/author(s).

DOI: N/A

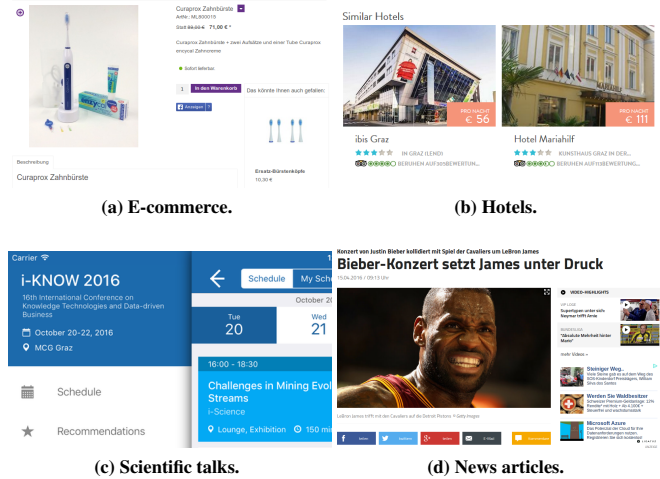


Figure 1: Providing recommendations in multiple domains requires supporting heterogeneous data structures, allow for domain-specific algorithm customization, as well as to support service isolation and fault tolerance.

from one or more source domains in order to generate predictions for a target domain. However, they also state that multi-domain approaches have mainly focused on the provision of cross-domain recommendations by jointly considering user preferences for items in various systems.

Other related works implicitly share this definition [2] and focus on cross-domain recommendations rather than on multi-domain ones. Thus, the main focus has been to tackle the data sparsity problem (e.g., via active transfer learning [14]) by utilizing Collaborative Filtering [4, 7, 10] and Content-based Filtering approaches [3, 11].

In the present paper, we build upon these works and we extend the scope of multi-domain recommender systems by introducing several guidelines concerning topics such as data heterogeneity and customization that should be taken into consideration. We base our findings on real-world applications (e.g., e-commerce¹, hotels², scientific talks³ or news articles⁴ – just to name a few) and propose a practical approach on how to support multi-domain recommendations on both the item and system level as described by [2].

To the best of our knowledge, this is the first work which addresses the question of what design decisions should be taken into consideration when building a recommender system for a multi-domain environment.

¹<http://www.mymanou.com/>

²<https://www.triprebel.com/>

³<http://uscn.me/r220>

⁴<http://www.clef-newsreel.org/>

```
{
  "id": "e12c4fb-ba85-46d5-896d-af65d1f3b48c",
  "item": "5a2bc423-15dc-47d3-8a8c-f543dc267a7c",
  "users_listened": [3907, 57017]
  "users_listened_count": 2
  "domain": "LastFM"
},
{
  "id": "8fe89bab-9631-4e87-81a0-a8dd3ffba774c",
  "user": "23861",
  "item": 4105,
  "rating": 1.0
  "domain": "MovieLens"
}
```

Listing 1: Example of different schema strategies to store data in the same place for item level multi-domain recommendations.

For example, as shown in Listing 1, we could easily derive different data structures to store and generate recommendations in the corresponding music and movie domain. Moreover, we found that the capability for horizontal scaling (i.e., creating shards and replicas) is extremely important when integrating new domains on an item-level (see Figure 2) as such a strategy easily increases the amount of data that needs to be stored and processed.

2.3 Customizing Recommendation Approaches

An important aspect of a multi-domain recommender system is customization. Typically, different application domains have different domain-specific data features, which means that, for example, a music recommender could solely use implicit user interactions (e.g., listened songs), whereas an e-commerce recommender could use explicit ones (e.g., ratings). On top of that, one would also need to separately determine and setup the correct algorithmic parameters for each domain. For example, in case of the Collaborative Filtering algorithm, the similarity function (e.g., Cosine or Jaccard similarity) and the neighbourhood size need to be determined for each domain. For other domains, one may need to setup custom filtering criteria (e.g., recommend items that are suited for minors or are part of a specific category). Thus, a multi-domain recommendation approach should be aware of the underlying data structures and domain-specific parameters.

As such, we propose to outsource the domain-specific algorithm setup into so-called recommender profiles. In our applications, we defined a customizer module (see Figure 2), which enables to set domain-specific algorithm configurations and to transfer it to modules utilized in a specific domain environment. This way, we can manage domain-specific configurations and dynamically integrate additional domains. An example of such recommender profiles for the music domain is given in Listings 2, 3 and 4.

Here, we define a configuration by a unique reference id, a reference to the algorithm implementation (e.g., class name) and the specific domain-relevant parameters. In case a recommender profile is created or updated for a particular domain at runtime, the changes need to be propagated throughout the whole system to every domain-dependant module. As such, each module from the same domain will be informed about the changes and the updated profile will be used as soon as a new recommendation request is received.

```
id: ktl_mp_lastfm
# reference for a MP implementation
algorithm: MostPopularGeneric

# algorithm specific parameters
parameters:
  # item-level domain?
  domain: lastFM
  # on what do I calculate the popularity?
  count_fields : [ users_listened_count ]
  user_action_fields : [ users_listened ]
```

Listing 2: A MostPopular recommender profile for the LastFM domain.

```
id: ktl_ub_cf_lastfm
# reference for a user-based-CF implementation
algorithm: GenericUBCF

# algorithm specific parameters
parameters:
  # item-level domain?
  domain: lastFM
  # How to calculate user similarity?
  similarity_function : OVERLAP # JACCARD, COSINE, etc.
  neighbourhood_size: 40
  user_action_fields : [ users_listened ]
```

Listing 3: A Collaborative Filtering recommender profile for the LastFM domain.

```
id: ktl_hybrid_cs_lastfm
# reference for a hybrid implementation
algorithm: CrossSourceHybrid

# algorithm specific parameters
parameters:
  # item-level domain given by combining profiles
  profile_ids : [ ktl_mp_lastfm , ktl_ub_cf_lastfm ]
  recommender_weights: [0.1, 0.9]
```

Listing 4: A Cross-Source Hybrid recommender profile for the LastFM domain.

2.4 Fault Tolerance

Deploying multiple modules in a distributed manner increases the probability of unexpected behaviour (e.g., hardware shutdown, I/O problems, software bugs, etc.). As we have proposed to use microservices, it is not necessary to cope with central node failures as it is the case with a master-slave architecture. In case a module fails, ZooKeeper, or any other orchestration service like Eureka or Consul, should remove the faulty module from its list of “live nodes” and no further requests will be redirected to it.

Thus, the module will not necessarily cause any major problems as long as there is another module of the same type available. When experiencing a high request load, it should be possible to deploy and register an additional module to ZooKeeper on the fly. To further improve the reliability of the system, multiple ZooKeeper instances can be used in a cluster in order to overcome the outage of single instances. In such a way, the runtime performance can be guaranteed for both item and system level multi-domain recommendations.

3 DOMAIN EXPERIMENTS

In order to demonstrate the application of our guidelines for providing recommendations in multiple domains, we performed several experiments on well-known datasets used in recommender systems research, (i.e., LastFM¹⁰, Foursquare¹¹ and MovieLens20M¹²). The LastFM dataset consists of 359,348 users, 268,736 artists and 17,559,530 implicit user interactions that denote the listening relationship between users and artists. Foursquare provides 2,809,581 ratings on a 5-star scale for different venues (i.e., restaurants) and contains 2,153,471 users as well as 1,143,092 venues in general. The MovieLens20M dataset has 138,493 users, 27,032 movies as well as 19,999,603 movie reviews on a 5-star rating scale with a step size of 0.5 (i.e., a 10-star scale).

We focused on providing recommendations for cold-start users and as such, we removed all users that interacted with more than 20 items. Next, we split the remaining data in two different sets (i.e., training and test sets) using a method similar to the one described in [9]. Thus, for each user, we withheld 10 items that were used for testing and the rest was used for training. This has resulted in an evaluation set of 2,409 users for LastFM, 41,628 users for FourSquare and 4,486 users for the MovieLens20M dataset. On these datasets, we evaluated a simple MostPopular (MP) approach (e.g., Listing 2), a user-based Collaborative Filtering (CF) approach (e.g., Listing 3) as well as a hybrid combination [1] of both (e.g., Listing 4). Additionally, we evaluated different neighborhood sizes N for the CF approach in order to optimize the results. For reasons of simplicity, we used a naive item overlap metric to measure the similarity between users (i.e., $OV(u_t, u_c) = |\Delta(u_t) \cap \Delta(u_c)|$, where $\Delta(u)$ corresponds to the set of items some user u has interacted with in the past). However, as shown in Listing 3 this is a domain-specific parameter that could be easily adapted.

The results of our evaluation are shown in Table 1 by means of the nDCG@10 and User Coverage (UC) metrics. The aim of this simple experiment was to show how different parameter setups can impact the performance in different domains. As shown, the neighbourhood size is one example of a parameter that needs to be optimized for a specific domain. By choosing the best parameter combination for the hybrid approach, we can provide more robust recommendations for all users in each domain.

4 CONCLUSION

In this work, we presented our approach on providing recommendations in a multi-domain environment. Specifically, we introduced the concept of recommender profiles in order to customize existing algorithms with domain-specific configuration. Apart from that, we provided guidelines with respect to service isolation, heterogeneous data and fault tolerance. Finally, we provided customization examples as well as evaluation results for the domains of (i) LastFM, (ii) FourSquare, and (iii) MovieLens. We believe that our findings and proposed guidelines are of use for developers and researchers of recommender systems to tailor and develop recommendations for multi-domain and distributed environments.

¹⁰<http://mtg.upf.edu/node/1671>

¹¹https://archive.org/details/201309_foursquare_dataset_umnn

¹²<https://grouplens.org/datasets/movielens/20m/>

	Approach	nDCG@10	UC
LastFM	MP	.0180	100%
	CF	N = 20	93.70%
		N = 30	
		N = 40	
		N = 50	
		N = 60	
	Hybrid	.1005	100%
Foursquare	MP	.0256	100%
	CF	N = 20	49.58%
		N = 30	
		N = 40	
		N = 50	
		N = 60	
	Hybrid	.0339	100%
MovieLens	MP	.0658	100%
	CF	N = 20	100%
		N = 30	
		N = 40	
		N = 50	
		N = 60	
	Hybrid	.0999	100%

Table 1: Evaluation results of our multi-domain experiment.

For future work, we plan to perform a more elaborate study using the proposed recommender profiles to study differences in domain-specific configurations (e.g., does a semantic relationship impact the choice of domain-specific parameters like the similarity metric or different filtering criteria?). Moreover, we plan to investigate datasets with textual contents in order to further explore how hybrid weightings may impact the performance in a specific domain with respect of not only accuracy but also diversity.

Acknowledgment. This work was funded by the Horizon 2020 project MoreGrasp (643955).

REFERENCES

- [1] S. Bostandjiev, J. O'Donovan, and T. Höllerer. Tasteweights: a visual interactive hybrid recommender system. In *Proc., RecSys '12*, pages 35–42. ACM, 2012.
- [2] I. Cantador, I. Fernández-Tobías, S. Berkovsky, and P. Cremonesi. Cross-domain recommender systems. In *Recommender Systems Handbook*. Springer, 2015.
- [3] A. M. Elkahky, Y. Song, and X. He. A multi-view deep learning approach for cross domain user modeling in recommendation systems. In *Proc. WWW'15*.
- [4] S. Gao, H. Luo, D. Chen, S. Li, P. Gallinari, and J. Guo. Cross-domain recommendation via cluster-level latent factor model. In *Proc. of ECML-PKDD'13*.
- [5] E. Lacić, D. Kowald, and C. Trattner. Socrecm: A scalable social recommender engine for online marketplaces. In *Proc. of the ACM Hypertext 2014*.
- [6] E. Lacić, M. Traub, D. Kowald, and E. Lex. Scar: Towards a real-time recommender framework following the microservices architecture. In *Proc. of LSRS'15*.
- [7] B. Loni, Y. Shi, M. Larson, and A. Hanjalic. Cross-domain collaborative filtering with factorization machines. In *ECIR*, pages 656–661. Springer, 2014.
- [8] A. McAfee and E. Brynjolfsson. Big Data: The management revolution. *Harvard Business Review*, 90(10):60–68, 2012.
- [9] D. Parra-Santander and P. Brusilovsky. Improving collaborative filtering in social tagging systems for the recommendation of scientific articles. In *Proc. of WI-IAT '10*, pages 136–142. IEEE Computer Society.
- [10] S. Sahebi and P. Brusilovsky. It takes two to tango: An exploration of domain pairs for cross-domain collaborative filtering. In *Proc. of ACM RecSys 2015*.
- [11] S. Sahebi and T. Walker. Content-based cross-domain recommendations using segmented models. In *CBRecSys@ RecSys*, pages 57–64, 2014.
- [12] M. Traub, D. Kowald, E. Lacić, P. Schoen, G. Supp, and E. Lex. Smart booking without looking: Providing hotel recommendations in the tripbel portal. In *Proc. of i-KNOW '15*.
- [13] S. Werner and A. Lommatzsch. Optimizing and evaluating stream-based news recommendation algorithms. In *CLEF (Working Notes)*, pages 813–824, 2014.
- [14] L. Zhao, S. J. Pan, E. W. Xiang, E. Zhong, Z. Lu, and Q. Yang. Active transfer learning for cross-system recommendation. In *AAAI*, 2013.