

Improving recommender system navigability through diversification: A case study of IMDb

Daniel Lamprecht
KTI, Graz University of
Technology
Graz, Austria
daniel.lamprecht@tugraz.at

Florian Geigl
KTI, Graz University of
Technology
Graz, Austria
florian.geigl@tugraz.at

Tomas Karas
KTI, Graz University of
Technology
Graz, Austria
karas@student.tugraz.at

Simon Walk
IICM, Graz University of
Technology
Graz, Austria
simon.walk@tugraz.at

Denis Helic
KTI, Graz University of
Technology
Graz, Austria
dhelic@tugraz.at

Markus Strohmaier
GESIS and University of
Koblenz-Landau
Cologne, Germany
strohmaier@uni-
koblenz.de

ABSTRACT

The Internet Movie Database (IMDb) is the world's largest collection of facts about movies and features large-scale recommendation systems connecting hundreds of thousands of items. In the past, the principal evaluation criterion for such recommender systems has been the rating accuracy prediction for recommendations within the immediate one-hop-neighborhood. Apart from a few isolated studies, the evaluation methodology for recommender systems has so far lacked approaches that quantify and measure the exposure to novel content while navigating a recommender system. As such, little is known about the support for navigation and browsing as methods to explore, browse and discover novel items within these systems. In this article, we study the navigability of IMDb's recommender systems over multiple hops. To this end, we analyze the recommendation networks of IMDb with a two-level approach: First, we study reachability in terms of components, path lengths and a bow-tie analysis. Second, we simulate practical browsing scenarios based on greedy decentralized search. Our results show that the IMDb recommendation networks are not very well-suited for navigation scenarios. To mitigate this, we apply a method for diversifying recommendations by specifically selecting recommendations which improve connectivity but do not compromise relevance. We demonstrate that this leads to improved reachability and navigability in both recommender systems. Our work underlines the importance of navigability and reachability as evaluation dimension of a large movie recommender system and shows up ways to increase navigational diversity.

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

i-KNOW '15, October 21-23, 2015, Graz, Austria

© 2015 ACM. ISBN 978-1-4503-3721-2/15/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2809563.2809603>

CCS Concepts

•Human-centered computing → Web-based interaction; Collaborative filtering; •Information systems → Recommender systems;

Keywords

Recommender Systems, IMDb, Navigation, Diversification

1. INTRODUCTION

Recommender systems support users in filtering information and selecting items among huge numbers of possible options. By connecting users with appropriate, relevant, or novel items, recommender systems also help to reduce information overload by filtering out unwanted items and reducing cognitive load on users [9, 10, 20]. By establishing connections between items, recommender systems enable users to browse and peruse a system. Users enjoy browsing a recommender system without the intention of making a purchase [9], which is especially relevant on systems where users immediately consume items (such as on YouTube [5]). Finally, recommendations are also important in the discovery of novel content [17].

In the past, the majority of research and development on recommender systems has focused on improving rating prediction accuracy. Spurred by the Netflix Prize challenge¹, where the evaluation criterion was the root mean squared error (RMSE) calculated on the rating predictions, researchers have found substantial improvements in terms of computing rating predictions [13].

So far, comparatively little attention has been paid to supporting, evaluating, or improving navigation and exploration properties of recommender systems. As a consequence, we still do not know much about how these scenarios are supported in state-of-the-art recommender systems. Learning more about the conditions of navigability in recommender systems is vital for researchers and practitioners who want to gain insight into how well these systems support navigation.

In this paper, we set out to analyze such properties in a real-world recommender system. To this end, we apply

¹<http://www.netflixprize.com>

a recently presented network-theoretic framework [15] that proposes a two-level approach:

1. The first step investigates the *reachability* of recommendation networks (i.e., the networks formed by items as nodes connected by recommendations as links) by analyzing the topological characteristics in terms of components, clustering, path lengths and partitions. This analysis quantifies what parts of the network are connected via links and how many hops it takes to reach them.
2. The second step analyzes the results of these findings in a more practical way by simulating browsing scenarios on these networks. This provides us with insight into how well these networks fare in real-world navigation scenarios.

We apply this approach to investigate the case of IMDb, the largest movie database in the world. In particular we are interesting in answering the following research questions:

- RQ 1** How well do the recommendations of IMDb support reachability and navigability?
- RQ 2** How can reachability and navigability on IMDb be improved?
- RQ 3** What are differences between collaborative filtering and content-based recommendations in terms of reachability and navigability?

In order to answer these questions, we analyze the two types of recommender systems present on IMDb in their entirety (see Figure 1 for an example of an IMDb page). Our results show that the recommendation networks on IMDb are split into a large number of disconnected components with large distances within components. As a result, the current state of IMDb recommendations does not support any kind of exploration scenario very well. As a remedy, we introduce recommendation diversification to better distribute the recommendations among items and show that two diversification approaches are able to substantially improve navigability.

2. RELATED WORK

The study of human navigation in networks was strongly influenced by Milgrams and Travers [19, 24], who performed a series of experiments on navigation in social networks. They found that even within very large social entities, such as the entire United States, humans were able to find connections to others through a very small number of intermediaries. This coined the term of *Six Degrees of Separation*. The notion of an efficiently navigable network was later formalized by Watts and Strogatz, who described high clustering and short path lengths as characteristics of highly navigable small-world networks [25, 26]. Kleinberg identified further properties that rendered networks efficiently navigable with decentralized search algorithms [11, 12]. The navigation model of greedy decentralized search was later used to analyze human navigation dynamics in networks [8, 16, 23].

West and Leskovec [27] studied human goal-oriented navigation in the information network of Wikipedia and found that humans took only a few clicks longer than the shortest possible paths. However, in contrast to the shortest paths,

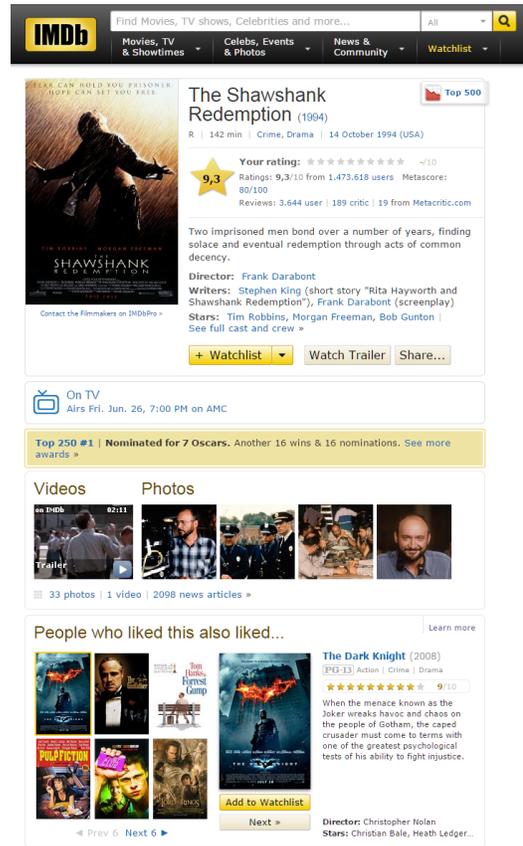


Figure 1: **IMDb page.** Example of an IMDb movie page, displaying facts, a voting score, links to videos and photos and collaborative filtering recommendations.

the resulting click trails exhibited a characteristic zoom-out phase (leading to more general concepts), followed by a phase of homing in to the target based on similarity.

Human navigation on recommender systems can occur in a range of uses cases. Recommendation browsing helps in discovering novel content [17], and the same has been found for search [28], where some users prefer navigation to search even when they know the target [22]. Generally, recommender systems help in learning and decision making [18, 20]. Users are more likely to follow links on movie recommendation sites than on factual websites such as Wikipedia [6]. On YouTube, recommendations fulfill the need for *unarticulated want* [5] and form a vital part of the user experience by connecting items.

A few studies have already investigated navigability on recommender systems. Music recommender systems were found to show heavy-tail degree distributions as well as small-world properties [2]. Several variations of IMDb recommendation networks have been found to exhibit long-tail degree distributions [7]. Celma and Herrera [3] found that collaborative filtering led to popularity bias and that a trade-off existed between accuracy and other evaluation metrics.

A simple method to improve navigability by selecting recommendations based on reachability was proposed by Seyerlehner et al [21]. We improve on this by taking the relevancy of recommendations as well as their directionality into account.

3. MATERIALS AND METHODS

3.1 Data Sets

The Internet Movie Database (IMDb) is a database of facts about movies and television shows. The website started out as a hobby project on Usenet and has since grown to be the largest movie website worldwide². The website presents facts and details about titles (movies, TV shows, short films and so forth), such as plot, cast, trailers and reviews, as well as information about actors and actresses, directors and crew. As of January 2015, the database contained facts about 3.1 million titles³.

Users on IMDb can contribute and edit facts, although changes are moderated before being entered into the database. Users can also rate movies, write reviews and participate in messaging forums.

IMDb offers two different recommender systems:

Collaborative Filtering Recommendations (CF).

IMDb uses non-personalized rating-based recommendations for its titles, listed as *People who liked this also liked...* on title pages. The interface shows a total of 12 CF recommendations, from which 6 are immediately visible (see Figure 1 for an example of this interface).

Content-based Recommendations (CB).

Up until a site redesign in 2010, IMDb used non-personalized content-based recommendations⁴. These recommendations were computed from a proprietary combination of facts such as title, keywords, genre and user votes. This interface including the recommendations is still available through a change in the user preferences. In the interface, 5 recommendations are visible initially, and a total of 10 are available by following a link.

The presence of two parallel recommendation engines enabled us to directly compare the navigability within two real-world recommender systems side-by-side. To obtain the data on the recommender systems, we performed an exhaustive search over the IMDb title IDs by enumerating the space of 10 million possible values. During our crawl in January 2015, we were able to obtain the entire database of about 3.1 million titles in this way. We then extracted facts, such as release date, plot, storyline and average rating, as well as all available recommendations of both types. In total, we obtained 785,019 nodes with content-based recommendations and 168,560 nodes with collaborative-filtering recommendations.

As the basis for the diversification approaches, we also inspected the reviews for each title and downloaded all ratings assigned as part of a review. After that, we visited the profile pages of all users who had written at least one review and additionally downloaded all of their ratings they had assigned without an associated review, if they were publicly available. To avoid problems with sparse data, we only used data from films with at least three ratings and users who had rated at least three titles. By combining the profile ratings with the reviews ratings, we obtained a total of 25,290,692 million ratings from 149,240 users for 168,078 titles.

²<http://http://www.imdb.com/pressroom>

³<http://http://www.imdb.com/stats>

⁴http://www.imdb.com/help/show_leaf?history

3.2 Recommendation Networks

We constructed unpersonalized top- N recommendation networks from the recommendations we obtained from IMDb. In each of these networks, the items were represented as nodes and recommendations formed directed edges. We constructed a total of four different networks: Two for collaborative filtering, with 6 and 12 recommendations per node (denoted as CF (6) and CF (12)), and two for content-based recommendations with 5 and 10 recommendations per node (denoted as CB (5) and CB (10)). The number of recommendations was therefore the same as in the user interfaces.

For the collaborative filtering networks, a fraction of nodes did not have any outgoing recommendations and were thus unreachable via recommendations but then constituted a dead end. These nodes made up 11% of the CF (6) and 21% of the CF (12) network.

3.3 Diversification

To improve navigability, we introduced diversity into the networks. User satisfaction with diversity for collaborative filtering has been found to peak between 30-40% diversity [29]. Based on this, we replaced recommendations as follows: For the immediately visible recommendations (5 for CB and 6 for CF), we replaced two recommendations. For the total recommendation list (10 for CB and 12 for CF) we replaced 4 recommendations. We use the following three approaches for diversification:

- **Random Recommendations.** The introduction of random links generally leads to well-connected networks with a small diameter. As such, introducing random recommendations effectively constituted an upper bound on the possible improvement through diversification.
- **Diversify.** Ziegler et al [29] proposed a method called *Diversify*. To apply it, we first build the recommendation network with the desired number of non-diversified recommendations (e.g., 4 for CF (6)). *Diversify* then introduces diverse recommendations for each node as the ones minimizing the similarity to the recommendations already present. We compute similarities between items by comparing their rating vectors.
- **Expanded Relevance (ExpRel)** Küçükünç et al. [14] proposed a method to take the location of recommendations in the network into account. We use a simplified version thereof: We first build the recommendation network $G = (V, E)$ with of the desired number of non-diversified recommendations (e.g., 4 for CF (6)). Based on this, for each node $n \in V$ we compute $\Gamma(n, 2)$, the set nodes reachable in the two-hop neighborhood of each node. We then rank potential diverse recommendations $d \in V$ based on the number of nodes in $\Gamma(d, 2) - \Gamma(n, 2)$, the number of additional nodes the recommendation would add.

For the random recommendations, we added the diverse recommendations to all nodes that had existing outgoing recommendations in the graph. For *Diversify* and *ExpRel*, we first computed the cosine similarities between all pairs of items for which co-ratings were present in our dataset. We then selected all items for which at least 100 similarities to other nodes could be computed and used the top 50 most similar nodes to select diversified recommendations from. This left us with 145,504 nodes for the CB and 118,691 for the CF networks.

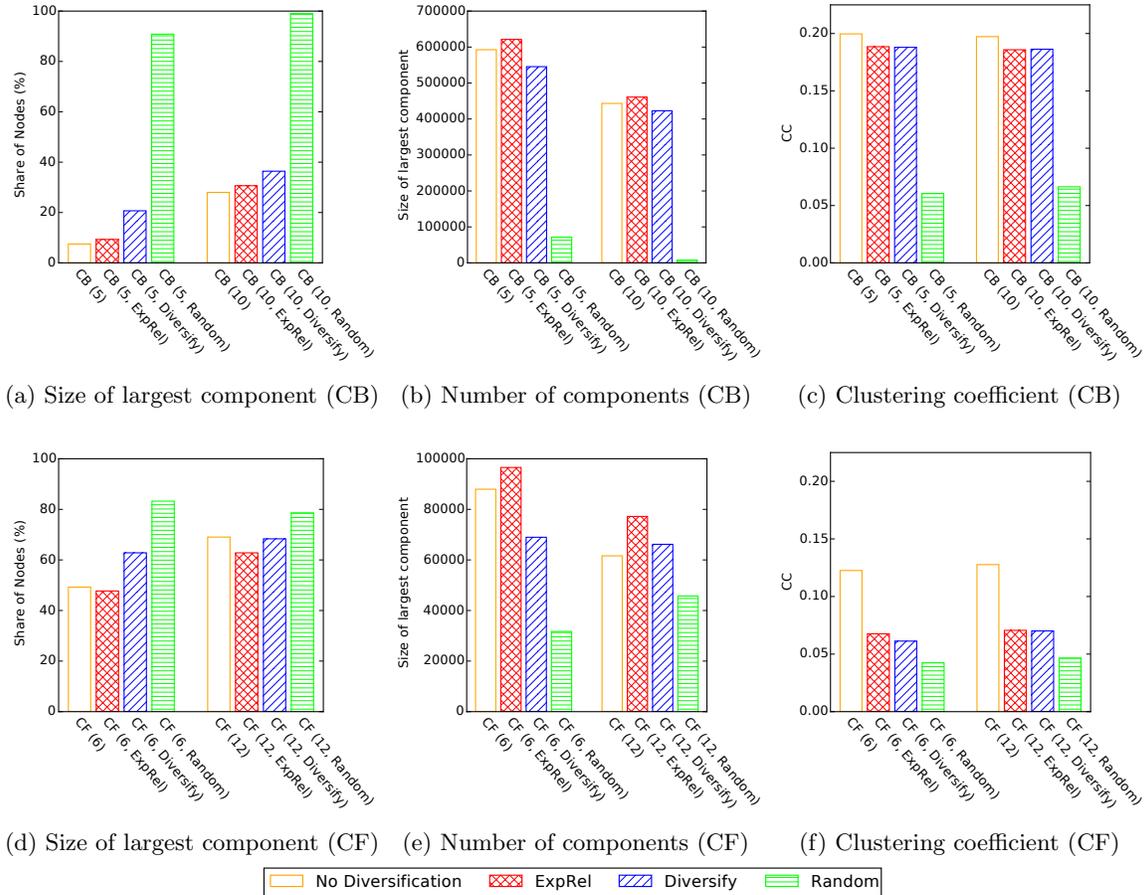


Figure 2: **Topology Analysis.** The figures show the sizes of the largest component, the numbers of components and the clustering coefficients. The unmodified recommendation networks, as present on IMDb, exhibit a comparatively small largest components and a high number of disconnected components. Diversification approaches change this and result in a larger component, while reducing clustering.

4. REACHABILITY

As the first part of our analysis, we study reachability of recommendation networks and analyze what parts of the graph are connected by paths of arbitrary lengths. This represents the basis for further analyses of efficient reachability and partition reachability, which permit us to gain more detailed insight into navigational dynamics.

4.1 Effective Reachability

As the first step, we investigate the fundamental problem of whether a connection between pairs of nodes exists at all.

Strongly connected components.

The largest component enables users to explore all of its items by following recommendation links and is a direct measure for the fraction of the network reachable via navigation. In addition to the largest component, the number of components present in the network shows the division into separate parts that are not interconnected by recommendation links. Figure 2 shows that in their unmodified versions, content-based recommendations led to substantially smaller largest components than collaborative filtering recommendations.

This confirms results from a previous study which found collaborative filtering to lead to larger components [15]. Possible contributing factors are the higher number of recommendations for collaborative filtering (6 and 12 versus 5 and 10 for content-based recommendations) as well as the higher total number of nodes in the content-based network. Diversification approaches were able to increase the size of the largest component substantially. The random diversification demonstrated the maximally achievable improvement, as random graphs are among the graphs with the highest possible reachability.

In terms of numbers of components, the results show that there exists a large number of disconnected components within the recommendation network. For the collaborative filtering recommendations, a major contributor to this is the fact that recommendations pointing away from items existed only for 79% of the nodes in CF (6) and 89% of nodes in CF (12). Those nodes therefore each formed a separate component with only one node in the directed graph. This problem was not present for the content-based recommendation network. However, the large number of disconnected components clearly hinders navigation in both networks. Diversification again mitigated this issue.

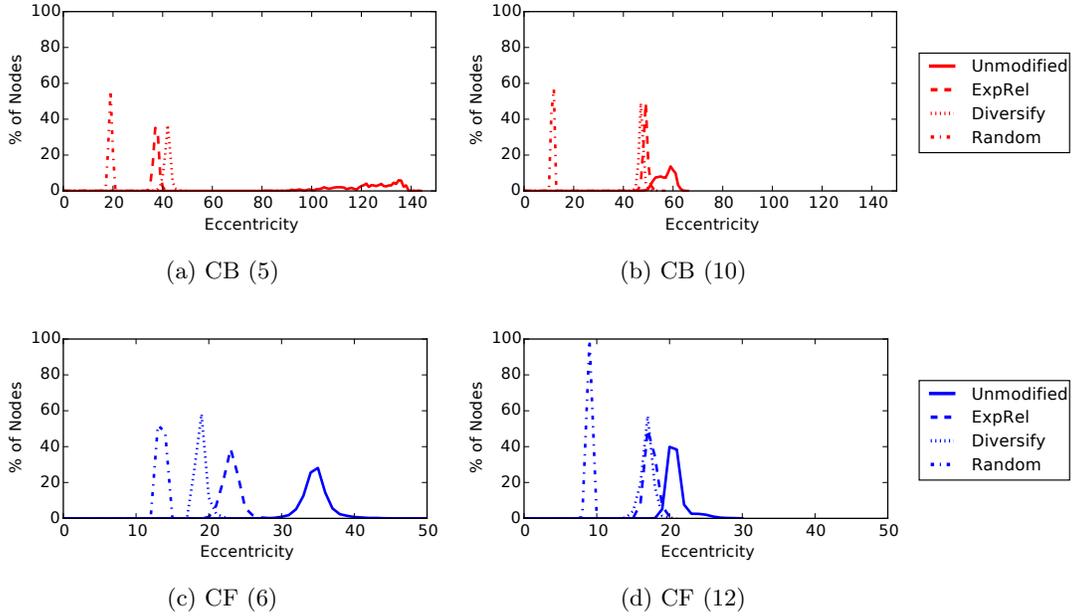


Figure 3: **Eccentricity Analysis.** This figure shows a sampled eccentricity distribution of both unmodified and diversified recommendation networks for a sample of 15% of the nodes in the largest strongly connected component (chosen uniformly at random). Eccentricity measures the longest shortest path from a node to any other node in the same component. With distances of up to 140 hops, the unmodified networks as present on IMDb do not lend themselves to navigation very well. Eccentricities can be reduced by introducing diversified recommendations.

Clustering Coefficient.

The clustering coefficient measures the fraction of neighbors that have a connection among themselves. High clustering implies more predictable browsing (with a large overlap of recommendations between related nodes) while low clustering increases the chance of being able to break out of the local context and follow a diverse or novel recommendation. Generally, high clustering with a few diverse links best supports navigation [26]. We define the clustering coefficient for recommendation networks as

$$C = \frac{1}{|V|} \sum_{i \in V} \frac{|\{(j, k) \in E | j, k \in \Gamma(i, 1)\}|}{|\Gamma(i, 1)| (|\Gamma(i, 1)| - 1)}, \quad (1)$$

where $\Gamma(i, 1)$ is the set of nodes reachable from i in one hop. The results show that the content-based networks exhibit higher clustering coefficients than the collaborative filtering networks. This indicates that content-based recommendations led to more redundancy in the resulting network. Together with the component sizes, it becomes apparent that a trade-off exists between reachability (i.e., the size of the largest component) and navigation predictability (i.e., higher clustering, which leads to better predictability of the area of a network a recommendation leads to).

4.2 Efficient Reachability

As the second step, we study the actual distance between pairs of nodes (given that there exists a path that connects them). This allows us to further investigate how well these networks support navigability and browsing. The probability that a user follows a link instead of typing in another URL or using the search function is around 65% [6] in movie recommender systems. This indicates that path lengths need

to be short to properly support browsing scenarios.

To assess the difficulty of navigation, we evaluate eccentricity distribution on the largest strongly connected component. The eccentricity of a node measures the longest shortest distance between the node and any other node of the same component, therefore allowing us to learn about distances in the recommendation network. For a node $i \in SCC(G)$,

$$ecc(i) = \max_{j \in SCC(G)} d(i, j), \quad (2)$$

where $SCC(G)$ is the largest strongly connected component in G and $d(i, j)$ is the geodesic distance between i and j . To evaluate eccentricity, we sampled the values for 15% of the nodes in the largest strongly connected component (between 8,000-112,000 nodes, chosen uniformly at random).

For the content-based network, eccentricities were comparatively large (cf. Figure 3), with distances reaching up to 140 hops. The collaborative filtering networks exhibited lower eccentricities, rendering them better suited for browsing. Diversification measures lowered eccentricities for both networks.

4.3 Partition Reachability

As the third step, we study reachability based on the bow-tie model. The bow-tie model is a partitioning of a graph into three major components: IN, SCC and OUT, as well as a few additional ones, with the disconnected nodes collected in *OTHER* [1] (see Figure 5 for details). A bow-tie analysis allows us to learn more about the navigational structures in recommendation networks beyond the largest strongly connected component.

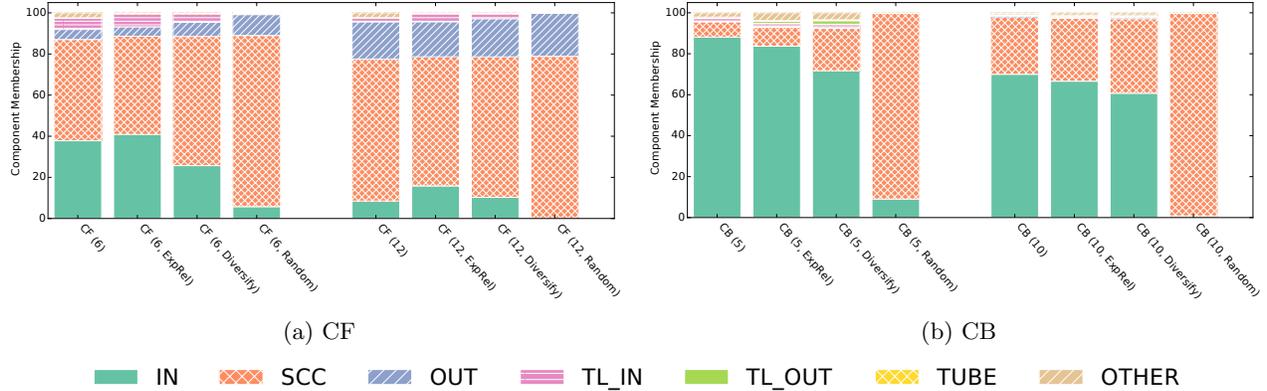


Figure 4: **Bow-Tie Analysis.** The figure shows the partition of the recommendation networks based on the bow-tie model (cf. Figure 5). The recommendation networks of IMDb consisted mainly of nodes in *IN*, *SCC* and *OUT*, implying that they were not in completely disconnected components. Diversification led to a larger share of nodes in the *SCC*.

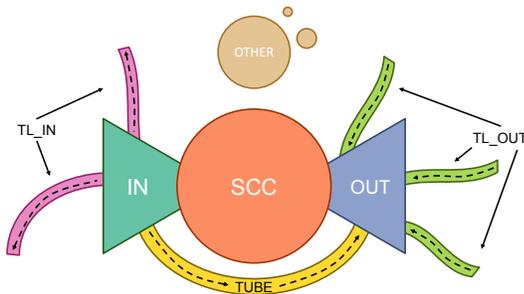


Figure 5: **Bow Tie Model.** The bow tie model [1] is a partitioning of a graph into a strongly connected component or core (*SCC*) as well as *OUT* which is reachable from it and *IN* which is able to reach it. Nodes in *TUBE* are on a detour from *IN* to *OUT*. The *TENDRILS* (*TL_IN*, *TL_OUT*) contain nodes pointing away from *IN* or pointing to *OUT*. Remaining nodes are collected in *OTHER*.

The bow-tie analysis confirmed that the size of the largest strongly connected component increased with diversification (cf. Figure 4). Moreover, it shows that most other nodes were in the *IN* and *OUT* components. From a navigation perspective, this is desirable as it implies that these nodes are either able to reach the largest component or are reachable from it. When following recommendations from a node contained in *IN*, it is likely that a user will be able to reach the *SCC*. Figure 6 depicts the changes in component membership from the unmodified network to a diversified one. Increasing the size of the *SCC* via diversification implies that some of the recommendations from items previously in *IN* now point to nodes in the *SCC* and therefore become themselves a part of it. Note that the number of nodes in *OTHER* components slightly increases due to the fact that diversifying recommendations removes some of the recommendations to sink nodes (that do not have any outgoing recommendations). Navigationwise, this implies that the number of dead-ends encountered by users browsing the recommendation network decreases.

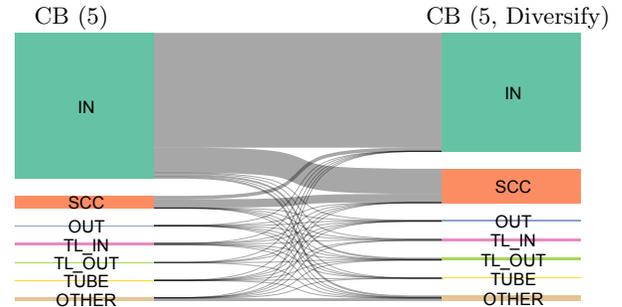


Figure 6: **Bow-Tie Membership Change Analysis for CB (5) to CB (5, Diversify).** Nodes were mostly part of *IN* in the unmodified recommendation network. Diversification moved items from *IN* to *SCC*.

5. NAVIGABILITY

As the first part of our analysis, we studied the reachability of recommendation networks. In the second part, we are now interested in how well the networks fare in terms of actual browsing scenarios. To this end, we simulate browsing in the networks and evaluate the results.

5.1 Start and target nodes

We evaluate browsing scenarios inspired by the desire to find a few movies relevant to certain genres. To this end, we take the genres (e.g., *Action*) as well as the genre combinations (e.g., *Action, Comedy*) as listed on IMDb⁵ for a total of 93 target genres. For each of these genres, we compute the 25 top-rated items with at least 1,000 ratings from our rating dataset and take them as target sets. We restrict our analysis to the largest strongly connected components (cf. Figure 2) and sample 100 start nodes for every target set, leaving us with 9,300 start-target missions to simulate.

5.2 Link selection strategy

Our simulation approach was based on *greedy decentralized search*, a method to analyze navigation dynamics in net-

⁵<http://www.imdb.com/genre/>

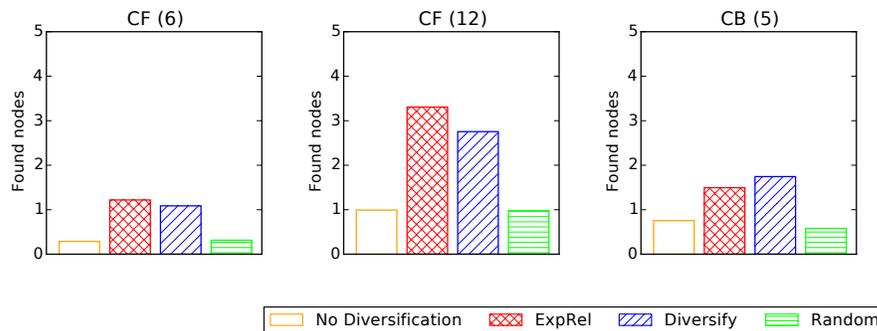


Figure 7: **Nodes found in navigation simulations.** The navigation scenarios were not very well-supported in the unmodified networks, where the simulations found between 0.2 and 1 node per run on average. With diversification approaches, the number of nodes found in the simulation of exploratory navigation scenarios increased by 100 – 300%.

works [8, 16, 23]. The simulation started at a target node and at each step greedily selected the outgoing recommendation with the highest similarity to the target set. The simulation kept track of visited nodes and explored each node only once. In case of a dead-end (no outgoing recommendations at all or no unvisited outgoing recommendations), the simulation backtracked to the previously visited node. We simulated each mission for a total of 50 clicks.

As the background knowledge to inform link selection, we computed the TF-IDF similarities between items by making use of the words contained in the title, plot and storyline descriptions. The value for a potential recommendation link was computed as the similarity between the current item and the average vector of the 25 target nodes. This is similar to the concept of information scent [4], where a link is thought to emanate a certain smell based on its usefulness with respect to the target.

We believe that these are plausible assumptions for users who have some idea where a recommendation could lead based on the information present with the recommendation in the interface (i.e., title and image).

5.3 Results

Figure 7 shows the results for the simulations. Overall, the navigation scenarios were not very well-supported in the unmodified networks, where the simulations found between 0.2 and 1 node per run on average. The diversification approaches were able to improve the outcomes compared to the unmodified recommendation networks substantially: both *ExpRel* and *Diversify* improved the number of found target nodes by 100 – 300%, thus strongly improving navigability in these networks.

Random diversification, however, did not lead to better results than the unmodified networks. Even though the injection of random links led to large components (cf. Figure 2), the resulting lower clustering meant that the similarity information was of little use in informing a navigation process.

6. DISCUSSION AND CONCLUSIONS

In this paper, we analyzed two recommendation networks from IMDb by applying a two-level evaluation approach for recommender systems to study reachability and navigability. In the following, we discuss the findings in the context of our research questions.

RQ 1 *How well do the recommendations of IMDb support reachability and navigability?*

The results of our analysis and our simulations show that with the unmodified recommendations present on IMDb, navigating the network (if at all possible) represents a very hard task for users. Within our navigation simulations, it was possible to retrieve only about one out of 25 target node within 50 steps, even though the target nodes were chosen to be the most popular items in terms of ratings and had a high number of votes.

RQ 2 *How can reachability and navigability on IMDb be improved?*

Applying two simple diversification measures led to improvement of reachability and navigability for both recommendation networks. The number of items the simulations was able to retrieve saw an up to threefold increase, thus making it more realistic for users to be able to gain useful knowledge from exploratory browsing in the network.

RQ 3 *What are differences between collaborative filtering and content-based recommendations in terms of reachability and navigability?*

The collaborative filtering recommendations (the approach currently in use by IMDb) led to a larger strongly connected component, resulting in a larger reachable share of the network than the network for content-based recommendations. However, in terms of the simulated navigation scenarios, content-based recommendation networks fared slightly better. This suggests that content-based recommendations make it easier to reach more popular nodes. Content-based recommendations led to networks with higher clustering—thus making link selection in them more predictable. Another possible explanation for this is that the information used for generating these recommendations overlapped with the background knowledge used in the simulations.

In the diversification measures we applied, we made the assumption that users would prefer between 30 – 40% diversification, based on a study by Ziegler et al. [29]. In future work, it would be interesting to conduct a usability study investigating the results of applying diversification on a live system and testing different fractions of diversified recommendations.

The results of this paper suggest that navigating and browsing recommendations are currently not very-well supported on IMDb. Our work shows a possible way of improving this via diversification measures.

7. ACKNOWLEDGMENTS

This research was supported by a grant from the Austrian Science Fund (FWF) [P24866].

8. REFERENCES

- [1] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. Graph structure in the web. *Computer Networks*, 33(1):309–320, 2000.
- [2] P. Cano, O. Celma, M. Koppenberger, and J. M. Buldu. Topology of music recommendation networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 16(1):013107, 2006.
- [3] Ò. Celma and P. Herrera. A new approach to evaluating novel recommendations. In *Proceedings of the 2nd ACM Conference on Recommender Systems*, pages 179–186. ACM, 2008.
- [4] E. H. Chi, P. Pirolli, K. Chen, and J. Pitkow. Using information scent to model user information needs and actions and the web. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 490–497. ACM, 2001.
- [5] J. Davidson, B. Liebald, J. Liu, P. Nandy, T. Van Vleet, U. Gargi, S. Gupta, Y. He, M. Lambert, B. Livingston, et al. The YouTube Video Recommendation System. In *Proceedings of the fourth ACM Conference on Recommender Systems*, pages 293–296. ACM, 2010.
- [6] D. F. Gleich, P. G. Constantine, A. D. Flaxman, and A. Gunawardana. Tracking the random surfer: Empirically measured teleportation parameters in pagerank. In *Proceedings of the 19th International Conference on World Wide Web*, pages 381–390. ACM, 2010.
- [7] J. Grujić. Movies recommendation networks as bipartite graphs. In *International Conference on Computational Science*, pages 576–583. Springer, 2008.
- [8] D. Helic, M. Strohmaier, M. Granitzer, and R. Scherer. Models of human navigation in information networks based on decentralized search. In *Proceedings of the 24th ACM Conference on Hypertext and Social Media*, pages 89–98. ACM, 2013.
- [9] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*, 22(1):5–53, 2004.
- [10] D. Jannach, M. Zanker, A. Felfernig, and G. Friedrich. *Recommender Systems: An Introduction*. Cambridge University Press, 2010.
- [11] J. Kleinberg. The small-world phenomenon: An algorithmic perspective. In *Proceedings of the thirty-second annual ACM Symposium on Theory of Computing*, pages 163–170. ACM, 2000.
- [12] J. M. Kleinberg. Navigation in a small world. *Nature*, 406(6798):845, August 2000.
- [13] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, (8):30–37, 2009.
- [14] O. Küçüküncü, E. Saule, K. Kaya, and Ü. V. Çatalyürek. Diversified recommendation on graphs: Pitfalls, measures, and algorithms. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 715–726, 2013.
- [15] D. Lamprecht, M. Strohmaier, and D. Helic. Improving Reachability and Navigability in Recommender Systems. *arXiv preprint arXiv:1507.08120*, 2015.
- [16] D. Lamprecht, M. Strohmaier, D. Helic, C. Nyulas, T. Tudorache, N. F. Noy, and M. A. Musen. Using ontologies to model human navigation behavior in information networks: A study based on wikipedia. *Semantic Web*, 6(4):403–422, 2015.
- [17] K. Lerman and L. Jones. Social Browsing on Flickr. In *Proceedings of 1st International Conference on Weblogs and Social Media*, 2007.
- [18] G. Marchionini. Exploratory search: from finding to understanding. *Communications of the ACM*, 49(4):41–46, 2006.
- [19] S. Milgram. The small-world problem. *Psychology Today*, 1:60–67, 1967.
- [20] P. Resnick and H. R. Varian. Recommender systems. *Communications of the ACM*, 40(3):56–58, 1997.
- [21] K. Seyerlehner, P. Knees, D. Schnitzer, and G. Widmer. Browsing music recommendation networks. In *Proceedings of the 10th International Society for Music Information Retrieval Conference*, 2009.
- [22] J. Teevan, C. Alvarado, M. S. Ackerman, and D. R. Karger. The perfect search engine is not enough: a study of orienteering behavior in directed search. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 415–422. ACM, 2004.
- [23] C. Trattner, P. Singer, D. Helic, and M. Strohmaier. Exploring the differences and similarities between hierarchical decentralized search and human navigation in information networks. In *Proceedings of the 12th International Conference on Knowledge Management and Knowledge Technologies*, page 14. ACM, 2012.
- [24] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32, 1969.
- [25] D. J. Watts. Networks, dynamics, and the small-world phenomenon. *American Journal of Sociology*, 105:493–527, 1999.
- [26] D. J. Watts and S. H. Strogatz. Collective dynamics of small-world networks. *Nature*, 393(6684):440–442, 1998.
- [27] R. West and J. Leskovec. Human wayfinding in information networks. In *Proceedings of the 21st International conference on World Wide Web*, pages 619–628. ACM, 2012.
- [28] R. W. White and J. Huang. Assessing the scenic route: measuring the value of search trails in web logs. In *Proceedings of the 33rd International ACM SIGIR conference on Research and development in information retrieval*, pages 587–594. ACM, 2010.
- [29] C.-N. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen. Improving recommendation lists through topic diversification. In *Proceedings of the 14th International Conference on World Wide Web*, pages 22–32. ACM, 2005.