

Assessing Impacts of a Power User Attack on a Matrix Factorization Collaborative Recommender System

Carlos E. Seminario and David C. Wilson

University of North Carolina at Charlotte
Charlotte, North Carolina, USA
cseminar@uncc.edu, davils@uncc.edu

Abstract

Collaborative Filtering (CF) Recommender Systems (RSs) help users deal with the information overload they face when browsing, searching, or shopping for products and services. Power users are those individuals that are able to exert substantial influence over the recommendations made to other users, and RS operators encourage the existence of power user communities and leverage them to help fellow users make informed purchase decisions, especially on new items. Attacks on RSs occur when malicious users attempt to bias recommendations by introducing fake reviews or ratings; these attacks remain a key problem area for system operators. Thus, the influence wielded by power users can be used for both positive (addressing the “new item” problem) or negative (attack) purposes. Our research is investigating the impact on RS predictions and top-N recommendation lists when attackers emulate power users to provide biased ratings for new items. Previously we showed that power user attacks are effective against user-based CF RSs and that item-based CF RSs are robust to this type of attack. This paper presents the next stage in our investigation: (1) an evaluation of heuristic approaches to power user selection, and (2) evaluation of power user attacks in the context of matrix-factorization (SVD) based recommenders. Results show that social measures of influence such as degree centrality are more effective for selection of power users, and that matrix-factorization approaches are susceptible to power user attacks.

1 Introduction

Robustness is one of the key problem areas in Recommender Systems (RSs). Although attacks on RSs have been researched in the past, users with malicious intent continue to find ways to bias predictions and disrupt the system. The problem with RS attacks is that, if left undetected or unmitigated, the system’s knowledge base becomes compromised, can generate biased recommendations for users, can cause users to waste time and money on inaccurate or false recommendations, and can also diminish the users’ trust in the overall system. In the Collaborative Filtering (CF) RS context, *power users* are those who can exert considerable influence over the recommendations presented to other users. Research has indicated that power users can have major

impacts on RS ratings predictions and top-N recommendations lists especially when the underlying RS algorithms are neighborhood-based (Lathia, Hailes, and Capra 2008) and that power user attacks on user-based systems are effective when power users are selected using techniques based on the underlying user-user relationships (Wilson and Seminario 2013). Furthermore, new items can sometimes encounter difficulty with market awareness/acceptance; to address this issue, marketers may rely on power users to help influence item recommendations to other users (Domingos and Richardson 2001; Anand and Griffiths 2011). And RS operators encourage the existence of power user communities, e.g., Amazon VineTM, to help their fellow users make informed purchase decisions¹. But it is possible for such influence to be used for malicious purposes as well as legitimate. And we are interested in studying what happens when power user influence is co-opted with biased ratings to game the system and sway recommendation results.

To address this issue, we investigate how power users can be identified and selected, as well as the impact of power user focused attacks on RSs. We adopt concepts of degree centrality from Social Network Analysis (Wasserman and Faust 1994) to select influential power users from the implied social graph of the RS user-user similarity matrix (Palau et al. 2004; Wilson and Seminario 2013), as well as other heuristic methods. We also use the “Power User Attack” (PUA) model (Wilson and Seminario 2013) to evaluate accuracy and robustness impacts of the PUA on a matrix factorization SVD-based CF recommender in order to determine the effectiveness of the attack on a model-based system.

This study provides results of experiments designed to answer the following research questions: (1) How are power users best identified in a RS? (2) What happens to the robustness of an SVD-based RS after the Power User Attack on new items? (3) How do the popular RS algorithms (user-based, item based, and SVD-based CF) compare in their robustness against power user attacks?

2 Related Work

Social Network Analysis affords well-known concepts that can be readily applied to RS power users and their influence

¹<http://www.amazon.com/gp/vine/help>

in the underlying RS social graph. The concept of Degree Centrality (Wasserman and Faust 1994) indicates that nodes (users) who have more edges (connections) to other nodes may have advantages; high in-degree refers to nodes that many other nodes connect to and corresponds to high prominence, prestige, or popularity and high out-degree refers to nodes that connect to many other nodes and corresponds to high expansiveness. Power users are of particular interest to RS operators and their client companies when launching a new item, because a positive endorsement (high rating) can result in making item recommendations to many other users. This “market-based” use of RS has been previously promoted as a solution to the “cold-start” or “new item” problem (Anand and Griffiths 2011). A viral marketing perspective to exploit the network value of customers was studied in (Domingos and Richardson 2001).

Attacks on RSs by providing false ratings are known as “shilling attacks” (Lam and Riedl 2004), or “profile injection attacks” (Mobasher et al. 2007; O’Mahony, Hurley, and Silvestre 2005). Research in attacks on recommender systems began in 2002 (O’Mahony, Hurley, and Silvestre 2002); a recent summary in (Burke, O’Mahony, and Hurley 2011) describes RS attack models, attack detection, and algorithm robustness. Previous work in this area has focused on the use of attack models based on hypothetical users that inject attack user profiles containing either random item ratings whose values are selected from a normal distribution around the mean rating of the dataset (this is not a very effective attack), item ratings whose values are selected from a normal distribution around the mean rating for each item (a more effective attack against neighborhood-based collaborative filtering algorithms), or a variant of these approaches (Lam and Riedl 2004; Burke et al. 2006; Mobasher, Burke, and Sandvig 2006; Mobasher et al. 2007; Williams et al. 2006). These hypothetical users are not representative of actual users, in fact, they are more like statistically “average” users.

Therefore, a gap in the research is that characteristics of real, more influential, “power” users have largely been ignored. Furthermore, research on attacks has revealed that knowledge of the underlying RS algorithms and dataset characteristics can enhance the effectiveness of an attack even though this knowledge may be difficult to obtain. For clarity, the power user attack envisioned in this research is not about having hundreds or thousands of actual power users colluding to mount an attack, rather, it is about an attacker being able to generate a set of synthetic power user profiles that, when stealthily injected into a RS, can effectively bias the recommendations. Ultimately, having the ability to generate a set of synthetic user profiles with power user properties can leave systems vulnerable to exploitation from more subtle, yet powerful, attacks by highly-influential power users. And this remains an open question in RS robustness research that we continue to explore in this study.

3 Singular Value Decomposition (SVD)

Recommender Systems, with thousands of users (rows) and items (columns), consist of a dataset with features that define a high-dimensional space and have sparse information

in that space. The data matrix is sparse because, typically, most of the users have only rated a small percentage of the items available. High dimensional data can be difficult to work with because adding more features can increase noise and error; also, there are not enough data points to get good estimates or predictions. To deal with this problem, dimensionality reduction techniques such as Singular Value Decomposition (SVD) have been used (Sarwar et al. 2000; Koren, Bell, and Volinsky 2009; Amatriain et al. 2011). SVD has been used as a tool to improve collaborative filtering by uncovering latent relationships between users and items and, once the matrix is factored, used to compute predictions.

The implementation of matrix factorization SVD we used was the Expectation Maximization (EM) algorithm (Dempster, Laird, and Rubin 1977) provided in the Apache Mahout platform². This algorithm requires two parameters: number of features and number of training steps. A sensitivity analysis was performed on these parameters to observe the impact on Mean Absolute Error (MAE) using the MovieLens ML100K dataset³. Based on these results, we found that when holding the number of training steps constant, MAE remains relatively flat as the number of features is varied. Conversely, when holding the number of features constant, MAE decreases to a minimum and then begins to increase. For 100 features, the minimum MAE occurs at 75 training steps; the differences in MAE between 25 and 75 steps and between 100 and 75 steps are significant ($p < 0.01$).

4 Power User Attack

In (Wilson and Seminario 2013), we defined a novel *Power User Attack* (PUA) model as a set of power user profiles with biased ratings that influence the results presented to other users. The PUA is distinct from previously studied types of RS attacks (O’Mahony, Hurley, and Silvestre 2005; Lam and Riedl 2004; Mobasher et al. 2007), e.g., “random”, “average”, “bandwagon”, etc., that rely on a set of carefully configured false user profiles which are injected into the dataset to mount the attack. The PUA relies critically on the method of power user identification/selection, so we also developed and evaluated a novel use of degree centrality concepts from social network analysis for identifying influential RS power users for attack purposes (Wilson and Seminario 2013). In that work, the PUA was evaluated using user-based and item-based CF recommender algorithms.

Like other attack models, PUA profiles contain the set of ratings a power user has made using the recommender system. The intent of an attack is to either promote (“push”) a target item by setting the rating to the maximum value or demote (“nuke”) a target item by setting the rating to the minimum value. The PUA consists of one or more user profiles containing item ratings (called attack user profiles) that push or nuke a specific item. However, unlike classic attack models (e.g., random, average, bandwagon) that employ straightforward statistical templates (e.g., average rating, popularity, and likability) to generate synthetic attack profile filler items

²<http://mahout.apache.org/>

³www.grouplens.org/ MovieLens dataset with 100,000 ratings, 1,682 movies, 943 users, 93.7% sparsity.

(Mobasher et al. 2007), very little is known about the profile characteristics of power users. And without this knowledge, it is difficult to generate synthetic power user profiles. Thus, our initial evaluation of this attack model employs selected power user profiles that already exist in the dataset to simulate injected attack user profiles. This is a limitation of the current study for PUAs in general, but serves to show the potential effectiveness of the PUA model for impacting impacting various RS algorithms. If an effective PUA could not be mounted by real power users, there would be little potential for PUAs based on modeling them. As with all attack model research, the potential difficulty or overhead cost for an attacker to acquire the necessary data is a consideration. And the degree of difficulty varies according to the methods employed. We believe PUAs to be well within the realm of possibility (e.g., insider, data breach, side-channel analytics) and focus here on their potential effectiveness.

To implement this attack, a group of power users are selected (see §4.1), the attack intent and target item(s) are specified, and the remainder of the profile for the PUA (the “filler”) remains unchanged for each power user in the attack. By keeping the power users’ profiles the same and then adding the target item rating for the attack, the power users’ connections to other users in the implied social graph remain essentially the same⁴. The number of power users participating in the attack defines the attack size; the larger the attack size, the larger the expected disruption in RS predictions and top-N recommendation lists.

4.1 Power User Selection

Power users in the RS context have been referred to as users with a large number of ratings (Herlocker et al. 2004) as well as those that are able to influence the largest number of other users (Domingos and Richardson 2001; Rashid, Karypis, and Riedl 2005; Anand and Griffiths 2011; Goyal and Lakshmanan 2012). To measure influence, (Rashid, Karypis, and Riedl 2005) used the number of prediction differences above a prediction threshold when a user is removed from the dataset, (Goyal and Lakshmanan 2012) used the number of users that had the prediction for a target item shifted sufficiently above a threshold so that the item appears in their top-N list, (Anand and Griffiths 2011) used MAE and coverage to evaluate various seed (influential user) selection algorithms, and (Domingos and Richardson 2001) used the expected lift in profit earned by influencing other users, recursively. Although maximizing the spread of influence through a social network is an NP-hard problem to solve optimally, several heuristics were analyzed by (Goyal and Lakshmanan 2012) to select groups of influential users including those with highest aggregate similarity to other users, highest positive average rating, and highest number of item ratings. The Number of Unique Prediction Differences algorithm (Rashid, Karypis, and Riedl 2005) was determined to be computationally inefficient and was not considered fur-

⁴In a few cases, power user profiles that already had a target item rating were updated in certain attack scenarios and, although the target item rating change might alter their neighborhoods, we believe the impact to this analysis is not an issue.

ther in our study.

We have developed an approach to power user selection for attack purposes (Wilson and Seminario 2013), based on social network analysis concepts of Degree Centrality (Wasserman and Faust 1994; Lathia, Hailes, and Capra 2008). Specifically, we use In-Degree Centrality (users who appear in the highest number of other users’ neighborhoods) with significance weighting (Herlocker et al. 1999) because when using similarity and neighborhood-based methods to select power users, significance weighting encourages strong connections between users who have rated many items in common. In an initial study, we found that both our approach and the Most Central heuristic (Goyal and Lakshmanan 2012) performed significantly better using significance weighting.

To evaluate the power user selection methods in this study, we use an ablation approach (Lathia, Hailes, and Capra 2008; Wilson and Seminario 2013), i.e., accuracy of the RS is measured as power users are removed from the dataset. If accuracy gets worse when power users are removed, the interpretation is that power users are impacting the RS recommendations; the power user selection method that is able to negatively impact the most is the better method.

5 Experimental Design

To address our research questions, we conducted an experiment using the MovieLens 100K dataset with an SVD-based recommender. Power users were selected from the dataset using three identification/selection methods. To simulate the PUA, power user profiles were converted to attack profiles by setting target items in those profiles to the maximum rating. Target items selected had no more than one rating in order to simulate a “new” item. Evaluations of accuracy and robustness were performed before and after the attack.

Evaluation Metrics: We use Mean Absolute Error (MAE) and prediction coverage for accuracy and coverage (Herlocker et al. 2004; Shani and Gunawardana 2011) using a holdout-partitioned 70/30 train/test dataset. We also use Hit Ratio, Prediction Shift, and Rank robustness measures (Mobasher et al. 2007; Burke, O’Mahony, and Hurley 2011) where a high Hit Ratio and a low Rank indicates that the attack was successful (from the attacker’s standpoint). Since the PUA being evaluated here is for new items (zero rating value), the Prediction Shift is expected to be close to the maximum rating as defined by the RS.

Datasets and Algorithms: We used the ML100K dataset with item ratings from 1 (did not like) to 5 (liked very much). For the SVD-based CF algorithm, we used the EM (see §3) algorithm as implemented in Mahout 0.4. Run-time parameters used for this algorithm were number of features (100) and number of training steps (75); settings were determined empirically as described in §3. The more traditional user-based and item-based CF algorithms were studied in a previous effort (Wilson and Seminario 2013) and those results will be used here for comparative purposes.

Power User Selection: The following methods were used, *InDegree*: Our method is based on the in-degree centrality concept from social network analysis, where power users

are those who participate in the highest number of neighborhoods. For each user i compute its similarity with every other user j applying significance weighting, then discard all but the top 50 neighbors for each user i . Count the number of similarity scores for each user j (# neighborhoods user j is in) and select the top 50 user j 's.

AggregatedSimilarity (AggSim): This is the Most Central heuristic from (Goyal and Lakshmanan 2012). The top 50 users with the highest aggregate similarity scores become the selected set of power users. This method requires at least 5 co-rated items between user i and user j and does not use significance weighting⁵.

NumberRatings (NumRatings): This method is based on (Herlocker et al. 2004) where “power user” refers to users with the highest number of ratings; it also is called the Most Active heuristic in (Goyal and Lakshmanan 2012). We selected the top 50 users based on the total number of ratings they have in their user profile.

Target Item Selection: For the ML100K dataset, 5 target items with no more than one rating, regardless of their rating value, were selected randomly, given our objective to attack only ‘new’ items. We recognize that 5 target items is a limitation in this study; however, new items are more vulnerable to attack than more popular items so this should provide a strong signal even with a small number of target items. We are considering a larger mix of new/existing target items as a future work.

Attack Parameter Selection: The Attack Intent is Push, i.e., target item rating is set to max (= 5). The Attack Size or number of power users in each attack is 50, 30, 10, 5, 3, 2, and 1. The maximum attack size (50) was selected based on previous research (Mobasher et al. 2007; Burke, O’Mahony, and Hurley 2011), where a 5-10% attack was shown to be effective; with ML100K, a 5% attack size is about 50 users. The attack profiles used were actual power user profiles and we added the target item rating. The Filler Size, or number on non-target items in each attack user profile, is determined by each power users’ profile size; therefore, filler size is not specified in this experiment.

Test Variations: One prediction algorithm, one dataset, three power user selection methods, and seven attack sizes. Each test variation was executed once for each of the 5 target items and data results were averaged over the 5 target items.

6 Results and Discussion

(1) *How are power users best identified in a RS?* Our assertion is that the amount of influence power users exerted on other users, before and after an attack, would indicate the best identification method. **Before the power user attack**, one measure of influence is the negative impact on RS accuracy (MAE) when removing power users (Lathia, Hailes, and Capra 2008). We removed from 0 to 50 (0% to 100%) of the identified power users from the dataset before any attacks took place for all three methods of power user selection; the most influential power users identified are removed first. The results for InDegree (Figure 1) show that as power




Figure 1: MAE impacts after removing Power Users using ML100K




Figure 2: MAE impacts after removing Power Users using ML100K

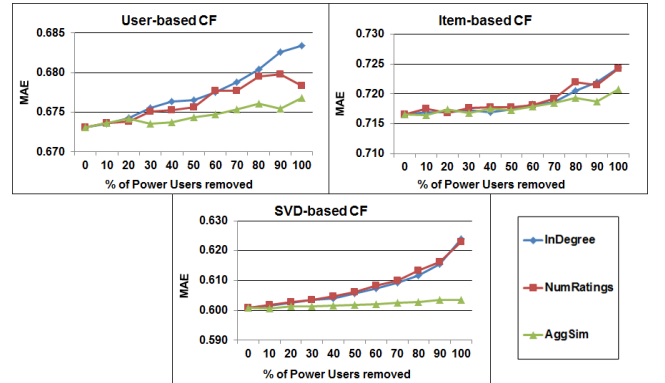


Figure 3: MAE impacts, after removing Power Users using ML10M

users are removed, accuracy impacts are significant on SVD-based recommenders ($p < 0.01$) when power users removed are $> 20\%$. Similar results occurred for the NumRatings method when power users removed are $> 60\%$, and influence of AggSim-selected power users remained flat. Furthermore, InDegree has significantly more impact on MAE than AggSim ($p < 0.01$) at all levels of power user removal and NumRatings ($p < 0.01$) when power users removed are $> 20\%$ and $< 80\%$. As a baseline, we removed users at random and found that the ablation curve for randomly-selected users is flat from 0% to 100% removed, i.e., their removal shows no significant impact on MAE. Coverage results (not shown) remained flat and at a high level ($> 99\%$) for all power user selection methods and number of power users

⁵Based on personal communication with the authors.

removed. The results obtained here are also consistent with those observed in our previous work (Wilson and Seminario 2013; Seminario 2013) using various CF algorithms: Figure 2 shows results using the ML100K dataset and Figure 3 shows results using the ML10M dataset⁶. **After the attack**, we expect for influence to be measured mainly by the impact on robustness metrics, i.e., the method that selects the most influential set of power users is the one producing the highest Hit Ratio and lowest Rank. Results show that all the power user selection methods were successful (from the attacker’s standpoint) at impacting the robustness metrics.

(2) *What happens to the robustness of an SVD-based RS after the Power User Attack on new items?* We found that the PUA was successful (from the attacker’s standpoint) at impacting RS robustness metrics across all three power user selection methods, as indicated by the Average Hit Ratio and Average Rank results shown in Figure 4; no significant differences were found between the three methods with 50 power user attack profiles. High levels of Average Hit Ratio and low levels of Average Rank were achieved with as few as 5 to 10 power users. Impacts to the robustness metrics indicate that a small number of power users⁷ can have significant effects on RS predictions and top-N recommendation lists for new items. With 50 power user attack profiles, the InDegree method showed a significantly lower (better) Average Rank than AggSim ($p < 0.01$) and a significantly higher (worse) Average Rank than NumRatings ($p < 0.01$). As expected, Prediction Shift (not shown) was high (> 4) given that the target items were “new” items.

This result is interesting given that SVD-based systems have been shown to be robust to attacks (Mehta and Nejdl 2009). In that work, the authors used clustering techniques to identify the attackers based on their statistical signatures, i.e., Random, Average, and Bandwagon attack models; the attack clusters were then eliminated from, or ignored during, the prediction calculation. In our experiment, the attackers were not eliminated from the dataset nor ignored during the prediction calculation, therefore, we see a more effective attack against the SVD algorithm.

(3) *How do the popular RS algorithms (user-based, item based, and SVD-based CF) compare in their robustness against power user attacks?* As noted above, robustness results (Figure 4) indicate that SVD-based recommenders are vulnerable to attack by power users with results comparable to user-based recommenders as shown on the left side of Figure 5 (Wilson and Seminario 2013), especially for the InDegree and NumRatings power user selection methods. The right side of Figure 5 indicates that item-based recommenders are less vulnerable to the impacts of the PUA. User-based CF is significantly more vulnerable to the power user attack than Item-based CF and is also consistent with previous findings (Lam and Riedl 2004; Mobasher et al. 2007; Burke, O’Mahony, and Hurley 2011) because the PUA, like the random and average attacks, are able to exploit the similarity between the attackers and non-attackers to favor the




Figure 4: ML100K – SVD-based Results

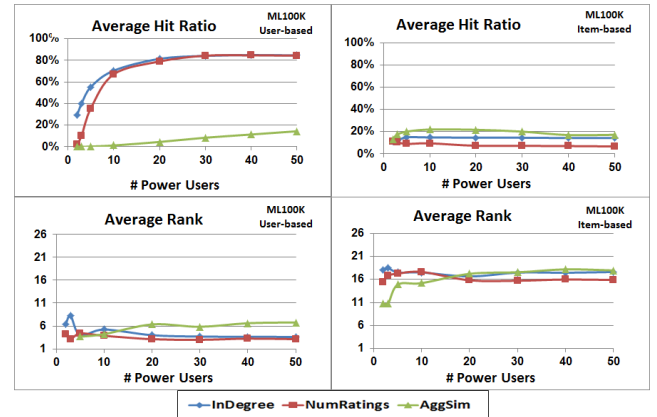


Figure 5: ML100K – User and Item-based Results

target item. For item-based CF, the AggSim method produced a more effective set of power users for the attack as compared to InDegree and NumRatings; however, the impact of the attack was weak, i.e. relatively low Hit Ratio and high Rank, compared to user-based CF.

Compared to user-based and item-based algorithms, we have shown a strong attack using an EM implementation of SVD although it appears insensitive to power user selection methods. Additional research is required to determine whether this is due to scale given ML100K’s size, use of the EM SVD algorithm vs. other SVD techniques, or the input parameters to the EM SVD algorithm.

⁶MovieLens dataset with 10,000,054 ratings, 10,676 movies, 69,878 users, 98.7% sparsity

⁷Note that 10 power users is $< 1\%$ of the ML100K user base.

7 Conclusion

This paper examined the use of social network based heuristics for identifying power users as part of attack vectors on recommender systems, as well as the impact of mounting a Power User Attack in the context of using an SVD-based recommender. We have shown that power users contribute to the improved prediction accuracy of SVD-based recommender systems and is consistent with our previous work using user-based and item-based recommenders and other datasets. We have also shown that a relatively small number of power users ($< 1\%$ of the user base) can have significant effects on RS predictions and top-N recommendation lists for new items using an EM implementation of a SVD-based recommender. Our work also indicates that the InDegree and NumberRatings methods of power user selection produce more influential set of power users than an Aggregated Similarity method. Our future work in this area will examine other contexts and approaches, including (1) additional matrix factorization approaches, (2) target item variations, (3) expanded domains/datasets, and (4) modeling power user profile characteristics for generating synthetic power user attack profiles to formally mount the PUA. The longer-term objective (future) is to develop a power user model so that synthetic power user attack profiles can be configured. Our conjecture is whether this research can result in a generalized model of power user characteristics that can be applied across different domains and RS algorithms.

References

- Amatriain, X.; Jaimes, A.; Oliver, N.; and Pujol, J. M. 2011. Data mining methods for recommender systems. In Ricci, F., et al., eds., *Recommender Systems Handbook*. Springer.
- Anand, S. S., and Griffiths, N. 2011. A market-based approach to address the new item problem. In *Proceedings of the 5th ACM Recommender Systems Conference*.
- Burke, R.; Mobasher, B.; Williams, C.; and Bhaumik, R. 2006. Classification features for attack detection in collaborative recommender systems. In *KDD '06: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM.
- Burke, R.; O'Mahony, M. P.; and Hurley, N. J. 2011. Robust collaborative recommendation. In Ricci, F., et al., eds., *Recommender Systems Handbook*. Springer.
- Dempster, A. P.; Laird, N. M.; and Rubin, D. B. 1977. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B* 39(1).
- Domingos, P., and Richardson, M. 2001. Mining the network value of customers. In *Proceedings of KDD'01*. ACM.
- Goyal, A., and Lakshmanan, L. V. S. 2012. Recmax: Exploiting recommender systems for fun and profit. In *Proceedings of the KDD'12 Conference*.
- Herlocker, J. L.; Konstan, J. A.; Borchers, A.; and Riedl, J. 1999. An algorithmic framework for performing collaborative filtering. In *Proceedings of the ACM SIGIR Conference*.
- Herlocker, J. L.; Konstan, J. A.; Terveen, L. G.; and Riedl, J. 2004. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems* 22(1).
- Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *Computer*.
- Lam, S. K., and Riedl, J. 2004. Shilling recommender systems for fun and profit. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*. ACM.
- Lathia, N.; Hailes, S.; and Capra, L. 2008. knn cf: A temporal social network. In *Proceedings of the 2nd ACM Recommender Systems Conference (RecSys '08)*.
- Mehta, B., and Nejdl, W. 2009. Unsupervised strategies for shilling detection and robust collaborative filtering. *User Modeling and User-Adapted Interaction* 19(1-2).
- Mobasher, B.; Burke, R.; Bhaumik, R.; and Williams, C. 2007. Toward trustworthy recommender systems: An analysis of attack models and algorithm robustness. *ACM Trans. Internet Technol.* 7(4):23.
- Mobasher, B.; Burke, R.; and Sandvig, J. 2006. Model-based collaborative filtering as a defense against profile injection attacks. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI'06)*.
- O'Mahony, M. P.; Hurley, N.; and Silvestre, G. C. M. 2002. Promoting recommendations: An attack on collaborative filtering. In *Proceedings of DEXA'02*.
- O'Mahony, M. P.; Hurley, N.; and Silvestre, G. C. M. 2005. Recommender systems: Attack types and strategies. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05)*.
- Palau, J.; Montaner, M.; Lopez, B.; and Rosa, J. L. D. L. 2004. Collaboration analysis in recommender systems using social networks. In *Cooperative Information Agents VIII: 8th International Workshop, CIA 2004*.
- Rashid, A.; Karypis, G.; and Riedl, J. 2005. Influence in ratings-based recommender systems: An algorithm-independent approach. In *Proceedings of the SIAM International conference on Data Mining, 2005*.
- Sarwar, B. M.; Karypis, G.; Konstan, J. A.; and Riedl, J. T. 2000. Application of dimensionality reduction in recommender system – a case study. In *ACM WEBKDD WORKSHOP*.
- Seminario, C. E. 2013. Accuracy and robustness impacts of power user attacks on collaborative recommender systems. In *Proceedings of the 7th ACM conference on Recommender Systems, RecSys '13*.
- Shani, G., and Gunawardana, A. 2011. Evaluating recommendation systems. In Ricci, F., et al., eds., *Recommender Systems Handbook*. Springer.
- Wasserman, S., and Faust, K. 1994. *Social Network Analysis: Methods and Applications*. New York, NY: Cambridge University Press.
- Williams, C.; Mobasher, B.; Burke, R.; Bhaumik, R.; and Sandvig, J. 2006. Detection of obfuscated attacks in collaborative recommender systems. In *Proceedings of the ECAI '06 Conference Workshop on Recommender Systems*.
- Wilson, D. C., and Seminario, C. E. 2013. When power users attack: assessing impacts in collaborative recommender systems. In *Proceedings of the 7th ACM conference on Recommender Systems, RecSys '13*.