

Automatically Building Research Reading Lists

Michael D. Ekstrand¹ Praveen Kanaan¹ James A. Stemper²
John T. Butler² Joseph A. Konstan¹ John T. Riedl¹
ekstrand@cs.umn.edu

¹GroupLens Research
Department of Computer Science and Engineering
University of Minnesota

²University of Minnesota Libraries

RecSys 2010



Recommending Research Papers

Way too many papers to read — need to find the right ones so we can get on to research!

Many papers, many reviewers — who is most qualified to judge the work?

Recommending Research Papers

Way too many papers to read — need to find the right ones so we can get on to research!

Many papers, many reviewers — who is most qualified to judge the work?

Working on a bibliography or reference list — who should have cited you?

Building a Reading List

Alice is a new grad student who took a class on adaptive web technologies. The professor assigned some readings, and she wants to learn more.

- Herlocker et al. An algorithmic framework for performing collaborative filtering. In *Proc. SIGIR 1999*.
- Hoffman. Latent semantic models for collaborative filtering. In *TOIS*, volume 22, issue 1 (January 2004).
- Guy et al. Personalized recommendation of social software items based on social relations. In *Proc. RecSys 2009*.
- Karypis. Evaluation of Item-Based Top- N Recommendation Algorithms. In *Proc. CIKM 2001*.

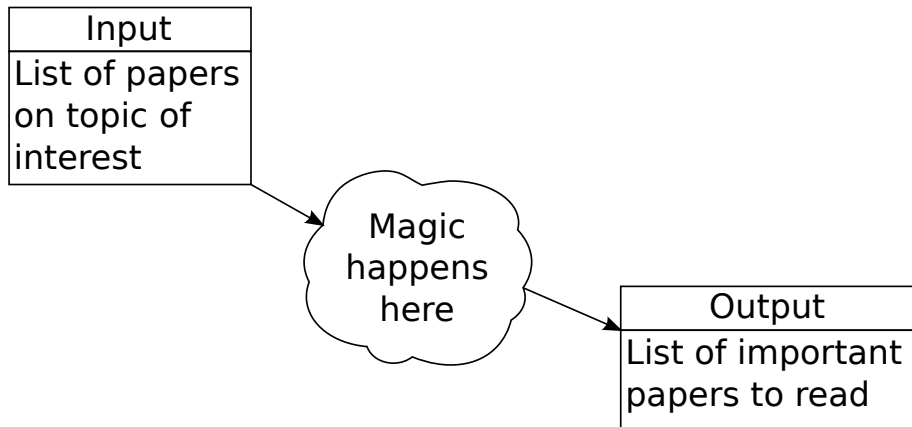
Building a Reading List

Alice is a new grad student who took a class on adaptive web technologies. The professor assigned some readings, and she wants to learn more.

- Herlocker et al. An algorithmic framework for performing collaborative filtering. In *Proc. SIGIR 1999*.
- Hoffman. Latent semantic models for collaborative filtering. In *TOIS*, volume 22, issue 1 (January 2004).
- Guy et al. Personalized recommendation of social software items based on social relations. In *Proc. RecSys 2009*.
- Karypis. Evaluation of Item-Based Top- N Recommendation Algorithms. In *Proc. CIKM 2001*.

How can we find important ones to read?

The *Reading List* Task



Domain Characteristics

Traditional CF methods are blind to item content or relationships.


The citation web has a defined structure.

Can we harness this to improve recommendation?


Questions

- Can we harness domain structure to improve recommendation of research papers?
- What algorithms perform well at building reading lists?
- How can we measure this performance?


Approach



Approach



Designing Candidate Algorithms



Algorithms

We adapt and combine off-the-shelf algorithms.

Collaborative filtering (CF)

Item-item CF over sets of items [Karypis, 2001]

Content-based filtering (CBF)

Lemur toolkit [Ogilvie and Callan, 2002] in BM25 mode with recommended baseline parameters

Graph ranking

- PageRank [Page et al., 1999]
- HITS authority scores [Kleinberg, 1999]
- SALSA [Lempel and Moran, 2000]
- Relative algorithms [White and Smyth, 2003]
 - ▶ Biased HITS
 - ▶ k -step Markov importance

Applying Item-Item CF

No users - papers are both users and items [McNee et al., 2002].

	A	B	C
A	✓		✓
B	✓	✓	
C		✓	✓

Becomes a form of co-citation analysis.

Rank-Weighting CF

Inspired by [Karypis, 2001]: user influence in item similarity does not need to be the same!

Karypis's approach: normalize user purchase vectors to unit vectors so users with few purchases influence item similarity more.

Our adaptation: weight citation vectors by graph rank. Highly-ranked papers have more influence.

$$\hat{\mathbf{a}} = r(a)\mathbf{a}$$


Use weighted citation vectors to compute cited paper similarity.

Using Weights with CBF

Two approaches

Subgraph ranking

Build subgraph from CBF results and rank its nodes.




Using Weights with CBF

Two approaches

Linear blending

Blend CBF scores with node weights in the global graph.



Learned coefficients with multivariate logistic regression (more on this later).

Hybridizing Recommenders [Burke, 2002]

Several ways of blending algorithms

- Use CF as input to CBF
- Use CBF as input to CF
- Blend output from CBF and CF to produce result


Basic algorithms + all hybrids = **177 algorithms**.

Hybridizing Recommenders [Burke, 2002]

Several ways of blending algorithms

- Use CF as input to CBF
- Use CBF as input to CF
- Blend output from CBF and CF to produce result

Basic algorithms + all hybrids = **177 algorithms**.




Hybridizing Recommenders [Burke, 2002]

Several ways of blending algorithms

- Use CF as input to CBF
- **Use CBF as input to CF**
- Blend output from CBF and CF to produce result

Basic algorithms + all hybrids = **177 algorithms**.




Hybridizing Recommenders [Burke, 2002]


Several ways of blending algorithms

- Use CF as input to CBF
- Use CBF as input to CF
- **Blend output from CBF and CF to produce result**

Basic algorithms + all hybrids = **177 algorithms**.



Pruning the Algorithm Pool



Designing an Evaluation Strategy

Goal: Measure algorithm performance at supporting a specific task.

Accuracy isn't enough [McNee et al., 2006].

Offline Evaluation

Use metadata from ACM Digital Library.

Simulate introductory reading list with hold-out test on articles in *ACM Computing Surveys*.

- Hold out 5 items from each citation list
- Attempt to recommend back the 5 items
- Skip surveys with less than 15 resolved citations

Result: 220 survey articles for training and testing.

Interlude - Learning CBF Blending Coefficients

Use half the articles to train the CBF blend.

Learned multivariate logistic regression with response of 1 for articles in the holdout set, 0 for other articles.

$$s(a) = \alpha L(a) + \beta r(a)$$

Since we're ranking, intercept and log don't matter.

Offline Evaluation (cont)

Measure using half-life utility metric [Breese et al., 1998]:


$$R_a = \sum_i \frac{u_{a,i}}{2^{(i-1)/(\alpha-1)}}$$

- $\alpha = 5$, the length of our reading lists
- $u_{a,i} = 1$ if article i is cited by survey a , 0 otherwise

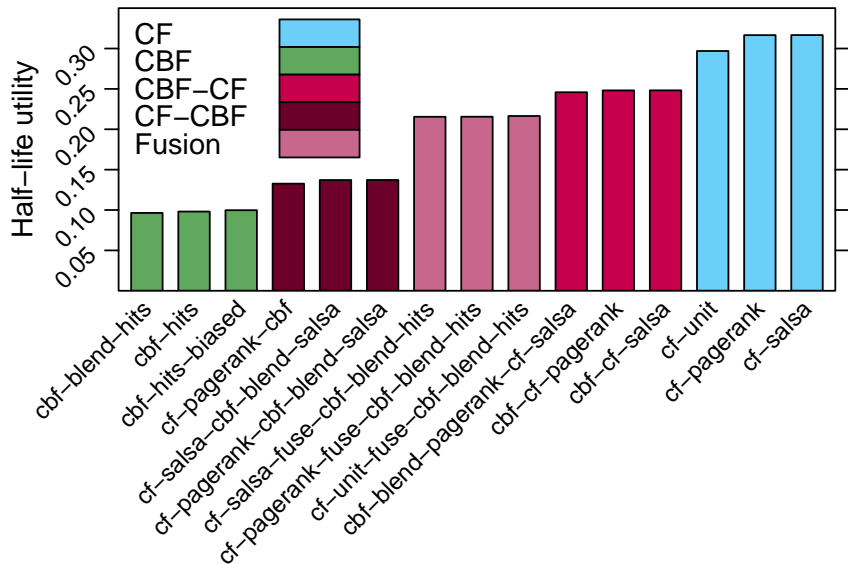
Aggregate to compute *fraction of potential utility achieved* (in range $[0, 1]$).

$$R = \frac{1}{nR_a^{\max}} \sum_a R_a$$


Offline Results — CF



Offline Results



Offline Results — CBF




Salsa performs like PageRank

Selected Algorithms

Chose 3 algorithms that performed well but have differing structures.

- CBF with Biased HITS in subgraph ranking configuration
- CBF fed into PageRank-weighted CF
- PageRank-weighted CF

Evaluating Performance



User-Based Evaluation

Asked graduate students to provide query sets of 5-10 papers and evaluate 3 5-item reading lists.

- Relevance of individual papers
- Importance of individual papers
- Quality of reading list as a whole
- Relative ranking of reading lists

All questions were set in the context of introducing a new researcher to the topic.



Use this page to provide us with 5-10 papers on a topic in computer science. We will use these papers as the starting point for building your reading lists. You can search for papers using the search box. Most papers in the ACM Digital Library, as well as some computer science papers from IEEE, Springer, and other publishers are available.

You may use multiple searches to build your list of papers. Once you have 5-10 selected papers, click "Continue" to proceed with the survey.

Gori, M., Pucci, A.. Research Paper Recommender Systems: A Random-Walk Based Approach. In Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence (Dec. 2006). WI '06.

McNee, S. M., Kapoor, N., Konstan, J. A.. Don't look stupid: avoiding pitfalls when recommending research papers. In Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work (Nov. 2006). CSCW '06.

Shimbo, M., Ito, T., Matsumoto, Y.. [Evaluation of kernel-based link analysis measures on research paper recommendation](#). In Proceedings of the 7th ACM/IEEE-CS joint conference on Digital libraries (Jun. 2007). JCDL '07.

McNee, S. M., Albert, I., Cosley, D., Gopalkrishnan, P., Lam, S. K., Rashid, A., Konstan, J. A., Riedl, J.. On the recommending of citations for research papers. In Proceedings of the 2002 ACM conference on Computer supported cooperative work (Nov. 2002). CSCW '02.

Instructions

1. Enter a search query to display papers in the Search Results pane
2. Move papers into your Selected Papers list with the buttons
3. Repeat steps 1 and 2 until you selected 5-10 papers, then click Continue

Selected List

- McNee, S. M., Kapoor, N., Konstan, J. A.. [Don't look stupid: avoiding pitfalls when recommending research papers](#). In Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work (Nov. 2006). CSCW '06.
- Gori, M., Pucci, A.. [Research Paper Recommender Systems: A Random-Walk Based Approach](#). In Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence (Dec. 2006). WI '06.
- McNee, S. M., Albert, I., Cosley, D., Gopalkrishnan, P., Lam, S. K., Rashid, A., Konstan, J. A., Riedl, J.. [On the recommending of citations for research papers](#). In Proceedings of the 2002 ACM conference on Computer supported cooperative work (Nov. 2002). CSCW '02.

In this portion of the survey, we will present you with up to 5-item reading lists. Evaluate each list as a whole with the questions provided at the end of the list.

Mishne, G. . [AutoTag: a collaborative approach to automated tag assignment for weblog posts](#). In *Proceedings of the 15th international conference on World Wide Web (May. 2006)*. WWW '06

Resnick, P.Iacovou, N.Suchak, M.Bergstrom, P.Riedl, J. . [GroupLens: an open architecture for collaborative filtering of netnews](#). In *Proceedings of the 1994 ACM conference on Computer supported cooperative work (Oct. 1994)*. CSCW '94

Fouss, F.Pirotte, A.Saerens, M. . [A Novel Way of Computing Similarities between Nodes of a Graph, with Application to Collaborative Recommendation](#). In *Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence (Sep. 2005)*. WI '05

Nie, L.Davison, B. D.Qi, X. . [Topical link analysis for web search](#). In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval (Aug. 2006)*. SIGIR '06

Sarwar, B.Karypis, G.Konstan, J.Reidl, J. . [Item-based collaborative filtering recommendation algorithms](#). In *Proceedings of the 10th international conference on World Wide Web (Apr. 2001)*. WWW '01

Feedback

If you were preparing a 5-item reading list to enable another researcher to understand the essentials of your topic, how many of the items above would you include?

0
 1
 2
 3
 4
 5

Overall, how good is this set of papers for introducing a new researcher to your topic?

Very Good
 Good
 Fair
 Poor
 Very Poor

You will find below a selection of papers that our algorithms have identified as candidates for a reading list. Beside each paper, we are asking you to rate it on the following three criteria:

How **relevant** is the paper to your topic of interest?

- Exactly captures the topic
- Mostly relevant
- Somewhat relevant
- Mostly irrelevant
- Entirely irrelevant

How **important** do you think the paper is to your topic? Measure this in terms of how long a reading list needs to be before you would include this paper. We will only ask you to rank the importance if the paper is at worst "Mostly irrelevant".

- Would include on a list of 5 papers
- Would include on a list of 10 papers
- Would include on a list of 25 papers
- Would include on a list of 50 papers
- Would omit or only include on very long lists

How **familiar** are you with this paper?

- I (co-)wrote it
- I have cited it
- I have read it
- I have heard of it (or am familiar with the authors)
- I have never heard of it

Evaluate Individual Articles:

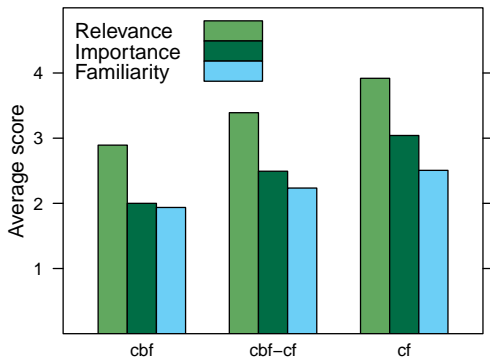

Fouss, F., Pirotte, A., Saerens, M. . A Novel Way of Computing Similarities between Nodes of a Graph, with Application to Collaborative Recommendation . In Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence (Sep. 2005). WI '05.	Relevance	Please select ▾
	Importance	Please select ▾
	Familiarity	Please select ▾
Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J. . GroupLens: an open architecture for collaborative filtering of netnews . In Proceedings of the 1994 ACM conference on Computer supported cooperative work (Oct. 1994). CSCW '94.	Relevance	Please select ▾
	Importance	Please select ▾
	Familiarity	Please select ▾
Nie, L., Davison, B. D., Qi, X. . Topical link analysis for web search . In Proceedings of the 29th annual international ACM SIGIR conference on Research and	Relevance	Please select ▾
	Importance	Please select ▾



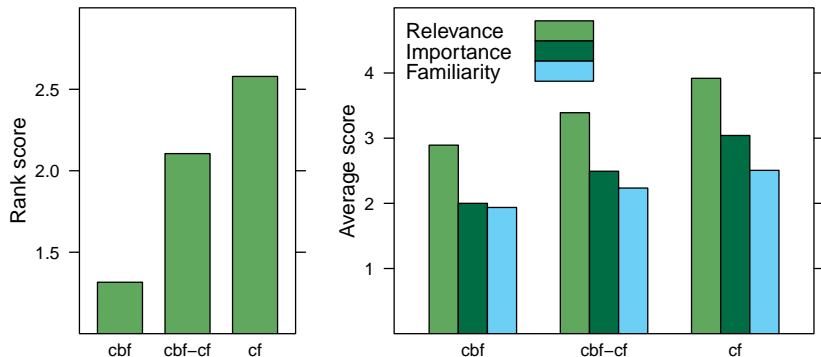
Below are summaries of the reading lists you have evaluated. Use the arrow buttons to put them in order from best to worst.

Reading list A >	< Reading list B >	< Reading list C
<p>Mishne, G. . AutoTag: a collaborative approach to automated tag assignment for weblog posts . In <i>Proceedings of the 15th international conference on World Wide Web (May. 2006)</i>. WWW '06.</p> <p>Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J. . GroupLens: an open architecture for collaborative filtering of netnews . In <i>Proceedings of the 1994 ACM conference on Computer supported cooperative work (Oct. 1994)</i>. CSCW '94.</p> <p>Fouss, F., Pirotte, A., Saerens, M. . A Novel Way of Computing Similarities between Nodes of a Graph, with Application to Collaborative Recommendation . In <i>Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence (Sep. 2005)</i>. WI '05.</p> <p>Nie, L., Davison, B. D., Qi, X. . Topical link analysis for web search . In <i>Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval (Aug. 2006)</i>. SIGIR '06.</p>	<p>Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J. . GroupLens: an open architecture for collaborative filtering of netnews . In <i>Proceedings of the 1994 ACM conference on Computer supported cooperative work (Oct. 1994)</i>. CSCW '94.</p> <p>Sarwar, B., Karypis, G., Konstan, J., Riedl, J. . Item-based collaborative filtering recommendation algorithms . In <i>Proceedings of the 10th International conference on World Wide Web (Apr. 2001)</i>. WWW '01.</p> <p>Shardanand, U., Maes, P. . Social information filtering: algorithms for automating "word of mouth" . In <i>Proceedings of the SIGCHI conference on Human factors in computing systems (May. 1995)</i>. CHI '95.</p> <p>Herlocker, J. L., Konstan, J. A., Borchers, A., Riedl, J. . An algorithmic framework for performing collaborative filtering . In <i>Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval (Aug. 1999)</i>. SIGIR '99.</p>	<p>Sarwar, B., Karypis, G., Konstan, J., Riedl, J. . Item-based collaborative filtering recommendation algorithms . In <i>Proceedings of the 10th international conference on World Wide Web (Apr. 2001)</i>. WWW '01.</p> <p>Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J. . GroupLens: an open architecture for collaborative filtering of netnews . In <i>Proceedings of the 1994 ACM conference on Computer supported cooperative work (Oct. 1994)</i>. CSCW '94.</p> <p>Herlocker, J. L., Konstan, J. A., Terveen, L. G., Riedl, J. T. . Evaluating collaborative filtering recommender systems . <i>ACM Trans. Inf. Syst.</i> (Jan. 2004).</p> <p>Herlocker, J. L., Konstan, J. A., Borchers, A., Riedl, J. . An algorithmic framework for performing collaborative filtering . In <i>Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval (Aug. 1999)</i>. SIGIR '99.</p> <p>Shardanand, U., Maes, P. . Social information filtering: algorithms for automating "word of mouth" . In</p>

User-Based Results



User-Based Results



CF performed the best (surprisingly)

Input Query (again)

- Herlocker et al. An algorithmic framework for performing collaborative filtering. In *Proc. SIGIR 1999*.
- Hoffman. Latent semantic models for collaborative filtering. In *TOIS*, volume 22, issue 1 (January 2004).
- Guy et al. Personalized recommendation of social software items based on social relations. In *Proc. RecSys 2009*.
- Karypis. Evaluation of Item-Based Top- N Recommendation Algorithms. In *Proc. CIKM 2001*.

How do we do?

CF's Recommender Systems Recommendations

- Sarwar et al. Item-based collaborative filtering recommendation algorithms. In *Proc. WWW 2001*.
- Resnick et al. GroupLens: An open architecture for collaborative filtering of netnews. In *Proc. CSCW 1994*.
- Shardanand and Maes. Social information filtering: algorithms for automating “word of mouth”. In *Proc. CHI 1995*.
- Deshpande and Karypis. Item-based top- N recommendation algorithms. In *TOIS*, volume 22, issue 1 (January 2004).
- Herlocker et al. Evaluating collaborative filtering recommender systems. In *TOIS*, volume 22, issue 1 (January 2004).

Task-Driven Design and Evaluation

Design recommenders to support human information needs, not just improve prediction error [McNee et al., 2006].

Task and context inform design of both recommenders and evaluation strategies.

Opportunity to harness unique characteristics (or “personalities”) of specific algorithms.

Conclusion

Contributions

- Design and evaluation of recommenders for reading lists
- Method for biasing CF with authority metrics
- CF works well for reading lists (surprising)

Conclusion

Contributions

- Design and evaluation of recommenders for reading lists
- Method for biasing CF with authority metrics
- CF works well for reading lists (surprising)

Open Questions

- How to tell user that they have an important paper?
- How to more accurately operate within topic scope?
- Is SALSA misused?

Conclusion

Contributions

- Design and evaluation of recommenders for reading lists
- Method for biasing CF with authority metrics
- CF works well for reading lists (surprising)

Open Questions

- How to tell user that they have an important paper?
- How to more accurately operate within topic scope?
- Is SALSA misused?

Questions?

ekstrand@cs.umn.edu

Funded by NSF grants IIS 05-34939 and IIS 08-08692