papers, though they probably consider other factors such as author and venue reputation. In these cases, the strength of the multi-layer approach is most evident. Whereas IQRA-TC performs the best, interestingly, TopCited (on the entire network) performs rather poorly. This finding indicates that our approach of selecting the documents similar to the query, followed by expansion using the citations, really helps in focusing the attention to relevant papers.

4.5 Timing Comparison

Table 4 shows the average times and the standard deviation across different queries for each of the query sets. We can see that IQRA-TC is among the fastest methods, and IQRA-ML also has good performance.

Query sets	Manual	Surveys	Citations
IQRA-TC	3.1 ± 1.5	**1.9 ± 0.7**	**2.4 ± 1.4**
IQRA-ML	7.9 ± 11.1	7.2 ± 2.4	6.1 ± 2.8
Okapi BM25	**1.8 ± 1.9**	4.7 ± 3.2	6.7 ± 4.0
TF-IDF	**1.8 ± 1.9**	4.7 ± 3.2	6.9 ± 4.0
TopCited	2.5 ± 3.4	5.2 ± 3.5	7.3 ± 4.2
CiteRank	20.8 ± 37.6	10.9 ± 3.2	12.7 ±3.8
PageRank (pre)	4.3 ± 5.2	13.2 ± 10.1	18.2 ± 12.6
PageRank (pos)	4.7 ± 1.9	7.6 ± 3.2	9.4 ± 3.8
PageRank (G_q)	3.1 ± 1.5	**1.9 ± 0.7**	**2.4 ± 1.4**

Table 4: Timing Comparison on CSX (in seconds)

4.6 Effect of Layers

We also investigated the effect of the different layers in our model. Table 5 shows the MAP@20 scores for various layers using the Manual, Surveys and Citations query sets on CSX. We observe that, as expected, the publications layer (P) plays a major role. However, adding authors (A), venues (V) and keywords (W) helps boost the performance even further. The PAWV model combines all of the layers and performs the best.

Layers	Manual Set	Surveys Set	Citations Set
P	0.188 ± 0.094	0.137 ± 0.129	0.077 ± 0.092
PA	0.205 ± 0.099	0.151 ± 0.144	0.086 ± 0.094
PV	0.187 ± 0.099	0.137 ± 0.130	0.077 ± 0.092
PW	0.205 ± 0.094	0.162 ± 0.151	0.093 ± 0.102
PAV	0.212 ± 0.099	0.151 ± 0.146	0.085 ± 0.094
PAW	0.216 ± 0.094	**0.171 ± 0.160**	0.098 ± 0.105
PWV	0.205 ± 0.093	0.163 ± 0.150	0.094 ± 0.104
PAWV	**0.223 ± 0.103**	**0.170 ± 0.161**	**0.097 ± 0.105**

Table 5: Effect of Different Layers (MAP@20): Publications (P), Authors (A), Venues (V), and KeyWords (W).

5. CONCLUSION

We have proposed a two-step approach for entity relevance and recommendation given a user-specified query. Instead of performing a query-independent search, we show that our strategy of staged query-dependent layer selection is much more effective. This is mainly due to two reasons, namely, fast pruning of irrelevant data, and query-dependent ranking propagation. Results on benchmark query sets show that our approach is more effective than existing methods. Our main conclusion is that for finding the most relevant citations for a paper, our top-cited method IQRA-TC serves well. The multi-layer approach IQRA-ML is a close second, but it has the potential for a more thorough literature survey by suggesting related entities like authors, venues and keywords. Showing the effectiveness of these extra layers is part of our ongoing work.

References

[1] C. Basu, H. Hirsh, and W. W. Cohen. Technical paper recommendation: A study in combining multiple information sources. *Journal of Artificial Intelligence Research*, 14:231–252, 2001.

[2] J. Beel, B. Gipp, S. Langer, and C. Breitinger. Research-paper recommender systems: a literature survey. *International Journal on Digital Libraries*, pages 1–34, 2015.

[3] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems*, 30(1):107–117, 1998.

[4] C. L. Giles, K. D. Bollacker, and S. Lawrence. Citeseer: An automatic citation indexing system. In *ACM Conference on Digital Libraries*, 1998.

[5] Q. He, J. Pei, D. Kifer, P. Mitra, and L. Giles. Context-aware citation recommendation. In *International Conference on World Wide Web*, 2010.

[6] K. S. Jones, S. Walker, and S. E. Robertson. A probabilistic model of information retrieval: development and comparative experiments: Part 1 & 2. *Information Processing & Management*, 36(6):779–808, 809–840, 2000.

[7] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of ACM*, 46(5):604–632, 1999.

[8] C. D. Manning, P. Raghavan, and H. SchÃijtze. *Introduction to Information Retrieval*. Cambridge University Press, Cambridge, UK, 2008.

[9] S. M. McNee, I. Albert, D. Cosley, P. Gopalkrishnan, S. K. Lam, A. M. Rashid, J. A. Konstan, and J. Riedl. On the recommending of citations for research papers. In *ACM Conference on Computer Supported Cooperative Work*, 2002.

[10] F. Meng, D. Gao, W. Li, X. Sun, and Y. Hou. A unified graph model for personalized query-oriented reference paper recommendation. In *ACM Conference on Information and Knowledge Management*, 2013.

[11] X. Ren, J. Liu, X. Yu, U. Khandelwal, Q. Gu, L. Wang, and J. Han. Cluscite: Effective citation recommendation by information network-based clustering. In *ACM International Conference on Knowledge Discovery and Data Mining*, 2014.

[12] A. Ritchie, S. Robertson, and S. Teufel. Comparing citation contexts for information retrieval. In *ACM Conference on Information and knowledge management*, 2008.

[13] T. Strohman, W. B. Croft, and D. Jensen. Recommending citations for academic papers. In *ACM Conference on Research and Development in Information Retrieval*, 2007.

[14] J. Tang and J. Zhang. A discriminative approach to topic-based citation recommendation. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 2009.

[15] J. Tang, J. Zhang, L. Yao, J. Li, L. Zhang, and Z. Su. Arnetminer: Extraction and mining of academic social networks. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '08, 2008.

[16] M. Valenzuela, V. Ha, and O. Etzioni. Identifying meaningful citations. In *AAAI Workshop on Scholarly Big Data*, 2015.

[17] D. Walker, H. Xie, K.-K. Yan, and S. Maslov. Ranking scientific publications using a simple model of network traffic. *Journal of Statistical Mechanics: Theory and Experiment*, (06):P06010, 2007.

[18] C. Wang and D. M. Blei. Collaborative topic modeling for recommending scientific articles. In *ACM International Conference on Knowledge Discovery and Data Mining*, 2011.

[19] N. Zheng and Q. Li. A recommender system based on tag and time information for social tagging systems. *Expert Systems with Applications*, 38(4):4575–4587, 2011.

[20] D. Zhou, S. Zhu, K. Yu, X. Song, B. L. Tseng, H. Zha, and C. L. Giles. Learning multiple graphs for document recommendations. In *International Conference on World Wide Web*, 2008.