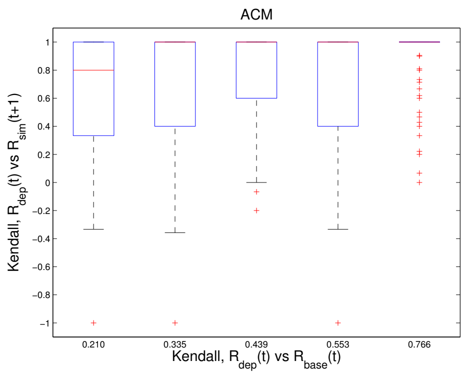

(a) Evaluating Dependency-based Ranking

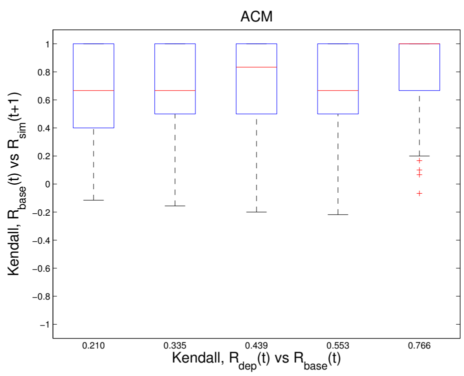

(b) Evaluating Baseline Co-authorship-based Ranking

Figure 11: ACM

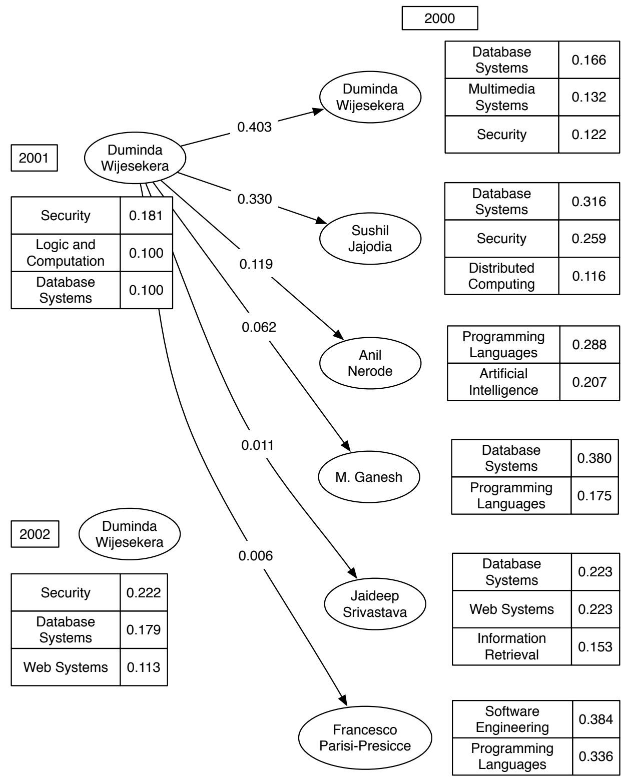

Figure 12: DBLP Case Study

4.5 Case Study Using DBLP, we provide a case study to help illustrate the workings of our proposed social dependency model. For this case study, we use the profile of Associate Professor Duminda Wijesekera. Figure 12 shows the social dependencies of Duminda Wijesekera for the year 2001. The directed edges show Duminda Wijesekera's dependencies on his co-authors who publish with him in the year 2001. Next to these co-authors are their respective topic distributions for year 2000. From the year 2000 to 2001, we observed that Duminda Wijesekera's topic in Security has increased from third position to first position [30]. Based on the dependencies, we observe that he depended on Sushil Jajodia most as compared to other co-authors (excluding himself). Based on the co-authors topic distribution, Sushil Jajodia's topic in Security is the highest which explains why Duminda Wijesekera's dependency on Sushil Jajodia is the highest [6]. In 2002, Duminda Wijesekera continues to increase his topic in Security [29, 31]. This illustrates how social dependency works based on the two components of interactions as well as content change.

5 Conclusion

In this paper, we address the problem of modeling the evolution of user interaction networks, in order to determine the social dependency weights among users at various time steps. We identify two primary factors to social dependency, namely: interactions between users, and temporal correlation between the users' topic distributions. We propose a *Decay Topic Model* to model a user's evolution of content at the topic level, as well as a *Social Dependency Metric* to determine the degree to which a user is dependent on another user. Comprehensive experiments on real-life co-authorship datasets DBLP and ACM show that our proposed models perform well against the baseline (co-authorship count) in two predictive tasks: predicting an author's ranking of co-authors by social dependency, as well as predicting the author's topic distribution in the next time step. This validates our hypothesis that we also need to take into account the changing topic preferences of users beyond just interactions (which the co-authorship baseline only models indirectly). For future work, we aim to incorporate additional factors to further improve the model. One is to learn the decay factor δ automatically. Another is to incorporate the magnitude (e.g., number of interactions) in addition to topic distribution in determining social dependency between users.

6 Acknowledgement

We acknowledge ACM for providing the ACM digital library data for this research. This work is supported by Singapore's National Research Foundation's research grant, NRF2008IDM-IDM004-036.

References

[1] Association for Computing Machinery. 2011. ACM Digital Library.

[2] Frank M. Bass. A new product growth for model consumer durables. *Manage. Sci.*, 50:1825–1832, December 2004.

[3] Chumki Basu, Haym Hirsh, and William Cohen. Recommendation as classification: using social and content-based information in recommendation. AAAI '98/IAAI '98, pages 714–720, 1998.

[4] David M. Blei and John D. Lafferty. Dynamic topic models. In *Proceedings of the 23rd international conference on Machine learning*, ICML '06, pages 113–120, 2006.

[5] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, March 2003.

[6] Alexander Brodsky, Csilla Farkas, and Sushil Jajodia. Secure databases: Constraints, inference channels, and monitoring disclosures. *IEEE Trans. Knowl. Data Eng.*, 12(6):900–919, 2000.

[7] Kevin R. Canini, Lei Shi, and Thomas L. Griffiths. Online Inference of Topics with Latent Dirichlet Allocation. In *Proceedings of AI Stats*, 2009.

[8] Dan Cosley, Daniel P. Huttenlocher, Jon M. Kleinberg, Xiangyang Lan, and Siddharth Suri. Sequential influence models in social networks. In *International Conference on Weblogs and Social Media*, 2010.

[9] David Crandall, Dan Cosley, Daniel Huttenlocher, Jon Kleinberg, and Siddharth Suri. Feedback effects between similarity and social influence in online communities. KDD '08, pages 160–168, 2008.

[10] Laura Dietz, Steffen Bickel, and Tobias Scheffer. Unsupervised prediction of citation influences. ICML '07, pages 233–240, 2007.

[11] Amit Goyal, Francesco Bonchi, and Laks V.S. Lakshmanan. Learning influence probabilities in social networks. WSDM '10, pages 241–250, 2010.

[12] M. G. Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):pp. 81–93, 1938.

[13] Timothy La Fond and Jennifer Neville. Randomization tests for distinguishing social influence and homophily effects. In *WWW*, pages 601–610, 2010.

[14] Jure Leskovec, Kevin J. Lang, Anirban Dasgupta, and Michael W. Mahoney. Statistical properties of community structure in large social and information networks. WWW '08, pages 695–704, 2008.

[15] Michael Ley. DBLP Computer Science Bibliography, 2005.

[16] Jianhua Lin. Divergence measures based on the shannon entropy. *IEEE Transactions on Information Theory*, 37(1):145–151, 1991.

[17] Lu Liu, Jie Tang, Jiawei Han, Meng Jiang, and Shiqiang Yang. Mining topic-level influence in heterogeneous networks. In *CIKM*, 2010.

[18] Hao Ma, Haixuan Yang, Michael R. Lyu, and Irwin King. Sorec: social recommendation using probabilistic matrix factorization. CIKM '08, pages 931–940, 2008.

[19] Miller McPherson, Lynn Smith-Lovin, and James M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, 2008.

[20] R. Nallapati and W. Cohen. Link-PLSA-LDA: A new unsupervised model for topics and influence of blogs. In *Proceedings of ICWSM*, 2008.

[21] Jennifer Neville and David Jensen. Relational dependency networks. *J. Mach. Learn. Res.*, 8:653–692, 2007.

[22] M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Phys. Rev. E*, 69(2):026113, Feb 2004.

[23] Huiming Qu, Jimeng Sun, and Hani T. Jamjoom. Scoop: Automated social recommendation in enterprise process management. In *ICSC*, 2008.

[24] Michal Rosen-Zvi, Thomas Griffiths, Mark Steyvers, and Padhraic Smyth. The author-topic model for authors and documents. UAI '04, pages 487–494, 2004.

[25] Upendra Shardanand and Pattie Maes. Social information filtering: algorithms for automating "word of mouth". CHI '95, pages 210–217, 1995.

[26] Parag Singla and Matthew Richardson. Yes, there is a correlation: - from social networks to personal behavior on the web. In *WWW*, pages 655–664, 2008.

[27] Jie Tang, Jimeng Sun, Chi Wang, and Zi Yang. Social influence analysis in large-scale networks. KDD '09, pages 807–816, 2009.

[28] Ken Wakita and Toshiyuki Tsurumi. Finding community structure in mega-scale social networks. WWW '07, pages 1275–1276, 2007.

[29] Lingyu Wang, Duminda Wijesekera, and Sushil Jajodia. Towards secure xml federations. In *DBSec*, 2002.

[30] Duminda Wijesekera and Sushil Jajodia. Policy algebras for access control: the propositional case. CCS, 2001.

[31] Duminda Wijesekera and Sushil Jajodia. Policy algebras for access control the predicate case. CCS, 2002.

[32] Xin Xin, Irwin King, Hongbo Deng, and Michael R. Lyu. A social recommendation framework based on multi-scale continuous conditional random fields. CIKM '09, pages 1247–1256, 2009.

[33] J. Yang and J. Leskovec. Modeling information diffusion in implicit networks. In *ICDM*, 2010.