ctdbd:07

Using Link Semantics to Recommend Collaborations in Academic Social Networks

Michele A. Brandão[1], **Mirella M. Moro (advisor)**[1]

[1]Departamento de Ciência da Computação, Universidade Federal de Minas Gerais

{micheleabrandao,mirella}@dcc.ufmg.br

Abstract. *Researchers have explored social network analysis (SNA) in many contexts with different goals. Here, we use concepts from SNA for recommending collaborators in academic networks. Specifically, we propose two metrics and verify how they influence on recommendation of new collaborations or intensification of existing ones. We also propose new algorithms for evaluating recommendations based on social concepts. Our experimental evaluation shows our new metrics improve the recommendations quality when compared to the state-of-the-art. Finally, we analyze the properties of the academic social networks in order to reinforce and justify the metrics' recommendation results.*

1. Introduction

Social networks (SN) aggregate different individuals according to their relationships. Here, we focus on people connected by research (or academic) ties. Now, Social Network Analysis (SNA) includes patterns and principles that are defined by social theories, such as homophily, proximity, and so on. Such patterns, principles and models can assist in exploring and predicting group and individuals' behavior. Indeed, many methods explore SNA for viral marketing, link prediction, and others. Moreover, link prediction may also be mapped to link recommendation; then, instead of *inferring* future connections, it also allows to *suggest* new ones. In such a research-oriented SN context, recommending or predicting new links may help a researcher to form new groups or teams, to improve the quality of communication in the network, and so on. Also, research groups with well connected co-authorship social network tend to be more prolific [Lopes et al 2011].

Discovering new links in this scenario is not trivial. When recommending friendships in an online SN, the number of friends in common can be used to estimate the social proximity between users [Lopes et al 2010]. On the other hand, in the academic context, social proximity has different interpretations in which the social connection between people and their academic background (e.g., geographic location and research area) must be considered. Specifically, our interest is in the social proximity between researchers that is defined by social theories. Therefore, we focus on discovering how the homophily and proximity principles increase the quality of the recommendations and (possibly) influence in the collaboration formation. Improving the SN link semantics allows to represent homophily as institutional affiliation and proximity as geographic location. In such a context, co-authors are similar if they are from the same institution and are close depending on the physical distance from their institution.

Another problem is how to evaluate the recommendations. Common metrics for evaluating recommendations, such as precision and recall, practically do not explore any particular feature of the SN. Therefore, we employ SNA-based concepts for evaluating the recommendations from the SN perspective (which makes sense because the recommendations were defined from the social perspective as well).

The contributions are summarized as follows: definition of two new metrics that consider social principles, called *Affin* and *GLI*, and two recommendation functions to recommend collaborators considering link semantics (Section 2); description of three metrics (*novelty*, *diversity* and *coverage*) for analyzing the quality of the recommendations (Section 3); an experimental evaluation using two real SN, an analysis of the academic social networks used in the experimentation and discussion on how the SN properties influence the new recommendation metrics (Section 4).

Additionally, after this dissertation defense, we (and our research group) have continued working in two of its open problems. First, the evaluation with real users, for which we have developed an online questionnaire with access to 213 Brazilian computer science researchers. Only 61 have answered (around 29%, which shows the difficult for evaluating the algorithms with real users), and the results showed that academic recommendations do not work both ways (i.e., A could agree with the recommendation for working with B, but B could find A not as suitable for working with). Second, the visualization of the recommendations, for which we have developed an online system that shows the recommendations considering social networks and different metrics (a preview is available at http://homepages.dcc.ufmg.br/∼mirella/Tools/CNARe/).

2. Defining the Link Semantics to Recommend Collaborators

Social Networks are formed by *actors* (people) and their *relational ties* (links). The importance of a relationship between its actors may be defined by a weight measure. Each weight is relevant because it reflects the link semantics, instead of just the network topological feature. Here, we use an academic SN in which two researchers are connected if they have co-authored a publication. The final goal is to recommend collaborations (new or intensification) over this network, which is mapped to predicting links in a SN.

Affin - Affiliation Metric. The Affin metric considers the homophily principle, which is derived from the researchers' affiliation. It considers the *affiliation weight* $Affin_{i,j}$: for any given pair of researchers $\langle i, j \rangle$, it is calculated by the number of papers of researcher i co-authored with people from j's institution divided by the total number of papers authored by i. *Affin* follows the intuition that an institution is more important to an author if he/she has already collaborated with someone from that institution; hence, this follows the idea that is likely to contact other researchers in the same institution. However, recommending based solely on the researchers' affiliations is not enough, because it disregards the history of the researchers' collaborations. Hence, we propose to combine Affin with existing metrics in order to be more useful for the recommendation function. Specifically, Affin is combined with *cooperation* ($Cp_{i,j}$, how much the two researchers have collaborated), *correlation* ($Cr_{i,j}$, how similar the areas of the researchers are) and *social closeness* ($Sc_{i,j}$, a normalized variant of the shortest path metric) [Lopes et al 2010].

GLI - Geographic Location Information Metric. The *GLI* metric follows the proximity principle. To measure the physical proximity between pairs of researchers, we introduce the *geographic location weight* $GLI_{i,j}$ that considers the geographic location information for any given pair of researchers $\langle i, j \rangle$ defined by the distance between geographical coordinates of the researchers *i* and *j* institutions. Here, a geographical coordinate is that of the city where a researcher's institution is. To define a qualitative scale, we consider the *travel time* that covers the distance (represented by $GLI_{i,j}$) between researchers. The travel time weight ($\Delta t_{i,j}$) allows to specify how far two researchers are from each other.

3. Evaluation Metrics

Evaluating the quality of recommendations and the effectiveness of recommendation functions are very difficult tasks, mainly for two reasons [Fouss and Saerens 2008]: *(i)* different algorithms may have different performance on different datasets, and *(ii)* the goals for which an evaluation is performed may differ. Also, having high accuracy is important, but *insufficient* to ensure the quality of the recommendations [Fouss and Saerens 2008, Shani and Gunawardana 2011]. Hence, it is important to consider a large number of metrics to analyze different aspects in the evaluation of the recommender systems. Next, we detail each metric and show how they are employed.

Accuracy. Here, accuracy is given by recall because: *(i)* the networks are sparse and the total number of possible links is large; *(ii) Affin* and *GLI* metrics aim to make networks more connected, as opposed to totally connected; and *(iii)* high recall indicates that the metrics provide correct recommendations.

Novelty. We adapt the idea from [Fouss and Saerens 2008] for the academic SN context. Given a set of target researchers $\mathcal{T}$, a recommendation list $\mathcal{L}$ and the total number of target researchers n. First, we calculate the frequency $\Im_r$ of each recommended researcher r, where $r \in \mathcal{L}$, to represent the popularity degree of the researchers: the less popular a recommended research, the most probable he/she is unknown to a target researcher. Then, we take the median f_m as a central tendency metric to represent the frequencies. Finally, the frequency median f_m of the recommended researchers is divided by the number of target researchers n. The resulting value represents the *novelty* in a recommendation list.

Diversity. The most explored method to measure diversity in a recommendation list is the intra-list similarity metric [Shani and Gunawardana 2011]. We use this method based on the approach presented in [Ziegler et al 2005]. Some changes have been made in this approach to evaluate collaboration recommendations: we use the correlation among researchers Cr (defined in [Lopes et al 2010]), which represents the semantics of the SN relations, to calculate this similarity. Finally, high values indicate low diversity.

Coverage. Coverage is given by a metric that computes how unequally different the recommended items are. Two different metrics of this distributional inequality are *Gini index* (GI) and *Shannon Entropy* (SE) [Shani and Gunawardana 2011].

4. Experimental Evaluation and Results

This section details the datasets employed in our experimental evaluation, presents the evaluation results and a graph analysis of the datasets.

Dataset Details. The experiments were performed using two real SN that were built from *CiênciaBrasil*[1] and DBLP[2] datasets. The social network of each dataset was divided in two parts: 90% of the data as validation set, and the remaining 10% for testing. Both parts also follow the time interval distribution, where the first part considers publications

[1]CiênciaBrasil: `http://pbct.inweb.org.br`
[2]DBLP: `http://www.dblp.org/db/`

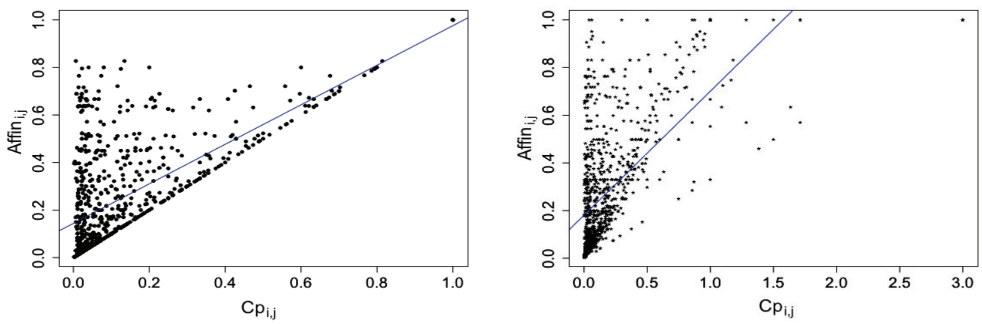

(a) *CiênciaBrasil*: CC = 0.573 (b) DBLP: CC = 0.601

Figure 1. The (clear) relation between *Affin* and Cooperation.

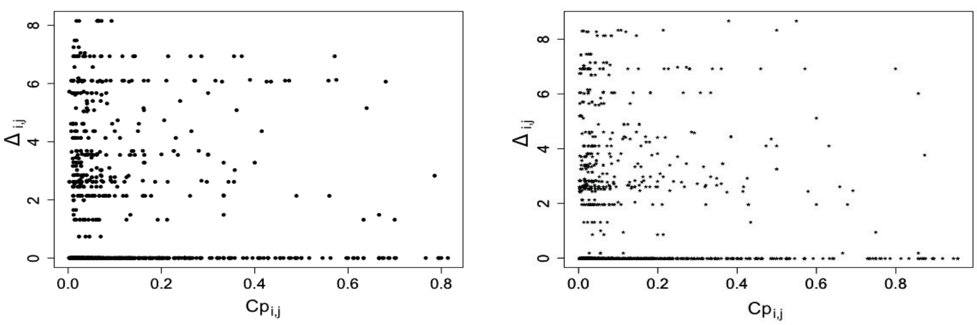

(a) *CiênciaBrasil*: CC = -0.047 (b) DBLP: CC = -0.090

Figure 2. The (non-existant) relation between travel time and Cooperation.

prior to the second part. Furthermore, we compare the results of *Affin* and *GLI* with *CORALS* [Lopes et al 2010]. In order to provide a better comparison, we have also combined *CORALS* and *Affin* in a new metric, called *CORALS+Affin*, that works on a network built as *CORALS* including all researchers correlated by research area, social closeness and affiliation. Note that we did not consider combining *CORALS* and *GLI* because, as shown next, there is no relation between cooperation and location.

Evaluation Results. Our evaluation results show that using institutional affiliation leads to an improvement in accuracy. Thus, in *CiênciaBrasil* and DBLP, *Affin* performs better than *GLI* and CORALS. The recall of *CORALS+Affin* is equal to *Affin*, because *Affin* adds affiliation to the original *CORALS*. A complementary result is that affiliation and cooperation ($Affin_{i,j}$ and $Cp_{i,j}$) are directly related, because the Pearson correlation coefficient (CC) in both SN is large (greater than 0.5) as shown in Figure 1. This fact explains why *Affin* provides more accurate recommendations. Regarding geographic location, *GLI* presents the worst accuracy results. For better understanding, we observe that the Pearson correlation coefficient between *travel time* and cooperation ($\Delta t_{i,j}$ and $Cp_{i,j}$) is in the range [-0.09; 0.0] which indicates lack of linear correlation. This is a common behavior in both *CiênciaBrasil* and DBLP datasets as shown in Figure 2.

In both social networks, *GLI* provides recommendations with more *novelty* and *diversity*. *Affin* presents the second best value for *diversity* and the same result as *CORALS+Affin* for *novelty*. Furthermore, using geographic location leads to an improvement in *coverage*. *GLI* generates a recommendation list with more unequally different researchers, and presents the best results for *Gini index* and *Shannon Entropy* in *CiênciaBrasil* and DBLP. *Affin* presents the second best result in *CiênciaBrasil* and the worst in DBLP. Moreover, *CORALS+Affin* provides better results than CORALS for *coverage* in both SN, because *CORALS+Affin* considers more researchers than CORALS in the recommendations, which increases the difference between them.

Besides recommending new collaborations, we also work on recommending (existing) collaborations that can be further intensified, i.e., the intensifiable collaborations. In order to evaluate such recommendations, we consider only accuracy. Note that other evaluations metrics do not apply for intensifiable collaborations, because *novelty*, *diversity* and *coverage* cannot be established for existing collaborations. Regarding the accuracy, *GLI* shows recommendations with the best recall (for both networks), which is justified because it distinguishes researchers with near and far *travel time*, increasing the number of relevant results. *Affin* presents the second best recall (for the two social networks). This shows that the affiliation can improve the accuracy of the recommendations. Finally, CORALS and *CORALS+Affin* present the same results.

Graph Analyses. We also study the interactions within the datasets and analyze the properties of their networks based on two different graph analyses. First, we infer the strength of ties between researchers that have collaborated. This study aims to show that our metrics recommend collaborations that will be weak ties, which are important for establishing bridges within the network. Then, we build and study social networks (using *CiênciaBrasil* and DBLP datasets) that represent the collaboration among researchers grouped by their institutions. These studies contribute to understand the results of the recommendation functions and provide further evidence that validates our metrics.

Regarding tie strength, it is important to know whether *Affin* and *GLI* can recommend collaborations that will also be weak ties. Hence, we analyze the relation between absolute cooperation and affiliation/*travel time*. We observe that affiliation and tie strength are directly related. Moreover, weak ties are common when the value of $Affin_{i,j}$ is a little larger than the threshold specified in Section 3.2 of the dissertation. This fact shows that *Affin* can recommend collaborations that will probably form weak ties. However, we observe that *travel time* and *tie strength* are not related. Hence, it is not possible to predict whether *GLI* will give recommendations that will form weak ties.

On the other hand, considering aggregated cooperation (we call $Cgroup_{i,j}$), one way to measure the cooperation between researchers from different institutions i and j is through dividing the number of papers of researchers from institution i co-authored with people from institution j by the total number of papers authored by researchers from institution i. We built graphs that represent the cooperation between researchers from pairs of institutions. In these graphs, the aggregated cooperation ($Cgroup_{i,j}$) defines the weight of the edges. We are interested in studying the properties of these networks and investigating if these networks follow the small-world phenomenon. In this study, we conclude that the networks that represent the cooperation between Brazilian computer scientists from different institutions have the small-word properties, which may help ensure that the physical distance alone does not influence in the cooperation between researchers (computer scientists) from different institutions. By contrast, this conclusion also validates the *Affin* metric, because *Affin* considers that researchers from different institution can cooperate.

5. Conclusion

This work introduced two new metrics for recommending collaborations in an academic SN. Given a recommendation system, the hardest part is to define which metric should the recommendation function rely upon when producing the results. The base of our work is to consider social aspects when recommending collaborations to researchers. Specifically, we consider the institutional affiliation aspect (*Affin*) and the geographic location information (*GLI*) of all researchers in the social network. Besides providing these metrics, we have also proposed new ways for evaluating the recommendation results. Furthermore, we have also taken a deeper look at the SN properties, which validated our conclusions. All those contributions advance the collaboration recommendation algorithms, as well as spread new insights on evaluating recommender systems strategies and social network analysis, which are prolific research areas. Finally, this work has generated the following publications: [Brandão and Moro 2012a], [Brandão and Moro 2012b], [Brandão et al. 2013b], [Brandão et al. 2013a], [Brandão et al. 2014]. Issues for future work include: refining the recommendation function with other link semantics; studying other geographic factors that influence in the cooperation; and considering other metrics to evaluate the quality of the recommendations.

Acknowledgments to CAPES, CNPq, Fapemig and InWeb – Brazil.

References

Brandão, M. A. and Moro, M. M. (2012a). Affiliation influence on recommendation in academic social networks. In *AMW*, pages 230–234, Ouro Preto, Brazil.

Brandão, M. A. and Moro, M. M. (2012b). Recomendação de colaboração em redes sociais acadêmicas baseada na afiliação dos pesquisadores. In *SBBD - short paper*, pages 73–80, São Paulo, Brazil.

Brandão, M. A., Moro, M. M., and Almeida, J. M. (2013a). Análise de Fatores Impactantes na Recomendação de Colaborações Acadêmicas Utilizando Projeto Fatorial. In *SBBD - short paper*, pages 1–6, Recife, Brazil.

Brandão, M. A., Moro, M. M., and Almeida, J. M. (2014). Experimental evaluation of academic collaboration recommendation using factorial design. *JIDM*, 5(1):52–63.

Brandão, M. A., Moro, M. M., Lopes, G. R., and de Oliveira, J. P. M. (2013b). Using Link Semantics to Recommend Collaborations in Academic Social Networks. In *WWW Workshops*, pages 833–840, Rio de Janeiro, Brazil.

Fouss, F. and Saerens, M. (2008). Evaluating Performance of Recommender Systems: An Experimental Comparison. In *Web Intelligence*, pages 735–738, Sydney, Australia.

Lopes et al, G. R. (2010). Collaboration Recommendation on Academic Social Networks. In *ER Workshops*, pages 190–199, Vancouver, Canada.

Lopes et al, G. R. (2011). Ranking Strategy for Graduate Programs Evaluation. In *ICITA*, pages 59–64, Sydney, Australia.

Shani, G. and Gunawardana, A. (2011). Evaluating Recommendation Systems. In *Recommender Systems Handbook*, pages 257–297. Boston, USA.

Ziegler et al, C.-N. (2005). Improving recommendation lists through topic diversification. In *WWW*, pages 22–32, Chiba, Japan.