SEMEF: A Taxonomy-Based Discovery of Experts, Expertise and Collaboration Networks

Delroy Cameron, Boanerges Aleman-Meza, Sheron L. Decker & I. Budak Arpinar

Large Scale Distributed Information Systems (LSDIS) Lab
Computer Science Department, University of Georgia
Athens, GA 30602-7404, USA
{cameron, boanerg, decker, budak}@cs.uga.edu

Abstract. Finding relevant experts in research is often critical and essential for collaboration. Semantics can be useful in this process. In particular, a taxonomy of Computer Science topics, together with an ontology of publications can be glued through explicit relationships from papers to one or more topics. These paper-to-topics relationships, extracted from paper abstracts and keywords are valuable for building an Expertise Profile for a researcher based on the aggregation of the topics of his/her publications. We describe an approach that finds experts, expertise and collaboration networks in the context of the peer-review process. We use DBLP bibliography data to determine different levels of collaboration based on degrees of separation. This helps in suggesting experts for PC membership and has the benefit of presenting potentially unknown experts to the PC Chair(s). We present our findings and evaluations in the context of expanding collaboration networks in a peer-review setting.

Keywords: C-Net, Collaboration Level, Collaboration Strength, Expert Finder, Expertise Profile, Expertise Rank, Semantic Association, Semantic Web, DBLP

1 Introduction

A method for identifying experts in a peer-review setting is important, especially because conferences, workshops, symposiums etc, necessitate that qualified reviewers assess the quality of research submissions. Since existing conference management applications such as Confious (confious.com) and OpenConf (openconf.org) provide minimal support for finding experts, the onus falls on conference organizers to compose program committees. Traditionally, creating a program committee is based largely on the organizers' knowledge of experts in the field. However, this approach has its limitations. First, due to various emerging online communities and diversification of research areas, it is possible that unknown experts may be overlooked [5]. Moreover the human effort required can be intensive and time consuming. We address the problems involved with finding experts, presenting a seamless semi-automated approach requiring minimal manual input.

One of the challenges in finding experts is that of obtaining data needed to reasonably assess and quantify expertise. Indeed, obtaining expertise data can be approached from many different perspectives [2]. Data could be extracted from Curriculum Vitas (CV) [4], intranet applications [9], Version Control Systems (VCS) [10] and from publications (white papers, technical reports, scientific papers etc) [13, 14]. In our approach, we use publications data to demonstrate the benefits of semantic in finding experts.

Semantic techniques play a critical role in determining expertise, particularly at finer levels of granularity. For example, a researcher with expertise in "Semantic Web Processes" may be a good match for a conference in "Web Services." Finding such information involves inexact matches of expertise that may not always be explicit. Hence, by using a taxonomy of topics to find expertise in subtopics of a given topic, we show that semantic techniques can be used to discover expertise at specific levels of granularity.

Publication data can also provide person-centric information about collaboration relationships between researchers that could be significant in determining which experts are invited for PC membership. For example, two researchers could have comparable knowledge on a certain topic, but a conference organizer may have a preference or bias to invite one over the other due to past collaboration or social relationship. The trivial and shortest collaboration distance is that of coauthors in a publication. In our approach, we address finding various degrees of separation between researchers, based on common-coauthors, same affiliation, etc.

The dataset for demonstrating our approach consists of three main parts. First, an ontology of expertise data extracted from various sources and linked to an ontology of Computer Science publications based on DBLP bibliography. Second, a taxonomy of Computer Science topics and third an ontology linking publications to topics in the taxonomy. These *papers-to-topics* relationships are used to classify authors as experts on certain topics. We integrate these datasets and apply various semantic web techniques on them to show that by discovering and exploiting semantic associations within a dataset, interesting relationships can be obtained and analyzed to provide meaningful results, otherwise difficult to obtain. The contributions of this paper are therefore as follows.

- We address the problem of finding experts by applying semantic technologies under the scenario of finding relevant reviewers for consideration for membership in a Program Committee for conferences (or workshops, etc). The main benefit of semantics is in achieving finer granularity in measuring expertise.
- We propose a solution to the problem of finding relevant experts that are potentially unknown to PC Chair(s). Our solution involves discovery of collaboration levels among experts groups to provide a second dimension in the selection process – a dimension indicating how experts are related among experts.

We demonstrate the effectiveness of this approach by comparing existing experts listed in PCs of past conferences with our recommended experts.

2 Expertise Dataset

In scientific research, the publications of a researcher can be viewed as representative of his/her expertise. Therefore a dataset replete with publication information, providing relationships between publications and areas of expertise (topics) is crucial. The benefits of relating papers to topics have been argued and demonstrated with a small dataset in [1]. To obtain a dataset, we focused initially on the DBLP bibliography, which is an excellent source of Computer Science publications. However, in spite of large amounts of bibliographical data, neither DBLP nor similar datasets explicitly relate researchers to areas of expertise. In parallel research we developed a methodology for identifying emerging trends in research areas. Such work uses a taxonomy of Computer Science topics to link publications to research areas [6]. We take advantage of such dataset to establish the relationships between papers to topics needed for computing expertise. In this section, we include a short overview of that dataset.

2.1. Topics Taxonomy

Many attempts that lead to the development of taxonomies across a variety of areas have been undertaken (e.g., [8, 11]). The taxonomy that we use in this work consists of 344 topics across different areas in Computer Science. The complete taxonomy is available online (http://cs.uga.edu/~cameron/swtopics/taxonomy).

2.2. Publications Data

For the purposes of this paper, we used a subset of DBLP containing publications in research areas including Databases, Web, Semantic Web, Information Retrieval, Data Mining and AI. Instead of the XML-version available at their site, we use an RDF representation of DBLP data called SwetoDblp (lsdis.cs.uga.edu/projects/semdis/swetodblp/). The methods described later in Section 3.3 require that data be represented in RDF. Table 1 describes the subset, which is roughly 10% the size of DBLP.

Table 1. Main classes for the subset of publications from DBLP data

Authors	67,366
Journal Articles	25,973
Articles in Proceedings	51,202

2.3. Papers-to-Topics Relationships

The dataset relating papers to topics used for this work related 3,970 papers to the taxonomy of topics. There were a total of 6,668 papers to topics relationships. A key aspect in obtaining such relationships involved the use of the ‘ee’ metadata value in DBLP to obtain terms for each paper. This attribute links from DBLP to the actual document (PDF, PS, etc) or other data sources such as ACM Data Library, IEEE Publications Database, and Science Direct containing additional publications data. These *ee* links were used in *focused crawling* of such data sources in which extractors retrieved terms from keywords and abstracts for many publications. These terms were then investigated in the taxonomy from which successful matches established relationships from the paper to the topic (s). A detailed description of the techniques involved in this process is covered in [6].

3 Approach

When finding experts, there are two fundamental considerations. First, determining '*Expertise Profiles'* and second, ranking *'Experts'* according to how their Expertise Profile matches a specific topic. The first task requires quantification of expertise, which we obtain by using publication impact data from Citeseer (citeseer.ist.psu.edu/impact.html). An advantage of using such data is that the publication venues listed were URLs that match those from DBLP. The disadvantage is that this data was last updated in 2003. Nonetheless, we consider it a credible source of publication impact statistics. The second consideration requires extracting and aggregating relevant expertise for identifying researchers in specific domains. Figure 2 shows the core architecture involved in this process.

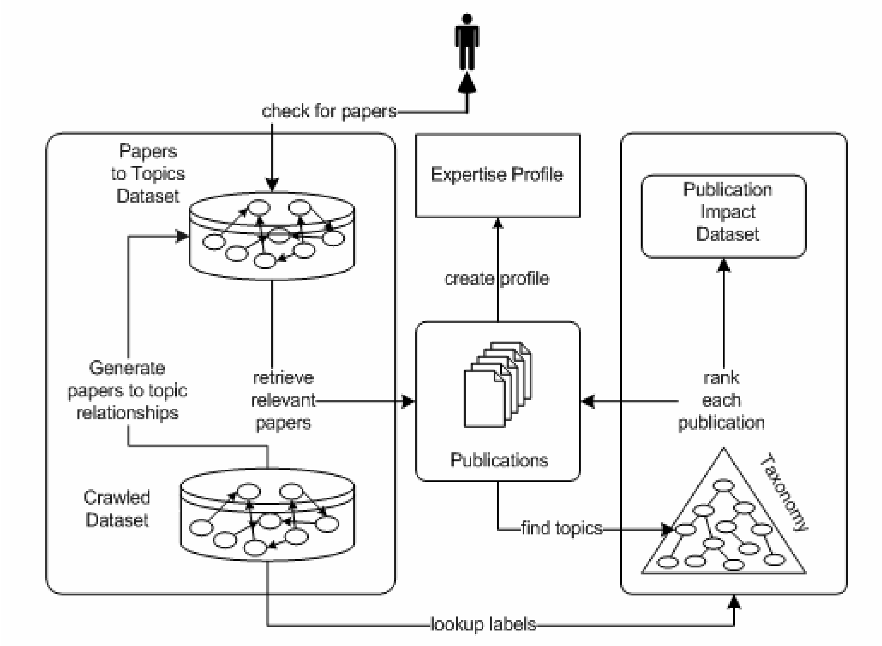

Figure 2. Core System Architecture

3.1. Expertise Profiles

We define the *Expertise Profile* of a researcher as the set of *topic-value* pairs (one pair for each topic), for which a researcher is considered knowledgeable. For example, if one paper has three topics, then three topic-value pairs will appear in the Expertise Profile. The aggregated topic-value pairs of the researcher reflect their expertise profile.

It is worthwhile to mention that every coauthor is implicitly connected to the topics of a publication. Therefore, for every paper we assume equal expertise among coauthors, and assign them equal values. We make this claim because the overall effect of all publications for an author alludes to an overall area or *'Areas of Expertise.'* Researchers with more publications would tend to accumulate higher values for expertise in given topics.

Figure 3 shows the algorithm for computing expertise profiles. The first step is to identify all of the publications of a researcher. From the papers-to-topics relationships, we determine the topic(s) for which each paper is related. Then we obtain the publication impact for each paper (if available) and update the Expertise Profile with the topic-value pair. For papers with multiple topics we simply find the aggregate sum of the publication impact for each topic of the paper. Although this may seem to be a simple algorithm, it provides a good enough measure of expertise across multiple topics.

```
Algorithm findExpertiseProfile(researcher, publications)
create 'empty expertise profile'
foreach paper of researcher do
    get 'topics' list of paper (using papers-to-topics dataset)
    foreach topic in topics list do
        get 'publication impact'
        if 'publication impact' is null
            'publication impact' ← default
        else
            'weight' ← 'publication impact' + existing 'weight' from expertise profile
        if 'expertise profile' contains 'topic'
            update 'expertise profile' with <'topic,' 'weight'>
        else
            add <'topic,' 'weight'> pair to 'expertise profile'
    end
end
return 'expertise profile'
```

Figure 3. Algorithm 1 for determining Expertise Profiles

Table 2 shows how this example applies to a researcher in our dataset. The researcher has publications in *World Wide Web (WWW)* and the *Conference on Information and Knowledge Management (CIKM)*. According to Citeseer the publication impact of these conferences is *1.54* and *0.73* respectively. Since the researcher has two publications on Search Engines, the total expertise for Search Engines is the sum of the two impact values *(2.27)*.

Table 2. Computing Expertise Profiles using Publication Impact

Publication	Topic	Publication Impact	Expertise Value
conf/www/FlakeGLG02	Search	1.54	1.54
	Search Engines	1.54	2.27
conf/cikm/GloverLBG99	Search Engines	0.73	
conf/cikm/KrugerGCGFLO00	Web Search	0.73	0.73
conf/www/GloverTLPF02	Classification	1.54	1.54
conf/cikm/GloverPLK02	Knowledge Discovery	0.73	0.73

3.2. Ranking Experts

Determining whether a researcher is an expert is a very broad and highly subjective subject. Ranking researchers could even be controversial. For example, a researcher may be an expert in *'Database Systems'* but not necessarily in *'Semantic Web.'* Identifying experts therefore requires a *specific context*. In any case, the approach we adopt involves finding exact as well as inexact matches of expertise, given a specific context.

Table 3, shows that *'Search'* is obviously a main area for the researcher described above. However, a close examination reveals that there is higher cumulative expertise in subtopics and related topics of Search. It is therefore implicit that expertise in subtopics is inclusive of the topic Search itself. Therefore in finding expertise, we contend that the weighted summation of the relevant topics in the expertise profile indicates the rank of the expert in the field. Note, that this measure excludes the expertise value for the topic *'Classification,'* as it is not a subtopic of Search. This example highlights the unique benefit of semantics in detecting inexact matches through the use of a taxonomy of topics.

Table 3. Computing Expertise Rank from Expertise Profiles

Topic	Expertise Value
Search	1.54
Search Engines	2.27
Web Search	0.73
~~Classification~~	~~1.54~~
Knowledge Discovery	0.73
Expertise Rank	4.54

Equation (1) expresses the formula for determining the expertise rank (E_{rank}), used to compute the values in Table 3. The variable m represents the total number of papers for a relevant topic, n represents the total number of relevant topics and T_i and $P_{(i,j)}$ refer to a topic and paper respectively. The summation of the publication impact for each paper $P_{(i,j)}$, for each relevant topic T_i, gives the total expertise rank (E_{rank}).

$$E_{rank} = \sum_{i=0}^{n} T_i \sum_{j=0}^{m} P_{(i,j)} \qquad (1)$$

3.3. Collaboration Networks

A semi-automated application for finding experts provides the important functionality of alleviating the job of the PC Chair(s). However, in the wider scope of finding experts, PC Chairs are themselves experts who in all likelihood are personally and/or professionally connected to other experts. We aim therefore not only to find experts, but also to analyze collaboration networks relative to the PC Chair(s). The goal is to not only suggest experts but also provide PC Chairs with information about who are the experts outside of their collaboration network. One benefit of this kind of analysis is to avoid possibly suggesting experts already known to PC Chairs i.e. "do not suggest to me experts whom I already know."

3.3.1 Discovery of Collaboration Levels

The notion of Semantic Associations [3] has been used to discover various paths that connect two entities in a populated ontology. This concept has applicability in a variety of areas, such as determining provenance and trust of data sources [7].

We implemented Semantic Association discovery to obtain multiple ways in which a given PC Chair is related to an expert. The goal is to automatically analyze each of these paths to determine the closest link between the two persons. Within publication data, researchers are implicitly related through coauthors, affiliation, conference proceedings, etc. Such relationships reflect different '*Collaboration Levels'* between various researchers. Table 4 describes these levels.

Table 4. Types of Collaboration Levels

Collaboration Levels	Description - w.r.t. PC Chair(s)	Degrees of Separation
STRONG	co-authors	One
MEDIUM	common co-authors	Two
WEAK	published in same proceedings	Unspecified
	co-authors with common co-authors	Three
	co-author related to editor of proceedings	Unspecified
EXTREMELY WEAK	co-author published in same proceedings	Three
UNKNOWN	no relationship based on dataset	Unknown

The strongest and most obvious association exists between coauthors. However, many weaker relationships may also exist. For example, two authors can be related through a common coauthor, or merely through publications in the same Proceedings. The least obvious relationships are classified as unknown. For such relationships, we make the important distinction that they do not necessarily imply lack of a relationship between the two entities, but rather no existing relationships according to our dataset. Indeed personal and/or professional relationships may exist elsewhere, as well as within publications data absent from our dataset.

3.3.2 C-Nets

In addition to determining levels of collaboration, a second technique for analyzing collaboration networks involves the creation of '*CollaborationNets'* or *'C-nets.'* The C-net paradigm follows a similar intuition to that of creating *"Relational ExpertiseNets"* from [14]. ExpertiseNets represent the correlation of all of a researcher's publications with respect to a cluster of cited publications on the same topic. In our work, C-Nets identify the ordering of experts within clusters of collaboration units obtained from the larger list of experts. For example, suppose that a Professor has three publications in a rather new topic and two of his students are coauthors in such publications. The goal is to use the measure of *'Collaboration Strength'* [12] with the cluster of experts, to distinguish which of these three persons might be the "boss," i.e., the Professor. In essence, finding the "expert among experts."

Therefore, we define a C-Net as a bidirectional Graph (G_k) consisting of a list of nodes, such that there exists a super node of maximum value (v_m), that is connected to every other node (v_i) in the bidirectional Graph (G_k) through an edge (e_j).

$$G_k = \{v_1, v_2, v_3 \ldots \ldots \ldots v_n\} \mid G_k \in G$$
$$e_j = \{v_m, v_i\} \text{ or } \{v_i, v_m\} \mid \forall i, \; i \neq m \tag{2}$$

In Equation (2), G represents the entire list of experts, while (G_k) is a C-Net for a cluster of experts within the expert list. The element (v_m) represents an expert whose expertise is higher than the others, and for whom there is a direct coauthorship relationship with each of the other experts.

In each C-Net the collaboration strength of a node is a measure that is computed using the method outlined in [12]. Table 5 shows a C-Net from our dataset, in which the super node (Expert A) has a significantly greater expertise than other nodes, and Coauthor4 has the highest collaboration strength with the super node, most likely as a result of having published more papers with him/her.

Table 5. C-Net Unit

Node	Expertise	Collaboration Strength	Chair1	Chair2
Expert A	14.80		WEAK	WEAK
Coauthor1	0.73	0.5	WEAK	WEAK
Coauthor2	0.73	0.5	WEAK	WEAK
Coauthor3	0.73	0.5	WEAK	EXTREMELY WEAK
Coauthor4	1.81	1.0	UNKNOWN	WEAK

4. Evaluation

We implemented a prototype application called SEMEF (SEMantic Expert Finder), which analyzes program committee lists of past conferences to evaluate our approach. We aimed to establish two things. First, we aimed to validate the efficacy of our system, as a plausible Expert Finder approach. Second, we leverage SEMEF to discover expert collaborations and C-Nets, to encourage collaboration networks expansion. Our evaluation therefore covers two areas: *Validation* and *Collaboration Network Expansion*.

4.1. Validation

To validate SEMEF, we consider World Wide Web Conference Tracks from the past three years. We obtain PC lists by manually looking up names in DBLP. Since SwetoDblp was derived from DBLP, person URIs in SwetoDblp match their DBLP entry. We determine the expertise profiles for each PC member using Algorithm 1 discussed in Section 3.1 and compare them against our expert finder list.

The input topics for which relevant expertise is needed for each track, was obtained from the Call for Papers (CFP) posting. Determining the relevant topics from the CFP might seem to be a challenging task. However, by using the taxonomy of topics, the closest topics to the CFP was selected in a relatively straightforward manner. Table 5 shows the topics we used for the WWW2006 Search Track based on the CFP and our taxonomy.

Table 5. WWW2006 Search Track Input Topics and Subtopics

Topics	Subtopics (included automatically)
Search	Fighting Search Spam, Search Engines, Search Engine Engineering, Search Engineering Improved Search Ranking, New Search Paradigms, Semantic Search, Search Technologies, Search and Querying, Similarity Search,
Ranking	Ranking and classification, Page Rank
Indexing	Indexing and querying, Search Querying and Indexing
Information Retrieval	Information Retrieval and Applications
Web Mining	Web mining with search engines, Mining the web
Web Search	None
Web Graph	
Link Analysis	

The list of subtopics shown is not exhaustive. The taxonomy of topics fully described in [6], has a maximum depth of three. Here we show a depth of two to avoid clutter.

The initial input topics allow us to do two things. First, it allows us to determine the *relevant* expertise profiles and expertise rank for each PC member (as described in Section 3.2). It should be noted that researchers may have expertise in a variety of areas, and thus considered experts in several fields. We require only the relevant expertise to determine whether a researcher is an expert given a set of input topics. The keen observer may be unsettled by the observation that a well-known researcher may have low expertise according to our list. However, we emphasize that our approach finds experts with respect to the input topics deemed relevant to the conference, workshop, etc. From these relevant expertise profiles, we achieve or second goal, which is that of finding a list of *relevant experts* for comparison against the PC-List.

To obtain the SEMEF list of experts given an array of topics, a second method outlined in Algorithm 2 (not shown), described by Equation (1) in Section 3.2 is used. This algorithm takes as input a list of topics and returns a list of researchers of highest expert rank for those topics.

Table 6. Program Committee List compared with SEMEF List

Percentage SEMEF List	Search Track (Number of PC Members in SEMEF List)				Cumulative Percentage in SEMEF
	Search2007	Search2006	Search2005	Average	
(top) 0 -10%	10	13	13	12	35%
10 – 20%	5	8	6	6	52%
20 – 30%	6	0	0	2	58%
30 – 40%	4	1	1	2	65%
40 – 50%	6	2	0	3	73%
50 – 60%	3	1	1	2	79%
60 – 70%	4	0	0	1	82%
70 – 80%	1	1	0	1	85%
80 – 90%	1	0	0	0	85%
90 – 100%	0	0	0	0	85%
Total	40/48	26/29	21/25	29/34	
	83%	89%	84%	85%	

Table 6 shows the relative distribution of experts in the PC-List compared with the SEMEF list. On average, SEMEF finds that 85% of the experts in the PC-List have some expertise in the topics of the track. Furthermore, of the average number of PC members per year (29), 35% are in the top 10% of our SEMEF list, while close to 60% is in the top 30% of our SEMEF list. This establishes both that the WWW Search Track has a good distribution of experts in its PC-Lists and that SEMEF is a plausible approach for finding them. Figure 4a shows a raw distribution of our results, while Figure 4b shows a cumulative distribution. Figures 4c and 4d show the average distributions.

a) 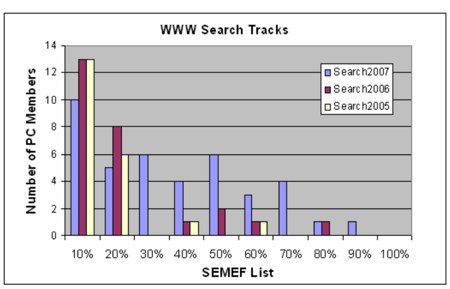

b) 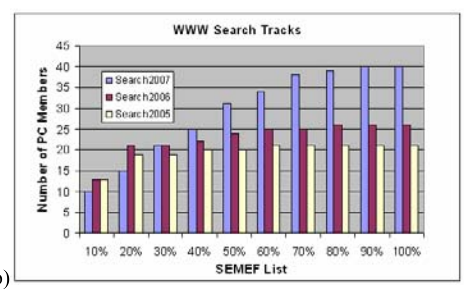

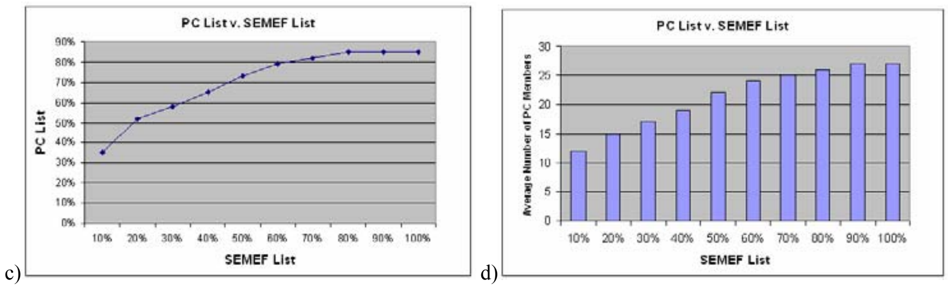

Figure 4. SEMEF Validation Results

4.2. Collaboration Network Expansion

Recommending experts for consideration on a program committee is not a straightforward task. To do this accurately, the closest relationships between PC Chairs and experts must be ascertained. As such, we find semantic associations between '*PC Chair-PC member'* pairs for each track. Table 7 shows that the majority of experts in the PC-List have a weak relationship to each Chair. This extends the argument that the WWW Search Track not only invites top experts, but more so experts with low collaborations levels relative to PC Chairs.

Table 7. PC Chair - PC Members Collaboration Relationships

Relationship	**PC List (Number of Expert Relationships)**						**Above Average Expertise (in PC)**
	Search2007		**Search2006**		**Search2005**		
	Chair1	Chair2	Chair1	Chair2	Chair1	Chair2	
STRONG	2	0	3	0	3	0	0
MEDIUM	10	7	6	2	7	8	4
WEAK	31	17	15	20	11	14	10
EXTREMELY WEAK	1	2	1	2	0	0	0

Table 8 shows that a larger number of experts, with expertise higher than the average expertise compared with the average expertise of the PC-List, but with equal collaboration levels relative to the PC Chair(s) exist. Through Collaboration Levels and C-Net clusters analysis, of such experts from the SEMEF list, more choices about possible experts to invite are available to the PC Chairs: a task otherwise undertaken manually. Hence the overall benefit of our semantic taxonomy-based approach.

Table 8. PC Chair - SEMEF List Collaboration Relationships

Relationship	SEMEF (Number of Expert Relationships)						Above Average Expertise (not in PC)
	Search2007		Search2006		Search2005		
	Chair1	Chair2	Chair1	Chair2	Chair1	Chair2	
STRONG	6	2	10	3	10	2	3
MEDIUM	106	53	88	55	88	76	16
WEAK	649	293	608	582	605	576	58
EXTREMELY WEAK	99	26	66	26	66	43	3

6. Related Work

In industrial settings, various approaches exist for finding experts. In [10], the concept of '*Expertise Atoms'* is used in software engineering VCS. The summation of expertise atoms (code changes) proves reliable in yielding expertise information. The obvious bottleneck in this approach is the extent of information sharing among the various software development companies, to whom such information might be important. Additionally, privacy and security concerns in many cases serve as a deterrent to information exchange. Thus, the key difference between this work and ours is that we exploit publicly available data for finding experts.

In [13] various algorithms for finding relevant reviewers are presented. Based on coauthorship graphs and relative-rank particle-swarm propagation, relationships between coauthors are given weights, which propagate through the coauthorship network using stochastic analysis on outgoing edges. States and energy levels of propagating nodes allow identification of the most qualified reviewers which are nodes with the highest energy levels within the network. This approach finds qualified reviewers for the bidding phase of the review process. We note that our approach is much more multifaceted. First, we make recommendations for not only discovering experts but also for expanding collaboration among them. Second, we provide the functionality of discovering close collaboration relationships, as well as analyzing the C-Nets of various researchers. This provides PC Chairs with the insight needed to form better PC-Lists. Furthermore, we extend our techniques to quantify expertise, in the form of expertise profiles for various experts. These are important examples of the extent and benefits of our approach.

The hierarchical ontological approach in [14], which classifies papers into expertise categories, bears some similarity to ours by using semantics on publication data. This similarity in approach is important in showing that there is good value in data derived from publications. Although this approach uses citation linkages and graph analysis to determine actual expertise values, we note the importance in using a taxonomy in relating publications to different topics. Obviously, the more well defined a taxonomy, the more well defined are expertise profiles, which is at the core of our work. We note one main difference between our work and this approach. In our work, we go one step further in using the taxonomy to find expertise in subtopics as well, to achieve finer granularity in expertise matching.

7. Conclusions & Future Work

The SEMEF approach presented in this paper is a new method for finding expertise, experts and expanding collaboration networks in the context of the peer-review process. We examined collaboration networks by discovering Semantic Associations between experts and PC Chairs using publications data. We also introduced the concept of CollaborationNets (C-Nets) for grouping experts.

In accomplishing these important tasks, a number of datasets were used, including a taxonomy of Computer Science topics. The taxonomy proved extremely useful in two key aspects. First, it was a central connection point for linking topics to papers, and subsequently obtaining experts on such topics. Second, the taxonomy allows us to find exact and inexact matches of expertise, which is significant if an expert finder application is to be meaningful.

We evaluated our methods by comparing experts found using our system with PC members from past conferences. We found that in general, our system is fairly accurate in corroborating the expertise of PC members. For instance, more than 50% of the PC list was found in the top 20% of our expert finder list. Furthermore, we found a significantly greater number of experts with similar relationships to PC Chairs and comparable (if not higher) expertise that could be considered for invitation for joining future program committees.

We realize that there is room for improving our methods. In fact, by merely improving techniques and sources for data collection, we stand to obtain additional data (not limited to publications) that could provide further information for collaboration level detection. A more sophisticated keyword Extractor Algorithm for obtaining topics from the Call for Papers will at a very minimum, reduce the manual input required in our system, and likely increase the probability of topic matches in our taxonomy. More complex expertise ranking techniques (e.g. [15]) could also improve the overall quality of our application. We investigate these in future work, with hopes of enhancing the SEMEF Semantic Web Expert Finder application.

References

1. Al-Sudani S., Alhulou R., Napoli A., Nauer E.: OntoBib: An Ontology-Based System for the Management of a bibliography. *17th European Conference on Artificial Intelligence,* Riva del Garda, Italy (Aug 28-Sept 3, 2006)
2. Aleman-Meza, B., Bojars, U., Boley, H., Breslin, J.G., Mochol, M., Nixon, L.J.B., Polleres, A., Zhdanova, A.V.: Combining RDF Vocabularies for Expert Finding, *4th European Semantic Web Conference*, Innsbruck, Austria, (June 3-7, 2007)
3. Anyanwu K., Maduko A., Sheth A.P.: SemRank: Ranking Complex Relationship Search Results on the Semantic Web. *14th International World Wide Web Conference*, Chiba Japan (May 10-14, 2005)
4. Bojārs, U., Breslin J.G.: ResumeRDF: Expressing skill information on the Semantic Web. *1st International ExpertFinder Workshop,* Berlin, Germany (January 16, 2007)
5. Cameron, D., Aleman-Meza, B., Arpinar, I.B.: Collecting Expertise of Researchers for Finding Relevant Experts in a Peer-Review Setting. *1st International ExpertFinder Workshop,* Berlin, Germany (January 16, 2007)
6. Decker, S.L., Aleman-Meza, B., Cameron, D., Arpinar, I. B.: Detection of Bursty and Emerging Trends for Identification of Researchers at the Early Stages of Research Areas, *(submitted for publication, http://lsdis.cs.uga.edu/~aleman/publications/decker07.pdf)*
7. Ding, L., Kolari, P., Finin, T., Joshi, A., Peng, Y., Yesha, Y.: On Homeland Security and the Semantic Web: A Provenance and Trust Aware Inference Framework. *AAAI Spring Symposium on AI Technologies for Homeland Security,* Stanford, California 2005
8. Gandon, F. Engineering an ontology for a multi-agent corporate memory system. Proc. ISMICK'01, 209-228.
9. Liu, P., Dew, P.: Using Semantic Web Technologies to Improve Expertise Matching within Academia, *2nd Int'l Conference on Knowledge Management,* Graz, Austria (June 2004)
10. Mockus, A., Herbsleb J. A.: Expertise Browser: A Quantative Approach to Indentifying Expertise. *International Conference on Software Engineering,* Orlando, Florida (May 19-25, 2002)
11. Mika, P.: Flink: Semantic Web Technology for the Extraction and Analysis of Social Networks, *Journal of Web Semantics,* 3 (2-3): 211-223, (2005)
12. Newman, M.E.J., Phys. Rev. e Stat. Phys. Plasmas Fluids Relate. Interdiscip. Top. 64, 016132, 2001
13. Rodriguez, M.A., Bollen, J.: An Algorithm to Determine Peer-Reviewers*, Los Alamos National Laboratory Technical Report,* LA-UR-06-2261, December 2005
14. Song, X., Tseng, B.L., Lin, C.-Y., Sun, M.-T.: ExpertiseNet: Relational and Evolutionary Expert Modeling, *10th Int'l Conference on User Modeling,* Edinburgh, Scotland (July 2005)
15. Zhang, J., Ackerman, M. S., Adamic, L.: Expertise Networks in Online Communities: Structure and Algorithms, *16th Int'l World Wide Web Conference*, Banff, Canada (May 8-12, 2007)