

Available online at www.sciencedirect.com



Journal of INFORMETRICS An International Journal

Journal of Informetrics 1 (2007) 68-82

www.elsevier.com/locate/joi

Mapping the bid behavior of conference referees

Marko A. Rodriguez^{a,*}, Johan Bollen^b, Herbert Van de Sompel^b

^a CCS-3: Knowledge & Information Systems Science Team, Los Alamos National Laboratory, United States ^b Digital Library Research & Prototyping Team, Los Alamos National Laboratory, United States

Received 7 July 2006; received in revised form 14 September 2006; accepted 19 September 2006

Abstract

The peer-review process, in its present form, has been repeatedly criticized. Of the many critiques ranging from publication delays to referee bias, this paper will focus specifically on the issue of how submitted manuscripts are distributed to qualified referees. Unqualified referees, without the proper knowledge of a manuscript's domain, may reject a perfectly valid study or potentially more damaging, unknowingly accept a faulty or fraudulent result. In this paper, referee competence is analyzed with respect to referee bid data collected from the 2005 Joint Conference on Digital Libraries (JCDL). The analysis of the referee bid behavior provides a validation of the intuition that referees are bidding on conference submissions with regards to the subject domain of the submission. Unfortunately, this relationship is not strong and therefore suggests that there may potentially exists other factors beyond subject domain that may be influencing referees to bid for particular submissions.

© 2006 Elsevier Ltd. All rights reserved.

Keywords: Peer-review; Referees; Bid behavior

1. Introduction

The peer-review process is the most widely accepted method for validating research results within the scientific community. However, its credibility as a valid certification mechanism has come under scrutiny. There exists a rich body of literature that points to many of the inadequacies of the current system (Evans, 1995; El-Munchid, 2001; Bence & Oppenheim, 2004), but of particular interest to this paper is the issue concerned with ensuring that referees are in fact reviewing manuscripts within their domain of expertise (Kassirer & Campion, 1994; Eisenhart, 2002). There exists a series of stages within the peer-review process that ultimately lead up to a referee review. One of the first and potentially most important stage is the one that attempts to distribute submitted manuscripts to competent referees. Unfortunately, it is difficult to study many of the stages of the peer-review process due to its confidential nature. Therefore, much of the peer-review process, including referee assignment, remains sheltered from the rigors of the scientific method. Fortunately, the program chairs and steering committee of the 2005 Joint Conference on Digital Libraries¹ (JCDL) has provided the Los Alamos National Laboratory (LANL) Digital Library Research and Prototyping team the referee bid data used for their 2005 conference peer-review process so that referee assignment could be analyzed for this study.

^{*} Corresponding author. Fax: +1 505 665 6452.

E-mail addresses: marko@lanl.gov (M.A. Rodriguez), jbollen@lanl.gov (J. Bollen), herbertv@lanl.gov (H. Van de Sompel).

¹ JCDL 2005 is located at: http://www.jcdl2005.org/.

^{1751-1577/\$ –} see front matter © 2006 Elsevier Ltd. All rights reserved. doi:10.1016/j.joi.2006.09.006



Fig. 1. Typical conference review stages.

In conference situations, where there exist a large number of submissions at one particular point in time (near the submission deadline date), conference organizers tend to rely on a pool of pre-selected referees to review the submission archive. The conference organizers require each referee to briefly look over each submission (e.g. read each submission abstract or ACM classification codes) and place submission bids. A referee bid states the referee's subjective opinion of their level of expertise with regards to a submission. Furthermore, conflict of interest situations are usually identified at this point. Once all the referee bids have been collected, the conference organizers can use any number of the many documented manuscript-to-referee matching algorithms to distribute each submission to a set of competent referees (Wei, Hartvigsen, & Czuchlewski, 1999). These stages are represented in Fig. 1. The data set provided by the 2005 JCDL program chair does not state which referees reviewed which submission, only the subjective opinion of the referee's level of expertise with respect to each submission.

Since conference organizers ask their referees to bid on submissions with regard to their domain of expertise, it is hypothesized that referee bidding is based on two factors: (1) the subject domain of the submission and (2) the expertise of the referee. The validity of this hypothesis is investigated using various statistical techniques that rely on a keyword analysis of submission abstracts and the location of each referee within the greater scientific community's co-authorship network. In short, the analysis demonstrates that the referees of the 2005 JCDL program committee are, in fact, bidding for submissions with respect to the subject domain of the submissions. Unfortunately, the strength of this relationship is not strong enough to conclude that submission subject domain is the only, or even the most significant, factor influencing referee bidding behavior.

2. The 2005 JCDL bid data set

The JCDL is an international forum that focuses on the technical, practical, and social issues concerning digital libraries. Each year, the JCDL hosts a conference to present technical papers, posters, demonstrations, tutorials, etc., that present recent developments in the digital library community. From June 7 to June 10 of 2005, the JCDL was held in Denver, Colorado in the United States (Sumner, 2005). The bid data provided by the 2005 JCDL program chair is considered extremely sensitive, therefore careful handling and analysis of this data was the first priority of this research endeavor. All information that is not publicly available from the JCDL website is, to the best of our knowledge, indeterminable from the presented results. Information, such as which submissions were rejected is not provided. The names of the referees have been anonymized by assigning each referee a unique random identifier. This section will discuss the bid data provided by the 2005 JCDL program chair as well as the various manipulations necessary to appropriately represent this information for analysis.

There were 264 submissions to the 2005 JCDL. Of those 264 submissions, 105 were full technical articles, 77 were short technical articles, 40 were posters, 17 were demonstrations, 4 were panel talks, 7 were tutorials, 7 were workshop talks, and 7 were doctoral presentations. The JCDL program committee provided the authors a table containing each submission's unique identification number, title, authors, type, and acceptance/rejection status. An example subset of this data is provided in Table 1. The submission titles and authors of those submission that were rejected by the committee have been replaced with the ### notation in order to protect the privacy of the submitters. Since accepted submissions are freely accessible, information pertaining to accepted publications is provided². Furthermore, note that the title and authors have been truncated to ensure that the table fits within the margins of this paper.

Each referee on the 2005 JCDL program committee was asked to bid on which submissions they wished to review in terms of their expertise in the subject domain of the submission. Therefore, accompanying the submission data table

² JCDL 2005 proceedings located at: http://www.informatik.uni-trier.de/ ley/db/conf/jcdl/jcdl2005.html.

Sub id	Submission title	Submission authors	Submission type	Submission status
13	###	###	Full Technical Article	Rejected
14	###	###	Full Technical Article	Rejected
15	Creating an Infrastructure for Collaboration	R. David Lankes,	Short Technical Article	Accepted
16	Graph-based Text Representation Model	Hidekazu Nakawatase,	Full Techinical Article	Accepted
17	An Evaluation of Automatic Ontologies	Aaron Krowne,	Full Techinical Article	Accepted

Table 1 Sample of the 2005 JCDL submission data

Table 2

Example bid matrix for each submission for each program committee referee, B

Sub/ref	1	2	3	4	5	
13	1	2	2	3	3	
14	2	3	2	3	3	
15	4	2	3	1	1	
16	3	3	1	2	0	
17	1	3	2	3	3	

Table 3

The meaning of the bid values within the bid matrix, **B**

Bid	Meaning of the bid value
0	Did not provide a bid
1	Expert in the domain of the submission and wants to review
2	Expert in the domain of the submission
3	Not an expert in the domain of the submission
4	Conflict of interest between referee and submission

Table 4

The meaning of the bid values within the modified bid matrix, \mathbf{B}'

Bid	Meaning of the bid value
0	Unknown expertise (wildcard)
1	Expert in the domain of the submission
2	Not an expert in the domain of the submission

there also exists an associated bid matrix, $\mathbf{B} \in \mathbb{B}^{|S| \times |R|}$, where *S* is the set of submissions, *R* is the set of referees, and $\mathbb{B} = \{0, 1, 2, 3, 4\}$. It is important to note that $|S| \gg |R|$. The rows of the bid matrix refer to the unique id of each of the submissions. The columns of the bid matrix refer to the referees of the program committee. The matrix entries are the bid values provided by each referee for each submission. Therefore, $b_{i,j}$ refers to referee *j*'s bid for submission *i*, where $0 \le b_{i,j} \le 4$. Table 2 is an artificial example of the supplied bid information. Note that the bid values for Table 2 were randomly generated and the referee names are not provided. The actual program committee for the JCDL is public information³, but their respective bid vectors are not.

The values of the bid matrix, **B**, are not on an interval scale, but instead are nominal (i.e. each value symbolizes a particular bid type). Table 3 provides the meaning for each of the bid values.

The bid matrix provided by the 2005 JCDL, **B**, contains extraneous information, such as 'wants to review' (b = 1) and 'conflict of interest' (b = 4). Since this study focuses specifically on referee expertise, this information will be discarded. Therefore, bid categories 1 and 2 will be considered the same and bid categories 0 and 4 will be considered wildcards. The modified bid matrix used throughout the remainder of this study has the properties of $\mathbf{B}' \in \mathbb{B}^{|S| \times |R|}$ where $\mathbb{B}' = \{0, 1, 2\}$. Table 4 has the bid meanings of the modified bid matrix.

The original artificial bid matrix provided in Table 2 is thus transformed into the one shown in Table 5.

Of the 264 submissions, only 118 of the submissions have actual bid data. This means that 146 submissions had bids of all 0. Therefore, only those submissions with a complete set of bid data will be analyzed for the remainder of

³ JCDL 2005 program committee available at: http://www.jcdl2005.org/progcomm.html.

Table 5 Example modified bid matrix for each submission for each program committee referee, \mathbf{B}'

Sub/ref	1	2	3	4	5
13	1	1	1	2	2
14	1	2	1	2	2
15	0	1	2	1	1
16	2	2	1	1	0
17	1	2	1	2	2

this study. In addition, of the 76 program committee members of the JCDL, 11 members gave no bid information. No bid information is defined as an individual whose bid vector is all 0s. These referees were removed from the analysis. Finally, since a portion of this analysis is based on co-authorship behavior, those referee committee members not located within the DBLP⁴ were not included in this study. Of the remaining 65 referees, 5 were not in the DBLP. Therefore, the bid matrix as defined for the remainder of this study has 118 rows (submissions), and 60 columns (referees), $\mathbf{B}' \in \mathbb{B}'^{118 \times 60}$.

3. The methodology

Intuitively, when ignoring conflict of interest situations, referee bidding should be based on two factors: (1) the domain of the submission and (2) the domain of expertise of the referee. Therefore, the referee bid matrix should be the result of each referees analysis of the submission abstracts and the referee's area of expertise (their location in the scientific community's co-authorship network). This idea, which is the hypothesis of this study, is represented by the arced dotted lines at the top of Fig. 2. To verify or falsify this hypothesis, a collection of statistical techniques are used to determine the relationship between referee bidding and submission subject domain. The two factors of the hypothesis are explored according to Tracks 1 and 2 of Fig. 2.

Track 1 provides a correlation between two submission similarity matrices. The first similarity matrix is constructed using referee bid data, S_b , and the second is constructed according to an submission abstract term analysis, S_t (Section 4). If referees are in fact bidding according to the subject domain of the submissions, then the correlation between S_b and S_t should be high. If the correlation is negative, or extremely low, then other factors that may not include submission subject domain are influencing referee bidding. Furthermore, it is possible to cluster submissions according to referee bid behavior. A intra-term analysis of these clusters provide an entropy value for each of the clusters. If the clusters created by referee bidding maintain a low entropy for their highest weighted terms and a low correlation between their term vectors, then it can be argued that referee bidding is driven by submission subject domain.

Track 2 provides the correlation between a referee similarity matrix created according to referee bidding behavior, \mathbf{R}_{b} , and a referee similarity matrix created using a relative-rank algorithm within a co-authorship network, \mathbf{R}_{g} (Section 5). A high correlation means that referees who are similar in expertise, as determined by their place in the co-authorship network, are also bidding similarly. A high correlation would be expected if referee bidding is based solely on submission subject domain. If this correlation is low, then other factors besides submission subject domain are influencing referee bidding. This paper will first explore Track 1 and then Track 2.

4. The bid matrix and submission similarity

This section will present the Track 1 analysis represented in Fig. 2. In order to determine the relationship between referee bidding and submission subject domain, the submissions are related according to the bid behavior of the program committee referees, S_b , and are related according to their abstract term-frequency inverse document-frequency (TFIDF) term weight distributions, S_t (Salton, 1998). In short, a TFIDF calculation determines the most descriptive words within a document (or document cluster) with respect to the entire document corpus. This section will first discuss the construction of S_b and then S_t .

Since the values of the bid matrix, \mathbf{B}' , refer to semantic categories and not a gradient scale, a Hamming distance function is used to determine the similarity of any two submissions (Hamming, 1950). Hamming distance is defined

⁴ Digital Bibliography and Library Project available at: http://www.informatik.uni-trier.de/ ley/db/.



Fig. 2. Experiment outline.

as the amount of characters that differ between two strings of equal length. For example, if there exists the strings "2212" and "1212", the Hamming distance is 1 since only their first characters differ. Given the Hamming distance between two bid vectors, $h(\vec{b'}_i, \vec{b'}_j)$, and the length of a vector, $l = |\vec{b'}_l|$, the similarity between any two submissions is calculated according to Eq. (1). To account for wildcard bids $(b'_{i,j} = 0)$, if any one of the two bid vectors being compared has an entry that contains a 0, that particular entry on both vectors is ignored and both their vector lengths, l, are reduced by 1. For example, when comparing the two bid vectors "0121" and "2120", their length, l, is 2 and their Hamming distance, h, is 0 because both their first and last entries are ignored and their second and third entries are equal. Therefore, their similarity is 1.

$$\mathbf{S}_{bi,j} = \mathbf{S}_{bj,i} = 1 - \frac{h(\vec{b'}_i, \vec{b'}_j)}{l}$$
(1)

Eq. (1) ensures a symmetrical submission similarity matrix, $\mathbf{S}_b \in \mathbb{R}^{|S| \times |S|}$, whose diagonal values are 1.0. According to the sample bid matrix presented in Table 5, the submission similarity matrix shown in Table 6 is constructed using Eq. (1).

 Table 6

 Submission similarity determined according to their Hamming distance

Id	13	14	15	16	17		
13	1.0	0.8	0.25	0.25	0.8		
14	0.8	1.0	0.0	0.5	1.0		
15	0.25	0.0	1.0	0.66	0.0		
16	0.25	0.5	0.66	1.0	0.5		
17	0.8	1.0	0.0	0.5	1.0		



Fig. 3. Submission similarity represented according to a hierarchical cluster.

4.1. Submission similarity and the dendrogram

Once a submission similarity matrix, S_b , has been constructed it is possible to hierarchically structure the submissions into a dendrogram in order to visualize the relationship between the various submissions by means of the complete linkage technique (Lance & Williams, 1967). The submission dendrogram constructed from S_b is presented in Fig. 3. Note that the titles of the rejected submissions have been left out. Accepted submission titles have been truncated to ensure readability. Furthermore, larger cluster patterns are represented as the eight boxed sections and are denoted C1 through C8. These clusters were extracted from the dendrogram by setting a threshold on the dendrogram tree height. The threshold, which is 1.1, was arbitrarily selected to expose enough clusters to make the following analysis interesting.

 Table 7

 Cluster feature vectors of the keywords in the submission abstracts

Cluster/term	Browser	Built	Bureau	Bush
3	3	7	3	1
4	4	3	2	0
5	1	0	1	0

A manual review of the clusters with respect to the submissions they contain demonstrates a congruency between submission topic and referee bidding. To validate this qualitative claim, three statistical techniques are used. The first involves analyzing the abstracts of the submissions of each of the clusters in order to determine cluster subject domain. The second involves determining the entropy value of each cluster. Clusters that are more strict with respect to a particular subject domain will tend to have a lower entropy. The third technique provides a correlation between a similarity matrix constructed from the cosine similarity of the TFIDF term weight vectors of each submission, S_t , and the matrix constructed from the referee bid behavior, S_b . This final correlation provides a single quantitative value expressing the relationship between submission subject domain and referee bidding.

4.2. Entropy in the submission clusters

The eight major clusters of the submission dendrogram derived according to referee bid data can be validated as meaningful categorizations by analyzing the terms of the submission abstracts. This requires that all submission abstracts be parsed to determine the full collection of keywords across all submission abstracts. Each abstract is processed by removing stop words and then applying the Porter stemming algorithm (Porter, 1980). These two processes remove overly frequent words (i.e. the, and, it) and perform suffix stripping (i.e. computer and computation stem to comput), respectively. Each time a particular term in the full collection of keywords is used in an abstract of one of the cluster submissions, the term frequency for that term in that cluster is incremented by 1. An example feature vector is provided in Table 7. For example, for all the submissions in cluster 3, the term *built* was used seven times.

For each cluster *i*, it is possible to determine how specific a particular term *j* is to that cluster according to Eq. (2) where freq(*i*, *j*) is the frequency of term *j* in cluster *i*, n(i) the total number of terms in cluster *i*, *N* the number of clusters (which is always 8 for this experiment), and $n_c(j)$ is the number of clusters for which term *j* appears (Salton, 1998).

$$\text{TFIDF}(i, j) = \frac{\text{freq}(i, j)}{n(i)} \times \log_{10}\left(\frac{N}{n_{\text{c}}(j)}\right) \tag{2}$$

The higher the TFIDF weight for term j in cluster i, the more specific term j is to the cluster i and therefore the more suited it is as a description of the cluster's subject domain. The following table presents the TFIDF calculations for the sample feature vector presented in Table 8.

In order to determine the subject domain of each of the eight clusters, the top 10 TFIDF weighted terms were extracted. Table 9 provides these terms ordered by their TFIDF weight where term 1 has a higher weight than term 2. The term weight distributions derived from the TFIDF calculation of the cluster abstracts can now be represented according to their internal cluster information content. Internal cluster information content can be calculated using the standard entropy equation as defined according to its information theoretic sense (Shannon, 1948). The lower the entropy, the more specialized, or focused, the cluster. The higher the entropy, the less specialized. Since clusters vary in size, the entropy for a cluster is calculated only for the top 10 term weights presented in Table 9. Furthermore, since an entropy calculation is defined for a probability distribution, Eq. (3) normalizes the top 10 term weights.

Cluster TFIDF term	weight vectors	of the keywords	in the	submission	abstracts

Table 8

Cluster/term	Browser	Built	Bureau	Bush
3	0.00	0.08	0.00	0.03
4	0.00	0.05	0.00	0.00
5	0.00	0.00	0.00	0.00

Table 9 Top 10 terms for the eight clusters defined in Fig. 3

	C1	C2	C3	C4	C5	C6	C7	C8
1	webcast	behavior	extract	patent	name	hidden	ecl	photo
2	cyberinfrastructur	drew	powerpoint	tobacco	surrog	crawler	morf	preserv
3	ncknow	note	handel	invent	music	subschema	relev	video
4	interview	overload	mainli	determin	disambigu	flora	item-level	european
5	teacher	engag	train	hidden	segment	ontos	network	alert
6	descriptor	factor	graph	control	candid	queri	citat	dark
7	faculti	gather	step	american	network	expans	circleview	region
8	lesson	school	weight	chemic	tempor	homepag	dlii	busi
9	survei	teamsearch	algebra	compani	citat	plant	extract	addit
10	transcript	visualis	basi	searchabl	genet	reusabl	meta-inform	mobil

$$\text{TFIDF}'(i, j) = \frac{\text{TFIDF}(i, j)}{\sum_{k=0}^{k<10} \text{TFIDF}(i, k)}$$
(3)

The entropy of a cluster is then calculated over the probability distribution as described by Eq. (4), where H(i) is the entropy for cluster *i*.

$$H(i) = -\sum_{j=0}^{j<10} \text{TFIDF}'(i, j) \log_2(\text{TFIDF}'(i, j))$$
(4)

The entropy values for the eight clusters of the dendrogram presented in Fig. 3 are presented in Table 10.

It is interesting to note that C5, the lowest entropy cluster, is composed mainly of submissions associated with name disambiguation and music in digital library research. On the other hand, the highest entropy cluster C2 has a mix of more unrelated submissions ranging from digital libraries in educational settings to infectious diseases and terrorism. Figs. 4 and 5 present the distribution of the term weights for the top 10 terms of the eight clusters. The steeper the distribution tail, the lower the cluster entropy and therefore the more focused the cluster is towards its higher weighted terms. The analysis of the terms for each cluster points to a qualitative relationship between referee bidding and submission subject domain.

A more quantitative validation can be determined when the TFIDF term weight vectors of the eight clusters are compared using a Spearman rank-order correlation. The correlations are performed on the cluster's TFIDF term weight vectors which contain the entire abstract dictionary, D, where |D| = 2121. Table 11 provides the Spearman rank-order correlations for each cluster comparison. What is noticeable from Table 11 is that all the correlations are less than ± 0.12 . The fact that the clusters, which are organized by referee bidding, yield very low correlations between their TFIDF term weight vectors means that the clusters are well separated according to their term distributions. If these correlations were high, then it would be difficult to claim that the clusters are organized according to subject domain and thus referee bidding would not be related to submission subject domain. Since the correlations are all less that ± 0.12 , this confirms the hypothesis that there does exist a relationship between the bidding behavior of the conference referees and the subject domain of the submission abstracts. The next section will further explore the strength of this relationship.

Table 10 Entropy values for the eight clusters defined in Fig. 3

Cluster	Entropy
1	3.2668
2	3.2840
3	3.2148
4	3.2213
5	3.2025
6	3.2281
7	3.2610
8	3.2442



Fig. 4. Normalized TFIDF weights for the 10 terms of Table 9 for the clusters 1 through 4 defined in Fig. 3.

4.3. Cosine similarity correlation

To further quantify the relationship between referee bidding and a submission's subject domain, it is possible to correlate the relationship between submissions based on referee bidding, on the one hand, and the relationship between submissions based on their TFIDF term weight vectors, on the other. This requires the construction of the similarity matrix S_t , which denotes the cosine similarity between every submission with respect to their complete TFIDF term weight vector. This means that each term in the abstracts of each submission is analyzed according to the TFIDF



Fig. 5. Normalized TFIDF weights for the 10 terms of Table 9 for the clusters 5 through 8 defined in Fig. 3.

Spearma	spearman rank-order correlations for the 2121 FFIDF term weights of the eight clusters defined in Fig. 5							
	C1	C2	C3	C4	C5	C6	C7	C8
C1	1.0	0.10008	0.04776	0.01666	0.02687	0.08817	0.07590	0.00653
C2	0.10008	1.0	0.07817	0.03279	0.03958	0.07983	0.08225	-0.06390
C3	0.04776	0.07817	1.0	0.04314	0.10796	0.11926	0.14579	0.02117
C4	0.01666	0.03279	0.04314	1.0	0.04393	0.07461	0.06287	-0.04580
C5	0.02687	0.03958	0.10796	0.04393	1.0	0.09844	0.11082	-0.05910
C6	0.08817	0.07983	0.11926	0.07461	0.09844	1.0	0.14143	-0.00126
C7	0.07590	0.08225	0.14579	0.06287	0.11082	0.14143	1.0	0.00549

Table 11 Spearman rank-order correlations for the 2121 TFIDE term weights of the eight clusters defined in Fig. 3

0.02117

equation presented in Eq. (2). This results in a matrix $\mathbf{T} \in \mathbb{R}^{|S| \times |D|}$ where |S| is the size of the submission archive and |D| is the size of the full collection of terms of all abstracts in the submission archive. For this particular experiment \mathbf{T} is therefore defined as $\mathbf{T} \in \mathbb{R}^{118 \times 2121}$. The TFIDF term weight vector of each submission can be compared against every other submission's TFIDF term weight vector using the standard cosine similarity function presented in Eq. (5), where \vec{t}_i is the TFIDF term weight vector for submission *i*. This equation guarantees a symmetrical matrix with a diagonal of 1.0.

-0.04580

-0.05910

-0.00126

0.00549

$$\mathbf{S}_{ti,j} = \mathbf{S}_{tj,i} = \frac{\vec{t}_i \cdot \vec{t}_j}{\|\vec{t}_i\| \cdot \|\vec{t}_j\|}$$
(5)

 S_b and S_t similarities can not be assumed to be parametric⁵, therefore the more robust Spearman rank-order correlation, ρ , was adopted to asses correlations. The Spearman ρ correlation between S_t and S_b was determined to be 0.153. This means that submissions categorized according to a TFIDF analysis of their abstracts and submissions categorized according to the referee bid behavior are in fact positively correlated, though not strongly.

$$\rho = 0.153, \ p < 2.2^{-16}$$

0.00653

C8

-0.06390

In literature, there exists arguments against judging research similarity according to manuscript term usage (Leydesdorff, 1989, 1997). While TFIDF weighting is an established method for determining document similarity (Baeza-Yates & Ribeiro-Neto, 1999), the construction of S_t occurs only at the single term level and therefore, a term may have several conceptual mappings (e.g. *network* in the graph theoretic sense and in the communication infrastructure sense). Thus, S_t is inevitably an approximation of the conceptual mapping between the various submissions. However, given that the analyzed abstracts were generated by a small, well-focused community of researchers, such homonyms may not be so prevalent (Leydesdorff, 1989).

5. The bid matrix and referee similarity

This section will overview the experiment as described by Track 2 of Fig. 2. If referees are deemed similar in expertise, as determined by their relative location to one another within the scientific community's co-authorship network, then similar referees should be bidding similarly. To test this hypothesis, two referee similarity matrices are created. The first referee similarity matrix, $\mathbf{R}_b \in \mathbb{R}^{|R| \times |R|}$, is constructed from the transpose of the modified bid matrix, \mathbf{B}'^{T} . Each referee is compared to each other referee with respect to their bidding behavior. Based on the transpose of the artificial data from Table 2, the same similarity equation used to construct the submission similarity matrix, Eq. (1), can be used to construct the referee similarity matrix presented in Table 12. The next section will present a dendrorgam of \mathbf{R}_b before discussing the second referee similarity matrix, \mathbf{R}_g .

5.1. Referee similarity and the dendrogram

Given the referee similarity matrix, \mathbf{R}_{b} , the dendrogram in Fig. 6 can be constructed. Unfortunately, due to privacy issues, the referee names are not provided. What is noticeable from the dendrogram is the collection of

1.0

⁵ Although $\mathbf{S}_{\mathbf{b}}$ follows a normal distribution (F = 0.858), but the similarities of $\mathbf{S}_{\mathbf{t}}$ follow an exponential distribution.

Table 12 Referee similarity determined according to their Hamming distance

Ref	1	2	3	4	5
1	1.0	0.5	0.75	0.0	0.0
2	0.5	1.0	0.2	0.4	0.75
3	0.75	0.2	1.0	0.2	0.0
4	0.0	0.4	0.2	1.0	1.0
5	0.0	0.75	0.0	1.0	1.0



Fig. 6. Referee similarity represented according to a hierarchical cluster.

nearly identical referees on the upper branch. When reviewing the modified bid matrix, \mathbf{B}' , it becomes apparent that 19 of the referees stated themselves to be expert in the domain of every submission (excluding their wildcard bids).

5.2. Relative-rank correlation

In order to provide a quantitative evaluation of the similarity of referees with respect to their bidding behavior and their domain of expertise, a relative-rank algorithm within a co-authorship network is computed to determine referee similarity. It has been widely accepted that co-authorship networks represent the relationship of individuals with respect to their domain of expertise (Newman, 2004). The relative-rank algorithm will determine the similarity of each referee with respect to each other referee as defined by their relative location to one another within the greater scientific community's co-authorship network. The similarity of the referees as determined by their relative-rank, \mathbf{R}_{g} , and their similarity as determined by their bid behavior, \mathbf{R}_{b} , can then be correlated. A high correlation means that referees of similar expertise are bidding in a similar manner. A low correlation means that referees of similar expertise are not bidding in a similar manner. The co-authorship network, *G*, used for this experiment was constructed from the DBLP database as of October 2005. The DBLP co-authorship network has 284,082 nodes (authors) and 2,167,018 edges (co-authorship relationships). This section will first formalize the co-authorship network data structure and relative-rank algorithm before discussing the results.

A co-authorship network is defined by a graph composed of nodes that represent authors and edges that represent a joint publication. Therefore, a co-authorship network is represented by G = (N, E, W), where N is the set of authors in the network, E is the set of edges relating the various authors, and W is the set of weights associated with the strength of tie between any two collaborating authors. Any edge, $e_{i,j}$, connects two authors, n_i and n_j , with a respective weight of $w_{i,j}$. Furthermore, $E \subseteq N \times N$ and |E| = |W|. The edge weight between any two authors is determined by Eq. (6), where the summation is over the set of all manuscripts registered with the DBLP, M, expressing a collaboration between authors n_i and n_j , and the function A(m) returns the total number of authors for manuscript m, where $m \in M$ and $w_{i,j} \in \mathbb{R}^+$ (Liu, Bollen, Nelson, & de Sompel, 2005; Newman, 2001).

$$w_{i,j} = w_{j,i} = \sum_{\forall m \in M \text{ authored by } i,j} \frac{1}{A(m) - 1}$$
(6)

To provide the reader with an understanding of the relationship between the 2005 JCDL program committee members, a subset of the DBLP co-authorship network which contains the program committee's co-authorship relationships is presented in Fig. 7. Note that this network was not constructed using referee bid data, but from information that is publicly available through the DBLP database. Furthermore, the co-authorship edge weights have been left out to improve readability.

The final analysis to be performed is to rank each of the 60 referees relative to one another so as to construct $\mathbf{R}_{g} \in \mathbb{R}^{|R| \times |R|}$, Eq. (7). For each referee in the JCDL program committee that provided valid bid data and is located in the DBLP, a similarity value to every other member in the committee was computed using a relative-rank algorithm (sometimes called a 'personalized' rank) (Rodriguez & Bollen, 2005; White & Smyth, 2003) within the DBLP co-authorship network. Since *G* is a weighted graph, the ranking algorithm actually used in this experiment is the weighted relative-rank implementation described in (Rodriguez & Bollen, 2005).

$$\mathbf{R}_{g} = \begin{pmatrix} \mathbf{R}_{g_{R_{1},R_{1}}} \cdots \mathbf{R}_{g_{R_{1},R_{|R|}}} \\ \vdots & \ddots & \vdots \\ \mathbf{R}_{g_{R_{|R|},R_{1}}} \cdots \mathbf{R}_{g_{R_{|R|},R_{|R|}}} \end{pmatrix}$$
(7)

An example of relative-ranking is as follows. Given a network, such as the one displayed in Fig. 7, the relative-rank algorithm would rank *FOX* more strongly to *NELSON* than to *RAY* since there exists a clique relationship between *FOX*, *NELSON*, and their co-authors. This network structure does not exist between *FOX* and *RAY*. Since co-authorship networks relate individuals with respect to similar domains of expertise, the conclusion to be drawn is that the stronger ranking of *FOX* to *NELSON* implies that *FOX* is more related by expertise to *NELSON* than he is to *RAY*. A simpli-



Fig. 7. Subset of the DBLP co-authorship network containing only connected JCDL referees.

fied version of the pseudo-code for constructing \mathbf{R}_{g} , Eq. (7), is presented in Algorithm 1. For a more indepth, and formal, review of relative-rank algorithms for network analysis, refer to (Rodriguez & Bollen, 2005; White & Smyth, 2003).

Algorithm 1. Constructing the referee similarity matrix \mathbf{R}_{g}



Fig. 8. Human-factor noise in review stages.

Because the similarity values of \mathbf{R}_g are not on an interval scale, the use of a non-parametric correlation test, such as the Spearman rank-order correlation, was warranted. The Spearman rank-order correlation between \mathbf{R}_b and \mathbf{R}_g values was calculated to be 0.101. The positive correlation indicates that referees are bidding with respect to their domain of expertise, but the low correlation may point to other factors.

$$\rho = 0.101, \quad p < 2.2^{-16}$$

Co-authorship networks are not the only data from which author similarity within a domain can be determined. For instance, many authors may publish single authored manuscripts and thus may find themselves not connected within the greater scientific community's co-authorship network. Furthermore, not all collaborations, and therefore similarity of expertise, lead to co-authored publications (Borgman & Furner, 2002) and therefore any assessment of author similarities based on collaboration relationships derived from co-authorship provides only a partial picture of the entire space of existing similarities. Nevertheless, those collaborations that have effectively resulted in co-authorships are positively identified by this method. Using a wider range of bibliographic data may result in an incrementally more complete assessment of author similarity. Unfortunately, such an effort was beyond the scope of this study, but can be included in follow-up studies.

6. Conclusion

This paper provided an exploration of the bidding behavior of the 2005 JCDL program committee. The various analysis techniques used demonstrate that the 2005 JCDL program committee did, in fact, bid for conference submissions with respect to the subject domain of the submission. On the other hand, the strength of this relationship is low and therefore demonstrates that other factors may be involved in referee bidding. One such factor seems to be referee fatigue. With 146 submissions having no bid data and with 19 referees stating themselves to be an expert in the domain of all submissions, human-driven referee bidding in conference settings may not be the most optimal technique for performing conference peer-review. Since bidding is the preliminary component of the manuscript-to-referee matching algorithm, sloppy bidding can have dramatic effects on which referees actually review which submissions, Fig. 8. In general, the stages that follow from the inclusion of noisy data in the peer-review chain can severely effect the quality of the peer-review process. It is speculated that referee fatigue not only influences the bidding and manuscript dissemination stages of the review cycle, but potentially more damaging, fatigued referees could be rejecting acceptable manuscripts or accepting fraudulent or faulty manuscripts in the review stage.

Future work in this area will focus on expanding this study's hypothesis in order to develop a mathematical model of the factors influencing referee bidding. Furthermore, an application of this methodology to bid data from other conferences can help to provide a broader perspective (and confirmation) of the factors influencing submission bidding in the peer-review process.

Acknowledgments

This research could only have been conducted with the help of the 2005 JCDL program chairs (Mary Marlino, Tamara Sumner, Frank Shipman) and steering committee (Erich Neuhold). Furthermore, the DBLP has once again provided the authors of this paper a thorough dataset for exploring the structure of science. The authors would like to thank the producers of R Statistics, GraphViz, Dia, and LATEX for their freely available software. Finally, Xiaoming Liu

and Karin Verspoor contributed by reviewing drafts of this manuscript. This research was financially supported by the Los Alamos National Laboratory.

References

Baeza-Yates, R. A., & Ribeiro-Neto, B. A. (1999). Modern Information Retrieval. ACM Press/Addison-Wesley.

- Bence, V., & Oppenheim, C. (2004). The influence of peer review on the research assessment exercise. Journal of Information Science, 30(4), 347–368.
- Borgman, C., & Furner, J. (2002). Scholarly communication and bibliometrics. Annual Review of Information Science and Technology, 36, 3–72.
- Eisenhart, M. (2002). The paradox of peer review: Admitting too much or allowing too little? Research in Science Education, 32(2), 241–255.
- El-Munchid, H. (2001). Evaluation of peer review in biomedical publications. Annals of Saudi Medicine, 21, 5-6.
- Evans, P. (1995). The peer-review process. Literati Newsline: Special Issue for Authors and Editors, 2.
- Hamming, R. W. (1950). Error-detecting and error-correcting codes. Bell System Technical Journal, 29(2), 147-160.
- Kassirer, J. P., & Campion, E. W. (1994). Peer review: Crude and understudied, but indispensable. *The Journal of the American Medical Association*, 272, 96–97.
- Lance, G., & Williams, W. (1967). A general theory of classificatory sorting strategies i. clustering systems. *The Computer Journal*, *10*(3), 271–277. Leydesdorff, L. (1989). Words and co-words as indiactors of intellectual organization. *Research Policy*, *18*, 209–223.
- Leydesdorff, L. (1997). Why words and co-words cannot map the development of science. Journal of the American Society for Information Science, 45(5), 418–427.
- Liu, X., Bollen, J., Nelson, M. L., & de Sompel, H. V. (2005). Co-authorship networks in the digital library research community. *Information Processing and Management*, 41, 1462–1480.
- Newman, M. E. J. (2001). Scientific collaboration networks. I. Network construction and fundamental results. Physical Review, 64.
- Newman, M. E. J. (2004). Coauthorship networks and patterns of scientific collaboration. Proceedings of the National Academy of Science, 5200– 5205.
- Porter, M. F. (1980). An algorithm for suffix stripping. Program: Automated Library and Information Systems, 14(3), 130-137.
- Rodriguez, M. A., & Bollen, J. (2005). Simulating network influence algorithms using particle-swarms: Pagerank and pagerank-priors. Technical report. Los Alamos National Laboratory [LAUR-05–6469].
- Salton, G. (1998). Term-weighting approaches in automatic text retrieval. Information Processing and Management, 24(5), 513-523.
- Shannon, C. E. (1948). A matematical theory of communication. Bell System Technical Journal, 27, 379-423.
- Sumner, T. (2005). Report on the fifth ACM/IEEE joint conference on digital libraries—cyberinfrastructure for research and education. *D-Lib* Magazine, 11(7/8).
- Wei, J. C., Hartvigsen, D., & Czuchlewski, R. (1999). The conference paper-reviewer assignment problem. Decision Science, 30(3).
- White, S., & Smyth, P. (2003). Algorithms for estimating relative importance in networks. In KDD '03: Proceedings of the ninth ACM SIGKDD international conference on knowledge discovery and data mining (pp. 266–275). New York, NY, USA: ACM Press.