

Investigations on Rating Computer Sciences Conferences

An Experiment with the Microsoft Academic Graph Dataset

Suhendry Effendy
National University of Singapore
effendy@comp.nus.edu.sg

Roland H. C. Yap
National University of Singapore
ryap@comp.nus.edu.sg

ABSTRACT

The rating of Computer Science (CS) conferences are important as it influences how papers published at the conferences and may also be used to evaluate research. In this paper, we proposed a method, *rSIT*, based on a small given set of top conference (*pivots*) and a relatedness measure based this set as well as basic baseline methods using citation count and field rating. We experimented with a snapshot dataset from Microsoft Academic Graph together with conference data from Microsoft Academic Search. We evaluated the conference ratings from our methods with the CCF conference rating list. We showed that *rSIT* correlates well with CCF rating and correlates better than ratings from using a baseline ranking with citation count or field rating.

Keywords

relatedness measure; classification; conference rating

1. INTRODUCTION

It is common to have informal or formal ratings of Computer Science (CS) conferences. By rating, we mean a quality/reputation measure for that conference. Typically conference ratings range from A, B or C (sometimes with 4 levels using A* as being higher than A). Such ratings are important since they also influence how papers published at the conference are perceived and also affect the paper submission process and the degree of competitiveness of the conference. Such ratings may also be used to evaluate research, e.g. a research granting agency may use this in the grant review process for a research grant.

Naturally a conference rating list is inherently subjective. It may also be tailored for a specific purpose. There are a number of public CS conference rating lists (e.g. CCF, see 3.1) which are in wide-spread use in particular contexts. There are also other lists which may not be publically available or are more informal, but that does not detract from their importance in a given context.

Rating conferences is subjective and typically a conference rating list is manually curated taking in various factors. Various lists may also not be in full agreement with each other, which is to be expected. However, we may expect also there to be substantial agreement.

Given the importance of conferences and conference rating in the research process within computer science, it is interesting to study what can be the underlying factors which may be important to the rating process. Although ultimately the rating process may need to be a manual process, we propose to study the possibility of computing some objective measures which give good correlation to the chosen ratings.

We highlight some terms used in the rest of the paper. When we say “conference rating” or simply rating, we mean the rating assigned to a conference which can come from a conference rating list or a rating computed using one of the methods in this paper, e.g. an A-rated conference. This is to be differentiated from a “conference ranking”, or simply ranking, which is some ordering defined on a list of conferences, e.g. rank by number of citations to papers in a particular conference.

In this paper, we extend previous work [6], which is perhaps one of the first initial works to investigate this question. Here, we propose a method based on a small given set of top conferences and a relatedness measure to this set which we call *rSIT*. Our method tries to take into account that the distribution of the relatedness measure differs between areas. We also propose two basic methods based on constructing a rank ordering based on citation count and field rating (a form of h-index). We evaluate all these methods, including the previous work, using the CCF conference rating list to serve as a source of independent conference ratings (CCF is widely-used in China). Our evaluation employs a snapshot of the Microsoft Academic Graph graph which contains scholarly big data released by Microsoft Research. The preliminary experimental results comparing our methods with CCF ratings suggest that citation count and field rating can give good correlation with the conference rating. However, our proposed method using our relatedness measure can give even better correlation and seems to give smaller differences in the level of disagreement versus CCF, i.e. turning an A-rated conference into C or vice versa.

1.1 Related Works

The problem on scientific collaboration networks have been well-studied. Most works investigate the structure of scientific research, focusing on the author collaboration work [1, 2, 3, 5, 8, 9]. However, there have not been much work on

investigating conference rating. Bird et al. [2] investigated the collaboration style between authors and areas in computer science. They also determined the areas of interest by manually selecting several representative conferences. Although they did not explicitly consider rating or conference reputation in their selection, we noticed that most of the selected conferences can be considered reputable.

Effendy et al. [4] proposed several conference relatedness measures. This was then applied in an initial investigation of conference rating prediction [6], which we will call rSTC. Another difference is that [6] experimented with bibliographic data from DBLP, here, we instead use the Microsoft Academic Graph dataset. In this paper, we propose a relatedness method, which we call rsIT. We also investigate some basic methods using citation count and field rating conference rankings.

2. RELATEDNESS MEASURE

The idea of conference relatedness measure is to have a metric which measures similarity between conferences. Three methods to measure relatedness between conferences were proposed in [4]: direct Jaccard, random walk, and pivot aggregation. We observed that among those three approaches, the random walk approach appears to give better similarity between conferences in related areas. In this work, we adapt the *random walk* (RW) approach.

Let C be the set of all conferences. Each conference $c \in C$ has a set (respectively, multiset)¹ of authors (c) . For any two conferences, c_i and c_j , we can compute their authors similarity using *Jaccard index*, which is defined as $w_{i,j} = \frac{|\alpha c_i \cap \alpha c_j|}{|\alpha c_i \cup \alpha c_j|}$, i.e. the size of intersection over the size of union of the two sets (or multisets).

The idea of relatedness measure is to relate conferences relative to a base conference which we call a *pivot*. Pivots are meant to be an influential conferences, e.g. highly rated conferences. In the RW method, relatedness between conferences is measured by the probability each conference c_i is visited in a random walk of length L originating from a pivot in the conference graph. A basic building block of the RW method is the transition matrix W' which is built by normalizing the Jaccard index as $W'_{i,j} = w_{i,j} / !_i$ and $!_i = \sum_j w_{i,j}$. Then $W'_{i,j}$ is the transition probability from c_i to c_j in the random walk. Let R^p be a vector of size $|C|$ (number of conferences) with p as the pivot conference; the i^{th} element of R^p (R_i^p is the relatedness score by RW from conference c_i w.r.t. p) is computed as follows. $E^p(s)$ is a vector of size $|C|$ representing the probability that a random walk of length s originating from p ends up at each conference.

$$E^p(0)_i = \begin{cases} 1; & \text{if } c_i = p \\ 0; & \text{otherwise} \end{cases}$$

$$E^p(s) = W' \times E^p(s-1) \text{ where } s > 0$$

$$R^p = E^p(L)$$

Thus, all scores in R^p are in the range of $[0::1]$ and sum up to 1. In our experiments, we chose $L = (v)$ where (v) is the *eccentricity* of node v , i.e. L is enough such that all conferences can be visited from the pivot, given all conferences are connected, which appear to be our case. In the rest

¹In our experiments, we use multisets.

Algorithm 1 Classify each conference into $[A, B, C]$ based on pivot thresholds (rsIT). P : set of pivots; C : set of all conferences; R is relatedness score of C to P ; $\#_A$ and $\#_B$: pivot thresholds parameter; Y : classified rating for all $c \in C$.

```

1: function rsIT( $R, P, C, \#_A, \#_B$ )
2:    $Y \leftarrow \{C \mid c \in C\}$ 
3:   for  $c \in C$  do
4:     for  $p \in P$  do
5:       if  $R_c^p \geq \#_A * R_p^p$  then
6:          $Y_c \leftarrow \max(R_c; A)$  .  $\max : A \succ B \succ C$ 
7:       else if  $R_c^p \geq \#_B * R_p^p$  then
8:          $Y_c \leftarrow \max(R_c; B)$ 
9:   return  $Y$ 

```

of this paper, we will refer the relatedness score R^p from a pivot by RW simply as *relatedness score*.

2.1 Rating Classification from Relatedness

In the previous work, Jahja et al. presented a method to classify conferences into ratings based on the relatedness score [6]. In this paper, we present a new approach also using relatedness which improves the resulting ratings as shown in the experiments. To differentiate both approaches, we will refer to the approach in [6] as rSTC (short for “rating by relatedness score with *threshold count*”), while the proposed approach in this paper is rsIT (short for “rating by relatedness score with *independent threshold*”). rsIT uses the relatedness score, but it treats the scores differently from rSTC.

We can think of the relatedness score as giving a certain ordering of conferences relative to the pivot. We observed that conferences with the same area to the pivot and more reputable are much closer to the pivot, i.e. have higher score. However, the strength of each pivot ordering varies from each other, e.g., conferences in cryptography (CRYPTO as pivot) are closer to each other than conferences in artificial intelligence (AAAI as pivot).² In part, this is due to AI being broader than cryptography, encompassing topics in vision, knowledge discovery, logics, search, etc. Based on these observations, we propose rsIT to classify the rating of conferences based on multiple pivot relatedness score by using two pivot thresholds, namely $\#_A$ and $\#_B$ (it can be generalized to more thresholds). Let P be the set of pivots. We classify a conference as **A**-rated if its relatedness score to any pivot $p \in P$ (i.e. R_c^p) is at least $\#_A * R_p^p$. If the score is between $(\#_A * R_p^p .. \#_B * R_p^p)$, we classify it as **B**; otherwise as **C**. Each pivot can be treated independently to each other, thus, this approach takes the differing strength of each pivot into consideration. The details of the rsIT algorithm is given in Algorithm 1. First, line 2 initializes all conferences as **C**-rated. Then, each conference examines its relatedness score to each pivot and determine their best possible rating based on thresholds $\#_A$ and $\#_B$ (line 3 - 8). Note that function *max* (line 6 and 8) returns the best rating out of the two parameters (the ordering is $A \succ B \succ C$). This algorithm can be easily extended to deal with more than three levels of conference ratings.

²Details on conference abbreviations and names can be found in the Appendix.

3. EVALUATING CONFERENCE RATINGS

3.1 Dataset

Recently (June 2015), Microsoft Research released a snapshot of its scholar data, namely the “Microsoft Academic Graph (MAG)” [7]. This dataset contains information on research publications across multiple disciplines, e.g., authors, journals, conferences, citations. We downloaded the dataset³ and extracted entries which correspond to publications in the CS conferences used in our evaluation. This data is then used to build a conference graph (using the Jaccard index as in [4, 6]) and compute the relatedness scores used in our rating classification experiments.

Next there is the question of how to evaluate the different rating techniques proposed in this paper. There are only a few public and widely used conference lists with ratings. In order to investigate and evaluate conference rating, we employed the China Computer Federation (CCF)⁴ list which rates each CS conference as A, B or C. CCF is widely used in China.

We highlight that conference rating is inherently subjective and such conference lists are typically the result of a manual curation process rather than an algorithmic process. Although conference rating is ultimately subjective, we have chosen for fairness to evaluate using an independent and widely used conference rating list. Such an evaluation is necessary to study what are the important factors which can make-up the construction of a conference rating list. We also remark that although CCF is a small list, we feel the choices in its curated ratings are, on a whole, quite reasonable.

In addition to rating CS conferences, CCF also classifies each conference by an area of research. As our method is pivot-based, we also employ the CCF area to choose the pivot conferences. CCF has a special category which they call *Inter* (Interdisciplinary and Emerging Topics). However, this is not really an area of research and appears more to be a miscellaneous category. *rSIT* assumes the chosen pivots are reputable and representative conferences from each area. As such, for this paper, we decided to remove all conferences listed in *Inter* from our dataset. We believe this removal is quite minor as there are only 7 conferences in *Inter*. Table 1 shows the number of conferences in CCF with A/B/C ratings by area of research and also the overall totals (excluding conferences in *Inter*).

Microsoft Academic Search⁵ (for short, we will call it “Libra”), is an experimental search engine service developed by Microsoft Research for academic publications. It also computes for their dataset, two conference lists which are sorted by citation count and field rating. The field rating is a measure defined by Libra, they explain that it is similar to h-index for authors⁶ extended to conferences. These two sorted lists can also be considered as giving a ranking (for their dataset) in terms of the citation count and field rating measures.

Given the citation count and field rating (from the Libra ranking lists), we can also generate a conference rating (Sec-

³<http://research.microsoft.com/en-us/projects/mag/>

⁴<http://www.ccf.org.cn/sites/ccf/paiming.jsp>

⁵<http://academic.research.microsoft.com/>

⁶An author with h-index of h has published h papers in which each has been cited by other papers at least h times.

Table 1: CCF Area and Rating Distribution

Area	CCF Rating			Σ
	A	B	C	
Ai-Pr	5	11	15	31
Cg-Mm	2	4	5	11
Db-DM-IR	5	9	11	25
Hci-Ubi	1	4	8	13
Inter *	(1)	(4)	(2)	(7)
Network	3	11	13	27
Se-Ss-PI	6	15	15	36
Sec	3	7	7	17
System	4	21	20	45
Theory	3	5	9	17
Σ - Inter	32	87	103	222

The CCF areas are: Ai-Pr: Artificial Intelligence & Pattern Recognition, Cg-Mm: Computer Graphic & Multimedia, Db-Dm-Ir: Database, Data Mining & Information Retrieval, Hci-Ubi: Human-Computer Interaction & Ubiquitous Computing, Inter: Interdisciplinary & Emerging Topics, Network: Computer Networks, Se-Ss-PI: Software Engineering, System Software & Programming Language, Sec: Network & Information Security, System: Computer System & High-Performance Computing, Theory: Theoretical CS.

Table 2: Tau-b Correlation between Libra Conference Ranking and CCF Rating

Libra Ranking	Tau-b to CCF
by citation count	0.619
by field rating	0.703

tion 3.2). Later, we evaluate with the rating by our *rSIT* algorithm.

3.2 Conference Ranking with Libra

We used the Kendall’s Tau correlation coefficient to analyze the correlation between Libra conference rankings and CCF rating, assuming that higher rating conferences should appear before lower rating conferences in the rank list. Specifically, we used Kendall’s Tau-b as obviously there are (many) ties in a rank list whose values are only A, B, and C. Two conferences are considered as being tied if both have the same rating and their appearance order in the list will not matter in Tau-b evaluation. The Tau-b statistic which ranges from -1 (perfect disagreement) to 1 (perfect agreement) can be used to measure the degree of agreement. Table 2 shows the Kendall’s Tau-b correlation coefficient between the Libra conference rankings and the CCF rating. As we can see, the Libra conference rankings have a strong correlation to the CCF rating with Tau-b correlation by field rating being higher. However, the correlations are still not so close to 1.0, thus, there are a number of pair of conferences which come out in reverse order of their CCF rating.

Figure 1 visualizes the citation count and field rating distribution in Libra against its CCF rating. As we can see, more A-rated conferences are on the left (have higher citation




Figure 1: Libra Citation Count and Field Rating Distribution and CCF Rating

count or field rating), while more C are to the right. Using a simple threshold, we can project each conference in a Libra ranking into a conference rating. Let N_A and N_B be defined as the desired number of A and B-rated conferences. We employed N_A and N_B as thresholds for the conference ranking as follows: the first N_A conferences in the ranking are classified as A, the next N_B are B, while the rest are C. We will refer to the rating derived from this process parameterized by N_A and N_B simply as “Libra rating”.

In our evaluation, we set N_A and N_B in accordance to the rating distribution in CCF. Table 3 shows the confusion matrix between projected rating from Libra conference ranking and CCF rating. The row corresponds to Libra rating, while the column corresponds to the CCF rating. For example, the entry at row-A column-C in the Citation Count table is 6; it means there are 6 conferences which are A-rated in Libra but C-rated in CCF. The diagonal entries correspond to the conferences where Libra and CCF agree. As N_A and N_B are chosen in accordance to CCF rating distribution, the row and column sums on Table 3 are identical. We can observe 3 kinds of disagreement in this confusion matrix: (i) AC (A-C or C-A); (ii) AB (A-B or B-A); and (iii) BC (B-C or C-B). We consider these kinds of disagreements as being at different levels and are prioritized by level. For example, a conference which is rated as A by one of its Libra ratings but rated as B in CCF (or vice versa), then it is an AB disagreement. We argue that, for the purposes of evaluating conference ratings, reducing AC disagreement is the most important as it has the largest rating difference, followed by AB. Depending on the situation, one may take into account or ignore BC disagreements. Note that a disagreement error may shift between levels so reducing AC may increase BC, hence, BC may be considered the least important.

The evaluation shows that the Libra ratings have good agreement with its CCF rating. As we can see in Table 3, the diagonal entries which correspond to the agreement are the largest in each row and column. However, we also observe some AC disagreements on some conferences, e.g. there are conferences which rated A by Libra rating, but C-rated in CCF. The field rating appears to have a better agreement to CCF than citation count as the diagonal sum is larger and it has fewer AC disagreement. This suggests that field rating is better in measuring conference rating reputation than citation count which is also supported by its higher Tau-b correlation in Table 2.

Table 3: Confusion Matrix of Libra Ranking and CCF

by Citation Count				
	A	B	C	Σ
A	19	7	6	32
B	13	52	22	87
C	0	28	75	103
Σ	32	87	103	222

by Field Rating				
	A	B	C	Σ
A	21	9	2	32
B	11	53	23	87
C	0	25	78	103
Σ	32	87	103	222

Row: Libra, Column: CCF

3.3 Evaluating conference ratings with rsIT

3.3.1 Ranking by Relatedness Measure

First, we show that conference ranking ordered by relatedness score has a strong correlation to its CCF rating. Figure 2 shows some representative results on relatedness score distribution over conferences in the same area with the pivots shown. The score shown for each conference is the highest score to the pivot set. The x-axis plot is normalized to the highest relatedness score in the data; in other words, $R_i^p = \max\{R\}$ where $\max\{R\}$ denotes the highest relatedness score among the respective pivots. Notice that the highest relatedness score will be for one of the pivot themselves (from the pivot to itself), thus, the left most data point in Figure 2 belongs to one of the pivots. The visualization has trends similar to Figure 1, more A rated conferences have higher score (closer to the pivot), while more C rated conferences are at the far end. However, we argue that this visualization shows that relatedness can give a clearer differentiation between conferences with different ratings for a set of conferences which motivates our rsIT method.

3.3.2 Conference Rating using Relatedness

We manually picked 17 A-rated conferences across various areas in CCF as pivots - the mapping to pivots is given in Table 4 and the conference abbreviations can be found in the Appendix. For a fair evaluation against CCF, we run the




Figure 2: Relatedness Score Distribution and CCF Rating

Table 4: 17 Selected Pivots

CCF Area	Pivots	Count
Ai-Pr	AAAI, ICML, CVPR	3
Cg-Mm	SIGGRAPH, MM	2
Db-Dm-Ir	SIGMOD, KDD, SIGIR	3
Hci-Ubi	CHI	1
Network	INFOCOM	1
Se-Ss-PI	ICSE, OSDI, POPL	3
Sec	CRYPTO, CCS	2
System	ASPLOS	1
Theory	STOC	1

Table 5: Confusion Matrix of rsIT ($\#_A = 0.339$, $\#_B = 0.109$) and CCF

	A	B	C	Σ
A	26	6	0	32
B	6	50	31	87
C	0	31	72	103
Σ	32	87	103	222

Row: rsIT, Column: CCF

rsIT (Algorithm 1) with $\#_A$ and $\#_B$ such that the number of A and B-rated conferences (N_A and N_B) correspond to the CCF rating distribution. Threshold parameters for $\#_A$ and $\#_B$ to satisfy N_A and N_B are computed by simple binary search ($\#_A$ and $\#_B$ are computed independently). We obtained values of $\#_A = 0.339$ and $\#_B = 0.109$. Table 5 shows the confusion matrix between the rating from rsIT with CCF. Observe that there are no CCF C-rated conferences classified as A by rsIT (i.e. no AC disagreement). However, such AC disagreements occurred in the Libra rating in Table 3. The AB disagreements by rsIT are also fewer than Libra rating. This suggests that rsIT rating may be better than the Libra rating.

We also compare our proposed relatedness-based algorithm, rsIT, with previously proposed rsTC[6]. rsTC directly uses N_A and N_B as upper bound parameters in rating conferences. Table 6 shows the confusion matrix between

rsTC and CCF. As we can see, only 17 conferences are classified as A by rsTC in our MAG-derived dataset, namely, the pivots themselves. The reason for this is because rsTC did not able to differentiate several closest conferences to the pivots and their count exceeds N_A . In this case, since N_A is an upper bound, then only the pivots are rated as A by rsTC.

Unlike rsTC, rsIT classifies each conferences based on the actual value of the relatedness scores (which are real numbers). Although the relatedness score value is not guaranteed to be unique, ties are less likely to happen in practice. This allows us to find a threshold $\#$ which can fit the count constraint N exactly. This is the case in Table 5 (notice that the row and column sums are identical). With rsTC, this is less likely to happen as it uses a discrete ranking from the relatedness score where there may be many ties.

Table 7 summarizes all the confusion matrices from Table 3, 5, and 6 in terms of agreement in A, B and C ratings (ordered by agreement priority) and disagreements in terms of AC, AB and AC (ordered by disagreement priority). As

Table 6: Confusion Matrix of rsTC [6] and CCF

	A	B	C	Σ
A	17	0	0	17
B	15	57	28	100
C	0	30	75	105
Σ	32	87	103	222

Row: rsTC, Column: CCF

Table 7: Confusion Matrix Summary

Total Agreement			
	A	B	C
LibraCC	19	52	75
LibraFR	21	53	78
rsTC	17	57	75
rsIT	26	50	72

Total Disagreement			
	AC	AB	BC
LibraCC	6	20	50
LibraFR	2	20	48
rsTC	0	15	58
rsIT	0	12	62

LibraCC: Libra rating by citation count
 LibraFR: Libra rating by field rating
 (Table 3)

we can see, rsIT gives better agreement in terms of the agreement priority. It also has the smallest disagreement in the term of the disagreement priority, i.e. lowest AC and AB disagreement.

4. CONCLUSION

In this paper, we proposed a method called rsIT to classify conferences into ratings and also methods based on ranking. The rsIT method uses a relatedness score to a set of pivots taking into account the different pivots have a different distribution of scores. We evaluated our methods using the Microsoft Academic Graph dataset and Libra results against the CCF conference ratings list. Experiments show that our rsIT method, which only uses bibliographic data, performs better (with respect to the CCF rating) than using ranking which based on citation count or field rating. Similarly, rsIT is better than a previous existing method (rsTC), which can suffer when there are many ties. Our experiments show that both relatedness score and citations (more so in the form of field rating) correlate well with conference rating, but relatedness score appears to deal better with errors which may occur when ranking by citation-based metrics. Possibly, this is because it is easier to manipulate citation counts, thus making it less robust. The field rating which is more difficult to manipulate than citation counts also seems to be better than citation counts. Some other experiments (beyond the scope of this paper) suggest that our relatedness measure is quite robust and consequently this also makes it hard to manipulate.

5. ACKNOWLEDGEMENTS

This work has been supported by Microsoft Research Asia.

6. REFERENCES

- [1] A. Barabási, H. Jeong, Z. Néda, E. Ravasz, A. Schubert, and T. Vicsek. Evolution of the social network of scientific collaborations. In *SIGKDD*, 2006.
- [2] C. Bird, E. Barr, A. Nash, P. Devanbu, V. Filkov, and Z. Su. Structure and dynamics of research collaboration in computer science. In *SDM*, 2009.
- [3] M. Biryukov and C. Dong. Analysis of computer science communities based on DBLP. In *ECDL*, 2010.
- [4] S. Effendy, I. Jahja, and R. H. C. Yap. Relatedness measures between conferences in computer science - a preliminary study based on DBLP. In *WWW*, Companion Volume, 2014.
- [5] R. Ichise, H. Takeda, and K. Ueyama. Community mining tool using bibliography data. In *Information Visualization*, 2005.
- [6] I. Jahja, S. Effendy, and R. H. C. Yap. Experiments on rating conferences with CORE and DBLP. *D-Lib Magazine*, 20(11/12), 2014.
- [7] Announcing the Microsoft Academic Graph: Let the research begin. https://blogs.msdn.microsoft.com/msr_er/2015/06/26/announcing-the-microsoft-academic-graph-let-the-research-begin.
- [8] F. Muhlenbach and S. Lallich. Discovering research communities by clustering bibliographical data. In *Intl. Conf. on Web Intelligence*, 2010.
- [9] O. R. Zaiane, J. Chen, and R. Goebel. Mining research communities in bibliographical data. In *WebKDD/SNA-KDD*, 2007.

APPENDIX

Conference Abbreviation & Name

Abbreviation	Conference Name
AAAI	AAAI Conf. on Artificial Intelligence
ASPLOS	Intl. Conf. on Architectural Support for Programming Lang. and Operating Systems
CCS	ACM Conf. on Computer and Communications Security
CHI	ACM Conf. on Human Factor in Computing Systems
CRYPTO	Advances in Cryptology
CVPR	IEEE Conf. on Computer Vision and Pattern Recognition
ICML	Intl. Conf. on Machine Learning
INFOCOM	IEEE Conf. on Computer Communication
ICSE	Intl. Conf. on Software Engineering
KDD	ACM Intl. Conf. on Knowledge Discovery and Data Mining
OSDI	USENIX Symp. on Operating Systems Design and Implementation
POPL	ACM Symp. on Principles of Programming Lang.
SIGGRAPH	ACM Intl. Conf. on Computer Graphics and Interactive Techniques
SIGIR	ACM Intl. Conf. on Research and Development in Information Retrieval
SIGMOD	ACM Intl. Conf. on Management of Data
STOC	ACM Symp. on Theory of Computing