

A Graph Analytics Framework for Ranking Authors, Papers and Venues

Arindam Pal
TCS Research
Tata Consultancy Services
Kolkata – 700156, India
arindam.pal1@tcs.com

Sushmita Ruj
Computer and System Sciences Division
Indian Statistical Institute
Kolkata – 700108, India
sush@isical.ac.in

ABSTRACT

A lot of scientific works are published in different areas of science, technology, engineering and mathematics. It is not easy, even for experts, to judge the quality of authors, papers and venues (conferences and journals). An objective measure to assign scores to these entities and to rank them is very useful. Although, several metrics and indexes have been proposed earlier, they suffer from various problems. In this paper, we propose a graph-based analytics framework to assign scores and to rank authors, papers and venues. Our algorithm considers only the link structures of the underlying graphs. It does not take into account other aspects, such as the associated texts and the reputation of these entities. In the limit of large number of iterations, the solution of the iterative equations gives the unique entity scores. This framework can be easily extended to other interdependent networks.

Keywords

CITATION INDEX, GRAPH ANALYTICS, INTERDEPENDENT NETWORKS, ITERATIVE ALGORITHMS, RANKING OBJECTS.

1. INTRODUCTION

An *author* is a person who publishes papers in some venues. A *paper* is a scientific publication, such as research/technical papers, books and book chapters. A *venue* is a place for publishing papers, like conferences, journals and workshops.

There are some general characteristics of good authors, papers and venues, which intuitively make sense. We list them below.

1.1 Characteristics of a Good Author

- Should have independent research capabilities.
- Should have many publications.
- Should have *important, high-quality* publications.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

International Workshop on Mining and Learning with Graphs, KDD 2016, San Francisco, California, USA

© 2016 ACM. ISBN 123-4567-24-567/08/06...\$15.00

DOI: 10.475/123_4

- Should have lot of citations for his papers.
- Should publish in *good* venues.
- Should work with other *good* authors.

1.2 Characteristics of a Good Paper

- Should be written by *good* authors.
- Should have many citations.
- Should be cited by other *good* papers.
- Should be published in *good* venues.

1.3 Characteristics of a Good Venue

- *Good* authors should publish there.
- *Good* papers should be published there.
- Should be cited by *good* papers published in other *good* venues.

There is no *perfect* citation index. Every index has some problems. A researcher can manipulate an index to his own advantage. We will explain some of the issues below.

1. **Self-citation:** An author can give citation to her own papers to boost her number of citations. Sometimes it's necessary, sometimes it is dishonesty. It can be easily detected.
2. **Mutual citation:** Researcher *A* gives citation to a paper of researcher *B*. Researcher *B* gives citation to a paper of researcher *A*. This is a mutual citation of length 2. It can be a longer cycle also. Now some people may form a cycle or even a clique of size *k*, and cite each others papers, thereby increasing their citation counts unfairly. This can be easily detected by a graph-based fraud detection algorithm.
3. **Big shots and the rich get richer phenomenon:** Suppose someone is a *big shot*, like a professor who has many PhD students. She is the center of a star graph. Any publication from her group will have her name, no matter whether she actually worked on that paper. Her paper and citation count will increase unfairly. This can be easily detected. This is also called the *rich get richer phenomenon*.

4. **Working in a large group:** Suppose a large group of ten people works on some problems. Any paper that comes out of that group has names of all the researchers. That increases the number of papers and citations artificially.
5. **Type of publication:** If you publish a book, book chapter or survey paper, that will have many more citations compared to a research paper. That does not mean that the book/survey has more research output.
6. **Area of research:** Some areas will have more citations than others. Papers in Natural Sciences like Physics or Biology have more citations than Computer Science. Within Computer Science, Machine Learning papers have more citations than Theoretical Computer Science papers.
7. **Difficulty of papers:** Generally, the more difficult and mathematical a paper is, the lesser the number of citations, because less people will understand and cite the paper. That does not mean that the paper is less important.
8. **Conference/journal type:** Some conferences typically have far more citations than others. ICASSP or ICC will have many more citations than STOC or KDD, because the first 2 accepts 1000+ papers, while the last 2 accepts 50-100 papers. That does not mean STOC/KDD is less important than ICASSP/ICC. In fact, people generally agree that it is much harder to get a paper accepted in STOC/KDD than ICASSP/ICC. This is also related to (6) and (7).
9. **Division of credits:** Some people get credit just because they are in a research group without giving any significant ideas. It is difficult to detect this.

2. RELATED WORKS

Many earlier works tried to rank authors and papers based on their importance. Some important ones are Hirsch's h -index [7], Individual h -index [1], Egghe's g -index [4], and Zhang's e -index [15]. One of the earliest work in this area is by Garfield [5], who defined a metric called *impact factor*. The impact factor of a journal is the average number of citations received by papers published in the journal in the past two years. Pinski and Narin [12] proposed a citation index by observing that not all citations are equally important. They noted that a journal is *important*, if it is cited by other *important* journals. Kleinberg [9] gave the first algorithm to rank web pages based on the hyperlink structure of the web graph. This algorithm is called Hypertext Induced Topic Search (HITS).

Brin and Page came up with the PAGERANK citation metric [2, 10] to use with the Google web search engine. This algorithm uses a *random walk* model on the web graph. They proved that the PAGERANK is the *stationary distribution* of this random walk. Gyöngyi et al. [6] gave an algorithm to detect spam web pages, which they called TRUSTRANK. The SIMRANK algorithm by Jeh and Widom [8] gives a measure of the similarity between two objects based on their relationships with other objects. Their basic idea is that two objects are similar if they are related to similar objects. Zhou et al. [16] proposed a method for co-ranking authors and their

publications using several networks associated with authors and papers. Their co-ranking framework is based on coupling two random walks that separately rank authors and documents using the PAGERANK algorithm.

Walker et al. [14] gave a new algorithm called CITERANK. The ranking of papers is based on a network traffic model, which uses a variation of the PAGERANK algorithm. Chen et al. [3] used a PAGERANK based algorithm to assess the relative importance of all publications. Their goal is to find some exceptional papers or *gems* that are universally familiar to physicists. Sun and Giles [13] proposed a popularity weighted ranking algorithm for academic digital libraries. They use the popularity of a publication venue and compare their method with the PAGERANK algorithm, citation counts and the HITS algorithm. Pal and Ruj [11] designed the CITEX citation index to assign scores to both authors and papers. They considered the structures of the publication graph and the citation graph.

3. PROBLEM DEFINITION AND MODEL

3.1 The Fundamental Graphs

- **Collaboration graph:** This graph is associated with the authors. The nodes of the graph are the authors. There is an undirected edge between two nodes, if the corresponding authors have written a paper together.
- **Citation graph:** This graph is associated with the papers. The nodes of the graph are the papers. There is a directed edge from a paper to another paper, if the first paper has cited the second paper.
- **Publication graph:** This graph is relating the authors, papers and venues. The nodes of the graph are the authors, papers and venues. There is an undirected edge between an author node and a paper node, if the author has written the paper. There is an edge between an author and a venue, if the author has published in that venue. There is an edge between a paper and a venue, if the paper was published in that venue.

We have a set of m authors $\mathcal{A} = \{a_1, \dots, a_m\}$, a set of n papers $\mathcal{P} = \{p_1, \dots, p_n\}$, and a set of r venues $\mathcal{V} = \{v_1, \dots, v_r\}$. A venue can either be a *journal*, *conference* or *workshop*. We represent this by a *publication graph* $G_P = (V_P, E_P)$, whose vertices are the set of authors, papers and venues, *i.e.*, $V_P = \mathcal{A} \cup \mathcal{P} \cup \mathcal{V}$. There are three types of edges in this graph.

- There is an undirected edge between author a_i and paper p_j , if author a_i has written paper p_j . Note that, this is a *symmetric* relation, so the edges are *undirected*. Associated with this, there is an $m \times n$ *author-paper matrix* M , whose rows and columns are a_1, \dots, a_m and p_1, \dots, p_n respectively, and whose $(i, j)^{th}$ entry $m_{ij} = 1$, if and only if author a_i has written paper p_j .
- There is an undirected edge between author a_i and venue v_k , if author a_i has published a paper at venue v_k . Associated with this, there is an $m \times r$ *author-venue matrix* N , whose rows and columns are a_1, \dots, a_m and v_1, \dots, v_r respectively, and whose $(i, k)^{th}$ entry $n_{ik} = 1$, if and only if author a_i has published a paper at venue v_k .

- There is an undirected edge between paper p_j and venue v_k , if paper p_j is published at venue v_k . Associated with this, there is an $n \times r$ *paper-venue matrix* L , whose rows and columns are p_1, \dots, p_n and v_1, \dots, v_r respectively, and whose $(j, k)^{th}$ entry $l_{jk} = 1$, if and only if paper p_j is published at venue v_k .

Since there are only edges between authors, papers and venues, the publication graph is a *tripartite graph*.

There is a *collaboration graph* $G_A = (V_A, E_A)$ associated with the authors, whose vertices are the set of authors, *i.e.*, $V_A = \mathcal{A}$. There is an *undirected* edge between author a_i and author a_j , if they have written a paper together. Note that collaboration (co-authorship) is not a *transitive* relation. Suppose authors a_i and a_j have written a paper p_r and authors a_j and a_k have written a paper p_s . It may be that a_i and a_k have not written a paper together. Associated with this, there is an $m \times m$ *collaboration matrix* Q , whose both rows and columns are a_1, \dots, a_m , and whose $(i, j)^{th}$ entry $q_{ij} = 1$, if and only if author a_i has written a paper with author a_j .

There is a *citation graph* $G_C = (V_C, E_C)$ associated with the papers, whose vertices are the set of papers, *i.e.*, $V_C = \mathcal{P}$. There is a *directed* edge from paper p_j to paper p_k , if paper p_j has cited paper p_k . Note that this is an *asymmetric* relation, so the edges are *directed*. Associated with this, there is an $n \times n$ *citation matrix* C , whose both rows and columns are p_1, \dots, p_n , and whose $(j, k)^{th}$ entry $c_{jk} = 1$, if and only if paper p_j has cited paper p_k . Note that the citation graph can't have any *directed cycle*. This is because a paper can only cite a previously published paper, so they are *totally ordered* in time. Moreover, the diagonal entries are all zero, because a paper can't cite itself.

We describe some important properties of the publication matrix and the citation matrix below.

- The *row sum* of the row i , $\sum_{j=1}^n m_{ij}$ is the number of papers published by author a_i .
- The *column sum* of the column j , $\sum_{i=1}^m m_{ij}$ is the number of authors who have written the paper p_j .
- There is no self-citation by any paper, *i.e.*, $c_{jj} = 0$ for $j = 1, \dots, n$.
- If the papers are numbered in *decreasing order of time* (newest first), the resulting citation matrix will be *upper-triangular*.

An example of a publication graph, a collaboration graph and a citation graph is given in Figure 1.

The following sets are important for further development.

1. For an author a , $AUTHORS(a)$ is defined as the set of co-authors of author a , *i.e.*, the set of authors who have written at least one paper with a . In other words, $AUTHORS(a) = \{b \in \mathcal{A} : (a, b) \in E_A\}$.
2. For an author a , $PAPERS(a)$ is defined as the set of papers written by author a . In other words, $PAPERS(a) = \{p \in \mathcal{P} : (a, p) \in E_P\}$.
3. For an author a , $VENUES(a)$ is defined as the set of venues in which author a has published. In other words, $VENUES(a) = \{v \in \mathcal{V} : (a, v) \in E_P\}$.




Figure 1: The publication graph, collaboration graph and citation graph showing authors, papers and venues with their edges labeled in black, blue and red colors respectively.

4. For a paper p , $AUTHORS(p)$ is defined as the set of authors who have written paper p . In other words, $AUTHORS(p) = \{a \in \mathcal{A} : (a, p) \in E_P\}$.
5. For a paper p , $CITE(p)$ is defined as the set of papers who have cited paper p . In other words, $CITE(p) = \{q \in \mathcal{P} : (q, p) \in E_C\}$.
6. For a paper p , $REF(p)$ is defined as the set of papers which have been given as reference (cited) by paper p . In other words, $REF(p) = \{q \in \mathcal{P} : (p, q) \in E_C\}$.
7. For a paper p , $VENUES(p)$ is defined as the set of venues in which paper p was published. In other words, $VENUES(p) = \{v \in \mathcal{V} : (p, v) \in E_P\}$. Note that, usually $VENUES(p)$ will have only one element v_k , where the paper p was published. But it may happen that the same paper may be published in a workshop, conference and journal. So, we keep the definition general.
8. For a venue v , $AUTHORS(v)$ is defined as the set of authors who have published in the venue v . In other words, $AUTHORS(v) = \{a \in \mathcal{A} : (a, v) \in E_P\}$.
9. For a venue v , $PAPERS(v)$ is defined as the set of papers published in the venue v . In other words, $PAPERS(v) = \{p \in \mathcal{P} : (p, v) \in E_P\}$.

4. OUR ALGORITHM

Our goal is to assign scores to authors and papers using the structure of the publication and citation graphs, so that important authors and papers get higher scores.

4.1 Computing Author, Paper and Venue Scores

For each author a_i , we have an *author score* (a -score) x_i . For each paper p_j , we have a *paper score* (p -score) y_j . For each venue v_k , we have a *venue score* (v -score) z_k . We represent the set of author scores as a column vector $\mathbf{x} = (x_1, \dots, x_m)^T$, the set of paper scores as a column vector $\mathbf{y} = (y_1, \dots, y_n)^T$, and the set of venue scores as a column

vector $\mathbf{z} = (z_1, \dots, z_r)^T$. We initialize all author, paper and venue scores to one, *i.e.*, $\mathbf{x} = \mathbf{1}_m, \mathbf{y} = \mathbf{1}_n$, and $\mathbf{z} = \mathbf{1}_r$. Then, we iteratively update the *a*-scores, *p*-scores and *v*-scores using the following rules.

1. For each author a_i , its *normalized a-score* \bar{x}_i is given by the *a*-score x_i divided by the sum of the number of co-authors, the number of papers written by him and the number of venues where his papers are published. In other words, $\bar{x}_i = \frac{x_i}{apv_i}$, for $i = 1, \dots, m$, where $apv_i = |AUTHORS(a_i)| + |PAPERS(a_i)| + |VENUES(a_i)|$ is the sum of the number of co-authors, the number of papers written and the number of venues where author a_i 's papers are published.
2. For each paper p_j , its *normalized p-score* \bar{y}_j is given by the *p*-score y_j divided by the sum of the number of authors who have written the paper, the number of papers who have cited the paper and the number of venues where the paper is published. In other words, $\bar{y}_j = \frac{y_j}{acv_j}$, for $j = 1, \dots, n$, where $acv_j = |AUTHORS(p_j)| + |CITE(p_j)| + |VENUES(p_j)|$ is the sum of the number of authors who have written the paper, the number of papers who have cited the paper and the number of venues where the paper is published.
3. For each venue v_k , its *normalized v-score* \bar{z}_k is given by the *v*-score z_k divided by the sum of the number of authors and papers published in venue v_k . In other words, $\bar{z}_k = \frac{z_k}{ap_k}$, for $k = 1, \dots, r$, where $ap_k = |AUTHORS(v_k)| + |PAPERS(v_k)|$ is the number of authors and papers published in venue v_k .

The iterative update rules are given below.

$$x_i = \sum_{l \in AUTHORS(i)} \bar{x}_l + \sum_{j \in PAPERS(i)} \bar{y}_j + \sum_{k \in VENUES(i)} \bar{z}_k \quad (1)$$

$$y_j = \sum_{i \in AUTHORS(j)} \bar{x}_i + \sum_{l \in CITE(j)} \bar{y}_l + \sum_{k \in VENUES(j)} \bar{z}_k \quad (2)$$

$$z_k = \sum_{i \in AUTHORS(k)} \bar{x}_i + \sum_{j \in PAPERS(k)} \bar{y}_j \quad (3)$$

The solution of these system of recursive, simultaneous equations gives the author, paper and venue scores. We have proved some results on the convergence of these equations. We are calculating the scores for real graphs from *Microsoft Academic Graph*, *Google Scholar* and *AMiner*. These results will be reported in the full version of this paper.

5. CONCLUSION AND FUTURE WORK

Here are some future directions to work on.

- Can we develop an axiomatic approach to citation analysis by starting from a set of basic axioms?
- In a real-world scenario, authors, papers and venues will be added over time. Dynamically modifying the scores from the current scores in an incremental manner is a challenging problem.

- Applying this framework to customer recommendation of products will be an important problem.
- This is an example of an *interdependent graph*. How to extend our method to other interdependent graphs?
- Using other parameters such as time of publication and age of authors can be useful.

6. REFERENCES

- [1] P. D. Batista, M. G. Campitelli, and O. Kinouchi. Is it possible to compare researchers with different scientific interests? *Scientometrics*, 68(1):179–189, 2006.
- [2] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems*, 30(1):107–117, 1998.
- [3] P. Chen, H. Xie, S. Maslov, and S. Redner. Finding scientific gems with google's pagerank algorithm. *Journal of Informetrics*, 1(1):8–15, 2007.
- [4] L. Egghe. Theory and practise of the *g*-index. *Scientometrics*, 69(1):131–152, 2006.
- [5] E. Garfield. Citation indexes for science. *Science*, 122(3159):108–111, 1955.
- [6] Z. Gyöngyi, H. Garcia-Molina, and J. Pedersen. Combating web spam with TrustRank. In *Proceedings of the Thirtieth international conference on Very large data bases-Volume 30*, pages 576–587. VLDB Endowment, 2004.
- [7] J. E. Hirsch. An index to quantify an individual's scientific research output. *Proceedings of the National academy of Sciences of the United States of America*, 102(46):16569–16572, 2005.
- [8] G. Jeh and J. Widom. Simrank: a measure of structural-context similarity. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 538–543. ACM, 2002.
- [9] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *J. ACM*, 46(5):604–632, 1999.
- [10] L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.
- [11] A. Pal and S. Ruj. CITE_X: A new citation index to measure the relative importance of authors and papers in scientific publications. In *IEEE International Conference on Communications, ICC 2015, London, UK, June 8-12, 2015*, pages 1256–1261, 2015.
- [12] G. Pinski and F. Narin. Citation influence for journal aggregates of scientific publications: Theory, with application to the literature of physics. *Inf. Process. Manage.*, 12(5):297–312, 1976.
- [13] Y. Sun and C. L. Giles. Popularity weighted ranking for academic digital libraries. In *Proceedings of the 29th European conference on IR research*, pages 605–612. Springer-Verlag, 2007.
- [14] D. Walker, H. Xie, K.-K. Yan, and S. Maslov. Ranking scientific publications using a model of network traffic. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(06):P06010, 2007.
- [15] C.-T. Zhang. The e-index, complementing the h-index for excess citations. *PLoS ONE*, 5(5):1–22, 2009.

- [16] D. Zhou, S. A. Orshanskiy, H. Zha, and C. L. Giles. Co-ranking authors and documents in a heterogeneous network. In *Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on*, pages 739–744. IEEE, 2007.