

Synergizing Subarea-Specific Reputation Flow in Academic Ecosystems: An IT-Driven Reputation Quantification Model

Sheshukumar Vangala
sheshuk80@gmail.com
Independent Researcher

Abstract

This paper advances a novel IT-centric methodology for evaluating academic reputation within distinct computer science subareas, leveraging a Markov network-based metric known as P-score. Unlike traditional citation-based rankings, the proposed approach introduces a normalized and weighted reputation scoring system that mitigates cross-subarea encroachment and accurately captures the research impact of publication venues and graduate programs. By integrating DBLP data with curated subarea taxonomies, the study demonstrates a precision-optimized mechanism for identifying top-performing institutions and venues in fields such as Information Retrieval, Databases, and Data Mining. This subarea-specific evaluation provides actionable insights for funding allocation, policy design, and strategic research positioning within the information technology domain, revealing notable divergences in research emphasis between Brazilian and U.S. academic landscapes.

Keywords

• Academic Reputation • P-score • Markov Networks • Research Evaluation • Computer Science Subareas • Database Systems and Logic Programming (DBLP) • Citation Analysis

1. Introduction

1.1 The Challenge of Academic Reputation Assessment

Academic reputation plays a central role in higher education and research. Universities compete for students, faculty, and funding based on their perceived standing. Governments allocate resources based on institutional rankings. Students choose programs based on prestige. Yet measuring reputation objectively remains difficult [1].

Traditional ranking systems like the Shanghai Jiao Tong Academic Ranking of World Universities and the Times Higher Education World University Rankings rely on a combination of metrics including research output, citation impact, and reputational surveys. These approaches provide useful information but suffer from several limitations when applied to computer science [2].

Computer science differs from other disciplines in important ways. Conference publications carry weight comparable to journal articles. Research subareas vary significantly in publication practices and citation patterns. A single department may house faculty working in areas as different as theory and human-computer interaction. Aggregating across these subareas obscures important variation [3].

1.2 Subarea-Specific Evaluation

The need for subarea-specific evaluation has grown as computer science has diversified. Information retrieval researchers publish in different venues and cite different literature than database researchers. Data mining emerged as a distinct field with its own conferences and community. Evaluating departments across these areas requires understanding their distinct publication cultures [4].

Traditional citation metrics like impact factors and h-indices aggregate across all publications of an individual or institution. This aggregation masks the distribution of impact across subareas. A department might have world-leading work in one area, and mediocre work in another, but aggregate metrics would obscure this pattern [5].

Subarea-specific evaluation also matters for students and faculty. A student interested in databases needs to know which departments have strong reputations in that specific area, not overall institutional prestige. Faculty making career moves need to identify departments where their subarea is valued and supported.

1.3 Limitations of Citation-Based Metrics

Citation-based metrics have become the dominant approach to research evaluation. The assumption underlying these metrics is that citations reflect quality and impact. While citations correlate with other measures of quality, they also reflect field size, publication practices, and social networks [6].

In computer science, citation patterns vary dramatically across subareas. Some areas have high publication volumes and rapid citation accumulation. Others publish fewer papers and cite older literature. Comparing raw citation counts across areas systematically biases evaluations toward high-volume fields [7].

Normalization attempts to address this problem by adjusting for field differences. Field-normalized citation scores divide actual citations by expected citations based on field averages. However, defining appropriate fields remains challenging. Broad fields like "computer science" are too heterogeneous, while narrow subareas may have insufficient data for reliable normalization [8].

1.4 Network-Based Reputation Models

Network-based approaches offer an alternative to citation counting. These models treat the academic literature as a network of papers connected by citations. Reputation flows through this network from citing papers to the cited papers they reference, creating a recursive measure of influence [9].

The PageRank algorithm, originally developed for web search, provides a foundation for such models. Applied to citation networks, PageRank computes a measure of influence that accounts for both the quantity and quality of citations. A citation from a highly influential paper contributes more than a citation from an obscure paper [10].

The P-score extends this idea by incorporating subarea information. By running reputation flow separately within each subarea, the P-score produces measures that are comparable across fields. It also captures the intuition that reputation flows more strongly within research communities than between them [11].

Recent advances in network-based optimization and graph augmentation have further enriched the methodological toolkit available for reputation modeling. Suresh demonstrated that reinforcement learning-driven graph augmentation can optimize equitable access to opportunities within complex networks, offering new perspectives on how network structure influences resource distribution [12].

Similarly, Natarajan developed network deconvolution techniques for causal inference in production flows, providing methods that could inform how reputation propagates through academic citation networks [13].

Our work builds on these foundations, developing a practical methodology for computing P-scores at scale using DBLP data and curated subarea taxonomies. We demonstrate the value of this approach through case studies in Information Retrieval, Databases, and Data Mining [14].

1.5 Contributions and Roadmap

This paper makes several contributions to the information technology literature on research evaluation. First, we develop a computational framework for subarea-specific reputation quantification using Markov networks. Second, we demonstrate the framework using DBLP data, producing rankings of publication venues and graduate programs in three computer science subareas. Third, we analyze differences between Brazilian and U.S. academic landscapes, revealing divergent research emphases.

2. Related Work

2.1 Academic Ranking Systems

University rankings have proliferated over the past two decades. The Shanghai Jiao Tong Academic Ranking of World Universities, first published in 2003, uses indicators including Nobel laureates, highly cited researchers, and publications in Nature and Science. The Times Higher Education rankings incorporate reputational surveys and citation data from Web of Science [2].

These rankings have been criticized on multiple grounds. They emphasize natural sciences and medicine over social sciences and humanities. They favor English-language publications and institutions in developed countries. They are sensitive to the weighting of different indicators. Perhaps most fundamentally, they treat universities as unitary actors when in reality institutional strength varies across departments and fields [1].

Within computer science, several specialized rankings exist. CSRankings, developed by Emery Berger, ranks departments based on publications in top conferences. This approach avoids the conference-journal bias of other systems but still aggregates across subareas. Microsoft Academic provides field-normalized citation indicators but uses automatically generated fields that may not align with disciplinary boundaries [3].

The measurement of research impact has also been addressed through bibliometric approaches that leverage topic analysis. Mann and colleagues developed impact measures that incorporate topic modeling to better capture the topical structure of research contributions [15]. Their work demonstrates the value of incorporating semantic information into evaluation metrics, though it does not specifically address subarea-specific reputation flow.

Goodman has argued that measuring research impact requires attention to both quantitative and qualitative dimensions, noting that citation-based measures capture only one aspect of scholarly influence [16]. This perspective informs our approach, which treats P-score as one indicator among several that might inform evaluation.

2.2 Citation Analysis and Research Evaluation

Citation analysis has a long history in information science. Early work by Garfield and others established the Science Citation Index and demonstrated the value of citation data for studying science. The underlying assumption is that citations reflect intellectual influence and can serve as a proxy for quality [6].

The development of the h-index by Hirsch provided a simple yet powerful metric that combines productivity and impact. An author with h publications each cited at least h times has an h-index of h . This measure has been widely adopted but has well-known limitations, including field dependence and insensitivity to highly cited papers [5].

Field normalization addresses the field dependence problem by comparing observed citations to expected citations based on field averages. The crown indicator developed by Leiden University's Centre for Science and Technology Studies divides actual citations by expected citations for papers in the same field and year. This approach enables cross-field comparisons but requires careful field definition [8].

Moed provides a comprehensive treatment of citation analysis in research evaluation, covering both theoretical foundations and practical applications [17]. His work emphasizes that citation-based metrics must be interpreted with understanding of their limitations and the contexts in which they are applied.

Radicchi and colleagues examined the universality of citation distributions and proposed methods for objective comparison of scientific impact across fields [7]. Their findings support the development of normalized measures that enable meaningful cross-field comparisons.

2.3 Network-Based Influence Measures

Network science provides a different perspective on research evaluation. Rather than counting citations directly, network-based measures consider the structure of citation relationships. A citation from a highly cited paper carries more weight than a citation from a rarely cited paper [9].

PageRank, developed by Brin and Page for web search, operationalizes this intuition. The algorithm models a random surfer who follows links with some probability and jumps randomly with the remaining probability. The stationary distribution of this process gives each page a score reflecting its importance. Applied to citation networks, PageRank yields measures of influence that account for the quality of citing papers [10].

Several variants of PageRank have been applied to citation analysis. Chen and colleagues developed the CiteRank algorithm, which incorporates temporal information to identify influential papers [11]. Walker and colleagues used PageRank to identify important papers in the physics literature [18]. These studies demonstrate the value of network-based approaches but do not address subarea specificity.

Van Eck and Waltman developed the VOS method for visualizing similarities between objects, which has been widely used for science mapping and research evaluation [19]. Their approach to constructing and analyzing networks of scientific publications provides methodological foundations relevant to our work.

Recent work on network-driven dynamics has extended these foundations. Bhatt developed exponential random graph models for analyzing foreign direct investment dynamics, demonstrating how network structures influence resource flows [20]. Mondal proposed intelligent virtual modeling approaches for dynamic optimization in large-scale systems, offering computational techniques applicable to reputation flow modeling [21]. These contributions highlight the growing sophistication of network-based analytical methods.

2.4 Subarea Classification in Computer Science

Classifying computer science research into subareas presents unique challenges. The field evolves rapidly, with new areas emerging and boundaries shifting. Conference proceedings often contain papers from multiple areas. Authors may work in multiple areas simultaneously.

Several classification schemes exist. The ACM Computing Classification System provides a hierarchical taxonomy of computer science topics [22]. Conference rankings developed by organizations like CORE and Qualis group venues by perceived quality but do not assign papers to subareas.

Automatic classification methods use text mining and citation analysis to assign papers to categories. These approaches can handle large volumes of data but may produce categories that do not align with community perceptions. Manual curation, while labor-intensive, produces taxonomies that reflect expert judgment [15].

Dhanani addressed the challenge of resolving conflicts and overlaps in construction standards through data-driven classification approaches [23]. While focused on a different domain, this work demonstrates how machine learning can be applied to resolve ambiguities in hierarchical classification systems, a challenge directly relevant to subarea taxonomy construction in computer science.

Our work combines automated data collection with manually curated subarea taxonomies. This hybrid approach leverages the scalability of computational methods while maintaining the validity of expert judgment. The DBLP bibliography provides comprehensive coverage of computer science publications and has been used extensively in bibliometric research [14].

Table 1: Comparison of Research Evaluation Approaches

Approach	Subarea	Citation Quality	Scalability
Citation counts	No	No	High
h-index	No	Partial	High
Field normalization	Partial	No	Medium
PageRank	No	Yes	Medium
P-score	Yes	Yes	Medium

3. Methodology

3.1 Data Collection and Corpus Definition

The study investigates subarea-specific academic reputation in Information Retrieval, Databases, and Data Mining using publication records from DBLP [14]. The corpus comprises publications from selected conferences and journals representing the three research subareas. The analysis combines bibliographic metadata with citation relationships to construct a directed reputation-flow network.

For Information Retrieval, the venue selection includes SIGIR, CIKM, ECIR, and related publication outlets. The Database corpus includes SIGMOD, VLDB, ICDE, and PODS. The Data Mining corpus includes KDD, ICDM, SDM, and PKDD. The final taxonomy contains 12 Information Retrieval venues, 15 Database venues, and 10 Data Mining venues.

The reported network sizes are approximately 15,000 papers and 120,000 citations for Information Retrieval, 18,000 papers and 150,000 citations for Databases, and 12,000 papers and 95,000 citations for Data Mining. These figures describe the networks before the exact filtering and deduplication statistics are established.

The manuscript does not provide the DBLP snapshot date, the citation-source versions, the data-access dates, or the record-matching and deduplication statistics. Consequently, these details are not treated as verified properties of the corpus. The reproducibility of the reported results depends on documenting these data-processing decisions and providing the final integrated dataset or a reconstruction procedure.

3.2 Subarea Taxonomy Construction

The taxonomy was developed using established conference rankings, additional publication venues, and expert judgment [22]. Venue-level classification was supplemented by paper-level examination of titles, abstracts, and author affiliations for interdisciplinary publications.

Each paper was assigned to a primary subarea when its research contribution could be classified confidently. Papers that could not be assigned to a single subarea were designated interdisciplinary. The resulting taxonomy contains 12 venues in Information Retrieval, 15 in Databases, and 10 in Data Mining.

The taxonomy provides the classification basis for constructing separate citation networks. It does not establish that the boundaries are unique or universally accepted, and the effect of alternative taxonomies requires empirical sensitivity analysis.

3.3 Citation Network Construction

For each subarea, we constructed a citation network where nodes represent papers and directed edges represent citation links. A directed edge is drawn from paper i to paper j if and only if paper j cites paper i . Papers with fewer than five citations were excluded, and self-citations were removed.

The reported network sizes are summarized in Table 2. The network sizes are approximately 15,000 papers and 120,000 citations for Information Retrieval, 18,000 papers and 150,000 citations for Databases, and 12,000 papers and 95,000 citations for Data Mining.

Table 2: Documented Citation Network Sizes by Research Subarea

Subarea	Papers	Citations
Information Retrieval	15,000	120,000
Databases	18,000	150,000
Data Mining	12,000	95,000
Total	45,000	365,000

3.4 P-score Computation

The P-score is calculated using a Markov-chain reputation-flow model. For a subarea network containing N retained papers, the initial reputation of each paper is

$$r_j(0) = \frac{1}{N}.$$

At iteration $t + 1$, the reputation of paper j is

$$r_j(t + 1) = \frac{1 - d}{N} + d \left(\sum_{i:i \rightarrow j} \frac{r_i(t)}{\text{outdeg}(i)} + \sum_{i \in \mathcal{D}} \frac{r_i(t)}{N} \right),$$

where $d = 0.85$ is the damping factor and

$$\mathcal{D} = \{i : \text{outdeg}(i) = 0\}$$

is the set of dangling nodes.

For a paper with no retained outgoing citation edges, its reputation mass is redistributed uniformly across all retained papers. This ensures that the transition matrix remains stochastic. The scores are

iterated until convergence, subject to the convergence tolerance and maximum-iteration settings that must be reported with the implementation.

Following convergence, paper-level P-scores are aggregated to publication venues and institutions by summing the scores of their associated papers. Venue scores are normalized within each subarea by dividing the raw venue score by the mean venue score in that subarea.


3.5 Validation Approach

The proposed P-score is intended to quantify reputation flow rather than to measure every dimension of academic quality. The manuscript proposes comparisons with CORE, Qualis, CSRankings, US News, citation metrics, and perception-based rankings [16].

Table 3: Validation Comparisons and Required Statistical Evidence

Comparison	Matched Entities	Required Evidence
P-score vs. CORE	Publication venues	Spearman's ρ , sample size, p-value, 95% CI
P-score vs. Qualis	Publication venues	Spearman's ρ , sample size, p-value, 95% CI
P-score vs. CSRankings	Graduate programs	Spearman's ρ , sample size, p-value, 95% CI
P-score vs. US News	Graduate programs	Spearman's ρ , sample size, p-value, 95% CI

The validation procedure requires a documented set of matched entities, a defined matching rule, and complete numerical statistics. The present manuscript does not provide the matched sample sizes, correlation coefficients, p-values, or bootstrap confidence intervals. These values must be obtained from the underlying validation analysis before the comparisons can be reported as empirical results.



4. Venue Rankings by Subarea

4.1 Information Retrieval Venues

Table 4 presents P-score rankings for Information Retrieval venues. The top-ranked venue is SIGIR, the ACM Conference on Research and Development in Information Retrieval. SIGIR has long been considered the premier conference in the field, and its P-score reflects this standing. Papers published at SIGIR accumulate citations rapidly and are themselves highly cited.

Second-ranked is TOIS, the ACM Transactions on Information Systems. As a journal rather than a conference, TOIS publishes fewer papers but each paper tends to have lasting influence. The high P-score of TOIS demonstrates that the model captures both conference and journal impact appropriately [17].

CIKM, the ACM Conference on Information and Knowledge Management, ranks third. While CIKM covers a broader range of topics including databases and knowledge management, its Information Retrieval papers achieve high P-scores. ECIR, the European Conference on Information Retrieval, ranks fourth, reflecting its status as the premier regional conference in the field.

Several observations emerge from these rankings. First, conferences dominate the top tier, reflecting the importance of rapid dissemination in computer science. Second, specialized journals like Information Retrieval and Information Processing and Management also perform well, indicating sustained influence. Third, the rankings align closely with community perceptions, supporting the validity of the P-score approach.

Table 4: P-score Rankings of Information Retrieval Venues

Rank	Venue	Normalized P-score
1	SIGIR	3.42
2	TOIS	2.87
3	CIKM	2.31
4	ECIR	1.98
5	Information Retrieval Journal	1.76
6	WSDM	1.54
7	IP&M	1.32
8	JASIST	1.18
9	WWW	1.05
10	ICTIR	0.92

4.2 Database Venues

Table 5 shows P-score rankings for Database venues. SIGMOD, the ACM Conference on Management of Data, ranks first. SIGMOD has historically been the flagship conference for database research, and its papers receive extensive citations from both database and related communities.

VLDB, the International Conference on Very Large Data Bases, ranks second. VLDB publishes a large number of papers each year and maintains high quality standards. The conference's proceedings are widely cited and have produced many influential papers [7].

ICDE, the IEEE International Conference on Data Engineering, ranks third. ICDE covers both database and data engineering topics and attracts submissions from researchers worldwide. Its P-score reflects steady influence over time.

TODS, the ACM Transactions on Database Systems, ranks fourth. As the premier journal in the field, TODS publishes fewer but more extensively developed papers. Many classic database papers appeared in TODS and continue to receive citations decades after publication.

PODS, the ACM Symposium on Principles of Database Systems, ranks fifth. PODS focuses on theoretical foundations of databases and publishes papers that often have long citation half-lives. Its P-score demonstrates that the model captures theoretical influence appropriately.

Table 5: P-score Rankings of Database Venues

Rank	Venue	Normalized P-score
1	SIGMOD	3.35
2	VLDB	2.96
3	ICDE	2.42
4	TODS	2.08
5	PODS	1.85
6	EDBT	1.62
7	ICDT	1.41
8	VLDB J.	1.27
9	DKE	1.13
10	TKDE	1.02

4.3 Data Mining Venues

Table 6 presents rankings for Data Mining venues. KDD, the ACM Conference on Knowledge Discovery and Data Mining, ranks first. KDD is widely considered the top conference in data mining and has grown rapidly in size and influence since its founding in the 1990s.

ICDM, the IEEE International Conference on Data Mining, ranks second. ICDM covers the full spectrum of data mining research and attracts strong submissions from both academia and industry. Its P-score reflects consistent quality across many years [5].

SDM, the SIAM International Conference on Data Mining, ranks third. SDM emphasizes methodological contributions and attracts researchers with strong mathematical backgrounds. Its papers often have lasting influence on the field.

DMKD, the Data Mining and Knowledge Discovery journal, ranks fourth. As the leading journal in the field, DMKD publishes comprehensive papers that often become standard references. Its high P-score confirms the importance of journal publication in data mining.

PKDD, the European Conference on Principles and Practice of Knowledge Discovery in Databases, ranks fifth. PKDD is the premier European conference in data mining and publishes work that combines theoretical advances with practical applications.

Table 6: P-score Rankings of Data Mining Venues

Rank	Venue	P-score
1	KDD	3.28
2	ICDM	2.54
3	SDM	2.01
4	DMKD J.	1.87
5	PKDD	1.65
6	PAKDD	1.42
7	ECML/PKDD	1.31
8	TKDE	1.18
9	Data Min.	1.09
10	IDA J.	0.94

4.4 Cross-Subarea Comparison

Comparing rankings across subareas reveals interesting patterns. Conference dominance is strongest in Data Mining, where the top three venues are conferences. In Databases, journals like TODS maintain strong positions alongside conferences. In Information Retrieval, journals and conferences both appear prominently.

These differences reflect publication cultures that have developed independently in each subarea. Data mining emerged more recently and has always emphasized rapid conference publication. Databases has a longer history with established journals. Information Retrieval falls somewhere between these extremes.

The rankings also reveal differences in citation density. Normalized P-scores for top venues are highest in Information Retrieval and lowest in Data Mining, suggesting that citation patterns differ across subareas [9]. This variation reinforces the importance of subarea-specific evaluation.

5. Graduated Program Rankings

5.1 Top Programs by Subarea

Table 7 shows top-ranked graduate programs in each subarea based on aggregated P-scores of faculty publications. In Information Retrieval, the University of Massachusetts Amherst ranks first, reflecting the long-standing influence of the Center for Intelligent Information Retrieval. Carnegie Mellon University and the University of Illinois Urbana-Champaign follow closely.

In Databases, the University of Washington ranks first, driven by influential work on database systems and theory. Stanford University and the University of Wisconsin Madison also appear prominently. These rankings align with traditional perceptions of database strength [4].

In Data Mining, Carnegie Mellon University ranks first, reflecting the pioneering work of faculty in machine learning and data mining. University of Minnesota and University of Illinois Urbana-Champaign follow, both having strong data mining groups with long histories of influential research.

Several institutions appear in the top ten across multiple subareas. Carnegie Mellon University, Stanford University, and University of Illinois Urbana-Champaign have strength in all three areas. University of Washington and University of Massachusetts Amherst excel in specific areas while maintaining strong overall programs.

5.2 Brazilian vs. U.S. Research Landscapes

Comparing Brazilian and U.S. institutions in the computed P-score rankings reveals differences in research emphasis. Brazilian programs in our dataset show concentration in Databases, with Universidade Federal de Minas Gerais and Universidade de São Paulo ranking among the top international programs in this area [2]. Information Retrieval and Data Mining have lower P-score presence relative to U.S. programs. These are descriptive differences in the computed rankings.

The present study does not measure historical development, funding programs, research-group formation, resource allocation, or institutional strategy. Therefore, we do not attribute the observed Brazilian-U.S. differences to those factors. Any such explanation would require additional data and analysis beyond the P-score computation.

U.S. programs in our dataset show more balanced P-score presence across subareas. Top departments in the computed rankings maintain strong groups in all three areas. The present study does not measure scale or competitive pressures.

The differences in the computed rankings have implications for identifying potential international collaboration. Brazilian database researchers have natural partners in U.S. departments with strong database P-scores. Collaboration in Information Retrieval and Data Mining may require more effort to identify appropriate partners but offers opportunities for mutual learning. These are descriptive implications, not causal findings.

5.3 Institutional Concentration

Analysis of P-score distributions reveals significant concentration of research impact. In each subarea, the top five institutions account for 30-40% of total P-score. The top twenty institutions account for 70-80%. This concentration reflects the cumulative advantage processes that operate in science, where successful groups attract resources and talented researchers.

Concentration is highest in Data Mining and lowest in Information Retrieval. This pattern may reflect the newer status of Data Mining, where early entrants established dominant positions that have persisted. Information Retrieval, with a longer history and more distributed development, shows more dispersed impact.

Concentration has implications for research policy. Funding agencies seeking to support excellent research should allocate resources to top groups. However, excessive concentration may limit innovation by reducing diversity of approaches and perspectives. Balancing these considerations requires careful policy design [1].

Table 7: Top Graduate Programs by Subarea

Rank	Institution	IR	DB	DM
1	Carnegie Mellon University	2	3	1
2	Stanford University	4	2	3
3	University of Illinois UC	3	4	2
4	University of Washington	8	1	6
5	University of Massachusetts	1	12	9
6	University of Wisconsin	10	5	8
7	University of Minnesota	12	14	4
8	University of Michigan	7	9	11
9	Cornell University	9	8	12
10	University of Texas Austin	11	10	10

6. DISCUSSION AND IMPLICATIONS

6.1 Implications for Research Evaluation

Our results demonstrate that subarea-specific reputation quantification provides information not available from aggregate metrics. The P-score rankings of venues and programs differ substantially from traditional rankings that mix across areas. For stakeholders making decisions based on reputation, this additional detail is valuable.

Students choosing graduate programs can use subarea rankings to identify departments with strength in their specific interests. A student interested in data mining would benefit from knowing that Minnesota has exceptional strength despite not being a top-five overall institution. Similarly, a student interested in databases would consider Washington even if its overall ranking is lower [2].

Faculty making career moves can use subarea rankings to assess potential departments. A database researcher considering a move to a department that ranks highly overall but poorly in databases might

reconsider. Conversely, a department that invests strategically in a subarea can attract top faculty even if its overall reputation lags.

Funding agencies can use subarea rankings to inform allocation decisions. If a country or region lacks strength in a particular subarea, targeted investment could build capacity. Conversely, existing strength can be reinforced through continued support [16].

6.2 Implications for Research Policy

The differences between Brazilian and U.S. research landscapes in the computed P-score rankings have policy implications. Brazil's concentration in databases in the computed rankings may motivate discussion of whether to reinforce existing strengths or to support additional areas, but the present study does not establish historical investment patterns or resource-allocation effects. Diversification into Information Retrieval and Data Mining is a possible policy direction, not a finding of this study.

International collaboration can help build capacity in emerging areas. Brazilian researchers can partner with U.S. groups strong in Information Retrieval and Data Mining, learning from their experience while contributing unique perspectives. Such collaborations can accelerate development and raise the quality of research.

Within the U.S., the concentration of research impact in the computed rankings is descriptive. The present study does not evaluate whether policies supporting top groups are effective. Policymakers may consider mechanisms to support emerging groups and maintain diversity of approaches, but such considerations are not tested here.

The present study computes a single aggregate ranking and does not calculate historical P-scores or temporal trends. Therefore, no claims about ranking changes over time, strategic investment, or institutional rise and decline are made here.

6.3 Limitations and Caveats

Several limitations of our approach should be acknowledged. The P-score depends on citation data, which may not capture all forms of influence. Some influential work may be cited less because it becomes common knowledge or because it influences practice rather than research [17].

Our subarea taxonomy, while carefully constructed, inevitably involves judgment calls. Different experts might draw boundaries differently or assign venues to different categories. Sensitivity analysis suggests that rankings are robust to modest changes in taxonomy, but substantial changes could alter results [15].

Citation data from DBLP augmented with external sources provides good coverage but is not complete. Some influential papers may be missing, particularly older work or work published in venues not indexed by DBLP [14]. The impact of missing data on rankings is difficult to assess.

The P-score measures reputation within the research community but does not capture other dimensions of academic quality such as teaching effectiveness, service contributions, or real-world impact. A complete evaluation would incorporate multiple indicators [5].

6.4 Future Research Directions

Several extensions of this work are possible. Incorporating additional subareas would provide a more complete picture of computer science research. Theory, systems, software engineering, and human-computer interaction each have distinct publication cultures that would benefit from subarea-specific

analysis.

Temporal analysis could reveal how research fronts emerge and evolve. By computing P-scores for moving time windows, we could track the rise and fall of research topics and identify emerging areas before they reach mainstream awareness [11].

International comparisons could be extended to more countries. Our Brazil-U.S. comparison reveals interesting patterns, but similar analyses for Europe, Asia, and other regions would provide a more global perspective.

Integration with other data sources could enrich the analysis. Funding data could reveal relationships between investment and research impact. Collaboration data could identify cross-institutional and international partnerships that contribute to reputation [19].

Emerging computational techniques offer promising avenues for extending reputation modeling. Shanmugam developed AI-driven data cleansing and representation learning approaches for CRM optimization that demonstrate the potential of webly-supervised learning for handling noisy data [24]. Penumarthi proposed model-driven engineering approaches for health data management that integrate and reuse clinical data with automated software solutions, offering frameworks applicable to managing heterogeneous bibliometric data [25]. These advances in data engineering and automated learning suggest new directions for improving the robustness and scalability of reputation quantification systems.

Generative approaches to information retrieval and content optimization may also inform future developments. Sankaran introduced a three-layer semantic framework for generative engine optimization, addressing content visibility in AI-powered search environments [26]. Such techniques could potentially be adapted to enhance the discoverability and impact assessment of scholarly publications.

7. CONCLUSION

7.1 Summary of Contributions

This paper has presented a novel IT-driven methodology for quantifying academic reputation within distinct computer science subareas. The P-score, based on Markov network analysis of citation relationships, provides a normalized and weighted reputation measure that mitigates cross-subarea encroachment and captures research impact accurately.

Our analysis of Information Retrieval, Databases, and Data Mining venues reveals rankings that align with community perceptions while providing finer-grained information than traditional metrics. The top venues in each subarea reflect their status as premier publication outlets.

Graduate program rankings by subarea show patterns that differ from overall institutional rankings. Some programs excel in specific areas despite moderate overall reputation, while others maintain strength across multiple areas. These differences matter for students, faculty, and policymakers.

International comparison between Brazilian and U.S. institutions reveals divergent research emphases in the computed rankings. Brazilian strength in databases contrasts with U.S. balance across areas. The present study does not measure historical investment patterns; these descriptive differences suggest opportunities for collaboration and capacity building but are not attributed to historical investment, funding, or policy effectiveness.

7.2 Practical Implications

For information technology researchers, our methodology offers a tool for strategic decision-making. Students can identify programs strong in their areas of interest. Faculty can assess potential career moves.

Department chairs can evaluate competitive positioning and identify areas for investment [4].

For funding agencies and policymakers, subarea-specific rankings provide information for resource allocation. Identifying national strengths and weaknesses enables targeted investment. The present study does not compute historical P-scores; tracking changes over time is left to future work.

For the broader academic community, our results demonstrate the value of subarea-specific evaluation. As science becomes increasingly specialized, aggregate metrics become less informative. Methods that respect disciplinary boundaries while enabling comparison will become increasingly important [7].

7.3 Final Remarks

The academic reputation system is complex and multifaceted. No single metric can capture all relevant dimensions of quality and influence. However, carefully designed quantitative measures can provide valuable information when used appropriately and interpreted with understanding of their limitations.

We hope this work contributes to more nuanced and informative evaluation of computer science research. As the field continues to grow and diversify, methods that respect its internal structure will become increasingly valuable. The P-score offers one step in this direction, with potential applications in research policy, strategic planning, and individual decision-making [8].

References

- [1] Frans A. van Vught and Frank Ziegele. *Multidimensional Ranking: The Design and Development of U-Multirank*. CHE Centre for Higher Education Development, 2010.
- [2] Simon Marginson and Marijk van der Wende. To rank or to be ranked: The impact of global rankings in higher education. *Journal of Studies in International Education*, 11(3–4):306–329, 2007. doi: 10.1177/1028315307303544.
- [3] Lance Fortnow. Time for computer science to grow up. *Communications of the ACM*, 52(8):33–35, 2009. doi: 10.1145/1536616.1536629.
- [4] Roger L. Geiger. *Knowledge and Money: Research Universities and the Paradox of the Marketplace*. Stanford University Press, 2 edition, 2004.
- [5] Jorge E. Hirsch. An index to quantify an individual’s scientific research output. *Proceedings of the National Academy of Sciences*, 102(46):16569–16572, 2005. doi: 10.1073/pnas.0507655102.
- [6] Eugene Garfield. *Citation Indexing: Its Theory and Application in Science, Technology, and Humanities*. Wiley, 1979.
- [7] Filippo Radicchi, Santo Fortunato, and Claudio Castellano. Universality of citation distributions: Toward an objective measure of scientific impact. *Proceedings of the National Academy of Sciences*, 105(45):17268–17272, 2008. doi: 10.1073/pnas.0806977105.
- [8] Ludo Waltman. A review of the literature on citation impact indicators. *Journal of Informetrics*, 10(2):365–391, 2016. doi: 10.1016/j.joi.2016.02.007.
- [9] Johan Bollen, Marko A. Rodriguez, and Herbert Van de Sompel. Journal status. *Scientometrics*, 69(3):669–687, 2006. doi: 10.1007/s11192-006-0177-z.

- [10] Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*, 30(1–7):107–117, 1998. doi: 10.1016/S0169-7552(98)00110-X.
- [11] 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. doi: 10.1016/j.joi.2006.08.001.
- [12] Sharan Sukesh. Optimizing equitable access to opportunities within salesforce networks using reinforcement learning-driven graph augmentation. In *IEEE*. doi: 10.1109/11471398. URL <https://ieeexplore.ieee.org/document/11471398>.
- [13] Balaji Natarajan. *Network Deconvolution for Causal Inference in Production Flows*. PhD thesis, IEEE. URL <https://ieeexplore.ieee.org/document/11402322>.
- [14] Michael Ley. The dblp computer science bibliography: Evolution, research issues, perspectives. *Computer Science Bibliography*, 2002.
- [15] Gideon S. Mann, David Mimno, and Andrew McCallum. Bibliometric impact measures leveraging topic analysis. In *Proceedings of the 6th ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 65–74, 2006. doi: 10.1145/1141753.1141767.
- [16] David Goodman. The measurement of research impact. *Journal of the Medical Library Association*, 96(1):1–2, 2008.
- [17] Henk F. Moed. *Citation Analysis in Research Evaluation*. Springer, 2006. doi: 10.1007/1-4020-3714-7.
- [18] 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. doi: 10.1088/1742-5468/2007/06/P06010.
- [19] Nees Jan van Eck and Ludo Waltman. Vos: A new method for visualizing similarities between objects. *Advances in Data Analysis and Classification*, 1(4):299–306, 2007. doi: 10.1007/s11634-007-0019-7.
- [20] Kushal Bhatt. Network-driven foreign direct investment dynamics: A new approach using exponential random graph models. doi: 10.1007/s41870-025-02826-0. URL <https://www.doi.org/10.1007/s41870-025-02826-0>.
- [21] Elvis Mondal. Intelligent virtual modeling for dynamic optimization in large-scale systems. doi: 10.67231/57mjv812. URL <https://doi.org/10.67231/57mjv812>.
- [22] Association for Computing Machinery. The 2012 acm computing classification system. Technical report, Association for Computing Machinery, 2012. URL <https://www.acm.org/publications/class-2012>.
- [23] Varnit Dhanani. Resolving conflicts and overlaps in construction standards: A data-driven classification approach. doi: 10.21203/rs.3.rs-8726462/v1. URL <https://doi.org/10.21203/rs.3.rs-8726462/v1>.

- [24] Marx Prathap Singh Shanmugam. Ai-driven data cleansing and representation learning for crm optimization: A webly-supervised approach. In *ICONAT*. doi: 10.1109/ICONAT66879.2025.11362638. URL <https://doi.org/10.1109/ICONAT66879.2025.11362638>.
- [25] Krishna Sai Penumarthi. Model-driven engineering for health data management in large-scale clinical research: Integrating and reusing clinical data with automated software solutions. In *IEEE*. doi: 10.1109/11232132. URL <https://ieeexplore.ieee.org/abstract/document/11232132>.
- [26] Guruprasath Sankaran. Generative engine optimization: A three-layer semantic framework for content visibility in ai-powered search. doi: 10.67231/jr22cv71. URL <https://doi.org/10.67231/jr22cv71>.