

Citation analysis of computer systems papers

Eitan Frachtenberg ^{Corresp. 1}

¹ Computer Science, Reed College, Portland, Oregon, United States

Corresponding Author: Eitan Frachtenberg
Email address: etc_26@yahoo.com

Citation analysis is used extensively in the bibliometrics literature to assess the impact of individual works, researchers, institutions, and even entire fields of study. In this paper, we analyze citations in one large and influential field within computer science, namely computer systems. Using citation data from a cross-sectional sample of 2,088 papers in 50 systems conferences from 2017, we examine four research areas of investigation: overall distribution of systems citations; their evolution over time; the differences between databases (Google Scholar and Scopus) for systems papers, and; the characteristics of self-citations in the field.

On citation distribution, we find that overall, systems papers were well cited, with the most cited subfields and conference areas within systems being security, databases, and computer architecture. Only 1.5% of papers remain uncited after five years, while 12.8% accrued at least 100 citations. For the second area, we find that most papers achieved their first citation within a year from publication, and the median citation count continued to grow at an almost linear rate over five years, with only a few papers peaking before that. We also find that early citations could be linked to papers with a freely available preprint, or may be primarily composed of self-citations. For the third area, it appears that the choice of citation database also makes little difference in relative citation comparisons, despite marked differences in absolute counts. On the fourth area, we find that the ratio of self-citations to total citations starts relatively high for most papers but appears to stabilize by 12--18 months, at which point highly cited papers revert to predominately external citations. Past self-citation count (taken from each paper's reference list) appears to bear little if any relationship with the future self-citation count of each paper.

The primary practical implication of these results is that the impact of systems papers, as measured in citations, tends to be high relative to comparable studies of other fields and that it takes at least five years to stabilize. A secondary implication is that at least for this field, Google Scholar appears to be a reliable source of citation data for relative comparisons.

Citation Analysis of Computer Systems Papers

Eitan Frachtenberg

Department of Computer Science, Reed College

Corresponding author:

Eitan Frachtenberg

Email address: eitan@reed.edu

ABSTRACT

Citation analysis is used extensively in the bibliometrics literature to assess the impact of individual works, researchers, institutions, and even entire fields of study. This paper analyzes citations in one large and influential field within computer science, namely computer systems. Using citation data from a cross-sectional sample of 2,088 papers in 50 systems conferences from 2017, the analysis investigates four research areas: overall distribution of systems citations; their evolution over time; the differences between databases (Google Scholar and Scopus) for systems papers, and; the characteristics of self-citations in the field.

On citation distribution, the data indicates that systems papers are generally well cited, with the most cited subfields and conference areas within systems being security, databases, and computer architecture. Only 1.5% of papers remain uncited after five years, while 12.8% of papers accrued at least 100 citations. For the second area, most papers achieved their first citation within a year from publication, and the median citation count continued to grow at an almost linear rate over five years, with only a few papers peaking before that. Additionally, early citations appear to be linked to papers with a freely available preprint, or may be primarily composed of self-citations. For the third area, the choice of citation database ostensibly makes little difference in relative citation comparisons, despite marked differences in absolute counts. On the fourth area, the data shows that the ratio of self-citations to total citations starts relatively high for most papers but appears to stabilize by 12–18 months, at which point highly cited papers revert to predominately external citations. Past self-citation count (taken from each paper's reference list) appears to bear little if any relationship with the future self-citation count of each paper.

The primary practical implication of these results is that the impact of systems papers, as measured in citations, tends to be high relative to comparable studies of other fields, and that it takes at least five years to stabilize. A secondary implication is that at least for computer systems, Google Scholar appears to be a reliable source of citation data for relative comparisons.

1 INTRODUCTION

Citation analysis plays a central role in bibliometric evaluation of journals, conferences, institutes, and individual researchers [Moed, 2006]. The advent of web-based citation databases has led to faster growth in their importance and use [Meho, 2007]. Nonetheless, citation analysis remains challenging when comparing citations across years [Varga, 2019], types of scholarly communication [Martins et al., 2010], or fields of study [Adam, 2002, Patience et al., 2017].

This study explores a multifaceted citation analysis that circumvents these challenges by focusing on papers from a single year, from conferences only, and from a single large field within computer science (CS), namely, computer systems (or just “systems” for short). Citation analysis can play an increasingly important role in judging and quantifying the importance of scientists and scientific research [Meho, 2007], as well as in policy making [Moed, 2006]. Concomitantly, the debate about citations' utility and interpretation appears to increase just as much [MacRoberts and MacRoberts, 1989, 2018]. In computer science and computer systems, citation analysis often focuses on conference proceedings in particular, since they are seen as more timely, more cutting-edge, and more strictly refereed than some journals [Goodrum et al., 2001].

The bibliometrics literature is rich with studies analyzing citations in various disciplines and fields,

including CS as a whole [Devarakonda et al., 2020, Hirst and Talent, 1977, Mattauch et al., 2020]. Even within CS, several fields, subfields, and specific venues have received bibliometric analyses [Broch, 2001, Frachtenberg, 2022b, Iqbal et al., 2019a,b, Lister and Box, 2008, Rahm and Thor, 2005, Wang et al., 2016]. The purpose of this article is to apply these methods to the CS field of computer systems and to understand the distribution and characteristics of its citations. To the best of the author's knowledge, this study is the first systematic analysis of a wide cross-section of computer systems research.

Systems is a large research field with numerous applications, used by some of the largest technology companies in the world. For the purpose of this study, we can define systems as the study and engineering of concrete computing systems, which includes research topics such as operating systems, computer architectures, data storage and management, compilers, parallel and distributed computing, and computer networks. The goal of this study is to characterize, for the first time, the citation behavior of this large and important CS field, while controlling for year, venue type, and research area. Specifically, this paper examines: total citations after five years and how they compare across subfields and against other disciplines and fields; the dynamics of citation counts over time; the effect of the source of the citation database; and the characteristics of self-citations in computer systems.

This study uses an observational, cross-sectional approach, analyzing 2,088 papers from a large subset of leading systems conferences. The study population came from a hand-curated collection of 50 peer-reviewed systems conferences from a single publication year (2017). Among other characteristics, the dataset includes paper citation counts at regular intervals for at least five years from the publication date, by which time many citation statistics stabilize [Larivière et al., 2009]. More details on the collection methodology and selection criteria can be found in previous work [Frachtenberg and Kaner, 2022]. Using this dataset, this study addresses the following high-level research questions:

RQ1: What is the distribution of paper citations in systems conferences after five years? Citations in this dataset exhibit the typical skewed distributions with most papers accruing a few citations, a handful of papers racking up thousands of citations, and a dearth of uncited papers. According to one comparison, citation counts of top-cited papers put the field of systems among the top ten scientific fields. However, citations are distributed unevenly across and within conferences and subfields, with security, databases, and computer architecture conferences ranking the highest median citations.

RQ2: How have citations evolved over this period? In this dataset, the median number of citations per paper grows at a nearly constant rate, with few papers peaking before the five-year mark, and few papers achieving “runaway” citation growth. Most papers are first cited within 9–12 months since publications, and the average time to have to wait to reach at least n citations grows almost linearly with n .

RQ3: How do Google Scholar and Scopus compare for these conferences? The data suggests, as have other studies before, that the Google's inclusive policy of counting many document types as potential citation sources does inflate the absolute citation counts compared to Scopus' counts. However, the two counts are nearly perfectly correlated and provide similar relative comparisons across papers, conferences, and years.

RQ4: How many citations are self-citations, and how do they evolve over time? The main finding emerging from the dataset is that self-citations are more prevalent among early citations and among papers that are less cited overall after five years. However, the average ratio of self-citations to total citations appears to stabilize for most papers about 12–18 months after publication. This dataset shows no clear relationship between a paper's eventual self-citations and the number of self-citations it contains in its own reference list.

In addition to answering these research questions, this study makes the following contributions:

- A critical view of H-index as a metric for conferences or journals.
- A discussion of the relationship between acceptance rates and citations.
- A characterization of papers that are cited relatively early.

As a final contribution, this study provides the dataset of papers and citations over time, tagged with rich metadata from multiple sources [Frachtenberg, 2021]. Since comprehensive data on papers and conferences with citations are not always readily available, owing to the significant manual data collection involved, this dataset can serve as the basis of additional studies [Saier et al., 2023].

The remainder of this paper is organized as follows. The next section (Section 2) describes the data collection and processing methodology in detail. The results section (Section 3) enumerates the findings,

101 organized by research question. Section 4 combines results on raw citation statistics to explore three
102 higher-level topics. Related work is surveyed in Section 5, and Section 6 summarizes the results and
103 suggests directions for future research.

104 2 MATERIALS AND METHODS

105 The most time-consuming aspect of this study was the collection and cleaning of data. This section
106 describes the data selection and cleaning process for conference, paper, and citation data.

107 Data were collected as previously described in [Frachtenberg, 2022a]. Specifically, the dataset
108 comes from a hand-curated collection of 50 peer-reviewed systems conferences from a single publication
109 year (2017) to reduce time-related variance. Conference papers were preferred over journal articles
110 because in CS, and in particular, in its more applied fields, such as systems, original scientific results are
111 typically first published in peer-reviewed conferences [Patterson et al., 1999, Patterson, 2004], and then
112 possibly in archival journals, sometimes years later [Vrettas and Sanderson, 2015]. These conferences
113 (detailed in Appendix A) were selected to represent a large cross-section of the field with different sizes,
114 competitiveness, and subfields. Such choices are necessarily subjective, based on the author’s experience
115 in the field. But they are aspirationally both spread enough to represent the field well and focused enough
116 to distinguish it from the rest of CS. For each conference, various statistics were gathered from its web
117 page, proceedings, or directly from its program chairs.

118 Since the main interest of this paper is in measuring the posthoc impact of each paper, as approximated
119 by its number of citations, citation data was regularly collected from two databases, Google Scholar (GS)
120 and Scopus. GS is an extensive database with excellent coverage of CS conferences that contains not
121 only peer-reviewed papers, but also preprints, patents, technical reports, and other sources of unverified
122 quality [Halevi et al., 2017]. Consequently, its citation counts tend to be higher than those of databases
123 such as Scopus and Web of Science, but not necessarily inferior when used in paper-to-paper comparisons
124 [Harzing and Alakangas, 2016, Martin-Martin et al., 2018]. Since we are mostly comparing relative
125 citation metrics, even if the GS metrics appear inflated compared to other databases, we should still be
126 able to examine the relationship between relative citation counts of papers and conferences. Nevertheless,
127 for papers covered in the Scopus database, both sources of citations are compared in RQ3 to ensure that
128 both metrics are in relative agreement with each other.

129 To collect citation statistics from GS, each paper’s citation count was collected from on
130 scholar.google.com at each monthly anniversary during the first year since publication. During the second
131 year, the statistics for each paper were collected every 3 months, and afterwards, every six months.
132 Scopus, on the other hand, offers retroactive citation statistics, so its citation and self-citation data were
133 collected for the end of each calendar year, as well as in each paper’s five-year anniversary.

134 Statistics

135 For hypothesis testing, group means were compared pairwise using Welch’s two-sample t-test and
136 group medians using the Wilcoxon signed-rank test; differences between distributions of two categorical
137 variables were tested with the χ^2 test; and correlations between two numerical variables were evaluated
138 using Pearson’s product-moment correlation coefficient. All statistical tests are reported with their p-
139 values. All computations were performed using the R programming language and can be found in the
140 source code accompanying this paper.

141 Ethics statement

142 All data for this study were collected from public online sources and therefore did not require the informed
143 consent of the authors. No funding was received to assist with the preparation of this manuscript.

144 Code and data availability

145 The complete dataset and metadata are available online [Frachtenberg, 2021]. For ease of reproducibility,
146 a Docker image with the source code and dataset is also available at <https://hub.docker.com/r/eitanf/sysconf> with the tag ‘citations’.

148 Limitations

149 To control for the effect of time on citations, additional data from more recent conference years in the
150 dataset was excluded. Undoubtedly, more data could strengthen the statistical validity of the observations;




Figure 1. citation distribution after five years of all cited papers. Also shown are the number of samples (papers), mean, and median citations per paper.

but it could also weaken any conclusions based on the inherent delays in the citation process and in variation over time. The methodology is also constrained by the manual collection of data, such as conference statistics; paper downloading and text conversion; cleanup and verification; etc. The effort involved in compiling all necessary data limits the scalability of this paper's approach to additional conferences or years.

The focus on citations as a primary metric of interest has also received significant criticism in the bibliometrics literature because of its own limitations. First, there are limitations stemming from the sources of citation data, such as inflated metrics in GS and partial coverage in Scopus, as previously mentioned. Second, even accurate citation counts present a variety of limitations, such as references expressing negative sentiments [Parthasarathy and Tomar, 2014] or researchers and venues gaming their own citation counts [Biagioli, 2016]. Nevertheless, citations are arguably the most popular method to evaluate research impact, either as raw counts or as inputs to compound metrics, such as the H-index. As such, this initial analysis of citation behavior in a previously unexplored and important field could further our understanding of both the particular field and bibliometrics as a whole.

3 RESULTS

3.1 RQ1: Citation distribution

Let us begin by describing the distribution of citations in the field. In addition to observing summary statistics, such as means, medians, and outliers, focusing our attention on distributions can increase transparency when comparing citation metrics because of their typical skewness [Larivière et al., 2016]. To better understand the distribution of 5-year citations, starting top-down by looking at the overall distribution and characteristics of all cited papers and then zoom in on distributions by conference and by paper topics. The missing piece is then filled by describing the distribution of uncited papers.

3.1.1 Overall distribution

Figure 1 shows the overall distribution of cited papers exactly five years from their publication date. It is indeed skewed and long-tailed (note the logarithmic scale) and resembles a log-normal distribution [Rahm and Thor, 2005, Redner, 1998, Wang and Barabási, 2021, Wu et al., 2009]. Consequently, the mean number of citations, 55, is much higher than the median, 23, and even higher than the 75th percentile, 53. The mean is pushed this high by a handful of outlier papers: there are just 7 papers with over 1,000 citations each, one as high as 5229! Similarly, we find as many as 267 papers (or 12.8% of the dataset) with at least 100 citations. This large group of papers would probably no longer be considered an outlier,

and 100 citations are likely a reliable signal of notable impact, again suggesting that the field as a whole may be quite influential.

If we follow Patience's method [Patience et al., 2017] to attenuate the outliers and compute the mean number of five-year citations among the top-cited 31–500 papers, we get an average of 120. That same study also compared this statistic across 236 science fields. So, using their data for comparison would rank systems as one of the top ten fields in citation count (but keep in mind that the original comparison used less-inflated citations counts from the Web of Science database.)

3.1.2 Citations by conference

Although overall citation counts for computer systems appear high, keep in mind that not all conferences within the field are equally well cited. Figure 2 shows the same kind of histogram, but this time broken down by conference and sorted by median citation per conference. We can observe the following relationships from the data:

- Citations vary widely between conferences, ranging from an average of about 4 for HCW papers to about 152 for SIGCOMM papers.
- Even within conferences, some outlier papers push the mean far from the median. For example, SP and ISCA exhibit long tails, with a few papers having thousands of citations, whereas NDSS, an equally well-cited conference, exhibits nearly identical mean and median.
- The likelihood of a conference containing an outlier paper with at least 500 citations increases with conference size. Papers with over 500 citations were published in conferences averaging 85.7 accepted papers, vs. an average of 58.6 for papers with fewer than 500 citations ($t = 2.47$, $p = 0.02$). Furthermore, a conference's size is weakly but positively correlated with its average paper citations ($r = 0.32$, $p = 0.03$), suggesting a random variable in the citation count of a paper.
- Median citations, on the other hand, do not appear to be significantly correlated with a conference's number of accepted papers ($r = 0.24$, $p = 0.09$). However, conferences with higher median citations are generally more competitive (i.e., lower acceptance rates). These two factors exhibit a strong negative correlation ($r = -0.64$, $p < 10^{-5}$). Likewise, mean citations per conference is also negatively correlated with the conference's acceptance rate ($r = -0.59$, $p < 10^{-5}$).
- Different conferences have different skews: the distribution can be wide or narrow; some have long right tails and some have none; and the modes appear at different locations, not always near the median. In other words, despite all conferences belonging to the same overall systems field, citation distributions still vary significantly with other factors including size, acceptance rate, and specific subfield.
- Related, some conferences sharing subfields of systems appear to be also clustered together in citation metrics. For example, NDSS, CCS, and SP—all focused on computer security—are similarly ranked in terms of median citations (although the means vary much more). Likewise, NSDI/SIGCOMM (focusing on networking), HPCA/MICRO and ASPLOS/ISCA (focusing on computer architecture) also pair with similar median citations.

This last observation naturally leads to the question of the relationship between a research subfield and its typical citations. To answer this question, let us next look at research topics at the individual paper level rather than at the conference's broad scope.

3.1.3 citations by Subfield

The definition of subfields in computer systems, like the definition of the field itself, is necessarily fluid and subjective. Based on research experience in the field, a best-effort attempt was made to categorize all papers by reading every single abstract and tagging each paper with one or more subfields. The list of selected tags is presented in Table 1. Naturally, experts would differ in their opinions on this list and possibly on tag assignments to papers, but the current assignment provides a starting point for comparison across systems subfields.

Papers in this dataset are generally focused on one or two of these topics, averaging 1.7 topics per paper. There are a total of 2,026 (97% of all papers) cited papers with at least one topic assigned. The remainder either had no citations or did not fit well with any of the tags. Figure 3 shows the distribution of citations by topic in these tagged papers (with multitopic papers repeated).

Again, there appears to be a positive relationship between quantity—the number of papers published on a topic—and “quality”—how well cited these papers generally are ($r = 0.57$, $p = 0.03$). This could




Figure 2. Citation distribution of all cited papers by conference. Also shown are the number of samples (papers), mean, and median citations per paper. Conferences ordered by median and showing acceptance rate in parenthesis next to their names.

Table 1. List of selects systems subfields

Tag	Subfield description
Architecture	Computer architecture
Benchmark	Workloads, bechnmarking, metrics, and methodology
Cloud	Cloud computing and related infrastructure
Compilers	Compilers, language and debugging sopport, runtime systems
Concurrency	Parallel and distributed computing
Data	Big data applications and infrastructure
DB	Databases, key-value stores, and database management systems
Energy	Power and energy efficiency, sustainable computing
GPGPU	Accelerator technologies and heterogeneous computing
HPC	High performance computing and supercomputing applications
Network	Networking algorithms, wireless networks, switching and routing
OS	Operating systems, scheduling, resource management
Security	Security, privacy, encryption, resilience
Storage	Persistent and ehpemeral storage: disk, flash, RAM, etc.
VM	Containers and virtual machines and, networks




Figure 3. Citation distribution for all cited papers by paper's topic. Papers with multiple topics appear in multiple histograms. Also shown are the number of samples (papers), mean, and median citations per paper. Topics ordered by median citations per paper.

Table 2. Uncited papers by conference, as paper count and percentage of all accepted papers.

Conference	Count	Percent
HPCC	6	7.79%
IPDPS	4	3.45%
EuroPar	3	6%
Cluster	2	3.08%
HiPC	2	4.88%
ICPE	2	6.9%
ISPASS	2	8.33%
CCGrid	1	1.39%
CLOUD	1	3.45%
HCW	1	14.29%
HotStorage	1	4.76%
HPCA	1	2%
HPDC	1	5.26%
ICPP	1	1.67%
IGSC	1	4.35%
IISWC	1	3.23%
PODC	1	2.63%
SYSTOR	1	6.25%

again indicate a random element leading to a higher probability of outlier papers in large subfields, or it could indicate that these topics were both sufficiently popular at the time to attract multiple submissions and popular over time to attract eventual citations.

Nevertheless, this correlation should also be considered with a grain of salt because of the subjective process of conference selection for this study. For example, if more conferences and workshops on virtualization had been chosen, the number of papers on the topic would have obviously increased, whereas the relatively large conferences on security skew topic counts from the other end. At any rate, in 2017, the broad research topics that resulted in the highest median citations appear to be security, databases and data management, computer architecture, and networks.

3.1.4 Uncited papers

Owing to the logarithmic citation scale, the preceding histograms omitted papers with zero citations. Some early studies claimed that, generally, most scientific papers are not cited at all [Hamilton, 1991, Jacques and Sebire, 2010, Meho, 2007]. More recent research found that the rate of uncited papers keeps decreasing, and estimates it to be less than 10% [Wu et al., 2009]. For example, one study computed the percentage of uncited papers in physics (11%), chemistry (8%), and biomedical sciences (4%) [Larivière et al., 2009]. Another large-scale estimate for the entire CS discipline found that 44.8% of CS papers remained uncited after five years [Chakraborty et al., 2015]. In this dataset of systems papers, only 32 papers remain uncited after five years (1.5%). Of these, the conference with the most uncited papers was HPCC, followed by IPDPS and EuroPar (Table 2). Both in absolute terms and in percentage terms, the number of uncited papers remains very low for most conferences.

The topic tags with the most uncited papers were Benchmark, Storage, and HPC, followed by GPGPU and Concurrency (Table 3). It is not surprising that the two distributions appear to be related. Many papers in the top six uncited conferences were tagged with some of the top six uncited topics.

Having examined the citation distribution at a fixed point in time, we can now examine how it evolved to this point from the date of publication.

3.2 RQ2: Citation dynamics

Observing the total citations of papers at a fixed time point offers only a static view of a metric that is inherently a moving target. Citations tend to follow different dynamics as different papers, disciplines, and fields exhibit very different aging curves [Pichappan and Ponnudurai, 1999, Wang, 2013]. After five years, all the papers in this dataset likely had a chance to be discovered by fellow researchers, as

Table 3. Uncited papers by topic tags

Topic	Count
Benchmark	8
HPC	7
OS	7
Storage	7
Concurrency	6
GPGPU	6
Data	5
Network	4
Energy	2
Architecture	1
Cloud	1
Compilers	1

evidenced by the fact that nearly all are cited by outside researchers. We can therefore concentrate next on citation dynamics.

For the second research question, three aspects of the time dimension are examined: general distribution over time, time to first citation, and citation velocity. Since different conferences publish in different months, causing nearly a year's gap between the first and the last, time is normalized for all three aspects by counting the number of months passed from the date it was published, rather than a fixed start day in 2017.

3.2.1 Citation distribution

Looking at the citation distribution at six-month intervals (Figure 4), three observations can be made. First, we can see that the spread—difference across papers—grows over time (note the logarithmic scale), as some papers accelerate at a faster rate than others, creating a larger range of citation values.

Most citations (as seen in the first, second, and third quartiles) grow at a rapid clip over time without showing signs of slowing down (note the logarithmic scale). In fact, if we fit the median with a simple linear regression model, shown as a gray line, we match the median citation values after the first year nearly perfectly (intercept: -2.33, slope: 0.43 citations per month). Linear models for the 25th and 75th percentiles would have different slopes, because of the increasing spreads, but still fit quite well.

In contrast, mean citation outpace medians' growth because of the disproportionate pull of outlier papers. Except for the first six months (when the difference between the mean and the median is less than half a citation), the presence of highly cited papers pushes the mean increasingly higher than the median. In the first six months, the presence of many (1,191) uncited papers is pulling the mean to ≈ 1.6 citations. However, as papers gain enough time to be discovered, cited, peer-reviewed, and then published, the mean quickly catches up. By 12 months, the mean number of uncited papers drops to 503, and by 18 months to 227.

This observation naturally leads us to the next aspect of time: how long does it take papers to be cited?

3.2.2 Time to first citation

The dataset includes citation data at monthly intervals during the first year from publication, then quarterly for the second year, and then every six months for the remaining three years. This resolution allows us to estimate with near-monthly accuracy the first time that GS detected a citation for each paper. Figure 5 shows the distribution of the approximate time in months it took GS to first detect any citations for each paper (compared to Figure 1, two papers briefly showed citations before reverting back to zero).

Assuming that the first citation is external, i.e., not a self-citation, we can think of this moment of first being cited as a paper's "discovery event." In this dataset, it averages about 9 months. (As we see later in RQ4, this assumption is invalid for many papers, which may postpone the discovery time by a few months.) Discovery time can be longer than subsequent citations because it requires the paper to be published and discovered by another researcher, who must then wait till their citing document is peer-reviewed and published, a process that can take several months.




Figure 4. Citation distribution for all cited papers over time. Center bars show median value; lower and upper hinges correspond to the 25th and 75th percentile; upper points show high-citation outliers; and diamonds show means for all citations, including uncited papers. Also shown in gray is a linear regression line for the median values.




Figure 5. Distribution of time to first citation for all cited papers. Also shown are the number of samples (papers), mean, and median times in months.




Figure 6. Average time to reach n citations for increasing values of n . Only papers with at least 10 citations after five years are included.

To illustrate this point, Figure 6 shows the mean time it takes for papers to reach n citations. Note that the average growth from the first citation to the next and then the next takes about the same time across the range (about 2–3 months for each additional citation, slowly decreasing). However, going from zero citations to the first citation takes on average about twice as long. This observation leads us to the next and final question of time, namely, how fast and how long do citations grow?

3.2.3 Citation velocity

As we have just observed, the mean time across papers to receive the first citation (6.24 months) is longer than the time to add the second citation (2.91 months). It continues to decrease slowly such that it only takes 1.86 months on average to add the 10th citation. Intuitively, this makes sense as papers are discovered by an increasingly larger network of researchers and pick up more and more citations during their growth phase. However, this phase cannot grow indefinitely, as the size of the potential network is bounded by the systems research community size, and the impact of systems papers is often limited to a few years until newer systems replace them. To visualize these growth patterns, we can plot the citation velocity of every paper (Figure 7).

The first observation we can make from the figure is that the citation growth for most papers is nominal, as evidenced by the thick near-flat bands of most papers at every conference. The growth rate hovers close to zero for conferences with low median citations and remains mostly below five citations per month even for better-cited conferences. The second observation is that a few papers do achieve runaway growth, sometimes exceeding the plot limit of 15 citations per month—even reaching 147 in one instance—but their overall number is small. Only 26 papers exceed ten citations per month at the end of the five years.

This brings us to the last observation and provides an answer to the first question posed. There were actually slightly more papers exceeding ten citations per month at the four-year mark (30 papers), meaning that a few of those peaked before reaching five years. Looking at the graphs (ignoring some sharp peaks and dips caused by noise from GS), we can observe the occasional paper that reached maximum velocity and then started to slow down. Again, the overall number of these appears small, suggesting that most papers in systems take longer than five years to peak.

3.3 RQ3: Comparison of GS and Scopus citations

The citation criteria employed by GS tend to be more inclusive than those of other databases, sometimes resulting in inflated citation counts [Harzing and Alakangas, 2016, Martin-Martin et al., 2018]. Moreover, different fields tend to be covered to different degrees by different databases, and not all databases cover conferences equally well (the systems conferences in this dataset are covered 100% by GS). Because of

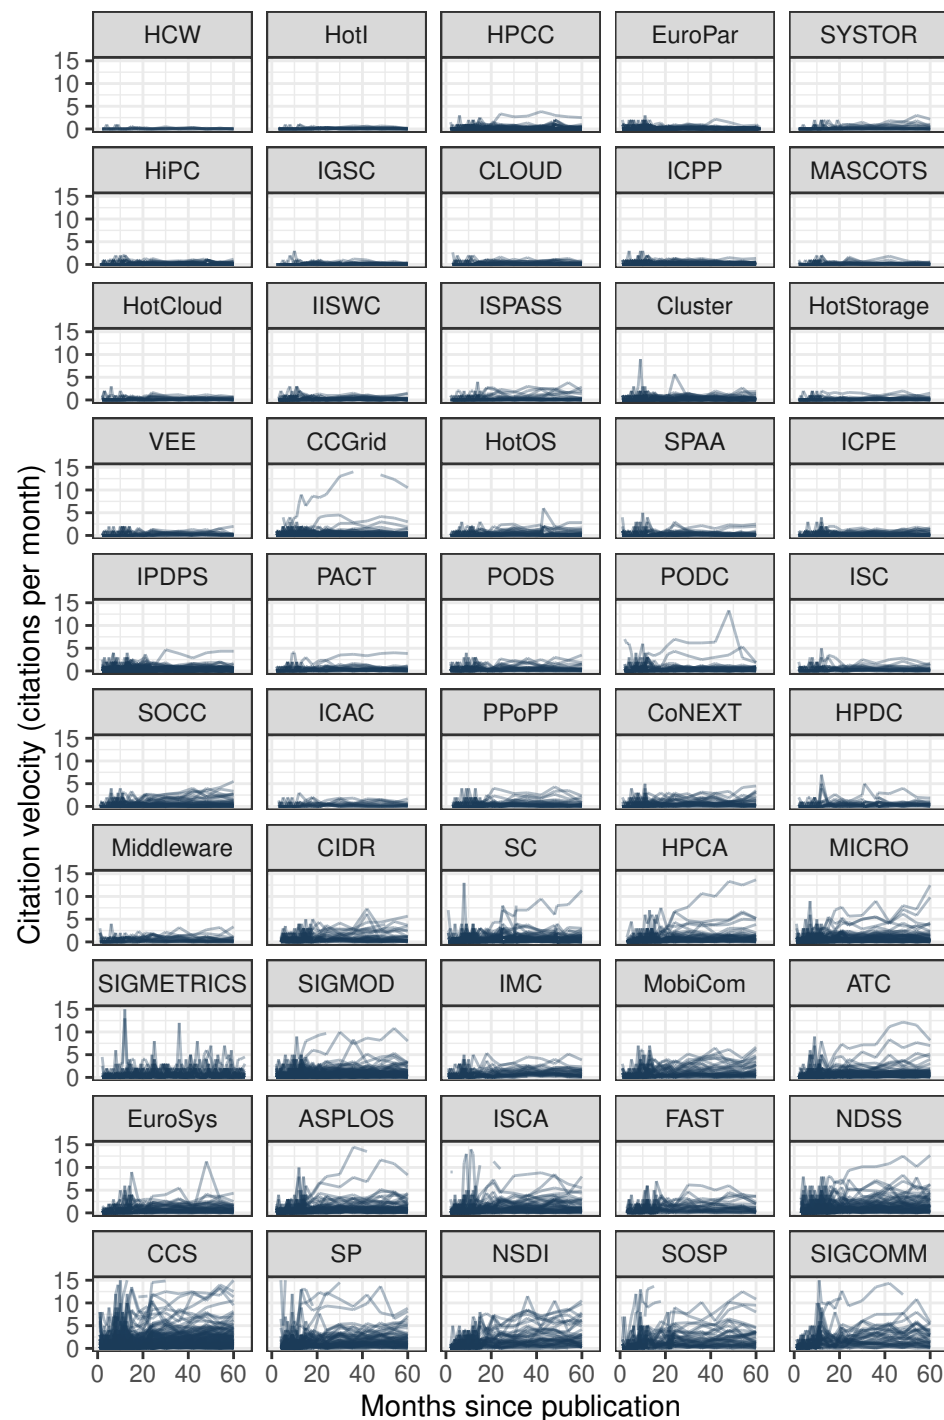


Figure 7. Citation velocity of every paper. Y-axis is clipped to show the bulk of the data with higher resolution, but some papers exceed 100 citations per month. Conferences ordered again by increasing median citations after five years.




Figure 8. Comparison of GS and Scopus citations after five years. Orange line denotes a 1:1 mapping.

this criticism of GS and as a way to add a measure of control to the findings, we can examine citation data from another database, Springer's Scopus.

Two types of statistics were collected from Scopus: total citations exactly five years from each paper's publication and total citations at the end of each year (with and without self-citations). We will first focus on the former statistic, which is directly comparable to the one collected from GS, and explore the latter in the next research question.

Note that Scopus coverage for these systems conferences is incomplete, missing a total of 123 papers relative to GS. For the remaining 1965 papers, Figure 8 shows a scatter plot of every paper's citation count in Scopus (x-axis) and GS (y-axis). The vast majority of points lie above the 1:1 intercept line, confirming that GS citation counts tend to be higher than those of Scopus. But the difference in magnitudes appears to be remarkably constant, with near-perfect correlation ($r = 0.99$, $p < 10^{-9}$, $R^2 = 0.98$). The implication here is that all the relative observations drawn so far on GS citations should generalize to Scopus citations as well, up to a constant factor. Other studies have found that Google Scholar citations are strongly correlated with those from Web of Science as well [Kousha and Thelwall, 2007].

3.4 RQ4: Effect of self-citations

Self-citations are fairly common in the sciences and have been estimated to comprise 10–40% of all scientific production, depending on the field [Aksnes, 2003, Snyder and Bonzi, 1998, Wolfgang et al., 2004]. On the one hand, self-citations represent a natural evolution of a research team's work, building upon their previous results, especially in systems projects that often involve incremental efforts of implementation, measurement, and analysis [Wolfgang et al., 2004]. On the other hand, self-citations can be problematic as a bibliometric measure of a work's impact because they obscure the external reception of the work and are prone to manipulation [Waltman, 2016].

For the final research question, let us quantify the degree to which self-citations affect the overall citation metrics in computer systems and how well self-citations can be predicted from the papers' references themselves. To this end, we will examine how self-citations evolve over time, how they relate to self-citations in the original papers, and how common they are overall.

3.4.1 Ratio of self-citations over time

We can start by computing the ratio of self-citations for every paper at every measurement time point (the end of each calendar year between 2017 and 2022, which represents different ages in months for each paper depending on the month they were published in). We can then average these ratios across all papers with the same measurement age, and divide those into four quartiles based on the total citations they had at the time. The results of this computation are shown in Figure 9. A few notable observations emerged from the data:




Figure 9. Mean ratio of self-citations out of total citations as a function of months since publication. Trend lines are smoothed using LOESS. Data are divided into four quartiles based on how many total citations the paper had after 5 years (Q4 means most cited papers).

- In the first few months after publication, most of the detected citations appear to be self-citations for all groups. A distinct transition to majority external citations starts about a year after publication.
- Highly cited papers (e.g., Q3 and Q4) are also cited earlier.
- The more a paper is cited overall, the lower the ratio of self-citations. Conversely, for Q1 and Q2, nearly all early citations were self-citations.
- That said, even in the fourth quartile, self-citations still comprise some 10% of all five-year citations on average.

Overall, it appears that self-citations are more characteristic of early citations, either because a paper has not had enough time to be well-known by almost anyone other than its authors, or as an attempt by the authors to increase its visibility [Chakraborty et al., 2015]. But as time passes and papers are evaluated by the community, highly cited papers accrue citations primarily from external researchers.

3.4.2 Backward and forward self-citations

These self-citation distributions inspire another curious question: Is there a pattern of specific self-citation behavior that affects both forward and backward citations? In other words, can we predict the number of self-citations a paper would receive after five years by counting self-citations in its own bibliography?

To answer this question, the total number of backward self-citations was estimated, as follows: First, the reference list from each PDF file was converted to a text file for every single paper and manually cleaned up conversion artifacts such as two-column papers breaking down the reference list. Then, the last names of all authors were searched for in every single reference. Any reference that matched one or more names was counted as a self-citation. Clearly, this method is imperfect because some last names may match different authors, and some last names may even represent a word in a paper's title. Nevertheless, manual inspection of several dozen papers has revealed very few inaccuracies.

The self-reference counts were converted to ratios by dividing by the total number of references in each paper and then compared to eventual (5-year) self-citations. Figure 10 shows the relationship between these two measures. We can make out several horizontal bands for low-cited papers with fixed ratios of eventual citations (0%, 50%, 100%, etc.), and one clear vertical band for papers with zero backward self-references. But overall, it shows little discernible patterns.

Likewise, the Pearson correlation between backward and forward self-citation ratios is nominal ($r = 0.16$, $p < 10^{-9}$). Removing some of the outliers (the top and bottom horizontal bands, as well as the leftmost band) does little to increase the correlation ($r = 0.19$, $p < 10^{-9}$). Even if we look only at the temporal subset of data when papers first exceed five citations (which typically comprise of a higher ratio




Figure 10. Relationship between five-year self-citations and self-citation count in each paper's own references. Orange line denotes a 1:1 mapping.

of self-citations), the observed correlation is not particularly strong ($r = 0.24$, $p < 10^{-9}$). It appears that the answer to this research question is mostly negative.

3.4.3 Overall self-citation ratios

One anecdotal observation we can make from Figure 10 is that most papers appear above the orange 1:1 line. In other words, a majority of papers (1,136 or 57.1% to be exact) had a higher self-citation ratio after five years than their backward self-citation ratio in their reference list.

Moreover, the mean self-citation ratio after five years, 24.1%, is on the one hand higher than the 15% rate found for physical sciences (and 9% over all sciences) in an earlier study on self-citation patterns across disciplines [Snyder and Bonzi, 1998]; but on the other hand, is in close agreement with the 24% rate found for CS papers in Norway [Aksnes, 2003]. The Norway study also found a high ratio of self-citing papers overall, agreeing with our data where 78.4% of papers had at least one self-citation after five years.

In terms of outliers, hidden in these ratios and averages are several papers whose total citations were predominantly self-citations. In their final citation count, 252 papers had over 50% self-citations, of which 11 papers were in Q4 with dozens of citations each. Four of the six most self-cited papers in relative terms were in architecture conferences, suggesting that research in this area often builds primarily on past research from the same group. That said, the overall correlation between a paper's total citations and self-citations after five years is low ($r = 0.19$, $p < 10^{-9}$), suggesting that for this dataset as a whole, self-citations are not typically a dominant component of overall citations.

4 DISCUSSION

This section integrates results from the previous section to explore three topics beyond raw citation counts: the use of H-index as a measure of conference quality, the relationship between a conference's acceptance rate and its eventual citations, and the characteristics of papers with early citations.

4.1 Conference H-index

One derivative aspect of citations that we can calculate is a conference's H-index. H-index was developed by Hirsch to measure the impact of individual researchers [Hirsch, 2005]. It is defined as the maximum integer n such that there exist n papers published in a given time window that received at least n citations. GS also reports the H-index and five-year H-index of conferences and journals and ranks them accordingly [Patience et al., 2017]. With the dataset of papers from 2017, we can compute a similar measure based on papers from that single year, if nothing else, to illustrate the weakness of this metric when applied to conferences.

Our partial H5-index metric ranges from 4 for HCW to 69 for CCS. It correlates strongly with the size of the conference ($r = 0.75$, $p < 10^{-9}$), suggesting that a conference's mere number of accepted papers, largely a policy decision, has a significant impact on its H-index metric. Moreover, recall the observation that the larger a conference, the more likely it is to have a "runaway" successful paper. This suggests that among all the factors that impact a conference's total citations (and consequently its H-index), there is also a non-negligible element of chance. In other words, a conference's steering committee could conceivably increase its H-index (and average citations) by simply increasing the size of the conference. Obviously, size is not the most important factor, and an overly permissive acceptance policy is likely to admit papers that would lower the overall citation average. As demonstrated in our dataset, the partial H5-index metric is clearly negatively correlated with acceptance rate ($r = -0.62$, $p < 10^{-5}$). Nevertheless, this finding weakens the case for the use of these metrics as reliable indicators of conference quality.

4.2 Conference acceptance rate

The strong correlation between a conference's acceptance rate and its median citation counts could suggest a causal relationship, that is, that competitive peer review selects for high-impact papers. However, the peer-review process is notoriously unreliable when it comes to selecting highly cited papers, which may not always even be the reviewers' goal [Coupé, 2013, Lee, 2019, Wainer et al., 2015]. Moreover, even if a causal relationship does exist, its direction is unclear. That is to say, researchers who predict that their paper will have a relatively low citation impact may self-select to submit it to a less competitive conference to increase its chances of acceptance. Since we regrettably cannot design a randomized controlled trial where some papers are randomly accepted in a given conference, and since we cannot accurately predict the citation impact of rejected papers, we do not have the necessary tools to evaluate the strength and direction of any such causal link. That said, even if the peer-review process has little predictive or selective power for high-impact papers, we could surmise that successfully publishing in competitive conferences increases the likelihood of eventual citations because intuitively, prestigious conferences increase the post-publication visibility and credibility of their papers.

4.3 Early cited papers

Although we can expect several months to pass before a paper is first discovered and cited by external scientists, some "early" papers were cited within six months of publication. Our dataset provides evidence for two possible explanations.

First, many of these citations could be self-citations, which do not require discovery. As we have seen for RQ4, more early citations are indeed self-citations. For example, the 255 "early" papers that were cited within six months of publication (Scopus data) averaged 55.61% self-citations. In contradistinction, in the 463 papers first cited in the following six-month period (months 7–12), this proportion was significantly lower at 41.98% ($t = 3.79$, $p < 10^{-3}$).

Second, the availability of preprints or other freely accessible versions of the paper *before* publication could accelerate the discovery process. The dataset includes the time (in months) it took GS to discover a freely available e-print version of each paper. For the "early" papers, this time averaged 3.76 months, compared to 5.55 months for the slow papers ($t = -3.19$, $p < 0.01$). For a frame of reference, this average was close to the overall average time to e-print across all papers (5.47 months). From this perspective, therefore, the papers cited within the second half-year are indistinguishable from all later-to-cite papers, as opposed to the distinct "early" papers cited in the first half-year.

5 RELATED WORK

Citation analysis is an active area of research in bibliometrics, with many studies looking at it both quantitatively—examining citation distributions, as this study does—or qualitatively, examining the strengths, weaknesses, and characteristics of citations as a metric of scholarly impact. This study does not aim to debate the merits of citation-based metrics, so the qualitative aspects will not be reviewed here. The interested reader is referred to the recent book "The Science of Science" for an overview of this debate [Wang and Barabási, 2021].

On the quantitative side, several studies exist that examine some of the metrics discussed in this paper, as well as others. Tsay and Shu analyzed citations in one journal and found that the most cited documents were journal articles, followed by books and book chapters, electronic resources, and conference proceedings [Tsay and Shu, 2011]. As mentioned in the introduction, in computer science,

conference papers take precedence to journal articles [Goodrum et al., 2001, Vrettas and Sanderson, 2015].

Citation analyses collect data from different databases, primarily Google Scholar, Scopus, and Web of Science. Other studies have evaluated these databases for their validity and reliability in citation analysis. Several of those concluded that Google Scholar both has higher coverage and more liberal policies for defining what constitutes a citation [Halevi et al., 2017], resulting in higher citation counts overall, as was also found here [Kousha and Thelwall, 2007]. This difference can make it difficult to argue about absolute numbers of citations since they vary significantly by database. However, as most of these studies also found—and ours agrees—the citation counts from the three databases appear to be strongly correlated [Harzing and Alakangas, 2016, Martin-Martin et al., 2018]. The implication is that comparing citations across papers, conferences, and years should produce similar conclusions, regardless of which reputable database is used for citation data.

Numerous studies examined specific aspects of citation analysis, such as the type of citing document or type of cited document. For example, Harter measured the citation counts and impact factor of early electronic journals to compare them with traditional journals [Harter, 1996]. A related study looked at how often electronic resources are cited in electronic journals [Herring, 2002]. Various other studies have found a potential link between the open sharing of research artifacts and increased citations [Frachtenberg, 2022a, Heumüller et al., 2020].

There exist many citation analyses for various disciplines and venues, too numerous to enumerate here. To the best of the author's knowledge, this study is the first to analyze citation characteristics across the entire field of computer systems research, and the first to cover the four research areas in one analysis. A few particularly relevant examples of similar studies follow, all focused on computer science fields.

In 2001, Broch analyzed the citations of one prestigious CS conference, SIGIR, over the years 1997-1999. Unlike this study, it covered three years, allowing the discover of some trends, but it focused narrowly on one conference only ($n=110$) and a relatively short post-publication window. It did find, as this study confirms, that papers that are electronically available tend to collect more citations. In another single-conference study, Frachtenberg analyzed the citations of papers in the SIGMETRICS'17 conference and found similar citation dynamics to those of the larger systems field in this more generalized study [Frachtenberg, 2022b]. Similarly focused, but taking a reverse perspective, Lister and Box looked at the outgoing citations of the SIGCSE'07 conference and examined in detail the venues and venue types of those cited papers [Lister and Box, 2008].

Some specific subfields of computer systems also received bibliometric analysis that covered more than a single venue. Iqbal et al. looked at five decades of the SIGCOMM conferences [Iqbal et al., 2019a], and more broadly, at 18 years of four venues focused on computer networks [Iqbal et al., 2019b]. The primary investigation of these longitudinal studies was aimed at changes over time, a dimension this cross-sectional study completely omits by limiting the data to a single year. Rahm and Thor looked at ten years of five primary database venues and looked at most cited papers, venues, authors, institutions, and countries, as well as citation skew and impact factor. In another subfield of computer systems, namely cloud computing, Khan et al. recently investigated the most-cited papers and the venues they appeared in [Khan et al., 2022]. Previously, Wang et al. also tried to identify important papers and research themes in the subfield using tools from social network analysis [Wang et al., 2016].

Most of these comparable studies did not go to the same level of detail on the questions of self citations, database selection, and citation dynamics described in this study. Moreover, although these studies share with the present one their attention on computer science (and some even on computer systems), they are all narrowly focused on a single conference or subfield, and do not cover computer systems broadly. However, we can find several studies on the other end, broadening their view either to all of computer science, or even multiple scientific disciplines at once.

For example, Chakraborty et al. looked at the growth dynamics of citations of some 1.5 million CS papers and identified several distinct patterns [Chakraborty et al., 2015]. Devarakonda et al. looked at some 8 million publications from 20 years and classified them into topics based on their citation patterns. Freyne et al. compared CS journals to conferences and found that conference papers are generally better cited. They also found that GS citations correlate strongly with Web of Science citations (when available), agreeing with our finding when comparing GS to Scopus. Most recently, Saier et al. [Saier et al., 2023] presented a very large machine-readable dataset of publication metadata (including citation networks) for all arXiv publications in all fields. Although the present work focuses on the data rather than the analysis,

and the papers are all preprints rather than peer-reviewed papers as in the other studies, this dataset should prove very useful for large-scale comparisons and bibliometric analyses of entire fields of study.

6 CONCLUSION AND FUTURE WORK

This work examined, for the first time, the citation characteristics of an important field of computer science, namely computer systems. It found that compared to other fields and disciplines, systems papers are very well cited: the top-cited papers rank the field among the highest scientific disciplines in citation counts, and only a few papers remain uncited after five years. From a practical perspective, these data serve as a grounding point for future citation analyses, either within computer systems or in comparison to other fields.

The overall ratio of five-year self-citations (24.1%) appears to be higher than in other scientific fields but agrees with the 24% rate found for all computer science papers in Norway. Interestingly, most papers exhibit some self-citations, especially within the first few months since publication (when the free availability of e-print versions of a paper also increases its early citation count). Over time, the self-citation ratio remains above 30% for papers with relatively few citations, while for the most cited 25% of papers, this ratio drops below 10%. Practically speaking, this observation suggests that analyzing self-citation statistics should only be performed from a vantage point of several years after publication.

Competitive conferences with a low acceptance rate tend to higher citations per paper on average. On the other hand, conferences that accept many papers overall (independent of the acceptance ratio) also exhibit more “runaway success” papers that lift their average citation counts, suggesting the possibility of a random factor. It also suggests that the H5-index metric for conferences is particularly prone to manipulation, as increasing the number of accepted papers is correlated with higher H5-index values on average. Whether the citation count comes from Google Scholar or Scopus did not matter much, as both were strongly correlated (Pearson’s $r=0.99$), leading to similar findings for all relative comparisons across papers and conferences.

This work can be extended in several directions. Two obvious extensions along the time axis are to continue collecting citation metrics for the same conferences and to look at additional conference years in order to verify the generalization of these results. A more intricate proposal is to move from the “how many”-type questions on citations in this study to the “why”-type questions. The goal of such an investigation would be to identify which among dozens of factors are most closely associated with increased citation counts in the field of computer systems and hypothesize potential explanations. The author plans to follow up with a study that analyzes the myriad factors that could be associated with citation counts: conference-related, paper-related, and author-related.

At a deeper level, the existing data could also help answer questions on the relationships between the text of each paper and its eventual citations. For example, by using topic modeling on the titles, abstracts, or full text of the papers, we could attempt to identify which topics were particularly impactful in 2017, and perhaps compare these findings to more recent systems papers to see how these topics have evolved. Yet another avenue for exploration of this dataset is the analysis of the various networks the data represents: in coauthorship collaborations, in citations, and in affiliations. These networks can be analyzed using tools from social network analysis to quantify the centrality of various papers, authors, or even topics, tying it back to topic modeling.

APPENDIX A. DETAILED CONFERENCE LIST

Each conference is described by its initialism, full name, commencement date, size (number of published papers), acceptance rate (if known), and the homepage containing the program.

1. ASPLOS: ACM International Conference on Architectural Support for Programming Languages and Operating Systems, 2017-04-08. 56 papers, acceptance rate: 17.5%. Homepage: <http://novel.ict.ac.cn/ASPLOS2017/>
2. ATC: USENIX Annual Technical Conference, 2017-07-12. 60 papers, acceptance rate: 21.7%. Homepage: <https://www.usenix.org/conference/atc17>
3. CCGrid: IEEE/ACM CCGrid, 2017-05-14. 72 papers, acceptance rate: 25.2%. Homepage: <https://www.arco.inf.uc3m.es/wp/ccgrid2017/>

- 583 4. CCS: ACM Conference on Computer and Communications Security, 2017-10-31. 151 papers,
584 acceptance rate: 18.1%. Homepage: <https://www.sigmac.org/ccs/CCS2017/>
- 585 5. CIDR: The biennial Conference on Innovative Data Systems Research, 2017-01-08. 32 papers,
586 acceptance rate: 41%. Homepage: <http://cidrdb.org/cidr2017/>
- 587 6. CLOUD: IEEE International Conference on Cloud Computing, 2017-06-25. 29 papers, acceptance
588 rate: 26.4%. Homepage: <http://www.thecloudcomputing.org/2017/>
- 589 7. Cluster: IEEE Cluster Conference, 2017-09-05. 65 papers, acceptance rate: 30%. Homepage:
590 <https://cluster17.github.io/>
- 591 8. CoNEXT: ACM International Conference on Emerging Networking Experiments and Technolo-
592 gies, 2017-12-13. 32 papers, acceptance rate: 18.7%. Homepage: [http://conferences2.
593 sigcomm.org/co-next/2017/#!/home](http://conferences2.sigcomm.org/co-next/2017/#!/home)
- 594 9. EuroPar: International European Conference on Parallel and Distributed Computing, 2017-08-30.
595 50 papers, acceptance rate: 28.4%. Homepage: <http://europar2017.usc.es/>
- 596 10. EuroSys: The European Conference on Computer Systems, 2017-04-23. 41 papers, acceptance
597 rate: 21.8%. Homepage: <https://eurosys2017.github.io/>
- 598 11. FAST: USENIX Conference on File and Storage Technologies, 2017-02-27. 27 papers, acceptance
599 rate: 23.3%. Homepage: <https://www.usenix.org/conference/fast17/>
- 600 12. HCW: IEEE International Heterogeneity in Computing Workshop, 2017-05-29. 7 papers, accep-
601 tance rate: 46.7%. Homepage: <http://hcw.eecs.wsu.edu/>
- 602 13. HiPC: IEEE International Conference on High Performance Computing, Data, and Analytics,
603 2017-12-18. 41 papers, acceptance rate: 22.3%. Homepage: <http://hipc.org/>
- 604 14. HotCloud: USENIX Workshop in Hot Topics in Cloud Computing, 2017-07-10. 19 papers, accep-
605 tance rate: 32.8%. Homepage: <https://www.usenix.org/conference/hotcloud17>
- 606 15. HotI: IEEE Annual Symposium on High-Performance Interconnects, 2017-08-28. 13 papers,
607 acceptance rate: 33.3%. Homepage: <http://www.hoti.org/hoti25/archives/>
- 608 16. HotOS: ACM Workshop on Hot Topics in Operating Systems, 2017-05-07. 29 papers, acceptance
609 rate: 30.9%. Homepage: <https://www.sigops.org/hotos/hotos17/>
- 610 17. HotStorage: USENIX Workshop on Hot Topics in Storage and File Systems, 2017-07-10. 21
611 papers, acceptance rate: 36.2%. Homepage: [https://www.usenix.org/conference/
612 hotstorage17](https://www.usenix.org/conference/hotstorage17)
- 613 18. HPCA: The IEEE Symposium on High Performance Computer Architecture, 2017-02-04. 50
614 papers, acceptance rate: 22.3%. Homepage: <http://hpca2017.org>
- 615 19. HPCC: IEEE International Conference on High Performance Computing and Communications,
616 2017-12-18. 77 papers, acceptance rate: 43.8%. Homepage: [http://hpcl.seas.gwu.edu/
617 hpcc2017/](http://hpcl.seas.gwu.edu/hpcc2017/)
- 618 20. HPDC: ACM International Symposium on High Performance Parallel and Distributed Computing,
619 2017-06-28. 19 papers, acceptance rate: 19%. Homepage: <http://www.hpdc.org/2017/>
- 620 21. ICAC: IEEE International Conference on Autonomic Computing, 2017-07-18. 14 papers, accep-
621 tance rate: 19.2%. Homepage: <http://icac2017.ece.ohio-state.edu/>
- 622 22. ICPE: ACM/SPEC International Conference on Performance Engineering, 2017-04-22. 29 papers,
623 acceptance rate: 34.9%. Homepage: <https://icpe2017.spec.org/>
- 624 23. ICPP: IEEE International Conference on Parallel Processing, 2017-08-14. 60 papers, acceptance
625 rate: 28.6%. Homepage: <http://www.icpp-conf.org/2017/index.php>

- 626 24. IGSC: IEEE International Green and Sustainable Computing Conference, 2017-10-23. 23 papers,
627 acceptance rate: unknown. Homepage: <http://igsc.eecs.wsu.edu/>
- 628 25. IISWC: IEEE International Symposium on Workload Characterization, 2017-10-02. 31 papers,
629 acceptance rate: 37.3%. Homepage: <http://www.iiswc.org/iiswc2017/index.html>
- 630 26. IMC: ACM Internet Measurement Conference, 2017-11-01. 28 papers, acceptance rate: 15.6%.
631 Homepage: <http://conferences.sigcomm.org/imc/2017/>
- 632 27. IPDPS: IEEE International Parallel and Distributed Processing Symposium, 2017-05-29. 116
633 papers, acceptance rate: 22.8%. Homepage: <http://www.ipdps.org/ipdps2017/>
- 634 28. ISC: ISC High Performance, 2017-06-18. 22 papers, acceptance rate: 33.3%. Homepage: <http://isc-hpc.com/id-2017.html>
635
- 636 29. ISCA: ACM/IEEE International Symposium on Computer Architecture, 2017-06-24. 54 papers,
637 acceptance rate: 16.8%. Homepage: <http://iscal7.ece.utoronto.ca/doku.php>
- 638 30. ISPASS: IEEE International Symposium on Performance Analysis of Systems and Software,
639 2017-04-24. 24 papers, acceptance rate: 29.6%. Homepage: <http://www.ispass.org/ispass2017/>
640
- 641 31. MASCOTS: IEEE International Symposium on the Modeling, Analysis, and Simulation of Com-
642 puter and Telecommunication Systems, 2017-09-20. 20 papers, acceptance rate: 23.8%. Homepage:
643 <http://mascots2017.cs.ucalgary.ca/>
- 644 32. MICRO: IEEE/ACM International Symposium on Microarchitecture, 2017-10-16. 61 papers,
645 acceptance rate: 18.7%. Homepage: <https://www.microarch.org/micro50/>
- 646 33. Middleware: The Annual Middleware Conference, 2017-12-11. 20 papers, acceptance rate: 26%.
647 Homepage: <http://2017.middleware-conference.org/>
- 648 34. MobiCom: ACM International Conference on Mobile Computing and Networking, 2017-10-17. 35
649 papers, acceptance rate: 18.8%. Homepage: <https://sigmobile.org/mobicom/2017/>
- 650 35. NDSS: The Network and Distributed System Security Symposium, 2017-02-26. 68 papers, accep-
651 tance rate: 16.1%. Homepage: <https://www.ndss-symposium.org/ndss2017/>
- 652 36. NSDI: USENIX Symposium on Networked Systems Design and Implementation, 2017-03-27. 42
653 papers, acceptance rate: 16.5%. Homepage: <https://www.usenix.org/conference/nsdi17/>
654
- 655 37. PACT: IEEE/ACM International Conference on Parallel Architectures and Compilation Techniques,
656 2017-09-11. 25 papers, acceptance rate: 23.1%. Homepage: <https://parasol.tamu.edu/pact17/>
657
- 658 38. PODC: ACM Symposium on Principles of Distributed Computing, 2017-07-25. 38 papers, accep-
659 tance rate: 24.7%. Homepage: <https://www.podc.org/podc2017/>
- 660 39. PODS: ACM Symposium on Principles of Database Systems, 2017-05-14. 29 papers, acceptance
661 rate: 28.7%. Homepage: <http://sigmod2017.org/pods-program/>
- 662 40. PPOPP: ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, 2017-02-
663 04. 29 papers, acceptance rate: 22%. Homepage: <http://ppopp17.sigplan.org/>
- 664 41. SC: The International Conference for High Performance Computing, Networking, Stor-
665 age and Analysis, 2017-11-13. 61 papers, acceptance rate: 18.7%. Homepage: <http://scl7.supercomputing.org/>
666
- 667 42. SIGCOMM: ACM SIGCOMM Conference, 2017-08-21. 36 papers, acceptance rate: 14.4%.
668 Homepage: <http://conferences.sigcomm.org/sigcomm/2017/>

- 669 43. SIGMETRICS: ACM SIGMETRICS, 2017-06-05. 27 papers, acceptance rate: 13.3%. Homepage:
670 <http://www.sigmetrics.org/sigmetrics2017>
- 671 44. SIGMOD: ACM International Conference on Management of Data, 2017-05-14. 96 papers,
672 acceptance rate: 19.6%. Homepage: <http://sigmod2017.org/>
- 673 45. SOCC: ACM Symposium on Cloud Computing, 2017-09-25. 45 papers, acceptance rate: unknown.
674 Homepage: <https://acmsocc.github.io/2017/>
- 675 46. SOSP: Symposium on Operating Systems Principles, 2017-10-29. 39 papers, acceptance rate:
676 16.8%. Homepage: <https://www.sigops.org/sosp/sosp17/>
- 677 47. SP: IEEE Security and Privacy, 2017-05-22. 60 papers, acceptance rate: 14.3%. Homepage:
678 <https://www.ieee-security.org/TC/SP2017/index.html>
- 679 48. SPAA: ACM Symposium on Parallelism in Algorithms and Architectures, 2017-07-24. 31 papers,
680 acceptance rate: 24.4%. Homepage: <http://spaa.acm.org/2017/index.html>
- 681 49. SYSTOR: ACM International Systems and Storage Conference, 2017-05-22. 16 papers, acceptance
682 rate: 34%. Homepage: <https://www.systor.org/2017/>
- 683 50. VEE: ACM International Conference on Virtual Execution Environments, 2017-04-09. 18 papers,
684 acceptance rate: 41.9%. Homepage: <http://conf.researchr.org/home/vee-2017>

685 REFERENCES

- 686 David Adam. Citation analysis: The counting house. *Nature*, 415(6873):726–730, 2002.
- 687 Dag W Aksnes. A macro study of self-citation. *Scientometrics*, 56(2):235–246, 2003. doi: 10.
688 1023/A:1021919228368. URL [https://link.springer.com/content/pdf/10.1023/](https://link.springer.com/content/pdf/10.1023/A:1021919228368.pdf)
689 [A:1021919228368.pdf](https://link.springer.com/content/pdf/10.1023/A:1021919228368.pdf).
- 690 Mario Biagioli. Watch out for cheats in citation game. *Nature*, 535(7611):201–201, 2016. doi: 10.1038/
691 535201a.
- 692 Elana Broch. Cite me, cite my references? (scholarly use of the acm sigir proceedings based on two
693 citation indexes). In *Proceedings of the 24th annual International ACM SIGIR Conference on Research*
694 *and Development in Information Retrieval*, pages 446–447, 2001. doi: 10.1145/383952.384090. URL
695 <https://dl.acm.org/doi/pdf/10.1145/383952.384090>.
- 696 Tanmoy Chakraborty, Suhansanu Kumar, Pawan Goyal, Niloy Ganguly, and Animesh Mukherjee. On
697 the categorization of scientific citation profiles in computer science. *Communications of the ACM*, 58
698 (9):82–90, August 2015. ISSN 0001-0782. doi: 10.1145/2701412. URL [https://doi.org/10.](https://doi.org/10.1145/2701412)
699 [1145/2701412](https://doi.org/10.1145/2701412).
- 700 Tom Coupé. Peer review versus citations—an analysis of best paper prizes. *Research Policy*, 42(1):
701 295–301, 2013. doi: 10.1016/j.respol.2012.05.004. URL [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S0048733312001412)
702 [science/article/pii/S0048733312001412](https://www.sciencedirect.com/science/article/pii/S0048733312001412).
- 703 Sitaram Devarakonda, Dmitriy Korobskiy, Tandy Warnow, and George Chacko. Viewing computer science
704 through citation analysis: Salton and bergmark redux. *Scientometrics*, 125(1):271–287, 2020. doi:
705 10.1007/s11192-020-03624-0. URL [https://www.ncbi.nlm.nih.gov/pmc/articles/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7968068/)
706 [PMC7968068/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7968068/).
- 707 Eitan Frachtenberg. Systems conferences analysis dataset. October 2021. doi: 10.5281/zenodo.5590574.
708 URL <http://github.com/eitanf/sysconf/>.
- 709 Eitan Frachtenberg. Research artifacts and citations in computer systems papers. *PeerJ Computer Science*,
710 8:e887, February 2022a. doi: 10.7717/peerj-cs.887. URL [https://peerj.com/articles/](https://peerj.com/articles/cs-887/)
711 [cs-887/](https://peerj.com/articles/cs-887/).

- 712 Eitan Frachtenberg. Multifactor citation analysis over five years: A case study of SIGMETRICS
713 papers. *Publications*, 10(4):47, December 2022b. doi: 10.3390/publications10040047. URL [https://](https://www.mdpi.com/2304-6775/10/4/47)
714 www.mdpi.com/2304-6775/10/4/47.
- 715 Eitan Frachtenberg and Rhody D. Kaner. Underrepresentation of women in computer systems research.
716 *PLOS ONE*, 17(4):e0266439, April 2022. doi: 10.1371/journal.pone.0266439. URL [https://](https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0266439)
717 journals.plos.org/plosone/article?id=10.1371/journal.pone.0266439.
- 718 Abby A Goodrum, Katherine W McCain, Steve Lawrence, and C Lee Giles. Scholarly publishing
719 in the internet age: a citation analysis of computer science literature. *Information Processing &*
720 *Management*, 37(5):661–675, 2001. URL [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/S0306457300000479)
721 [article/pii/S0306457300000479](https://www.sciencedirect.com/science/article/pii/S0306457300000479).
- 722 Gali Halevi, Henk Moed, and Judit Bar-Ilan. Suitability of google scholar as a source of scientific
723 information and as a source of data for scientific evaluation—review of the literature. *Journal of*
724 *informetrics*, 11(3):823–834, 2017.
- 725 David P Hamilton. Who’s uncited now? *Science*, 251(4989):25–25, 1991. doi: 10.1126/science.1986409.
726 URL <https://www.science.org/doi/10.1126/science.1986409>.
- 727 Stephen P Harter. The impact of electronic journals on scholarly communication: A citation analysis. *The*
728 *Public-Access Computer Systems Review*, 7(5):5–34, 1996.
- 729 Anne-Wil Harzing and Satu Alakangas. Google scholar, scopus and the web of science: a longitudinal
730 and cross-disciplinary comparison. *Scientometrics*, 106(2):787–804, 2016.
- 731 Susan Davis Herring. Use of electronic resources in scholarly electronic journals: a citation analysis.
732 *College & Research Libraries*, 63(4):334–340, 2002. URL [https://crl.acrl.org/index.](https://crl.acrl.org/index.php/crl/article/viewFile/15538/16984)
733 [php/crl/article/viewFile/15538/16984](https://crl.acrl.org/index.php/crl/article/viewFile/15538/16984).
- 734 Robert Heumüller, Sebastian Nielebock, Jacob Krüger, and Frank Ortmeier. Publish or perish, but do not
735 forget your software artifacts. *Empirical Software Engineering*, 25(6):4585–4616, November 2020.
736 doi: 10.1007/s10664-020-09851-6.
- 737 Jorge E Hirsch. An index to quantify an individual’s scientific research output. *Proceedings of the National*
738 *academy of Sciences*, 102(46):16569–16572, November 2005. doi: 10.1073/pnas.0507655102.
- 739 Graeme Hirst and Nadia Talent. Computer science journals—an iterated citation analysis. *IEEE Transac-*
740 *tions on Professional Communication*, PC-20(4):233–238, December 1977. doi: 10.1109/TPC.1977.
741 6591956.
- 742 Waleed Iqbal, Junaid Qadir, Saeed-Ul Hassan, Rana Tallal Javed, Adnan Noor Mian, Jon Crowcroft, and
743 Gareth Tyson. Five decades of the acm special interest group on data communications (sigcomm) a
744 bibliometric perspective. *ACM SIGCOMM Computer Communication Review*, 49(5):29–37, 2019a.
745 doi: 10.1145/3371934.3371948. URL [https://dl.acm.org/doi/pdf/10.1145/3371934.](https://dl.acm.org/doi/pdf/10.1145/3371934.3371948)
746 [3371948](https://dl.acm.org/doi/pdf/10.1145/3371934.3371948).
- 747 Waleed Iqbal, Junaid Qadir, Gareth Tyson, Adnan Noor Mian, Saeed-ul Hassan, and Jon Crowcroft.
748 A bibliometric analysis of publications in computer networking research. *Scientometrics*, 119(2):
749 1121–1155, 2019b. doi: 10.1007/s11192-019-03086-z. URL [https://link.springer.com/](https://link.springer.com/article/10.1007/s11192-019-03086-z)
750 [article/10.1007/s11192-019-03086-z](https://link.springer.com/article/10.1007/s11192-019-03086-z).
- 751 Thomas S Jacques and Neil J Sebire. The impact of article titles on citation hits: an analysis of general
752 and specialist medical journals. *JRSM short reports*, 1(1):1–5, 2010. doi: 10.1258/shorts.2009.100020.
753 URL <https://journals.sagepub.com/doi/pdf/10.1258/shorts.2009.100020>.
- 754 Daud Khan, Masoumeh Khalil Arjmandi, and Mayank Yuvaraj. Most cited works on cloud computing:
755 The ‘citation classics’ as viewed through dimensions.ai. *Science & Technology Libraries*, 41(1):42–55,
756 2022. doi: 10.1080/0194262X.2021.1951424. URL [https://doi.org/10.1080/0194262X.](https://doi.org/10.1080/0194262X.2021.1951424)
757 [2021.1951424](https://doi.org/10.1080/0194262X.2021.1951424).

- 758 Kayvan Kousha and Mike Thelwall. Google scholar citations and google web/url citations: A multi-
759 discipline exploratory analysis. *Journal of the American Society for Information Science and Technol-*
760 *ogy*, 58(7):1055–1065, 2007.
- 761 Vincent Larivière, Yves Gingras, and Éric Archambault. The decline in the concentration of citations,
762 1900–2007. *Journal of the American Society for Information Science and Technology*, 60(4):858–862,
763 2009. URL <https://arxiv.org/ftp/arxiv/papers/0809/0809.5250.pdf>.
- 764 Vincent Larivière, Veronique Kiermer, Catriona J MacCallum, Marcia McNutt, Mark Patterson, Bernd Pul-
765 verer, Sowmya Swaminathan, Stuart Taylor, and Stephen Curry. A simple proposal for the publication
766 of journal citation distributions. *BioRxiv*, page 062109, 2016.
- 767 Danielle H Lee. Predictive power of conference-related factors on citation rates of conference papers.
768 *Scientometrics*, 118(1):281–304, 2019. doi: 10.1007/s11192-018-2943-z. URL [https://link.](https://link.springer.com/article/10.1007/s11192-018-2943-z)
769 [springer.com/article/10.1007/s11192-018-2943-z](https://link.springer.com/article/10.1007/s11192-018-2943-z).
- 770 Raymond Lister and Ilona Box. A citation analysis of the sigcse 2007 proceedings. In *Proceedings of the*
771 *39th SIGCSE technical symposium on Computer science education*, pages 476–480, 2008.
- 772 Michael H MacRoberts and Barbara R MacRoberts. Problems of citation analysis: A critical review.
773 *Journal of the American Society for information Science*, 40(5):342–349, 1989.
- 774 Michael H MacRoberts and Barbara R MacRoberts. The mismeasure of science: Citation analysis.
775 *Journal of the Association for Information Science and Technology*, 69(3):474–482, 2018.
- 776 Alberto Martin-Martin, Enrique Orduna-Malea, Mike Thelwall, and Emilio Delgado Lopez-Cozar. Google
777 scholar, web of science, and scopus: A systematic comparison of citations in 252 subject categories.
778 *Journal of informetrics*, 12(4):1160–1177, 2018.
- 779 Waister Martins, Marcos Gonçalves, Alberto Laender, and Nivio Ziviani. Assessing the quality of
780 scientific conferences based on bibliographic citations. *Scientometrics*, 83(1):133–155, 2010.
- 781 Sandra Mattauch, Katja Lohmann, Frank Hannig, Daniel Lohmann, and Jürgen Teich. A bibliometric
782 approach for detecting the gender gap in computer science. *Communications of the ACM*, 63:74–80,
783 2020. doi: 10.1145/3376901.
- 784 Lokman I Meho. The rise and rise of citation analysis. *Physics World*, 20(1):32, 2007. URL [https :](https://arxiv.org/pdf/physics/0701012)
785 [//arxiv.org/pdf/physics/0701012,](https://arxiv.org/pdf/physics/0701012).
- 786 Henk F Moed. *Citation analysis in research evaluation*. Springer Science & Business Media, 2006.
- 787 G Parthasarathy and DC Tomar. Sentiment analyzer: Analysis of journal citations from citation databases.
788 In *5th international conference-confluence the next generation information technology summit (conflu-*
789 *ence)*, pages 923–928. IEEE, November 2014. doi: 10.1109/CONFLUENCE.2014.6949321.
- 790 Gregory S. Patience, Christian A. Patience, Bruno Blais, and Francois Bertrand. Citation analy-
791 sis of scientific categories. *Heliyon*, 3(5):e00300, 2017. ISSN 2405-8440. doi: [https://doi.](https://doi.org/10.1016/j.heliyon.2017.e00300)
792 [org/10.1016/j.heliyon.2017.e00300](https://doi.org/10.1016/j.heliyon.2017.e00300). URL [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/S2405844016322800)
793 [article/pii/S2405844016322800](https://www.sciencedirect.com/science/article/pii/S2405844016322800).
- 794 David A Patterson. The health of research conferences and the dearth of big idea papers. *Communications*
795 *of the ACM*, 47(12):23–24, December 2004. doi: 10.1145/1035134.1035153.
- 796 David A Patterson, Lawrence Snyder, and Jeffrey Ullman. Evaluating computer scien-
797 tists and engineers for promotion and tenure. *Computing Research News*, September
798 1999. URL [http://www.cra.org/resources/bp-view/evaluating_computer_](http://www.cra.org/resources/bp-view/evaluating_computer_scientists_and_engineers_for_promotion_and_tenure/)
799 [scientists_and_engineers_for_promotion_and_tenure/](http://www.cra.org/resources/bp-view/evaluating_computer_scientists_and_engineers_for_promotion_and_tenure/).
- 800 P Pichappan and R Ponnudurai. Skewness in citation peak. In *Proceedings of the 7th Conference of the*
801 *International Society for Scientometrics and Informetrics*, pages 395–407, 1999.

- 802 Erhard Rahm and Andreas Thor. Citation analysis of database publications. *ACM Sigmod Record*, 34
803 (4):48–53, 2005. doi: 10.1145/1107499.1107505. URL <https://dl.acm.org/doi/pdf/10.1145/1107499.1107505>.
804
- 805 Sidney Redner. How popular is your paper? an empirical study of the citation distribution. *The*
806 *European Physical Journal B-Condensed Matter and Complex Systems*, 4(2):131–134, 1998. URL
807 <https://link.springer.com/content/pdf/10.1007/s100510050359.pdf>.
- 808 Tarek Saier, Johan Krause, and Michael Färber. unarxiv 2022: All arxiv publications pre-processed for
809 nlp, including structured full-text and citation network. *arXiv preprint arXiv:2303.14957*, 2023. URL
810 <https://arxiv.org/abs/2303.14957>.
- 811 Herbert Snyder and Susan Bonzi. Patterns of self-citation across disciplines (1980-1989). *Journal*
812 *of Information Science*, 24(6):431–435, 1998. doi: 10.1177/016555159802400606. URL <https://doi.org/10.1177/016555159802400606>.
813
- 814 Ming-yueh Tsay and Zhu-yee Shu. Journal bibliometric analysis: a case study on the jour-
815 nal of documentation. *Journal of Documentation*, 67(5):806–822, 2011. doi: 10.1108/
816 00220411111164682. URL <https://www.emerald.com/insight/content/doi/10.1108/002204111111164682/full/html>.
817
- 818 Attila Varga. Shorter distances between papers over time are due to more cross-field references and
819 increased citation rate to higher-impact papers. *Proceedings of the National Academy of Sciences*, 116
820 (44):22094–22099, 2019. doi: 10.1073/pnas.1905819116. URL <https://www.pnas.org/doi/full/10.1073/pnas.1905819116>.
821
- 822 George Vrettas and Mark Sanderson. Conferences versus journals in computer science. *Journal of the*
823 *Association for Information Science and Technology*, 66(12):2674–2684, April 2015. doi: 10.1002/asi.
824 23349.
- 825 Jacques Wainer, Michael Eckmann, and Anderson Rocha. Peer-selected “best papers”—are they really
826 that “good”? *PLOS ONE*, 10(3):e0118446, 2015. doi: 10.1371/journal.pone.0118446. URL <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0118446>.
827
- 828 Ludo Waltman. A review of the literature on citation impact indicators. *Journal of informetrics*, 10(2):
829 365–391, 2016. doi: /10.1016/j.joi.2016.02.007. URL <https://www.sciencedirect.com/science/article/pii/S1751157715300900>.
830
- 831 Dashun Wang and Albert-László Barabási. *The science of science*. Cambridge University Press, Cam-
832 bridge, UK, 2021.
- 833 Jian Wang. Citation time window choice for research impact evaluation. *Scientometrics*, 94(3):851–872,
834 2013.
- 835 Nianxin Wang, Huigang Liang, Yu Jia, Shilun Ge, Yajiong Xue, and Zhining Wang. Cloud computing
836 research in the is discipline: A citation/co-citation analysis. *Decision Support Systems*, 86:35–47, 2016.
- 837 Glänzel Wolfgang, Thijs Bart, and Schlemmer Balázs. A bibliometric approach to the role of author
838 self-citations in scientific communication. *Scientometrics*, 59(1):63–77, 2004. doi: 10.1023/b:scie.
839 0000013299.38210.74. URL <https://link.springer.com/content/pdf/10.1023/B:SCIE.0000013299.38210.74.pdf>.
840
- 841 Ling-Ling Wu, Luesak Luesukprasert, and Lynne Lee. Research and the long tail: A large-scale citation
842 analysis. In *42nd Hawaii International Conference on System Sciences*, pages 1–10. IEEE, 2009. URL
843 <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=4755756>.