

Conference acceptance rate signals future impact of published conference papers.

BY JILIN CHEN AND JOSEPH A. KONSTAN

Conference Paper Selectivity and Impact

STUDYING THE METADATA of the ACM Digital Library (<http://www.acm.org/dl>), we found that papers in low-acceptance-rate conferences have higher impact than those in high-acceptance-rate conferences within ACM, where impact is measured by the number of citations received. We also found that highly selective conferences—those that accept 30% or less of submissions—are cited at a rate comparable to or greater than ACM Transactions and journals.

In addition, the higher impact of selective conferences cannot be explained solely by a more strict filtering process; selectivity signals authors and/or readers of the quality of a venue and thus invites higher-quality submissions from authors and/or more citations from other authors.

Low-acceptance-rate conferences with selective peer-review processes distinguish computer science from other academic fields where only journal publication carries real weight. Focus on conferences challenges the

» key insights


- **Papers published in highly selective CS conferences are cited more often on average than papers published in Transactions and journals.**
- **Conference selectivity serves two purposes: pick the best submitted papers and signal prospective authors and readers about conference quality.**
- **Below a certain acceptance rate, selectivity can backfire; conferences rejecting 85% or more of their submissions risk discouraging overall submissions and inadvertently filtering out high-impact research.**

field in two ways: how to assess the importance of conference publication, particularly compared to journal publication, and how to manage conferences to maximize the impact of the papers they publish. “Impact factor” (average citation rate) is the commonly used measure of the influence of a journal on its field. While nearly all computer scientists have strong intuition about the link between conference acceptance rate and a paper’s impact, we are aware of no systematic studies examining that link or comparing conference and journal papers in terms of impact.


This article addresses three main questions: How does a conference’s acceptance rate correlate with the impact of its papers? How much impact do conference papers have compared to journal papers? To what extent does the impact of a highly selective conference derive from filtering (the selectivity of the review process) vs. signaling (the message the conference sends to both authors and readers by being selective)? Our results offer guidance to conference organizers, since acceptance rate is one of the few parameters they can control to maximize the impact of their conferences. In addition, our results inform the process of evaluating researchers, since we know that computer scientists often defend the primary publication of results in conferences, particularly when being evaluated by those outside the field (such as in tenure evaluations).² Finally, we hope these results will help guide individual researchers in understanding the expected impact of publishing their papers in the various venues.

Data and Methodology

We based our study on ACM Digital Library metadata for all ACM conference and journal papers as of May 2007, as well as on selected other papers in the *ACM Guide to Computing Literature* for which metadata was available. Since there is no established metric for measuring the scientific influence of published papers, we chose to estimate a paper’s influence as the number of times it was cited in the two years following publication, referred to as citation count or simply as impact. We excluded from this count “self-citation” in subsequent papers by the authors



Overall, the conference papers had an average two-year citation count of 2.15, and the journal papers had an average two-year citation count of 1.53.



of the original work. Using citations as a measure of scientific influence has a long tradition, including the journal-impact factor.¹ We chose two years as a compromise between measuring long-term impact and the practical importance of measuring impact of more recent work.¹ Less than two years might be too short for the field to recognize the worth of a paper and cite it. More than two years would have excluded more recently published papers from our analysis due to insufficient time after publication, so would not have allowed us to include the current era of widespread use of the Digital Library.^a

For conferences, we counted only full papers, since they represent their attempt to publish high-impact work, rather than posters and other less-rigorously reviewed material that might also appear in conference proceedings. Conferences where acceptance rates were not available were excluded as well. For journals, we included only titled ACM Transactions and journals; only these categories are generally viewed as archival research venues of lasting value.

Finally, since our data source was limited to the metadata in the *ACM Guide*, our analysis considered only citations from within that collection and ignored all citations from conferences and journals outside of it; this was a pragmatic constraint because, in part, other indexing services do not comprehensively index conference proceedings. While it means that all our numbers were underestimates and that the nature of the underestimates varied by field (we expected to significantly underestimate artificial intelligence and numerical-computation papers due to the large number of papers published by SIAM and AAAI outside our collection), such underestimates were not biased toward any particular acceptance rate in our data set.^b

a To ensure that the two-year citation count was reasonable, we repeated this analysis using four- and eight-year citation counts; the distributions and graphs were similar, and the conclusions were unchanged.

b We hand-checked 50 randomly selected conference papers receiving at least one citation in our data set, comparing citation count in the data set against citation count according to Google scholar (<http://scholar.google.com/>). When trying to predict Google scholar citation count from ACM citation count in a linear regression, we found an adjusted *R*-square of

Therefore, this limitation did not invalidate our results.

Our analysis included 600 conferences consisting of 14,017 full papers and 1,508 issues of journals consisting of 10,277 articles published from 1970 to 2005. Their citation counts were based on our full data set consisting of 4,119,899 listed references from 790,726 paper records, of which 1,536,923 references were resolved within the data set itself and can be used toward citation count. Overall, the conference papers had an average two-year citation count of 2.15 and the journal papers an average two-year citation count of 1.53. These counts follow a highly skewed distribution (see Figure 1), with over 70% of papers receiving no more than two citations. Note that while the average two-year citation count for conferences was higher than journals, the average four-year citation count for articles published before 2003 was 3.16 for conferences vs. 4.03 for journals; that is, on average, journals come out a little ahead of conference proceedings over the longer term.

Results

We addressed the first question—on how a conference’s acceptance rate correlates with the impact of its papers—by correlating citation count with acceptance rate; Figure 2 shows a scatterplot of average citation counts of ACM conferences (y-axis) by their acceptance rates (x-axis). Citation count differs substantially from the spectrum of acceptance rates, with a clear trend toward more citations for low acceptance rates; we observed a statistically significant correlation between the two values (each paper treated as a sample, $F[1, 14015] = 970.5, p < .001^c$)

0.852, showing that overall ACM citation count is proportional to Google scholar citation count with a small variation. When added as an additional parameter to the regression, acceptance rate had a nonsignificant coefficient, showing that acceptance rate does not have a significant effect on the difference between ACM citation count and Google scholar citation count. We also hand-checked 50 randomly selected conference papers receiving no citations in our data set, finding no correlation between acceptance rate and Google scholar citation count.

c This F-statistic shows how well a linear relationship between acceptance rate and citation count explains the variance within citation count. The notation $F[1, 14015] = 970.5, p < .001$ signifies one degree of freedom for model (from using only acceptance rate to ex-

Figure 1. Citation count distribution within two years of publication.

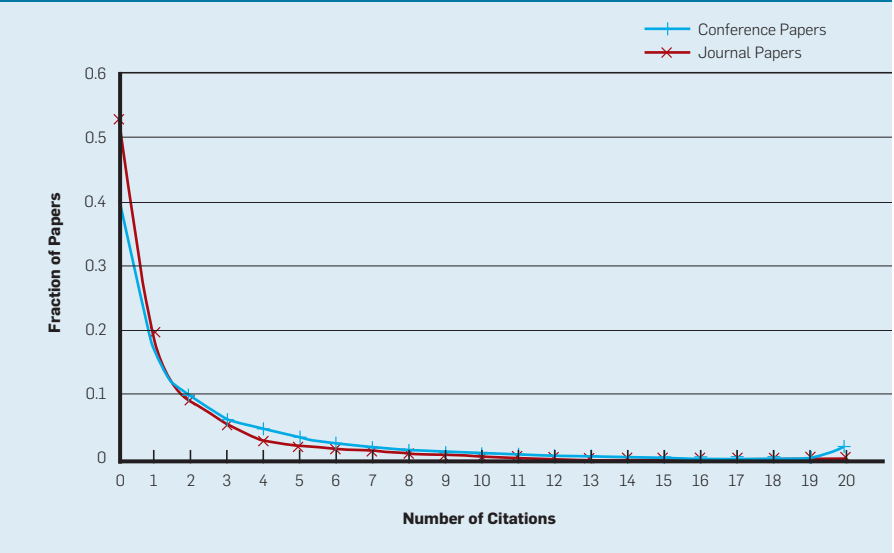
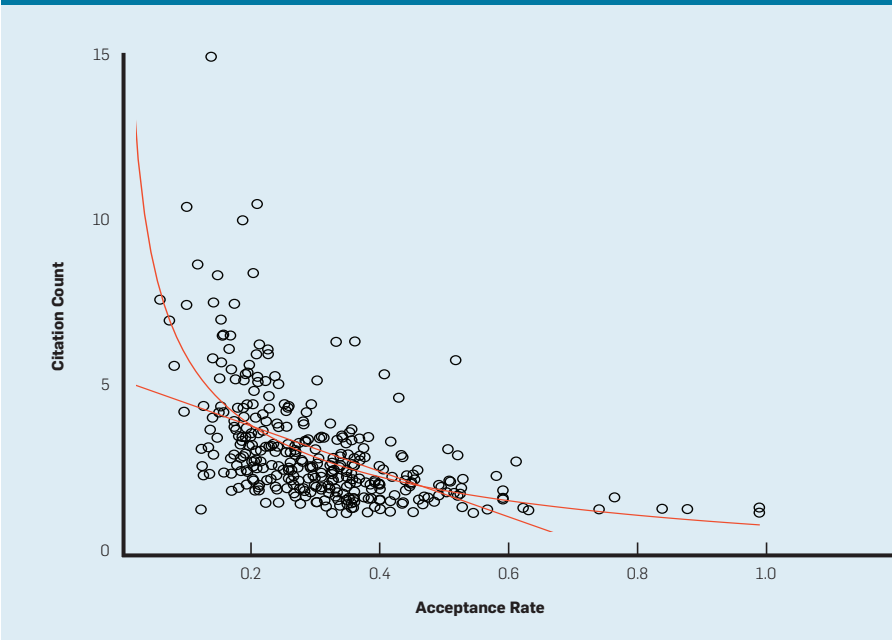


Figure 2. Average citation count vs. acceptance rate for ACM conferences.



and computed both a linear regression line (each conference weighted by its size, adjusted R -square: 0.258, weighted residual sum-of-squares: 35311) and a nonlinear regression curve in the form of $y = a + bx^{-c}$ (each conference weighted by its size, pseudo R -square: 0.325, weighted residual sum-of-squares: 32222), as shown in Figure 2.

Figure 3 is an aggregate view of the data, where we grouped conferences

plain citation counts), 14,015 degrees of freedom for error (from the more than 14,000 conference papers in our analysis), an F-statistic of 970.5, and probability less than 0.001 that the correlation between acceptance rate and citation count is the result of random chance.

into bins according to acceptance rates and computed the average citation counts of each bin.^d Citation counts for journal articles are shown as a dashed line for comparison. Conferences with rates less than 20% enjoyed an average citation count as high as 3.5. Less-selective conferences yielded fewer citations per paper, with the least-selective conferences (>55% acceptance rate) averaging less than $\frac{1}{2}$ citation per paper.

d We excluded conferences with an acceptance rate less than 10% and an acceptance rate over 60%, as there were too few conferences in these categories for meaningful analysis.

Figure 4 shows the percentages of papers within each group where citation count was above a certain threshold. The bottom bands (reflecting papers cited more than 10, 15, or 20 times in the following two years) show high-acceptance-rate conferences have few papers with high impact. Also notable is the fact that about 75% of papers published in >55%-acceptance-rate conferences were not cited at all in the following two years.

Addressing the second question—on how much impact conference papers have compared to journal papers—in Figures 3 and 4, we found that overall, journals did not outperform conferences in terms of citation count; they were, in fact, similar to conferences with acceptance rates around 30%, far behind conferences with acceptance rates below 25% (T -test, $T[7603] = 24.8$, $p < .001$). Similarly, journals published as many papers receiving no cita-

tions in the next two years as conferences accepting 35%–40% of submissions, a much higher low-impact percentage than for highly selective conferences.

The same analyses over four- and eight-year periods yielded results consistent with the two-year period; journal papers received significantly fewer citations than conferences where the acceptance rate was below 25%.

Low-acceptance-rate conferences in computer science have a greater impact than the average ACM journal. The fact that some journal papers are expanded versions of already-published (and cited) conference papers is a confounding factor here. We do not have data that effectively tracks research contributions through multiple publications to assess the cumulative impact of ideas published more than once.

Pondering why citation count correlates with acceptance rate brings us to

the third question—on the extent the impact of a highly selective conference derives from filtering vs. signaling—as this correlation can be attributed to two mechanisms:

Filtering. A selective review process filters out low-quality papers from the submission pool, lowering the acceptance rate and increasing the average impact of published papers; and

Signaling. A low acceptance rate signals high quality, thus attracting better submissions and more future citations, because researchers simply prefer submitting papers to reading the proceedings of and citing publications from better conferences. While filtering is commonly viewed as the whole point of a review process and thus likely explains the correlation to some extent, it is unclear whether signaling is also a factor. As a result, to address the third question, we clarified the existence of signaling by separating its potential effect from filtering.

We performed this separation by normalizing the selectivity of filtering to the same level for different conferences. For example, for a conference accepting 90 papers at a 30% acceptance rate, the best potential average citation count the conference could have achieved by lowering the acceptance rate to, say, 10% for the same submission pool would be the average citation count of the top 30 most-cited papers of the 90 accepted (presumably the 30 best papers of the original 300 submitted). We treated these 30 papers as the top 10% best submissions in the pool; other submissions were either filtered out during the actual review or later received fewer citations. Their citation count was thus an upperbound estimate of what might be achieved through stricter filtering, assuming conference program committees were able to pick exactly the submissions that would ultimately be the most highly cited. Using this normalization, we compared the same top portions of submission pools of all conferences and evaluated the effect of signaling without the influence of filtering. We normalized all ACM conferences in Figure 3 to a 10% acceptance rate and compared the citation counts of their top 10% best submissions; Figure 5 (same format as Figure 3) outlines the results. We excluded transactions and journals, as we were unable to get actual

Figure 3. Average citation count by acceptance rate within two years of publication.





Figure 4. Citation count distribution by acceptance rate within two years of publication.



acceptance-rate data, which might also be less meaningful, given the multi-review cycle common in journals.

Figure 5 suggests that citation count for the top 10% of submitted papers follows a trend similar to that of the full proceedings ($F[1, 5165] = 149.5, p < .001$), with generally higher count for low acceptance rates. This correlation indicates that filtering alone does not fully explain the correlation between citation count and acceptance rate; other factors (such as signaling) play a role.

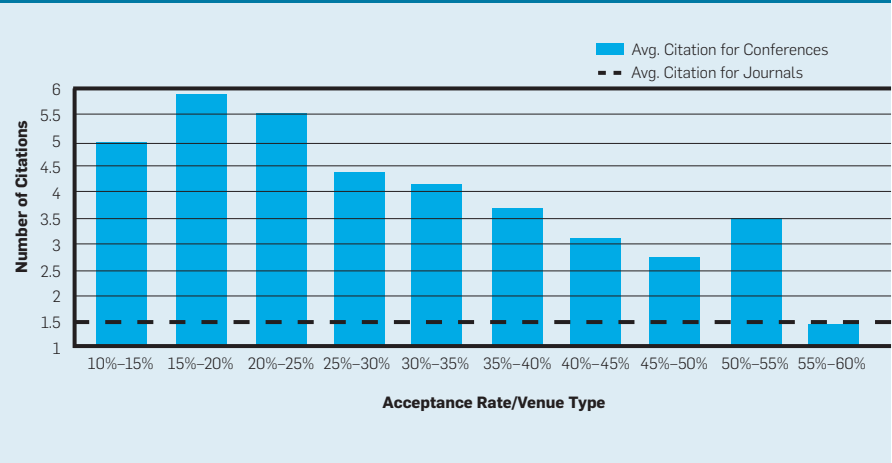
Discussion

Combining the results in Figures 3 and 5 provides further insight into the relationship between acceptance rate and citation count. For conferences with acceptance rates over 20%, the citation numbers in the figures almost consistently drop as the acceptance rate increases, suggesting that in this range, a higher acceptance rate makes conferences lose out on citation count not only for the conference but for its best submitted papers. Either higher-quality papers are not submitted to higher-acceptance-rate conferences as frequently or those submitted are not cited because readers do not explore the conferences as often as they explore lower-acceptance-rate conferences to find them.

The case for conferences with acceptance rates below 20% is more intriguing. Note that the lower impact of the 10%–15% group compared with the 15%–20% group in Figure 5 is statistically significant ($T[1198] = 3.21, p < .002$). That is, the top-cited papers from 15%–20%-acceptance-rate conferences are cited *more* often than those from 10%–15% conferences. We hypothesize that an extremely selective but imperfect (as review processes always are) review process has filtered-out submissions that would deliver impact if published. This hypothesis matches the common speculation, including from former ACM President David Patterson, that highly selective conferences too often choose incremental work at the expense of innovative breakthrough work.³

Alternatively, extremely low acceptance rates might discourage submissions by authors who dislike and avoid competition or the perception of there being a “lottery” among good papers for

Figure 5. Average citation count vs. acceptance rate within two years of publication, top 10% of submissions.



a few coveted publication slots. A third explanation suggests that extremely low acceptance rates have caused a conference proceedings to be of such limited focus that other researchers stop checking it regularly and thus never cite it. We consider all three to be plausible explanations; intuitively, all would hurt the impact of lower-acceptance-rate conferences more than they would hurt higher-acceptance-rate conferences.

Conclusion

Our results have several implications: First and foremost, computing researchers are right to view conferences as an important archival venue and use acceptance rate as an indicator of future impact. Papers in highly selective conferences—acceptance rates of 30% or less—should continue to be treated as first-class research contributions with impact comparable to, or better than, journal papers.

Second, we hope to bring to the attention of conference organizers and program committees the insight that conference selectivity does have a signaling value beyond simply separating good work from bad. Adopting the right selectivity level helps attract better submissions and more citations. Acceptance rates of 15%–20% seem optimal for generating the highest number of future citations for both the proceedings as a whole and the top papers submitted, though we caution that this guideline is based on ACM-wide data, and individual conferences should consider their goals and the norms of their sub-disciplines in setting target acceptance rates. Furthermore, many conferences

have goals separate from generating citations, and many high-acceptance-rate conferences might do a better job getting feedback to early ideas, supporting networking among attendees, and bringing together different specialties.

Given the link between acceptance rate and future impact, further research is warranted in the degree to which a conference’s reputation interacts over time with changes in its acceptance rate. Though a number of highly selective conferences have become more or less selective over time, we still lack enough data to clarify the effect of such changes. We hope that understanding them will yield new insight for conference organizers tuning their selectivity in the future.

Acknowledgments

This work was supported by National Science Foundation grant IIS-0534939. We thank our colleague John Riedl of the University of Minnesota for his valuable insights and suggestions.

References

1. Garfield, E. Citation analysis as a tool in journal evaluation. *Science* 178, 60 (Nov. 1972), 471–479.
2. National Research Council. *Academic Careers for Experimental Computer Scientists and Engineers*. U.S. National Academy of Sciences Report, Washington, D.C., 1994; http://www.nap.edu/catalog.php?record_id=2236
3. Patterson, D.A. The health of research conferences and the dearth of big idea papers. *Commun. ACM* 47, 12 (Dec. 2004), 23–24.

Jilin Chen (jilin@cs.umn.edu) is a doctoral student in the Department of Computer Science and Engineering at the University of Minnesota, Twin Cities.

Joseph A. Konstan (konstan@cs.umn.edu) is Distinguished McKnight Professor and Distinguished University Teaching Professor in the Department of Computer Science and Engineering at the University of Minnesota, Twin Cities.