

Paper Rating vs. Paper Ranking

John R. Douceur
Microsoft Research

Abstract

Within the computer-science community, submitted conference papers are typically evaluated by means of rating, in two respects: First, individual reviewers are asked to provide their evaluations of papers by assigning a rating to each paper's overall quality. Second, program committees collectively rate each paper as being either worthy or unworthy of acceptance, according to the aggregate judgment of the committee members. This paper proposes an alternative approach to these two processes, based on rankings rather than ratings.

1. Introduction

When an academic journal receives a submission, the journal asks reviewers to assess the quality of the paper relative to an established quality bar for the journal. The bar is determined by the selection of papers that have appeared in previous volumes of the journal. Once the reviewers have judged a submission to be above bar, the manuscript will be published, either in the next issue or – in the event that a particularly large set of high-quality submissions is received in a brief interval – in a shortly following issue. If, over time, the backlog of accepted-but-not-yet-published papers continues to grow, the journal's editors may ask future reviewers to raise their standards for subsequent papers. However, submission quality need have no immediate effect on the acceptance bar.

By contrast, academic conferences typically have a target number of papers to accept, or at least a target range. Therefore, the quality bar is at least somewhat dependent on the quality of the submissions to that particular year's conference, rather than strictly by the conference's history. Conferences have no freedom to delay the effect of current submission quality on the acceptance bar. Decisions must be made about whether to accept or reject each submitted paper, in light of the space budget of the conference.

Since a reasonable decision about each paper cannot be made in isolation from decisions about other papers under consideration, two common practices in program selection are highly suspect:

- in the reviewing process, asking reviewers to render a judgment about whether a submitted paper meets the conference's quality bar
- in the PC meeting, making accept/reject decisions individually for each paper

Neither of these practices is reasonable given that the bar is not known a priori. Moreover, by employing these common practices, conference organizers incur two significant problems:

- conflating reviewers' standards of stringency and leniency with the reviewers' judgment of the merits and weaknesses of each paper
- psychologically entrenching early acceptance decisions based upon insufficient information

Herein, we propose that both of these problems could be avoided by employing rankings rather than ratings for both individual reviewer assessments and program-committee discussion.

1.1. Assumed goal

We are assuming that goal of a program committee is to ensure that every accepted paper is higher quality than every rejected paper.¹ Though ideal, this goal is absurd in at least three respects: First, no objective standard of quality exists, so the goal is not well formed. Second, even if we assume that the opinions of PC members serve as an acceptable metric for evaluating paper quality, there may be differences of opinion among members regarding judgments of quality. And third, it is not generally possible to eliminate cases in which two reviewers disagree about which of two papers should exclusively be in the accepted set [3].

Despite the impossibility of the idealized goal, there are efforts we could take toward minimizing violations of

¹ This may not be strictly true, insofar as PCs may wish for balance among multiple subject areas and may thus tolerate lower-quality papers on subjects with lesser representation. In such cases, the recommendations of this paper could be applied within each subject area.

this goal. In this, we are aided by the fact that there is often a sizeable set of papers that could plausibly end up on either side of the acceptance decision.

1.2. Scope of proposal

The key issue this paper addresses is that, under the present system, subjective judgment of a paper's quality is bound up with an additional subjective determination of whether that quality is above or below the bar for acceptance.

This paper does not attempt to address any of the underlying causes of reviewer subjectivity, such as:

- emphasis preference – Reviewers may differ on the importance of aspects of a paper, such as novelty, completeness, extent of evaluation, currency, conference topicality, and clarity.
- topic interest – Reviewers have differing areas of interest; what is boring to one is engrossing to another.
- qualifications – Reviewers differ in their technical ability to adequately assess papers on various topics.
- defaults – Reviewers differ in their judgments of what to do with a paper they don't fully understand; whereas some are inclined to be charitable, others tend to be ruthless.

These sources of subjectivity present challenging problems, but they are beyond the scope of this paper.

2. Processes and problems in rating

Typically, reviewers are asked to rate each paper with ratings such as “strong accept”, “weak accept”, “weak reject”, and “strong reject”. Then, in the PC meeting, the committee collectively assigns a rating to each paper, based largely on the ratings provided by individual reviewers. The rating categories are similar, although they are characterized differently, such as “accept”, “accept if room”, “accept as poster”, and “reject”. Such ratings cause problems in both phases of the review process.

2.1. Rating-based reviews

As described above in §1.2, reviewers may have many axes of difference in the way they evaluate papers. But even if two reviewers happen to have the exact same emphasis preferences, topic interests, qualifications, and defaults, they might give drastically different ratings to a paper, because of differences in how

stringent or lenient they tend to be. In practice, this means that a paper reviewed by a stringent reviewer will receive a less favorable rating than a paper of comparable quality reviewed by a lenient reviewer.

Some program chairs have attempted to neutralize these tendencies by tagging each rating with a percentile range, such as “strong accept (top 10%)”, “weak accept (top 25% but not top 10%)”, etc. However, anecdotal evidence suggests that many reviewers discard these prescriptions in favor of the direct interpretations of each rating.

It might be possible to enforce a curve on ratings with sophisticated conference-management software that evaluates how well a reviewer's ratings fit the curve intended by the program chair. Imagine a dialog box that tells a reviewer, “You have strongly accepted 30% of the papers you reviewed. The overall acceptance rate for this conference will be approximately 12%. For randomly assigned papers, there is less than a 2% probability that the selection of your papers is skewed enough to warrant this discrepancy. Are you confident of your recommendations?”

If we were to take such an approach, we would have to answer the question of what to do when a reviewer insists on submitting off-curve ratings. If the software allows this to happen, then willful reviewers will easily circumvent this hypothetical safeguard. But if the software does not allow off-curve ratings to be submitted, we risk annoying reviewers, who might then decide not to submit any review because they feel themselves over-constrained.

2.2. Rating-driven PC meetings

The focus of a PC meeting (whether electronic or in-person) is to judge each submitted paper as either above or below the bar for acceptance. However, conferences typically have both a limit on the number of accepted papers and a (not necessarily official) quota to fill. For the count of accepted papers to fall within this target range, the quality bar must be set according to the quality distribution of submitted papers. However, this distribution is unknown until the committee has had the opportunity to discuss a significant fraction of papers.

This presents a Catch 21²: One cannot discuss whether to accept a paper without first determining where the bar is, but one cannot determine where the bar is

² Almost, but not quite, a Catch 22.

without first discussing a representative sample of papers. However, this is exactly what is called for by a process of sequential discussions on the acceptability of each paper in turn.

This situation gives rise to a dynamic that is likely to be familiar to anyone who has ever served on a PC: Early in the PC meeting, members maintain a very high standard for papers, rejecting good ones for fairly minor reasons. Later on, as it becomes clear that the quota will not be met, members start becoming looser about what they consider acceptable. Eventually, someone notes that the committee is accepting papers that are notably weaker than papers it had earlier rejected. This observation prompts earlier rejections to be revisited in light of the revised bar.

However, strong empirical evidence from psychology [2] shows that once a person renders a judgment on the desirability of an item, his opinion becomes reinforced, which strongly biases future judgments about the same item. Thus, even though a prematurely rejected paper may be brought up for reconsideration, it will generally not receive as much leniency as a paper that had not been tarnished with an early negative judgment. As Triesthof famously quipped, “You never get a second chance to make a first impression.”

Note that this problem occurs irrespective of the order in which papers are discussed. Therefore, it cannot be fixed by modifications to the paper-discussion order, such as discussing high-variance papers first.

3. Proposal: ranking

We propose that the problems enumerated in §2 could be avoided by basing reviews and PC discussion on rankings instead of ratings. Although rankings could be applied to reviews without applying them to PC meetings, or vice versa, the full benefits of ranking are only obtained when implemented together.

3.1. Ranking-based reviews

For many years, college admission boards have faced the problem of varying stringency among high schools in judgments of students’ grades. The widely adopted solution to this problem is for colleges to judge students by their class rank instead of by their GPA. In fact, the recent trend among some high schools of not reporting class rank has led college boards to complain that this reduces their objective information on students’ academic performance [4].

Class rank is immune to grade inflation. Analogously, a reviewer’s ranking of a set of papers is immune to the reviewer’s standards for acceptance. Papers reviewed by a stringent reviewer will not suffer unfairly in comparison to those reviewed by a lenient reviewer.

Some PC chairs have attempted to circumvent differing standards by normalizing reviewers’ ratings. However, if the number of rating choices is too small, they may contain too little information to discern the reviewer’s relative opinion of papers. On the other hand, if the number of rating choices is very large, this is really just a poor way of collecting rankings, since psychological evidence suggests that experts are better at rendering comparative judgments than absolute ones [5, 6]. In addition, if ratings are explicitly bound up with decision intentions (such as “accept”, “weak reject”, and so forth), this may still incur the problems described in § 2 above.

It may not be necessary to restrict reviewers to a total ordering. Perhaps a reviewer could be allowed to indicate that two or more papers are of equal rank, which would provide some additional reviewer flexibility. Empirical evaluation could help determine whether this freedom would tend to be abused by reviewers who hate (or love) every paper they read.

3.2. Ranking-driven PC meetings

A ranking-driven PC meeting can cleanly separate two processes that are tightly coupled – and conflated – in a rating-driven PC meeting: the judgment of paper quality (relative to other papers) and the determination of where to set the bar for acceptance.

Prior to the meeting, the chairs establish a straw-man global ranking. A simple method for producing such a ranking is, for each paper, convert each reviewer’s rank to a numerical score, and average the scores of all reviewers. Then, sort the papers according to their average scores. It remains to be seen what function would be best for the rank-to-score conversion, but it seems likely that the function should be nonlinear: There is probably more quality difference between papers ranked 1 and 2 than there is between papers ranked 10 and 11. The function should perhaps also account for the reviewer’s self-assessed confidence rating. A similar procedure is currently used in many PC meetings for determining a rough ranking for paper discussion order; however, since discussion is focused on individual accept/reject decisions, rather than changes to the rank order, the problems enumerated in § 2.2 remain.

The PC meeting then proceeds in two phases. In the first phase, the committee debates the relative ranks of papers. A typical instigating comment might be, “I thought paper 384 was far better than paper 721, but it is ranked three slots lower.” For pairs of papers which no single reviewer has reviewed, it is still possible to have an intelligent discussion among the reviewers of each paper: “Paper 219 has a really solid evaluation. Would the reviewers of the papers ranked above it please comment on the quality of the evaluation sections?”

Since the lower and upper bounds of the acceptable paper count are usually established beforehand, it is straightforward to avoid debating the relative ranks of any set of papers that are all well above or well below the cutoff range. The rankings of such papers, relative to each other, will not affect their ultimate acceptance or rejection. Papers whose ranks are within or near the cutoff range are the best candidates for intense debate, and so should be discussed first.

Divergent opinions could give rise to ordering cycles, but such cycles highlight papers that are important to thoroughly debate and/or to solicit additional reviews.

In the second phase, the committee establishes the cutoff point for paper acceptances. This decision could be based on a number of factors:

- The quality of papers within the target range might suggest that the bar should gravitate toward the upper or lower end of the range.
- The PC might be inclined to be generally lenient, or to be generally stringent.
- If the papers within the target range have some particularly desirable property, such as a fresh topic, the bar should perhaps go below them.
- A large gap in the assessed quality of adjacent papers may indicate a good cutoff point.
- If a short-distance cycle remains in the final ranking, the cutoff point should be positioned so that no cycle spans the cutoff, if possible.

The main benefit of a ranking-based discussion is to avoid prejudicing the committee’s judgments, but it has another benefit as well. A particular PC member may be especially dominant or persuasive, and in a rating-based system, he can intimidate or cajole the few other reviewers of a paper into accepting his view. However, in a ranking-based system, if any member wishes to significantly raise or lower the rank of a paper, he will have to argue against the reviewers of many other

papers. This will decrease the unwarranted influence of fearsome or charismatic members.

4. Challenges

A basic implementation challenge is that existing conference management software is designed to operate on ratings. Modifying this software to operate on rankings could require substantial reworking.

The ranking-driven PC meeting begins with a straw-man global ranking. If this ranking is poorly established, it may result in a very inefficient meeting. In the general case, it is impossible to convert a set of individual rankings into a global ranking [1]. However, we have no need for an optimal – or even consistent – global ranking. The initial ranking need only be good enough to avoid wasting time comparing papers of wildly different quality.

It is possible that there may be insufficient information for the PC to determine the relative ranking of certain pairs of papers. We suspect that, in practice, this will not be a common occurrence, because the transitivity of partial ordering will inform the relative ranks of most papers unless their levels of judged quality are very similar. If there are two papers that seem to settle near each other as the ranking is adjusted, and if no single reviewer is familiar with both papers, and if the papers lie near a likely cutoff point for acceptance, then this is a strong indication that a reviewer of each paper should be appointed to read the other paper and make a solid comparison.

Perhaps the biggest challenge with the ranking-driven PC meeting is that it complicates anonymous reviewing for papers submitted by PC members. In rating-driven meetings, a single paper is discussed at a time, so any members with conflicts-of-interest can step out of the room. In a ranking-driven meeting, multiple papers will be in discussion concurrently, since their relative merits are under consideration. Although a discussion of the merits and demerits of anonymous reviewing is beyond the scope of this paper, a simple way of addressing conflicts of interest is not to have PC members leave the room during discussion. We are aware of at least one top-tier conference’s PC meeting that merely required conflicted members to refrain from discussion, rather than leaving the room.

A final challenge is finding a program chair willing to try one or both parts of our proposal. The best venue might be a small workshop with few submissions. Such a venue might be more willing to allow conflicted

reviewers to silently observe paper discussions. Also, with a small number of submissions, there is less of a danger that reviewer coverage will be insufficient for making paper comparisons.

Acknowledgements

The author extends thanks to the WOWCS 2008 PC members who provided detailed reviews and many helpful comments on the submitted version of this paper, to the other WOWCS 2008 authors for the inspiring ideas in their submitted papers, to the various program chairs who have been daring enough to invite the author onto their program committees, and to the steering committees who have been even more daring in asking the author to serve as chair.

References

- [1] K. J. Arrow. "A difficulty in the concept of social welfare," *Journal of Political Economy* 58(4), August 1950, pp. 328–346.
- [2] J. Brehm. "Post-decision changes in desirability of alternatives," *Journal of Abnormal and Social Psychology* 52, 1956, pp. 384–389.
- [3] J. Duggan and T. Schwartz. "Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized," *Social Choice and Welfare* 17, 2000, pp. 85–93.
- [4] A. Finder, "Schools avoid class ranking, vexing colleges," *New York Times*, March 8, 2006.
- [5] C. Spetzler and C.-A. Stäel von Hostein, "Probability encoding in decision analysis," *Management Science*, 22, 340–358, 1975.
- [6] H. Wang, D. H. Dash, and M. J. Druzdel. "A method for evaluating elicitation schemes for probabilistic models," *IEEE Trans. on Systems, Man, and Cybernetics – Part B: Cybernetics*, 32(1):38–43, February, 2002.