The Task of the Referee*

Alan Jay Smith
CS Division, EECS Department
University of California
Berkeley, California 94720, USA

Abstract

The task of the referee is to evaluate a (research) paper in a timely manner for publication in a specific journal or conference proceedings. This involves determining if the work presented is correct, if the problem studied and the results obtained are new and significant, if the quality of the presentation is satisfactory or can be made so, and what revisions and changes to the paper are necessary and/or desirable. The evaluation must be relative to the degree of selectivity of the specific publication.

In this article, we discuss the problem of how to evaluate a paper for publication, and by inference, how to write one. The role of the editor, and rules and procedures used by most computer science journals are discussed. Brief discussions of refereeing proposals and survey and tutorial papers are also given.

1. Introduction

There is a constant stream of papers written and submitted for publication to conferences, journals, newsletters, anthologies, annuals, trade journals and newspapers, and other periodicals. Many such publications use referees as impartial, external experts to evaluate papers. This approach is often called "peer review", and is used to separate the wheat from the chaff. Refereeing is a public service, and is one of the professional obligations of a computer science professional. Typically, referees learn to produce referee reports by practice, by feedback from editors, by seeing referee reports for their own papers, and by reading referee reports written by others - i.e. solely by experience.

This article is an attempt to provide guidelines to referees on how to evaluate a paper, how to write a referee report, and how to apply commonly used standards and procedures. It is intended to replace the use of [Fors65], which does not reflect the procedures used in Computer Science and Engineering. The focus of this article is research in applied areas of computer science, such as systems, architecture, hardware, communications, and performance evaluation, but most of the discussion is generally applicable. Some discussion is provided as well on refereeing proposals, and also survey and tutorial papers. Authors should also find this material useful in preparing papers and submitting them for publication.

---

* Prof. Smith's research (regarding which he has received many referee reports) is supported in part by the National Science Foundation under grant MIP-8713274, by NASA under consortium agreement NCA2-128, by the State of California under the MICRO program, and by the International Business Machines Corporation, Digital Equipment Corporation, Hewlett Packard, Apple Computer, and Signetics/Philips Research Laboratories.

May, 1989

2. The Task of the Referee

The two major components of a referee report are:

(a) A recommendation for or against publication in a specific publication or presentation at a specific forum. An equivocal recommendation is acceptable if adequate discussion is provided for the guidance of the editor or program committee. If rejection is recommended, and if the paper does contain some publishable research, the report can suggest another place to publish. In all cases, sufficient discussion must be provided to justify the recommendation.

(b) A list of necessary and recommended changes and revisions. (A recommendation to reject the paper does not excuse the referee from suggesting changes that might permit the paper to be published elsewhere, or after resubmission.) The extent of necessary revisions, for journal publication, is largely separate from the recommendation for (eventual) publication; for a conference, the time factor generally will preclude substantial revisions.

It is very important that the referee walk the fine line between being overly permissive ("publish everything") and overly restrictive ("nothing is good enough to publish"). If one is insufficiently critical, one is encouraging bad research, causing trees to be cut down to make paper to publish the work, giving recognition (of a sort) and honors (of a sort) to those that don't deserve it, misleading the naive and inexperienced reader, misleading the author as to what is publishable, encouraging disrespect for the field, distorting commercial development, distorting hiring, promotion and tenure decisions, and perhaps subtracting from the general store of knowledge. In [Thom84], the problem of an excess of mediocre (or worse) papers in the literature is noted and discussed.

If one is overly critical of research, one blocks good research from publication, or causes it to be delayed in publication, wastes the time of authors, damages careers, and perhaps leaves journals with nothing to publish and conferences with nothing to present. It is particularly important not to reject new and significant work which runs counter to the prevailing wisdom or current fashions.

Basically, it is important for a referee who wants to be taken seriously to have a middle of the road view, to be able to distinguish good from bad work, and major from minor from negative contributions to the literature. A referee who always says "yes" or always says "no" is not helpful.

3. The Referee Report

A good referee report should have several parts. First, it is helpful if the referee summarizes the point of the paper in 1-5 sentences, both for the use of the editor, and to ensure that the referee actually knows what the point of the paper is. Second, the referee should evaluate the goal of the work both with respect to its validity and to its significance. Third, the referee should evaluate the quality of the work (methodology, techniques, accuracy, errors), and finally, the referee must provide an overall recommendation as to publication. If the recommendation is negative, the referee should always be clear about why s/he is recommending rejection. The referee should also be clear about the strength of his/her opinions; an equivocal ("maybe") recommendation is acceptable if the reasons for it are clearly documented. In any case, the referee report must contain enough discussion and information to justify the recommendation.

If the recommendation is favorable, it is essential that the referee provide as long a list as appropriate of both necessary and suggested changes. If the recommendation is negative, but the paper can be salvaged and either submitted elsewhere for publication or resubmitted to this publication, then a similar (but perhaps less detailed) list should also be provided. Suggestions for alternate places to publish are always welcome.

Typically, the text of the referee report is given to the author, stripped of all (surrounding) material identifying the referee. Therefore, while is important that the referee report be clear and explicit, it should not be insulting. Words such as "fool" and "idiot" should not be used to refer to the author, nor terms such as "trash" for the paper. The review of a research paper, in particular, should be directed at the paper, and not be a personal attack on the author. (The review of a proposal may be different, as is discussed below.)

Last but not least, the referee must make sure that his report reaches the editor in a timely manner. Computer science journals are notorious for having long delays between submission and publication; the two major components of that delay are the referees and the publication queue for the journal itself. Imagine if it were your paper! Likewise, referee reports must reach program chairs sufficiently before the program committee meeting that the material can be assembled and prepared for discussion.

4. Issues in Evaluating Research

What is the author trying to do?

What is the problem being considered? Is it clearly stated? Do you know what the author is trying to do? Does the author know? Does the author make clear what the important issues are? Does the author tell you (early in the paper) what he has accomplished? (E.g. if this is a system description, has the system been implemented or is this just a design?)

Is this work appropriate for this forum? One does not submit queueing theory papers to Datamation, or discussions of the latest release of MVS to JACM or Proceedings of the IEEE. Does this paper have anything to do with computer science?

Is the goal of this paper significant?

For that matter, is the problem real? Does it contradict any known physical laws (e.g. perpetual motion machines) or widely reported measurements?

Keep in mind what the Walrus said [Caro65]:

> "'The time has come,'
> the Walrus said,
> 'To talk of many things:
> Of shoes - and ships - and sealing wax -
> Of cabbages - and kings -
> And why the sea is boiling hot -
> And whether pigs have wings'"

Is this a careful analysis of how the sea got to be boiling hot, or an elegant study of the flight characteristics of pigs?

Is there any reason to care about the results of this paper, assuming (for the moment) that they are correct? Is the problem or goal major, minor, trivial or non-existent? Is the problem now obsolete? (E.g. reliability studies for vacuum tube mainframe computers.) Is the problem so specific or so applied as to have no general applicability and thus not be research?

Is the problem, goal, or intended result *new*? Has it been built before? Has it been solved before? Is this a *trivial* variation on or extension of previous results? Is the author aware of related and previous work, both recent and old? Does s/he *cite* that work? Are distinctions between this and previous work given and are they specific? If this work describes an implementation, are there any new ideas?

Is the method of approach valid?

Is there something about the approach to this problem that invalidates the results? Can you tell what the method is, or do you have to ferret it out from the middle of the mathematical formulas? What are the assumptions? How realistic are they? If they aren't realistic, does it matter? How sensitive are the results to the assumptions?

Is the method of approach significantly less than the state of the art? E.g. data is available, but the author is using a random number driven simulation, with unrealistic parameters. Does it matter? - back of the envelope calculations are often sufficient.

If this is a presentation of a new *idea* - e.g. a new design, or a new concept, is the *appropriate* amount of discussion or analysis presented? (Neither too much nor too little.)

Is the actual execution of the research correct?

Are the mathematics correct? Are the results consistent with the assumptions and/or with observed facts or measurements? Is the simulation methodology described in sufficient detail to convince the reader that the results are valid? If the simulation is stochastic, are confidence intervals on the results given? Are the results plausible, or even possible? Did the author do what he claims? (E.g. did he simulate the original system or a plausible model of it, or did he simulate his approximate mathematical model of the system?)

Are the correct conclusions being drawn from the results?

Are any conclusions being drawn from the results? What are the applications or implications of the results? Is there an adequate discussion of "why" these results were obtained?

Is the presentation satisfactory?

The first question is whether the paper is well enough written that it is possible to given it a technical evaluation. A paper which is incomprehensible is not publishable. Presuming that the paper is readable at all, an evaluation of the presentation is needed.

Is there too much or too little detail? Are the grammar and syntax correct? Is it well written? Does the abstract describe the paper? Does the introduction adequately explain the problem and the framework for the research? Are the figures and tables well labeled? Are the figures legible? meaningful? Are explanations poor or nonsensical? Is the author too verbose? too terse? Can the reader follow without reading all previous papers by this author? If the author refers the reader to other papers for crucial details, do you believe him? If sections of this paper are missing (due to a deadline), do you believe that they will be filled in as promised? Is the paper too colloquial in style? Are there many typographical errors? Does the paper contain too much material? (The paper should be long enough to present the necessary material, and no longer. Let the editor or program chair worry about page limits.)

What did you learn?

What did you (or what should the reader) learn from this paper? If you didn't learn anything, and/or if the intended reader won't learn anything, the paper is not publishable.

5. Overall, how good is it? What do you recommend?

Can you put the paper into one of these categories?

1. Major results - very significant. (less than 1% of the papers written.)
2. Good, solid, interesting work; a definite contribution. (less than 10% of the papers you will see.)
3. Minor, but positive, contribution to knowledge. (perhaps 10% to 30% of the papers submitted.)
4. Elegant and technically correct but useless. This category includes sophisticated analyses of flying pigs (see above).
5. Neither elegant nor useful, but not actually wrong.
6. Wrong and misleading.
7. The paper is so badly written that a technical evaluation is impossible.

In putting a paper into one of the above categories, it is important to normalize to the standards of the field as a whole, not to your own standards (which may be high or low) or to the average of the papers that you yourself write, or to the average of the papers that you find worth reading.

After categorizing the paper, the question is: what are the standards of this journal or conference? Is this the Proceedings of the IEEE or the ACM Transactions on Computer Systems or the ACM Symposium on Operating Systems Principles (all quite selective) or the Tahiti Conference on Beach Ball and Computer Systems? (Fictional, but a presumed boondoggle.) You should compare this paper not with the best or worst in that journal or conference, but with the average.

You should then make a recommendation. It can be favorable ("publish") or unfavorable ("reject"). The strength of the recommendation should be clearly stated to the editor (e.g. "wonderful paper, definitely accept"; "useful paper, probably accept"; "complete trash; definitely reject"). It is okay, but not desirable, to say "maybe"; if you feel that the paper has something worthwhile to say, but you're not sure if it is good enough to publish in this journal or conference proceedings, you can give an equivocal response.

You can also recommend that a paper be rejected for being inappropriate for this journal or conference.

If the paper is inappropriate, or marginal in quality for this forum, you can also suggest someplace else to submit the paper.

In any case, you *must* provide sufficient discussion and justification for whatever recommendation you make. A recommendation without sufficient justification will be given very little weight by the editor or program committee.

A recommendation to accept after revision will usually result in the editor asking you to review a revised version of the paper. It is important to ensure that the paper has been revised in a manner that is satisfactory, but the referee should be very cautious in finding new problems with the paper, when those "problems" were in the first version as well. Genuine problems should be cause for further revision or rejection, but one should not harass the author by proposing or requiring revision after revision unnecessarily.

6. Surveys and Tutorials

Surveys and tutorials constitute a distinct class of articles from research papers. Most or all of the work reported in such papers is not new, and is not expected to be new, although such a paper may be a convenient place for an author to include a variety of minor research results which would not stand on their own in separate papers.

The categories of survey and tutorial are overlapping but not identical. A pure tutorial is intended for the non-expert (generally novice) reader, and is intended to explain some body of material. The tutorial may not be comprehensive, and may be oriented toward a specific view of a field. On the other hand, the survey should provide broad and thorough coverage of some field or body of knowledge, and may be aimed at a reader ranging from the novice to the almost-expert.

In reviewing a tutorial paper, there are some specific issues to address: Does the paper cover the material promised by the title or abstract? Is this a reasonable body of knowledge to be covered by a tutorial article? (Is the scope too wide, too narrow, or too bizarre to be useful?) Is the material in the article correct? Is the level of coverage excessively simple-minded or excessively sophisticated, given the likely audience? Is the paper well written and clear? (This is a crucial issue for tutorials, but journals that publish tutorials, such as IEEE Computer and ACM Computing Surveys, often have editors that help with revisions.)

For a survey paper, many of the same issues apply. Does the paper cover the material promised by the title or abstract, and is this a reasonable body of knowledge to be surveyed at one time? Is the material in the article correct, and is the author sufficiently expert on the subject that he is able to correctly interpret results and provide perspective on the field? Has the author provided a balanced and comprehensive view of his topic? (Does s/he cite the important relevant literature, or does s/he omit mention of his competitors or enemies?) To the extent that the survey includes new research results, do those results meet the criteria given above for validity and correctness? (A survey does not have to stand on its own as a research paper, and so any research presented does not have to be very significant.) Finally, is the paper well written and clear?

7. Proposals

A proposal is a request to a funding agency for financial support, supposedly to do the research described in the proposal. Reviewing proposals is quite different from reviewing papers, and some special considerations apply. We provide some brief comments on this subject here.

The primary difficulty with reviewing a proposal is that the investigator is supposed to be telling you what s/he *plans* to do, rather than what has been done. The questions to be asked, then are: (a) Is the topic (or topics) of research significant? (b) Is the method of approach described (briefly) and is it reasonable? (c) Does the investigator appear to have sufficient expertise to produce useful results? (d) Is the budget reasonable, given the proposed research, the qualifications of the investigator (and his/her students) and the typical level of funding provided by the agency in question?

The easiest way to write a detailed and specific proposal is to propose to do research that is already complete, or at least substantially underway; this approach is quite common for an established researcher. Unfortunately, that isn't the purpose of a research proposal, and requiring a high level of detail and specificity in the proposal discriminates against newcomers to the field, and also against those that propose new work. Thus, the most serious difficulty faced by the reviewer of proposals is to evaluate a *proposal* and not to expect to see a research *report*. Similarly, a proposal may include a larger scope of work than can be reasonably accomplished with the time and effort specified. If the investigator clearly recognizes this, and indicates that s/he will pick and choose sub-topics within the area of research, depending on their interest and the availability of someone (e.g. a graduate student) to work on them, this is not a negative factor.

A major difference between research proposals and papers is that a proposal is speculative, and the reviewer has to evaluate what is likely to result. Such an evaluation, as noted below, should rely strongly on the reputation of the investigator, *when that reputation exists*. People with a history of good research will probably do good work, no matter how sloppy or brief their proposal. People with a history of low quality research (no matter how voluminous, nor how hot the topic) will probably continue in a consistent manner, no matter how exciting the proposal. Therefore, a very large fraction of the evaluation of a proposal by a *well known* investigator should depend on the reputation of that investigator. Is is very important, however, *to not discriminate against newcomers to the field*, who have no reputation, either good or bad. In the latter case, one must rely much more heavily on the text of the proposal, and on other information such as the investigator's Ph.D. dissertation, his academic record, his host institution, his Ph.D. institution, comments by his advisor or others who know him, etc.

Proposal reviewers are asked to comment on the proposed budget. Keep in mind that many factors affect the size of the budget other than the proposed scope of research, such as the agency providing the funding, the availability of facilities and staff, etc. Note particularly that for a new investigator, there is a major difference between no funding and minimal funding (2 months summer salary, amounts for travel, supplies and computer time). Funding a new investigator at a low level is often a good gamble - two or three years later, the investigator will have a track record, and if it is a good one, higher levels of funding can be justified.

8. Other Issues

8.1. Simultaneous Submission, Prior Publication, Unrevised Retries

If a paper is submitted simultaneously to two or more places, all editors and/or program chairs should be advised of (and approve of) this, and all referees should be notified. It is unethical to submit a paper simultaneously without notification, and that is a sufficient basis for rejecting the paper. There is a very good chance that if a paper is simultaneously submitted, the simultaneous submission will be detected.

If a paper has already been published (e.g. in the proceedings of a conference) and is then submitted for republication (e.g. in an archival journal), it is essential that the editor and referees be so notified. Some associations (e.g. IEEE, ACM) permit republication in their journals, but implicitly the paper must meet a higher standard for republication than if it had never been published. The paper must be important enough to merit republication. Significant extensions or major revisions are often a sufficient reason for republication. If the first version of the paper was published by a commercial (for profit) publisher, then it may be illegal (due to copyright) to republish the paper without explicit permission.

It is not uncommon to receive a paper to referee which you have previously recommended be rejected by some other publication. If the paper has not been significantly changed to comply with your previous review, it is appropriate to return a copy of the previous review, along with a blunt note suggesting that the author might try making revisions in accordance with referee reports.

8.2. Acknowledgements and Plagiarism

It is important that papers not plagiarize, and that joint work and contributions of others be fully acknowledged. Referees should explicitly point out any such problems discovered.

8.3. Timely Response, and Returning a Paper

It is important that referees respond reasonably promptly. Conferences have deadlines, and reports received after the program committee has met are useless. Journals do not generally have deadlines, but taking a long time to review a paper, and preventing its consideration for publication through delay, is professionally unethical. If you can't read the paper in a reasonable amount of time (typically 4-8 weeks), send it back to the editor, or at least get the editor's agreement to the delay.

Keep in mind that if you expect to have your own papers published, you have a responsibility to referee a reasonable number of papers. It is part of your job as a researcher, and the option of sending it back to the editor should not be abused. (Any editor who finds that an individual refuses to referee papers, either explicitly or through inaction, should feel free to send submissions by that individual back to him/her without further consideration.)

If you are sent a paper which you are not qualified to referee, you may also send it back to the editor (or program chair). Note, however, that you may have been specifically selected to provide an "outside" view of the field, and that fact may qualify you to provide a (limited) opinion.

If you are going to send a paper back without refereeing it, please do so *immediately*. Be sure to return the manuscript.

8.4. Does the Author's Reputation Matter?

The question is whether the reputation of the author should be allowed to influence the evaluation of the paper. This is an issue on which there is no consensus, and here I present my personal opinion: *to a small extent, and only in some circumstances.*

A research paper must stand on its own. The only time to take into account the reputation of the author is with regard to ambiguities, points that aren't clear, and reference to work that isn't presented. If the author is well regarded, one may be inclined to assume that any problems will be corrected (and must be corrected) on revision. If the author is poorly regarded, through a well earned bad reputation, then one can reasonably assume that omissions and ambiguities probably represent concealed (deliberately or otherwise) errors. It is very important that if the author is unknown to you, you not assume that the author is in the latter category.

A special case is when you are reviewing a proposal, as noted above. For a proposal, the reputation of the author should be as much as 75% and the proposal itself should be as little as 25%. The best predictor of future good work is a record of past good work, and conversely. A well written proposal by a researcher with a well deserved reputation for poor work should not be funded - it is very easy to make promises. It is again very important to not discriminate against people unknown to you - everyone has to start somewhere.

8.5. Confidentiality

It is the practice in Computer Science for the editor to transmit the referee reports verbatim (i.e. a photocopy) to the author, but without the referee's name, and without surrounding identifying information (e.g. institutional letterhead). If you don't want to be identified, you should not put identifying information in the text of your report.

Papers that are submitted for publication are not necessarily public. You should not use the material in a paper you have refereed, nor distribute copies of the paper, unless you have knowledge that the paper has indeed been made public, e.g. by being distributed as a technical report.

8.6. Conflicts of Interest

If you have a conflict of interest, you should make it known to the editor. If the conflict is severe, you should not referee the paper, but should instead return it to the editor. In particular, if you have a feud with an author, or a significant personal disagreement, it would be wise to send the paper back. If you are competing with the author for funding, and this is a proposal, you should make that known to the program officer.

The opposite type of conflict also occurs - you are being asked to referee a paper by a friend, colleague, former or current student, boss or subordinate, or former advisor. If you feel that you cannot provide an objective review, then you should return the paper to the editor.

8.7. Role of the Editor or Program Chairman

The editor has several tasks. (Here I am jointly referring to the editor in chief, who typically has the authority to decide whether to accept a paper, and the associate editor(s), who solicits the referee reports and recommends to the editor in chief whether to publish.) The editor receives the paper from the author and maintains correspondence with the author. The editor selects the referees, sends them each a copy of the paper with suitable instructions, and awaits their results. The editor *should* remind tardy referees, and find new referees after a certain period of no response has elapsed.

The editor should select referees that are knowledgeable in the subject matter of the paper, and can be relied upon to provide a fair and objective evaluation. Unfortunately, it is not always possible to do this - there are too many papers to be reviewed, and too few people known to be sufficiently expert and responsible. There is also another problem - by definition, people in area X believe that work in area X is worthwhile. A report received from someone in area X will evaluate the paper in area X by the standards of area X, but will seldom, if ever, say that work in area X is pointless and should be discontinued. It is, however, quite possible that such a response is appropriate. If you receive a paper to referee which is outside your area, you should consider whether it has been sent to you deliberately, and for that reason. Someone has to say that the emperor has no clothes.

After the editor has received a sufficient number of referee reports, the editor must decide whether to accept the paper, and if so, to what extent revisions are required. *The editor does not simply count the referee reports as votes.* The editor must read the referee report recommendations, and their reasons, and must decide, using his own judgement, whether to accept the paper. An editor, in theory, can overrule the unanimous recommendation of the referees; in practice, the editor can and sometimes does side with a minority of the referees. It is important that the referees state the reasons for their recommendations and justify them - those count as heavily or more heavily than the recommendations themselves.

The editor must also resolve conflicting recommendations, and tell the authors to what extent they must comply with the referee comments in making changes. A wise editor will transmit copies of all referee reports to all referees, both to educate the referees, and to be fair to the author in the case of conflicting reviews.

In the case of a conference, the program chairman (and his secretary) handles the mechanics of selecting referees and collecting and tallying their reports. Typically, the program committee, in a meeting or conference call, will decide which papers to accept by majority vote. The program chair may or may not have a vote that is larger than that of the others on the committee, but s/he seldom has the authority to accept or reject papers over the opposition of a majority of the program committee. Due to the large number of papers to be handled in a very short time, referees and authors are not usually given the personal attention provided by an editor who handles only one paper per month.

9. When You Are The Author

This article has been directed at the referee, but instructions to the referee are also instructions to the author. When starting research, when writing a paper, when finishing the paper, and when deciding where to submit it, ask yourself: how will this paper do when refereed according to the criteria given here?

Some specific things to think about are: Are you submitting the paper to the right place? Some journals and conferences will not consider material outside a specific scope; why waste 3-12 months to find out that your paper wasn't appropriate. Likewise, if you know that your paper is minor, why send it to a highly selective forum - send it somewhere where it has a reasonable chance of being accepted. If you suspect that further work is needed before publication, do that work; it may turn an unpublishable paper into a publishable one, without the 3-12 month extra delay. In general, a look at an issue of the publication to which you are considering submission will answer many of these questions; it is also helpful to look over the information provided by the journal to prospective authors; e.g. [CACM89, IEEE84].

Keep in mind that a good referee report is immensely valuable, even if it tears your paper apart. Consider - this report was gratuitously prepared by someone whose time you could not buy. All the errors they find, all the mistaken interpretations they make are things that you can correct before publication. Appreciate referee reports, and make use of them. Some folks become insulted, and ignore them; that is a waste of an invaluable resource.

An author receiving a negative referee report often suspects that the editor / program committee / program chair / referees are incompetent, biased, or otherwise unfair. While this sometimes happens, it is the exception; individual referee reports are often wrong, but a *set* of negative referee reports is a 90% accurate indicator that your paper has a problem, and needs to be either rewritten (or redone) before resubmission, or discarded as unpublishable or embarrassing.

Authors are particularly referred to [Levi83], [Day77], [Wegm86] and [Mano81], which provide discussions of how to write technical papers. Refereeing is also a good way to learn to write better papers; evaluating the work of others gives one insight into one's own.

10. Conclusions

Scientific progress rests on a foundation of peer review - evaluation of research for publication and funding by the researchers in the area, or by researchers qualified to evaluate the work. Good quality reviews - referee reports - are essential to this process, to separate the wheat from the chaff, and guidelines for refereeing have been presented.

Acknowledgements

I'd like to thank Peter Denning, Domenico Ferrari, Susan Graham, Anita Jones, Edward Lazowska, and Ken Sevcik for their comments on drafts of this article, particularly Anita Jones; some of her comments have been incorporated verbatim. That is not to say that these people all agree with this the material here; the opinions expressed are the author's own.

Bibliography

[CACM89] Communications of the ACM, ''Information for Authors'', CACM, 32, 3, March, 1989, pp. 411-414.

[Caro65] Lewis Caroll, Alice Through the Looking Glass, Chap. 4, Walrus and the Carpenter, Stanza 11, 1865.

[Day77] Robert Day, ''How to Write a Scientific Paper'', IEEE Transactions on Professional Communication'', PC20, June, 1977, pp. 32-37.

[Fors85] Bernard Forscher, ''Rules for Referees'', Science, October 15, 1965, pp. 319-321.

[IEEE84] IEEE, ''Guidelines for Authors'', IEEE Software, 1, 1, January, 1984, pp. 7-8.

[Levi83] Roy Levin and David Redell, ''An Evaluation of the Ninth SOSP Submissions'', ACM Operating Systems Review, 17, 3, July, 1983, pp. 35-40.

[Mano81] Frank Manola, ''How to Get Even with Database Conference Program Committees'', IEEE TC newsletter Database Engineering, 4, 1, September, 1981, pp. 30-36.

[Thom84] Keith Stewart Thompson, ''Marginalia / The Literature of Science'', American Scientist, 72, March-April, 1984, pp. 185-187.

[Wegm86] Mark N. Wegman, ''What It's Like to be a POPL Referee, or How to write an extended abstract so that it is more likely to be accepted'', Sigact News, 17, 4, Spring, 1986, pp. 50-51.