

How to Review HCI/Visualization Papers



My guidelines on writing good and helpful reviews for papers in the human-computer interaction and visualization research fields.

Editor's note: *I created this document in February 2009 as a guide for teaching my students how to review scientific papers. A recent discussion (December 2015) in the InfoVis community is challenging the field to improve its reviewing process; see for example [this post](#) by the UW IDL research group. In response, I just updated my guide and I am now considering to propose it as a reviewing guide for the upcoming IEEE InfoVis 2016 conference (for which I am papers co-chair with [Bongshin Lee](#) and [Kwan-Liu Ma](#)).*

For this reason, comments on this guide are welcome and even encouraged. I also have a vague suspicion that I have borrowed material from others without attributing them, but since I first wrote this as an internal document in February 2009, these sources have been lost. Please let me know if I am using any material without proper attribution (and if this is the case, please accept my apologies in advance). Also, please feel free to give additional links to other reviewing guides.

Note (December 20, 2015): *This is in no way an official document for the IEEE InfoVis conference. It represents my own personal opinion.*

How to Review HCI/Visualization Papers

Niklas Elmqvist, University of Maryland, College Park

Scientists communicate their results through scientific publications in workshops, conferences, and journals. Of course, not all results may be useful, correct, or even interesting. Therefore, science needs to determine whether a specific paper has a meaningful contribution to the field or not. However, because we all (supposedly, at least) work on the frontiers of knowledge, there is no higher authority to appeal to in order to judge the quality of a paper. That is why science uses peer review, where your colleagues (peers) are in charge of reading and evaluating your work in order to decide whether it should be published in a specific workshop, conference, or journal.

The concept of *peer review* means that the job of a scientist is not restricted to producing new results on his or her own, but also to review and evaluate results produced by others. Because a single person may have abnormal or anomalous opinions on a specific paper, editors and conference organizers usually enlist several reviewers for every submitted paper. Typically, high-quality publication venues such as prestigious conferences (in our fields, InfoVis, CHI, UIST, etc) and top journals (IEEE TVCG and ACM ToCHI) require at least three or more reviewers for each paper. The editor then makes his or her decision based on the cumulative scoring of all of the involved reviewers.

Remember this (shamelessly adapted from Winston Churchill and someone else):

Peer review is the absolute worst method for paper evaluation, except all other methods.

The imperfections of peer review aside, this document will offer some guidelines on good reviewing. I will discuss some basic approaches to how to review a paper efficiently, and I will also give some links to external reviewing guides I have found that come in useful.

A Responsibility, Not a Right

Because the scientific community is based on peer review, every active scientist needs to contribute and do their part by helping to peer review other people's papers. If you think about it, given that each paper that is submitted to a conference or a journal gets a minimum of three external reviews, a basic rule of thumb is that for every paper you submit, you should review three papers yourself.

Remember also that being assigned to be an external reviewer is an honor, and it shows that you enjoy the trust of the scientific community. It is the first step in becoming actively involved in your research domain, and if you continue doing a good job, this will be recognized by other, more senior, members of the community. Do well in your job as a reviewer, and sooner or later you will find yourself being assigned more responsibility, such as helping to organize conferences, serving on panels, being part of the technical program committee of a conference, or even serving on the editorial board of a journal.

Confidentiality

Being a reviewer for a paper is a job with a lot of responsibility, including being fair and objective. However, perhaps one of the most important of your responsibilities as a reviewer is to maintain confidentiality of your assigned papers. This means that you must do the following:

1. Never share ideas, concepts, or results from papers that you have reviewed with others.
2. Never use ideas, concepts, or results from papers that you have reviewed in your own work.
3. Never share the details or the results of the peer-review process with anyone else (including the authors, if you happen to know them).

The first two guidelines no longer apply once a paper has been accepted and has been published in the conference or journal (because by then, the authors' work is public knowledge). However, the third guideline is still in effect even after the paper has been accepted and published. Do not break the trust you have been given as a reviewer. (The only exception is for some special situations, like for CHI paper shepherding, where breaking your reviewer anonymity is often conducive to more effective communication with the authors.)

Conflicts of Interest

As already stated, reviewers have significant influence on whether a submission is accepted or rejected to a conference or journal. Such outcomes may make or break a career in certain

cases, which is why you should take this job very seriously. It also means that the reviewers for each submission must be fair and objective. However, if a reviewer has worked closely with one or several of the authors for a submission, they may not be expected to act impartially. More importantly, even if the reviewer does act impartially, they may be *perceived* not to be impartial by others.

Having a deep personal or professional relation with, being related to, or working at the same institution as one or several of the authors of a submission is called having a *conflict of interest* with the submission. The [VGTC ethics guidelines](#) give careful rules on what constitutes a conflict of interest or not, but the rule of thumb is simple: if you do not feel that you can review a paper impartially (or if someone might *perceive* that you cannot review it impartially), you have a conflict of interest. This includes both positive and negative relations; perhaps you have a strong personal relationship with one of the authors, or perhaps one of the authors slighted you in the past. Avoid having to fight your own bias.

Roles in the Review Process

There are several roles for people involved in the review process; the exact ones depends on whether you are submitting to a conference/workshop or a journal, but they can be summarized as follows:

- **Papers chair** (conference)/**Editor-in-Chief** (journal): The papers chair (conference) or EIC (journal) is ultimately responsible for accepting and rejecting papers, but often only makes final decisions based on recommendations by the program committee (conferences) or editorial board (journals). This person (or persons) also often appoints the members of the committee or editorial board.
- **Program committee member** (conference)/**Associate Editor**(journal): The PC member (conference) or AE (journal) is directly responsible for reviewing a specific submission, and typically recruits the external reviewers for the paper. Many conferences involve two PC members for each submission; they are then often called “primary reviewer” and “secondary reviewer”. CHI and related conferences tend to use the name “associate chair” (AC) for its PC members, and 1AC and 2AC for primary and secondary reviewers, respectively. Depending on the conference, a PC member/AE/AC may or may not write their own review. Regardless, it is this person’s job to remind reviewers, engage them in discussion, and make a final recommendation on the paper.
- **External reviewers** (both): Also called a “tertiary reviewer” (cf. primary and secondary reviewer), the external reviewer is one of the typically two to three persons who reads the

paper and writes an individual review. This is typically the first role you will be asked to serve in as you become part of a research area. Don't be fooled by the fact that there are several external reviewers to each paper; each external reviewer is extremely important and plays a significant role in getting a paper accepted or rejected. Put differently, most journals or conferences make it very difficult for a program committee member or associate editor (or even a papers chair) to make a decision outside the mandate given to them by the external reviewers. In other words, your work in this position is vital.

Critical Reading

Learning to read critically is of key importance when performing a review, but it is also the best way to learn how to write well—if you are able to turn the same critical eye to your own work, you will be much better at fixing potential problems with your own work before you submit it.

When reviewing a paper, I find it almost always necessary to print out the paper so that you can make notes and comments on the document as you read it. Reviews are also something you can work on when you do not have access to your computer (either because you cannot use a computer, or because you don't want to), such as on the plane, on the bus, or in your favorite couch, so having a printout makes sense. It also helps if you currently lack internet connectivity.

When reading a paper for a review, plan on reading the entire paper at least once, probably twice (especially if you are new to reviewing). If you have not read the full paper, you are not doing your job. In addition, you will most likely have to re-read sections of the paper (such as the technical description) more than twice to make sure that you understand all aspects. Do not start writing your review before you have finished this initial read phase (do make notes, however).

Bear in mind that reading critically also means getting your priorities right. Some errors in a paper are of minor importance to the general state of the work, whereas other errors are not so minor. It is important to make this distinction and not punish an otherwise good paper with sound and clear ideas for, e.g., typographical errors or small methodological flaws in the evaluation. It is possible, even easy, to be too critical. Most reviewers read papers looking for reasons to reject the paper. This is not necessarily good, because this approach may very well punish good but non-traditional ideas that have a potential for revolutionizing a field. Try to look for reasons to accept a paper instead.

Another often occurring phenomenon for reviewers (especially new reviewers) is to get a feeling of jealousy when reading a paper, especially one that introduces a novel idea that is conceptually simple yet has not been proposed in the past. The tendency in these cases is to strike down on minor problems in the paper or to make sweeping statements (such as “this is a trivial contribution that surely must have been done by someone in the past”; never write this without finding an explicit reference). After a few times of getting back reviews of this kind yourself for your own work, you will start to realize how frustrating this is for the authors. Remember to be as fair as you can be, and put yourself in the authors’ position—if you received the review you are writing for a paper of your own, would you be upset and feel unfairly treated? This is the litmus test for a good review, regardless of whether it is critical or supportive of a paper.

The Review Form

When you receive a paper to review, you typically also get a review form to fill out as you conduct your review. Review forms look different depending on the conference or journal; use the specific format given to you by the organizers. These days, conferences and journals alike tend to have web-based submission systems, so filling out a review form is usually done in your web browser. I always recommend writing your review on your own local hard disk and then uploading the finished form when you are done. This is so that you have your own copy of the review (for potential issues down the road), as well as to guard against issues with your web browser crashing and losing all your typed text, or even the submission system losing your review (it can happen).

Giving Ratings

Most review forms ask you to give numerical values on different categories, including the technical contribution, the significance, the originality, and the technical soundness. Often you also get to give an overall rating, which is typically the most important measure and states whether you think the paper should be accepted or not. Ratings have different scales depending on the conference, but a typical 5-scale conference rating has the following format:

- Strong accept – you strongly support this paper being accepted for the conference.
- Accept – you support this paper being accepted.
- Borderline/neutral – could go either way (try to avoid this rating because it does not help the editor or coordinator making the decision on the paper’s fate).

- Weak reject – this paper should not be accepted, but there is a small chance it could get in.
- Strong reject – this paper should absolutely not be accepted.

For journals, the ratings typically have the following format:

- Accept with no changes – the paper should be accepted as is.
- Accept with minor changes – the paper needs some small changes, but after these have been conducted, the paper can often be accepted for publication with no additional review cycle.
- Major revisions needed – the paper needs major revisions, so a new review cycle is needed.
- Revise and resubmit as new – there are major problems with the paper; should the authors decide to submit a new (hopefully totally rewritten) version of the paper, it will be treated as a new submission.
- Reject – the paper will have no chance of being accepted for publication.

Expertise

In addition to ratings on the paper, you are often also asked to give a self-evaluation of your expertise in the topic of the paper. This is your chance to tell the committee how confident you are in your review. For example, if you are an expert and know most of the literature in a specific field, your rating will carry more weight than if you only have passing knowledge about the field. Be fair to yourself and to the authors by trying to be as accurate as possible here.

A common mistake among junior reviewers is that they rate themselves too highly in their expertise. While it may certainly feel like you are an expert in a certain domain if you just submitted a paper and you have surveyed the literature, this may not be totally accurate. An expert is someone who has published papers, preferably several times, in the particular domain. If this is true for you, go ahead and mark yourself as an expert, but if not, you and the authors may be better served if you are a little more humble.

Finally, another issue to consider is whether you have sufficient expertise to even review a given paper. Try to avoid accepting reviews for papers that you are only passingly familiar with. In other words, if you feel yourself using the expertise “passing knowledge”, you may have overextended yourself. Do your best to spot these situations when you are invited to

review a paper so that you don't have to go back and decline with only a week or a couple of days left until the deadline.

Writing the Actual Review

Most, if not all, review forms include a free-text area where you can write your review (sometimes several such areas where you are asked to answer different aspects of the paper). This is perhaps the most important part of the review, because a review with only numerical ratings (even if they are high and supportive of the paper) is worthless if they are not backed up by text that justifies the ratings.

It is difficult to give an exact amount of text needed in a proper review, and it depends on the type of paper, but a rule of thumb is that you want at least three or four paragraphs of text. Very good or very bad papers sometimes warrant less than that if it is clearly obvious which category the paper belongs to. Borderline cases usually require the most text, especially if you have to develop a specific argument in favor of or against the paper.

There are several ways to structure your free-text review. You can decide which format you want to use yourself, but I often base my reviews on the CHI review format, which looks like this:

- **Summary.** Summarize the contribution of the paper in a single sentence. Follow up with a single sentence describing the potential benefit. Use your own words to show that you have understood the message of the paper. Sometimes a paper does less than what the authors claim—this is the place to take them down to earth by stating the contribution of the paper in your own words. However, try to be neutral here, and don't let your opinions shine through (too much).
- **Verdict.** After the neutral summary, I try to give my general feelings about the paper. I summarize my review in two or three sentences, describing both positive and negative points. This paragraph is a roadmap for the rest of the review, and you often end up writing it last. Often, you don't really know how you feel about a paper until you have finished the review proper and write the verdict paragraph.
- **Related work.** Discuss all of the specific related work that is relevant to this paper (even if the authors have cited this in their own work). List all of the important works (give complete references if you can, or at least author and year) on the topic. Try to resist the temptation to only reference your own work; the resulting literature survey not only becomes relatively one-sided, it also makes it easy for the paper authors to figure out your identity!

- **Originality.** Often related to the previous paragraph, this is your comment on how original this idea is. Is it an entirely new approach to a problem hitherto unseen in the literature, or is this incremental work building on existing knowledge in the domain? Entirely original ideas are naturally better, but solid incremental work also has its place in research. Sometimes this paragraph is combined with the related work.
- **Significance.** How important is the contribution over the existing work in the literature? What are the potential benefits to science and to society? Not all problems are worth solving and not all questions are worth answering. Discuss what this paper actually contributes to the research field.
- **Validity.** This section discusses the methodology and the approach taken to the research, and the corresponding confidence other researchers and practitioners can have in the results. This is the part where you evaluate and discuss the validation mechanisms the authors used to prove the significance of their work. In case they provide an implementation, discuss the feeling you have about the implementation (feelings you get from reading the technical description as well as watching any video). For a theoretical argument, analyze and discuss the internal and external validity of their argument. For a user study, discuss whether the study was conducted competently, if the tasks were ecologically valid, if the data was analyzed correctly, etc.
- **Readability.** I usually offer some comments on the state of the manuscript in terms of its grammatical correctness, its structure, and how easily it reads. Note that you are allowed to reject a paper due to poor writing or layout. Sad as it may sound, science is as much about communicating your results as it is about getting results in the first place. Readers are generally busy and pressed for time, so the very least we can expect is that writers invest the necessary time to make the reader's job as easy as possible.
- **Direct feedback.** I often find myself spotting errors, typos, and grammatical problems as well as issues with layout, figures, and tables when I am reading a paper. As a favor to the author, you can always provide a bullet list of direct comments on the paper at the end of your review.

Comments to the Committee

Sometimes you will find yourself having things to say about a paper that you don't want to divulge to the authors. Examples could be the following:

- You feel strongly about a paper but do not want to offend the authors.

- You were involved in a previous peer review of this paper and have some knowledge about how the paper has been changed since then. (Note that sometimes you may want to divulge this information to the authors in the review body itself; this is usually fine.)
- You suspect foul play (plagiarism, bogus data, etc) in the paper.
- You want to communicate specific pieces of information missing in the paper that could reveal your own identity (such as citations to your own work).

In all of these cases, you can use the specific field in the review form for comments that go only to the committee and are not divulged to the authors. In the example of a missing citation to your own work, you could for instance ask that the PC member in charge of the paper adds your citation to his/her meta-review so that it comes from a more neutral source that is harder to trace.

Bear in mind, however, that sometimes your comments to the committee are not hidden from other reviewers. Never write anything in the private comments that you would not be willing to stand for if they were (accidentally) made public.

Engage in the Discussion

Many paper reviewing systems, particularly for conferences, include mechanisms for internal discussions between reviewers. Sometimes you will be asked by the primary or secondary reviewer to give comments on the paper after reading other reviews. This is often done in lieu of a physical program committee meeting, where reviewers can otherwise discuss in the real world about aspects of a specific paper (even at a physical meeting, only the primary and secondary reviewers tend to be invited). Remember to take this opportunity because it helps the primary reviewers greatly in making a final decision. Avoid the temptation to be inactive, and instead use the discussion phase to give more feedback on the quality of the paper.

Writing a Good Review for a Good Paper

It is often very easy to write a review for a bad paper, because you will have a long list of complaints that you can easily put into words. It will not be hard to fill your 4-5 paragraphs, and your biggest challenge will be to be as polite as possible while still criticizing the work of the authors. However, it is often much more challenging to write a good review for a good paper.

If you feel strongly positive about a paper, the biggest mistake you can ever make is to give a high rating and then write a very short review. It often feels difficult to write a long a positive

review, because the paper basically did everything you expected from it, right? There is nothing more to add. However, as a meta-reviewer or program committee member, positive but short reviews are probably the least useful after rude and negative ones, because they basically give you no ammunition on how you can argue for this paper to be accepted. In other words, it is difficult to justify a paper being accepted if all your reviewers give very short reviews with no details.

To put this another way, it is equally important that you give many and good reasons for a paper to be accepted as it is for when you are recommending rejection. A verbose and lucid description of all of the strong points of the paper will help convince the powers that be that you know what you are talking about as a reviewer. A short review does nobody any favors, especially authors of papers you want to accept. Never forget that.

Finally, even good papers have points of improvement. It is easy to fall in to the temptation of merely filling the review body with a long list of weaknesses that should be improved. This may cause the editors or program committee members to think that you gave a high rating by mistake. Be sure to also include language that speaks to the strengths of the paper.

Conclusion

Reviewing is complex business, and there is a lot more to it than can be covered in basic guide like this. Be sure to read the below guides that other people have written on this topic to get more than my viewpoint. Furthermore, the best way to learn how to review is to actually review papers; while this can sound counter-intuitive, one way to get started is to read already published papers with a critical eye and write a review for them. As soon as you feel ready, you should also volunteer to review for conferences like CHI, UIST, and InfoVis; the reviewing system often allows even new researchers to volunteer for reviewing service. With some practice, you will begin to learn your own method, style, and voice. Good luck!

External Guides

- [CHI Reviewing Guide](#) – the best reference for how to review for CHI
- [Guide to Reviewing CHI Papers and Notes](#) – another guide for CHI reviewing
- [IEEE VGTC Reviewing Ethics](#) – ethics guidelines for IEEE VGTC conferences
- [The Task of the Referee](#) – by Alan Jay Smith (from Silvia Miksch's website)
- [A Guide for the New Referee in Theoretical Computer Science](#) – by Ian Parberry (from Silvia Miksch's website)