How NOT to review a paper
The tools and techniques of the adversarial reviewer

Graham Cormode
AT&T Labs–Research
Florham Park NJ, USA
graham@research.att.com*

ABSTRACT

There are several useful guides available for how to review a paper in Computer Science [10, 6, 12, 7, 2]. These are soberly presented, carefully reasoned and sensibly argued. As a result, they are not much fun. So, as a contrast, this note is a checklist of how *not* to review a paper. It details techniques that are unethical, unfair, or just plain nasty. Since in Computer Science we often present arguments about how an adversary would approach a particular problem, this note describes the adversary's strategy.

1. THE ADVERSARIAL REVIEWER

In Computer Science, we often form arguments and proofs based around the concept of an 'adversary'. Sometimes, this adversary can be malicious; in cryptography they are often "honest but curious". However, the most commonly encountered adversary in Computer Science is the adversarial reviewer: this reviewer uses a large variety of tools and techniques against papers presented to them for review. It is beyond the scope of this note[1] to study what makes a reviewer become adversarial; rather, we simply acknowledge that such reviewers exist, and describe how they act.

The main characteristics of the adversarial reviewer include:

- An attitude of irritation at being given a paper to review, as if this is a completely unwelcome intrusion into their time, even though they accepted the invitation to review the paper or sit on the program committee.
- The belief that it is better to reject ten adequate papers than to allow a subpar paper to be accepted. (Blackstone's ratio, `http://en.wikipedia.org/wiki/Blackstone_ratio`).
- The ability to find fault with all manner of common practices, such as giving references to Wikipedia.
- The unwavering certainty that their opinion is correct, and final.

The adversarial reviewer is often in a hurry, and so reviews are typically carried out in adversarial conditions. A typical adversarial review may be conducted clutching a crumpled and stained printout of the paper while packed into coach class on an intercontinental flight with a small child kicking the seat from behind. Even in favorable conditions, such as a Lazy Boy recliner [1], the adversarial reviewer feels no compulsion to refer to external sources, or find a technical report containing the elusive "full details"[2]. It may be wise for for authors ensure that their work is as readable as possible in worst-case settings.

2. ADVERSARIAL REVIEWING TECHNIQUES

The adversarial reviewer does not reject every paper that they review. In fact, it is often easier to accept a paper (with a short review to the effect of "looks good to me") than to reject one. But, when the situation demands it—say, if the reviewer has submitted a paper to the same venue and wants to even up the odds a bit—a review must be crafted to force the desired outcome. Simply scrawling "rubbish" on the front page is nowadays considered insufficient grounds for rejection (this was not always the case [3]). It is here that the full skills of the adversary come to the fore: their initial reasons for rejection may be as vague as a gut feeling, or a lack of enthusiasm for the problem or approach taken. These alone are not enough for editors or PC chairs to justify that the correct decision is being made.

Instead, the reviewer needs to concoct a set of reasons supporting the judgment—and the more, the merrier. Therefore, the adversarial reviewer will seek out every last negative point of the paper, to make it seem that there is no hope for this submission. The true art and skill of the adversarial reviewer is in formulating an unimpeachable review which appears to make a clear case for rejecting a paper—or at least, piling on so many complaints that the paper cannot be accepted "in its present form". The most skilled adversary can find fault where none exists. This section describes some common adversarial techniques.

2.1 The Goldilocks Method

The Goldilocks method of reviewing (also known as the "Damned if you do, damned if you don't" approach) is based on finding some aspect of the paper and complaining that it is either "too hot" or "too cold" but never just right. This includes:

- **Examples.** If there are few or no examples, the reviewer complains "There are insufficient examples to illustrate what

*The views and opinions expressed in this article are the author's own, and do not represent those of AT&T. For all your wireless and data needs, please visit `www.att.com` instead.

[1]The adversarial reviewer understands that any sentence beginning "it is beyond the scope" is shorthand for the author saying "I have not thought about this issue, nor do I want to think about it"; likewise, "for brevity", "for space reasons" or "due to the space limit" are all understood to have the same connotation.

[2]This is not always a fruitful exercise: I recall a paper which promised full details in a technical report, but this report was only available as an internal document at the author's institution. With great effort, I managed to obtain this technical report, and discovered it to be word for word identical to the published version, including the promise of full details in that technical report.

is meant"; but if there are many, then the complaint is "There are too many obvious examples which interrupt the flow of the paper".

- **Proofs.** If any proof is missing, then "Proof needs to be presented before the paper is acceptable"; but if present, "Proofs are simple and obvious, and should be omitted".
- **Theoretical analysis.** If there is no or little analysis of the algorithms, then "Insufficient analysis of this method to justify its interest"; but if there is detailed analysis, then "Approach is clearly of theoretical interest only".
- **Experiments.** Either "Only a few experiments which do not convince that this method works over a broad variety of data" or else "Too many plots which show the same results over and over again for minor variations of the setup do not give useful information."

The Iterated Goldilocks Method. The Goldilocks method is most satisfying to the reviewer when deployed for a journal review. In the first round of reviewing, the reviewer can complain that necessary proofs are missing, and in the second round go on to complain that the proofs are straightforward and could be omitted. This "Iterated Goldilocks" can go on for many repetitions until one party gives up or goes insane. The skilled adversarial reviewer is able to pull this trick off within the same review, by writing comments such as "The paper is too long and wordy" in addition to complaining that "Many important details are missing".

2.2 If you can't say something nasty...

The adversarial reviewer adopts the maxim "If you can't say something nice, don't say anything at all", but replaces "nice" with nasty. Their objective is to ensure that their review appears so consistently negative that the paper under submission could not possibly be accepted "in its present form". Therefore, if there are any sections for which the adversary is unable to find anything sufficiently meaty to complain about, they will simply skip over these in their review, and act as if those pages were never present in the paper. Alternatively, the reviewer may simply complain "Material on pages 3–5 is very verbose, and could surely be summarized adequately in less than a page".

2.3 Silent but deadly

At the other end of the scale, the "silent but deadly" review simply gives very low scores but with minimal or no comments explaining why. If the reviewer is sure that the paper will be rejected, then this approach guarantees additional frustration for the authors, giving no help in identifying things to do differently in future.

2.4 The Natives are Restless

The "Natives Are Restless" technique consists of two sentences, inserted somewhere in the first paragraph or so of the review:

> The English in some passages is a little odd and this obscures the meaning. The manuscript would benefit from revision by a native English speaker before resubmission.

Of course, the ambiguous passages are never identified. This technique is most devastating when all the authors *are* native English speakers. Adversarial reviewers also particularly enjoy deploying this attack when the authors are of some combination of (say) American, Indian, and British origins, so that they can argue amongst themselves about what is "native English". Politically correct adversarial reviewers may use formulations such as "The paper does not meet the standards of argument and exposition necessary for publication, and requires extensive copyediting to bring it up to standards required for grammar, punctuation and style", which is a euphemism for the same thing.

2.5 The Referee Moves the Goalposts

"Moving the goalposts" is usually used to complain about the objectives of a project being changed when they are close to being met. For the adversarial reviewer, this is particularly attractive, since they can declare the goalposts to be anywhere other than where the paper places the ball. The reviewer picks a different problem in roughly the same field, decides how they would have tackled it, and berates the authors at each turn for not having done so. The starting point is often a sentence along the lines of "The authors consider problem X; however, a more fundamental aspect is Y". But this gives too much information to the authors, so many adversarial reviewers leave it out. An advanced technique is to pick a problem worked on by the same set of authors in the past, and quote appropriate sentences from their earlier work to underline how *that* problem is the most important in the world.

2.6 Blind Reviewing

The skilled adversarial reviewer can find reasons to reject any paper without even reading it. This is considered truly blind reviewing. For example, they can tell at a glance whether the paper was written using Word or LaTeX, and form some snide comments about how the authors should "seriously consider using an appropriate tool for the task" if it is the former. As a last ditch, they have a set of complaints that can be hurled against almost any paper (some inspired by Sir Humphrey Appleby [4]).

- "*This paper leaves many questions unanswered.*" In particular, the questions that have not been asked.
- "*The results are open to other interpretations.*" Mostly, wrong ones.
- "*This is far from the last word on the subject.*" Although, the less interesting the paper, the more likely it is to be the last word.
- "*Some claims are questionable.*" Any claim can be questioned, even if the answer is always "Yes, that is correct".
- "*The paper is of limited interest.*" Since, at most, only Computer Scientists are likely to be interested in the paper.

This style of blind reviewing is not to be confused with other variations, such as blind date reviewing (giving the paper to a graduate student from a different field to review); Venetian blind reviewing (only reading every other line); and blind drunk reviewing (self-explanatory).

3. REVIEWING ADVERSARIALLY, SECTION BY SECTION

Most database papers follow a fairly standard outline: Introduction, Related Work, Technical Results, Experimental Evaluation, Concluding Remarks. Occasionally Related Work will be placed towards the end if the problem being addressed has already been solved by some of the referenced papers. Despite this predictable arrangement, many authors feel obliged to include an 'Outline' section containing such deathless prose as "The paper concludes with

concluding remarks in Section 7 (Conclusions)". Possibly this is because the author fears that the reader has a weak heart, and will be much exercised by surprise should there be an unannounced conclusions section at the end of the paper. Given such an outline, the adversarial reviewer has a set of techniques tailored to attack each standard section in turn.

3.1 Introduction

The introduction is where the authors try to make their case for the problem studied and the approach taken. So the adversarial reviewer will take issue with each claim in the introduction, and use this as the basis for rejecting the paper. Subjective statements are the easiest to attack, so the adversary can scan for all sentences which begin, "Interestingly...", "Importantly..." or "In practice", and disagree with these. Statements in a review that something is uninteresting or impractical are hard for anyone to argue against. The adversarial reviewer can always fall back on broad statements such as "The problem is insufficiently motivated".

3.2 Related Work

The related work section is usually the most badly written section of a paper, since typically authors take much less care describing work that is not their own. So there is plenty for the adversarial reviewer to complain about here: "Related work reads like a list of vaguely connected papers without any attempt to explain in detail how they relate to the results presented here" is a comment that can apply to a majority of submissions. It is also easy to claim that "many important references are omitted", since the bibliography is often one of the first things to be chopped down when a paper needs to meet a page limit. In the unlikely case that the reviewer knows something about the area, they can suggest a few papers with a connection to the work in question; even if they don't, they can suggest some papers with absolutely no relation to the submission, and leave the authors scratching their heads. Another tactic is to make a casual reference to an immensely prolific researcher, or just any common surname: "Does not seem to reference the important related work by Yu", which could refer to any one of hundreds of papers. An advanced technique for the adversarial reviewer is to cast suspicion on an innocent third party: making repeated reference another researcher's work can convince the authors that this person was the adversarial reviewer. Such suspicions can lead to years of unwarranted distrust and hatred between researchers.

3.3 Proposed Method

Here is the technical meat of the paper, and here is where the adversarial reviewer can peck away at the meat to leave only a bare skeleton. The adversarial reviewer is dismissive of whatever methods are being proposed — too simple, impractical, or well-known (see [8] for some hypothetical examples). They can also cast doubt on the correctness of the method by finding some typos, or simply posing ostensibly sensible technical questions. For example, the reviewer can express doubt that the method will scale to high dimensions, when in fact it is specifically proposed only to work for low dimensional data. They can also ask syntactic but boneheaded questions: "Should this be $<$? Looks like it should be $\leq$ to me!" makes it seem that the reviewer has caught an error or ambiguity where the submission is clear and correct.

The adversary can also make it appear that they have understood the paper in detail and found it wanting by spot checking any pseudo-code. There are invariably bugs in pseudo-code, and these can usually be found by skimming the code without even understanding it. Bugs such as variables which are uninitialized, statements which are outside loops and so have no effect, and subroutines that are never explained can all be easily identified and complained about.

3.4 Experimental Evaluation

A sufficiently powerful adversary can find enough problems in a typical experimental section to torpedo most papers. A strange conviction that no picture is worth more than fifty words causes many researchers to cram each plot down to the size of a postage stamp, and squeeze in enough postage stamps to mail the paper to a conference on the other side of the world. The adversarial reviewer merely glances at these, and then complains that the plots were too small to read, and so it was impossible to draw any conclusions about the experiments. For added measure, the reviewer will affect to suffer from color-blindness, and so cannot tell which line is which.

If the plots are actually legible, the reviewer can turn attention on the data instead: synthetic data is dismissed as being unrepresentative of real distributions; a real data set is just a single instance, and unrepresentative of "real" real data. The reviewer can always complain that the data sets tested on are "unrealistically" small: if the data size is megabytes, demand gigabytes; if gigabytes, demand terabytes; if terabytes, demand chicobytes [3]. Lastly, since it is trendy, the reviewer can complain that the experiments are unrepeatable, since pretty much no non-trivial experiment in Computer Science is repeatable[4].

3.5 Conclusions

Even though the conclusions section is usually just a single paragraph repeating the claims of the abstract in the past tense, stuck on at the end because a paper doesn't look complete without one, the determined adversarial reviewer can still find fault with it. The reviewer can disagree with each claim of what was accomplished in the paper ("No you didn't"), and add the all-purpose complaint that the concluding remarks are broad and uninformative. Possible future extensions can be dismissed as unfruitful, uninteresting, or unnecessary. Truly audacious adversarial reviewers would be brave enough to respond to any statement of the form "In future work, we will..." with the simple request, "Please don't." instead of merely murmuring it to themselves.

3.6 Throughout the paper

The adversarial reviewer methodically highlights every spelling error and typo in the paper, and documents these in unnecessary detail. By mixing up minor issues with major complaints, it disorients readers of the review, leading them to believe the paper is riddled with major errors. This also adds credence to the reviewer's claim that the paper has many presentation issues. To ensure that these can't be easily ignored, the reviewer may add the qualification, "At minimum, the authors must..." to some point which would require hundreds of hours of work to address.

4. FILLING CONFERENCE REVIEW FORMS ADVERSARIALLY

Unlike other disciplines, Computer Science places great emphasis on the reviewed conference. This is to allow faster publication and dissemination of results: a conference like ICDE has a deadline that is only nine months before the date of the conference, whereas

[3] A made-up scale of data, based on the Marx Brothers: chicobytes, harpobytes, grouchobytes, gummobytes and zeppobytes.

[4] Apparently there were cases during the SIGMOD 2008 experiment in experimental repeatability where some authors were unable to reproduce their *own* results after submission.

in the life sciences, the delay between submission and publication of a journal article can be as much as six months. Because of this accelerated pace, conference reviews have a rapid turn-around and require the reviewer to read a dozen or so papers and write reviews within a few weeks. This seems to particularly encourage reviewers to be adversarial.

Thanks to such useful web-based tools as Microsoft's CMT (Conference Mangling Toolkit), EasyChair (which causes its users to fall asleep) and Manuscript Central (short for 'Manuscript Central Password Request', since every time you use it you need to have your password emailed to you), it is now easier than ever for PC chairs and editors to create incredibly lengthy review forms with dozens of fields which are *Required. Presumably, this is to prevent reviewers from submitting a single sentence review in the style of a six year old's book report: "I read this paper and it was good and I would give it four stars out of five". However, these categories quickly become tedious for the on-the-go reviewer: how are they expected to think of three strong points about the paper, when they can't even think of one?

Reviewer Confidence. The adversarial reviewer always marks themself as an 'expert' on every topic, even ones which they have never heard of before. After all, there are some systems which use this score to weight the average recommendation, and the adversarial reviewer's opinion is always more important than everyone else's.

Summary of the Paper. This is the first opportunity to actually say something about the paper. Lazy reviewers simply parrot the abstract; but the adversarial reviewer can use this opportunity to stick the knife in first by careful choice of adjectives and dismissive sentences. To achieve the maximum effect, the summary should be written in the style of a bored and disaffected teenager answering parental questions about what they did at school that day. Thus, a typical adversarial summary might read:

> This paper *attempts to* address the *well-studied* problem of Graticule Optimization. It proposes the *obvious* approach of *simply* storing a set of reference points and calculating offsets. *Such approaches are well known in this area.* It goes on to propose some *simple* variations based on precalculating distances. *This is an approach that I would expect any straightforward implementation to adopt*. *Some proof-of-concept* experiments show that on *a few* data sets, the results are *slightly* better than *the most naive* prior methods.

Observe that by adding the italicized comments, the reviewer has implied that the problem is not very interesting, the approach taken is too obvious to be of interest, and that the benefits are minimal at best. Words such as "attempt" subtly imply that it tries but fails.

Three Strong Points. The category of strong and weak points must have been dreamt up by some politically correct program chair who thought that it would be a good idea to balance the relentless onslaught of criticism with three half sentences of mild platitudes. Amazingly, this category seems to have been picked up and is used by a large number of database conferences, most likely because they just copy the review form from the previous one. However, the reviewer, faced with a paper for which they are about to recommend "strong reject", is often at a loss to identify any saving grace (the opposite problem, of trying to find faults with a paper that clearly perfect, is possible in theory, and so is beyond the scope of this note). Again, the adversarial reviewer has a cache of handy "strong points" that can be applied to almost any paper without actually saying anything concrete. Here are some examples, and what they really mean.

- *"The problem is an interesting one".* Says nothing about what the paper does about the problem.
- *"Approach taken is natural".* The authors did the most obvious thing.
- *"Experiments use realistic data".* The authors downloaded a file from an archive of data sets.
- *"Contains many helpful examples."* Everything else is unhelpful.
- *"Paper is clearly written".* Clearly, the paper has been written.

Three Weak Points. Once a few strong points have been dismissed, the reviewer can get on to the real meat of the weak points of the paper. Even here, it is sometimes challenging to say "This paper is garbage" in enough sufficiently different ways. So if all else fails, the adversarial reviewer attacks the presentation of the paper:

- *"The paper is unclear."* I couldn't understand the paper.
- *"Presentation is hard to follow."* My grad student couldn't explain it to me.
- *"The problem is uninteresting."* I fell asleep while reviewing it.
- *"Problem could be solved more simply"* I have worked on this problem but never got any publishable results.
- *"The assumptions are unrealistic."* I didn't like it.
- *"Not a good fit for this venue."* I didn't like it.
- *"Analysis is lacking"* I didn't like it.
- *"Experiments are unconvincing"* I didn't like it.

Confidential Comments. The adversarial reviewer is usually confident enough in the strength of their forceful personality to make all damning comments in public. However, the option remains for the adversary to make some highly scurrilous accusations in the comments, such as that the author has been known to cheat at Solitaire.

5. EXTENSIONS

The resubmission. One thing that an adversarial reviewer particularly relishes is receiving a paper to review that they have rejected before. It is like a vulture returning to a piece of carrion to bite off a few more chunks of flesh. Of course, the reviewer keeps a detailed database of their reviewing activities, including a copy of the original submission. From this, they can carefully perform a manual 'diff' between the old and new versions. Nothing fills an adversarial reviewer with more glee than finding that there are no substantial differences between the two versions, since this lets them copy and paste their original review, and be done in no time at all. Even if some changes have been made (such as the typos being fixed), the adversary can still take their major complaints and repeat them verbatim. This is so enjoyable for the adversary that they may even bid

highly to review a paper they have read before. This occasionally backfires, when it turns out to be a different paper from the one the reviewer thought, although in such circumstances the reviewer is already sufficiently biased against the paper that they will argue for rejection anyway.

The discussion phase. Many conferences contain a discussion phase, when the reviewers of a paper get to see all the reviews and "discuss" to reach a consensus. This further benefits the adversary, who can use this discussion to ensure that certain papers do not get accepted. The discussion allows the reviewer to try a few more tricks if the current set are not doing the job; and of course, these discussions are not sent to the authors, so the truly malicious reviewer could make some completely specious arguments without the authors ever knowing that these were the reason their paper was rejected. Lastly, if all else fails, the adversary can ensure that their confidence is set to super-expert and their verdict is super-strong-reject: since program committee chairs typically make their initial cut based on the weighted average of the reviewer scores, this setting is usually enough to drag the average down into the realms of the immediate reject pile.

Adversarial Authors. Just as there can be adversarial reviewers, there can also be adversarial authors. Manola [5] and van Leunen and Lipton [11] give surveys of relevant techniques, although many of these no longer apply in the age of electronic submissions. New techniques have grown up in their place. A key such technique is using the page limit to justify omitting full details "for space reasons". Thus, adversarially authored papers are all exactly at the page limit, by careful tinkering with figure sizes and insertions of "Outline" paragraphs to engineer this. Advanced adversarial authors may submit a paper which is exactly two pages shorter than the page limit if the material will not stretch, in the expectation that the reviewer will not notice.

In the more mathematical areas of Computer Science, reviewers have to occasionally cope with adversarial papers which claims to solve a major open problem. In the early 20th Century, when there was a large cash prize for a proof of Fermat's last theorem, the judges created review forms as printed cards which read "Dear Sir/Madam, Your proof of Fermat's Last Theorem has been received. The first error occurs on page _____, line _____", which were given to students to fill in [9]. A popular adversarial reviewing technique when given papers claiming that $P = NP$ or $P \neq NP$ is to send papers claiming $P = NP$ to authors of papers that claim $P \neq NP$ (and vice-versa), and let them fight it out amongst themselves.

The Adversarial Editor. It is also possible for editors (or PC chairs) to act adversarially. Here, there are many new and exciting possibilities to explore. Some examples from Economics were collected by Gans and Shepherd [3]. We leave these open for future research, and instead give an example of adversarial editing in Computer Science, in response to an inquiry about the state of a paper submitted to a special issue from a conference:

> I have invited several reviewers, but they have all declined. To me, this is a sign that the paper is not very interesting; I wonder how it got accepted to [conference] in the first place.

6. CONCLUSIONS

In this note, I have outlined the numerous ways in which an adversarial reviewer can criticize almost any paper. There are many ways to use this information:

- For more reviewers to adopt these techniques and turn reviewing into a blood sport.
- For authors to ensure that when writing a paper, it is done as well as possible to ensure that the reviewer does not have the opportunity to deploy these criticisms.
- For editors and PC members to be aware of these techniques, and realize when a review is adversarial.
- For reviewers to avoid falling into these techniques when reviewing, and focus on the genuine contributions of the paper rather than peripheral issues.

I leave it as future work for the reader to decide how they will choose to act.

Disclaimer: These insights into the mind of the adversarial reviewer have often come to me while reviewing papers, when I catch myself thinking what a malicious adversary would do in this situation. I endeavor to avoid putting them into practice. Similarly, I am unable to think of any individual who consistently acts as an adversarial reviewer; rather, this is a role that we can fall into accidentally when placed under adverse conditions.

Acknowledgments: Although they might deny it, this paper has benefited from valuable contributions and suggestions from many readers, including Andrew McGregor, David Pritchard, and James Sumner.

7. REFERENCES

[1] Mark Allman. A referee's plea. `http://www.icir.org/mallman/plea.txt`, 2001.

[2] Mark Allman. Thoughts on reviewing. *ACM SIGCOMM Computer Communication Review (CCR)*, 38(2), April 2008.

[3] Joshua S. Gans and George B. Shepherd. How are the mighty fallen: Rejected classic articles by leading economists. *The Journal of Economic Perspectives*, 8(1):165–179, 1994.

[4] Anthony Jay and Jonathan Lynn. *The Complete 'Yes Minister'*. BBC Books, 1984.

[5] Frank Manola. How to get even with database conferece program committees. *IEEE Data Engineering Bulletin*, 4(1):30–36, September 1981.

[6] Ian Parberry. A guide for new referees in theoretical computer science. *Information and Computation*, 112(1):96–116, 1994.

[7] Timothy Roscoe. Writing reviews for systems conferences. `http://people.inf.ethz.ch/troscoe/pubs/review-writing.pdf`, 2007.

[8] Simone Santini. We are sorry to inform you... *Computer*, 38(12):128–127, 2005.

[9] Simon Singh. *Fermat's Last Theorem*. Fourth Estate, 1997.

[10] Alan Jay Smith. The task of the referee. *IEEE Computer*, 23(4):65–71, 1990.

[11] Mary-Claire van Leunen and Richard Liption. How to have your abstract rejected. *SIGACT News*, 8(3):21–24, 1976.

[12] Toby Walsh. How to write a review. `http://www.labunix.uqam.ca/~jpmf/int-mgmt/walsh1.pdf`, 2001.