

Rebuttal How-to: Strategies, Tactics, and the Big Picture in Research

Danfeng (Daphne) Yao
Department of Computer Science
Virginia Tech

Accepted for publication by Communications of ACM (May 2023)

Abstract

Rebuttals are not published, thus, it is difficult for junior researchers to read successful rebuttals and improve. This article demystifies rebuttal writing by showing the arm-the-champion strategy and a few key tactics. More importantly, we also discuss the conformity nature of conference reviewing and why researchers should not be defeated by paper rejections.

Introduction

At this year's Individualized Cybersecurity Research Mentoring Workshop (iMentor) [1], I plan to give a 45-minute presentation on writing conference rebuttals. The insights would benefit others beyond the cybersecurity research community, so I decide to organize my thoughts and share them with our broader computing community. Submitting rebuttals, also known as the author response period, is a common practice in computing conferences. After reading the reviews, the authors can submit a succinct itemized response, e.g., around 700 words within 4-5 days. The program committee then discusses the work further to reach the final verdict for the manuscript.

Writing rebuttals energizes me, even though most of my rebuttals do not change the rejection fate of my papers. I genuinely enjoy the opportunity of communicating with experts in the field. This process of intensely reading others' opinions about my work also helps solidify my own research style. Discussing rebuttal plans is a wonderful team bonding activity, so be sure to analyze reviews together with coauthors. Unconstructive reviews, however, induce impostor syndrome [2]. So, plan rebuttals *after* you have overcome the initial shock.

However, for junior researchers under job and graduation pressure, this rebuttal process can be rather stressful and confusing. Hope this discussion on rebuttal strategies, tactics, and the big picture of research would be useful. The principles mentioned here are useful for writing itemized author responses for journals, too.

Joan of Arc said, "All battles are first won or lost, in the mind". Have faith in your work, regardless of the outcome. Authors need to think beyond the outcomes of a particular paper. I have also shared some rebuttals of mine to help others get familiar with this type of writing [3].

Why bother submitting a rebuttal?

Researchers show 1-4.4% of papers were positively impacted by rebuttals in five recent conferences [4]. The low statistical impact on papers' outcomes is also observed by others [5]. I

briefly describe three scenarios where writing rebuttals could be immensely useful: *i)* major revision competency, *ii)* refutation to set the record straight, and *iii)* training critical thinking and question answering.

For conferences that have a major revision (e.g., USENIX Security Symposium) or conditional accept option, authors should submit author responses. For USENIX Security, the acceptance rates of major revised manuscripts are very high: 85.7% in 2020 [6] and 86% in 2021 [7]. However, before handing out the precious major revision verdict, reviewers examine the rebuttal for indicators showing the team's revision competency. Are the authors willing to conduct the necessary new experiments? Can authors adequately address the requested revision items in a few months? Poorly written rebuttals may indicate low-quality revision down the road. Thus, the rebuttal needs to show strong revision competency and commitment.

I sometimes submit rebuttals even when my paper receives entirely hopeless ratings. The purpose is to refute. This refutation scenario is where authors need to clarify factual errors and serious misunderstandings about their work, e.g., regarding novelty, significance, or correctness. You may need to strongly refute reviewers to set the record straight, preparing for future submissions. The paper review circle is small -- resubmissions may be assigned to the same, possibly biased, reviewers at later conferences. I experienced this situation in our CryptoGuard work [8], with re-occurring negative talking points incorrectly insisting that earlier prototypes had already solved the problem of deployment-grade cryptographic API misuse detection. It is unnatural not to feel upset and frustrated -- the work is your brainchild, but some experts think it is unworthy. Tap into that energy when writing rebuttals. Such situations also indicate that your research vision and style differ from others -- an advantage in the long run.

Lastly, writing rebuttals improves one's ability to brainstorm and answer questions in a straightforward manner. Mastering these essential skills requires practice. Therefore, even with heartbreaking ratings, completing the rebuttal process has training values.

Dynamics of the paper review process

Strategically, the key mission of rebuttals is two-fold: *i)* to solidify the champion's support and *ii)* to help the champion defend your work or "embolden" the champion [9]. The latter is also known as "arm the champion" [10], i.e., equipping the champion with additional evidence, so they can argue for your acceptance. Champions are reviewers who vocally advocate for your work [11]. Some conference assigns a champion to each paper [12].

The intuitive idea that a rebuttal does not need to carefully address champions' comments [10] is overly optimistic. Upon seeing weak rebuttals, a champion may declare ending their championing position and lower their rating, which was also reported by other researchers (e.g., [16]). After all, most champions do see your work's flaws and want to hear your defense and justification. Thus, seriously addressing champions' comments and concerns is absolutely necessary.

Social psychology experiments have repetitively shown that people have a tendency to conform [13]. Herd mentality also continues in online spaces [14]. Most reviewers tend to converge to the majority opinion, as it is the safest. Consider this scenario. There are four initial reviews, including one *reject*, one *weak reject*, one *weak accept*, and one *accept*. However, the *weak accept* reviewer -- upon seeing the assertive negative reviews -- immediately lowers their rating to *weak reject*. At this point, what does it take for the only positive reviewer to champion this submission? Championing takes energy. It also takes courage, patience, and persistence. The positive reviewer needs to read all reviews, re-read parts of the paper, and initiate the discussion to admit the current deficiencies in the paper, but continue to explain why these deficiencies are somewhat fixable or tolerable. In addition, the champion would need to ask strong detractors whether or not the rebuttal alleviates their concerns. Occasionally, this negotiation may also give the superficial impression that champions may have a lower scientific standard, which is not true at all. Being a champion is exhausting. Therefore, if the champion's concerns about your work are inadequately resolved, it would be difficult for them to support your work wholeheartedly.

The above example involves mixed diverging reviews. If planned well, such controversial cases may present promising outcomes. Another common situation is where all reviewers sound lukewarm and ratings are mostly borderline-level, e.g., two weak rejects and one weak accept. The excitement is clearly low. This scenario is more thorny, as the paper is likely to be rejected if there is no further discussion. As authors, you still have hope, as you can organize and interpret the reviews as part of your rebuttal, possibly delivering some degree of enthusiasm via your words. The rebuttal is your last resort to energize reviewers. For most conferences, PC chairs or area chairs also read your rebuttals and attempt to energize the discussion, as others also pointed out [15]. Oftentimes, I found reviewers give low ratings without any serious issues or have unrealistic expectations for research prototypes. For cybersecurity conferences, occasionally my rebuttals had to remind reviewers that security is relative and it is impossible to achieve absolute security.

Next, let me highlight some key rebuttal writing tactics with a simple question.

Some rebuttal writing tactics

When addressing critical comments, the response needs to be thoughtful and have depth. Brainstorm a little. Share your expert opinions. Let me illustrate using one seemingly innocuous comment.

Will your solution handle situation A?

This is not a simple yes or no question, if A is a complex scenario and you have not thoroughly discussed it. Trivializing reviewers' critical comments is a common rebuttal pitfall. Recognize key concerns and address them with thoughtfulness. For cybersecurity work, typically there are several possible options.

- *To defend your work:* Simply stating yes is insufficient. One needs to provide evidence, reasoning, or justification, as well as possible challenges associated with porting your solution to handle A.
- *To point out that the concern is universal:* A is a common challenge (e.g., lack of security ground truth in real-world code), thus not a dealbreaker. However, do thoughtfully discuss how A impacts your work, e.g., on precision and recall values and security guarantees.
- *To point out it is an open problem:* No one knows how to handle A well, but you plan to offer good insight in the revision.
- *To admit true limitations:* This response still needs to be thoughtfully worded to provide good insights. For example, A is out of our threat model, but other solutions can be borrowed to handle A. However, the integration may present complications in our setting. After all, it is unreasonable to expect one paper to solve everything everywhere all at once.
- *A combination of all:* For example, A may represent an open problem that is universal to some approaches that your work can potentially handle, but with limitations. Capturing these subtleties helps convey your technical depth.

Paraphrase the reviews and reorder them. You have full control over what goes into the rebuttal. You can reorder comments by importance. You can paraphrase the comments. Describe their concern in your own words. This process also shows your deep understanding and appreciation of their comments. Most importantly, using consistent terminology in both the question and the answer makes it easy to follow.

Avoid excessively referencing existing sections. What was already said in the paper clearly is unconvincing; repeating it would be unwise. The rebuttal is a precious space to provide new information, new perspective, new content, and new results (if permitted). Virtually all my rebuttals include new numbers and new experimental data. Running experiments takes time. Thus, the author team needs to meet as soon as possible to devise a rebuttal plan.

Convey your strong willingness to revise and do new work. A common pitfall in writing rebuttals is the lack of commitment to revision. However, depending on whether the conference offers a major revision or conditional acceptance option, you would write differently. Most reviewers do not feel comfortable conditionally accepting a paper, if the new version will likely look drastically different. Several blog articles also give other great suggestions [9, 10, 15, 16, 17].

Appeal a decision

My recent *Communications Medicine* work on AI digital health fairness [18] was rejected in the initial round of review. Luckily, *Nature* journals have a streamlined appeal process. Eventually, the editor decided to reverse their decision and send my revised version out again to reviewers. Journal reviewers operate in isolation and thus do not influence each other's opinion, at least in the initial round of review.

In contrast, appealing is less useful in our conference review systems based on my experience, as outlier reviewers tend to quickly give up and conform to the majority verdict. Thus, even if the PC chair agrees to add a couple more reviewers who like your work, the outcome rarely changes. Here is why -- the minute new reviewers see the existing negative ratings, they will likely swiftly reduce their ratings and declare alignment with the existing votes. This behavior is understandable, as championing such a paper is clearly an uphill battle. The authors can always improve their work and resubmit later, so why rush? Regardless of the outcomes, PC chairs and area chairs are instrumental in creating thorough post-response discussion, by asking "Reviewers, does the rebuttal address your concerns? Why not?"

The big picture in research

In *The Art of War*, Sun Tzu said that "Victorious warriors win first and then go to war, while defeated warriors go to war first and then seek to win." For researchers, doing good work in the first place is the ultimate key. The most brilliant rebuttal could not rescue ill-formed research.

The peer review system has randomness due to the unpredictable nature of the discussion (e.g., your champion may be unavailable for discussion due to a medical emergency). The peer review system also has implicit biases. For example, the importance of a research direction is heavily influenced by the reviewers' own research interests. In my experience with cybersecurity conferences, papers in hot areas tend to receive longer reviews with more excitement, as they address "important and timely problems". In contrast, work addressing older-but-still-unsolved problems (e.g., false positives in deploying anomaly-based intrusion detection) is likely met with dampened enthusiasm and perfunctory lackluster reviews. Non-expert reviewers may also have the incorrect impression that the problem has long been solved. Continuing to diversify program committees, as many conferences are doing, would help reduce such implicit biases.

For researchers, understanding the review mechanism and the dynamics among reviewers would help them navigate the publication process. Keep submitting. Win the battle in your mind.

References

[1] ACM SIGSAC Individualized Cybersecurity Research Mentoring Workshop (iMentor), co-located with ACM Conference on Computer and Communications Security (CCS) 2023.

Available at: <https://sites.google.com/vt.edu/imentor>

[2] Danfeng (Daphne) Yao. Depth and Persistence: What Researchers Need to Know About Impostor Syndrome. *Communications of the ACM*. 64(6). June 2021.

[3] Danfeng (Daphne) Yao. Samples of Computer Science Conference Author Responses.

Available at: <https://people.cs.vt.edu/~danfeng/Yao-rebuttal-samples.pdf>

[4] Nachum Dershowitz and Rakesh M. Verma. Rebutting Rebuttals. *Communications of the ACM*. 2023. Available at: <http://www.cs.tau.ac.il/~nachumd/papers/Rebuttals.pdf>

- [5] Joanna McGrenere, Andy Cockburn, and Sandy Gould. CHI 2020 – The effect of rebuttals. 2019. Available at: <https://chi2020.acm.org/blog/chi-2020-the-effect-of-rebuttals/>
- [6] Srdjan Čapkun and Franziska Roesner. Message from the USENIX Security '20 Program Co-Chairs. https://www.usenix.org/sites/default/files/sec20_message.pdf
- [7] Rachel Greenstadt and Michael Bailey. Message from the USENIX Security '21 Program Co-Chairs. 2021. https://www.usenix.org/sites/default/files/sec21_message.pdf
- [8] Sazzadur Rahaman, Ya Xiao, Sharmin Afrose, Fahad Shaon, Ke Tian, Miles Frantz, Murat Kantarcioglu, and Danfeng (Daphne) Yao. CryptoGuard: High Precision Detection of Cryptographic Vulnerabilities in Massive-sized Java Projects. ACM CCS 2019.
- [9] Michael Hicks. Advice on writing an author response. Available at: <http://www.pl-enthusiast.net/2014/09/17/advice-writing-author-response/>
- [10] Andreas Zeller. Patterns for writing good rebuttals. 2022. Available at: <https://andreas-zeller.info/2012/10/01/patterns-for-writing-good-rebuttals.html>
- [11] Oscar Nierstrasz. Identify the Champion. Available at: <https://scg.unibe.ch/download/champion/>
- [12] Andreas Zeller. How do different fields review papers? Experiences from ICSE, PLDI, and CCS. 2021. Available at: <https://andreas-zeller.info/2021/07/27/Reviewing-across-fields-ICSE-PLDI-CCS.html>
- [13] Conformity. Available at: <https://en.wikipedia.org/wiki/Conformity>
- [14] Herd mentality. Available at: https://en.wikipedia.org/wiki/Herd_mentality
- [15] Devi Parikh, Dhruv Batra, and Stefan Lee. How we write rebuttals. Available at: <https://deviparikh.medium.com/how-we-write-rebuttals-dc84742fece1>
- [16] Meredith Ringel Morris. My Rebuttal-Writing Process for HCI Venues. Available at: https://cs.stanford.edu/~merrie/merrie_rebuttal_tips.pdf
- [17] César Soto Valero. How to Write a Good Paper Rebuttal, a practical guide. 2022. <https://www.cesarsotovalero.net/blog/how-to-write-a-good-paper-rebuttal.html>
- [18] Sharmin Afrose, Wenjia Song, Charles B. Nemeroff, Chang Lu, Danfeng (Daphne) Yao. Subpopulation-specific Machine Learning Prognosis for Underrepresented Patients with Double Prioritized Bias Correction. Communications Medicine. 2 (111). 2022

Biosketch:

Daphne Yao is a Professor of Computer Science at Virginia Tech. She is an ACM Distinguished Scientist and currently serves as the Vice Chair of ACM SIGSAC. She can be reached at danfeng@vt.edu.