Report of the CRA Working Group on Research Integrity

August 2023

Working Group Members

Alex Aiken* (co-chair), Professor, Stanford University

Nancy M. Amato* (co-chair), Professor and Head, University of Illinois Urbana Champaign

Kivmars Bowling, Publications Director, Society for Industrial and Applied Mathematics

Leila De Floriani, Professor, University of Maryland

Eric de Sturler, Professor, Virginia Tech

Maria Gini,* Professor, University of Minnesota

Vicki Hanson, CEO, Association for Computing Machinery

Arvind Krishnamurthy,* Professor, University of Washington

Kate Larson,* Professor, University of Waterloo and CS-Can|Info-Can

Wei Li, Professor, Texas Southern University

Michael Littman, Division Director, National Science Foundation

Fatma Ozcan, Principal Software Engineer, Google

Melissa Russell, Executive Director, IEEE Computer Society

Vivek Sarkar, Professor, Georgia Institute of Technology

Ann Schwartz, Director of Research Computing Initiatives, Computing Research Association

Eugene H. Spafford,* Professor, Purdue University

Divesh Srivastava,* Head of Database Research, AT&T

*CRA Board of Directors Member

Table of Contents

Executive Summary

We enumerate current threats to research integrity in the field of computing. We evaluate whether these threats have significantly changed in recent years and whether the research community's mitigation efforts appear to be sufficient. We recommend a set of best practices to address specific issues with paper reviewing:

- To the extent possible, data collected to make decisions that can affect the reviewing of papers should not be self-reported and preferably should be gathered automatically.
- The process of assigning reviewers to papers should ensure that reviewers cannot effectively assign themselves or otherwise manipulate assignments of specific papers.
- A minimum of three reviews and preferably more, using geographically and organizationally diverse assignments, should always be used to dilute the effect of any single malicious reviewer.
- To the extent possible, universities, professional societies, and funding agencies should regularly provide the community with general information on the quantity and nature of violations and the penalties levied against offenders.

We also recommend that professional societies develop more rigorous standards for sponsoring special topic meetings to reduce the instances of such meetings where the proceedings have very low or no standards for publication.

Finally, we recommend the development of training materials for new members of the research community that convey the standards and expectations for authors and reviewers as well as the consequences for violations. These materials could be disseminated, or potentially even mandated, by professional societies, universities and/or funding agencies.

Charge and Background

The computing research community relies on the integrity of authors, reviewers, conference chairs, and journal editors to function properly. By *integrity* we mean the observation of formal rules and norms about what is acceptable in the preparation, submission, and peer review of research publications.

The CRA Working Group on Research Integrity was formed to address concerns that the frequency of violations of these rules and norms has increased in recent years. The charter of the group was to enumerate the current threats to research integrity and to make recommendations on best practices to mitigate those threats (where there is some consensus) or suggest directions for further study (where no such consensus exists).

The potential audiences for this report include authors, reviewers, program chairs, editorial board members, developers of conference management software, professional societies, department chairs and other university administrators, research managers in industry, and funding agencies. Not everything in the report will be of interest to everyone; where appropriate we identify specific audiences.

The purpose of any field of research is to advance knowledge. Given that a lack of integrity in research cannot improve the research itself, it is worth asking what motivates someone to break the rules and norms of research integrity. At least part of the answer seems to be that rewards are not always based on the quality and impact of research, but rather on proxies: the number of papers published, the number of citations, and/or fundraising success. Increases in violations of research integrity are likely driven by more organizations adopting such proxies and consequently more people responding to the resulting incentives. As long as organizations reward researchers based on metrics rather than an assessment of research quality done by independent experts, we will be managing symptoms rather than addressing underlying causes. We refer interested readers to the CRA memo *Incentivizing Quality and Impact: Evaluating Scholarship in Hiring, Tenure, and Promotion*.

Threats to Research Integrity

The current specific threats to research integrity discussed by the working group are listed below.

- *Subversion of peer review* occurs when reviewers work towards predetermined outcomes for papers under review.
- *Abuse of meeting sponsorship* involves creating conferences or workshops that do not meet community quality standards.
- *Fabrication and falsification of results* is knowingly attempting to publish results that are not scientifically valid.
- *Plagiarism* is taking credit for previous work (by others or even oneself), most often by appropriating ideas or text from past publications without attribution.
- *Manipulation of citation counts* is inflating the number of references to a paper beyond what is justified by the scholarship.
- *Ignoring the standards of authorship* primarily involves including authors on papers who have not contributed significantly to the work or, conversely, excluding people who have contributed significantly from authorship.
- *Abuse of generative AI in the publication process* involves using automated tools to substitute for human intellectual contributions in the submission or review process. It is a very recent potential threat. At this time it is too early to judge what its impact may be; this topic will need to be revisited when the community has additional experience with the use of generative AI in the publication process.

A number of practices have been found to mitigate efforts to undermine research integrity in the research and publication process. We stress that there are no foolproof practices. We do not believe it is possible or even desirable to attempt to eliminate all abuses, as overzealous efforts to root out violations of research integrity could well end up harming the work of the great majority who play by the rules. Our recommendations should be applied with the costs as well as the benefits in mind.

In what follows we consider each of the threats to research integrity listed above. We make specific recommendations for changes in current practices for peer review and the approval of

sponsored meetings. We also recommend continuing the community's current efforts to encourage reproducibility as a way of discouraging fabrication and falsification of results. We provide a brief summary of the other areas of concern without making specific recommendations.

Our primary overarching recommendation is to produce and disseminate training materials for new members of the computing research community that explain and illustrate the standards and expectations for authors and reviewers.

Subversion of Peer Review

We believe that the subversion of peer review is the most immediate and serious of the current threats to research integrity. The working group heard evidence that abuses of the peer review system have significantly increased and evolved to exploit the very different reviewing environment that exists today compared to a decade ago. The scale of peer reviewing for refereed conferences in computing has grown substantially in recent years and it is now clear that these complex review processes can be vulnerable to subversion in various ways.

The primary threat to peer review is a *malicious reviewer* (or reviewers) who attempts to influence the reviewing process to achieve a particular outcome for a paper independent of the paper's merits. We divide the recommendations for best practices for peer review into several categories: *data collection* to make decisions that may impact the reviewing process, *assigning reviewers* to papers, and *investigating and sanctioning* those who violate standards. We reiterate that these recommendations are specifically focused on supporting the integrity of peer review, though we also believe that these recommendations are consistent with achieving the best reviewing outcomes generally.

Data Collection

Data that could affect the reviewing process should not be self-reported. Given the scale of many computing-related conferences, this data should be collected automatically if possible. Where automatic collection is not possible, measures should be taken to minimize opportunities and the potential consequences of being given incorrect self-reported information by malicious reviewers.

The identities of reviewers and authors should be established through a persistent unique ID, such as ORCID [Haak, et al., 2012], OpenReview ID [Soergel et al., 2013] or ResearcherID [ResearcherID, 2023], rather than through self-reporting. Persistent identities can be linked to professional credentials, making it more difficult for someone to use multiple credible identities. While these credentials are not generally verified today, they could (and should) be in the future.

To the extent possible, conflicts of interest should be collected automatically rather than be self-reported. By a conflict of interest (COI) we mean any relationship between reviewers and authors that could reasonably influence a review. Current tools for detecting COIs (e.g., CLOSET [Bhowmick, 2023]) partially automate the process by using publicly available information to identify, for example, same-institution and co-author conflicts. These tools have been used in large conferences, such as PVLDB and SIGMOD, though use is not yet

widespread. Adoption of persistent identifiers with linked information about institutions, publications, grants, and students will increase the accuracy and coverage of automated COI detection.

Reliance on bidding by reviewers for the papers they will review should be limited to ensure that reviewers cannot effectively assign themselves specific papers. Many computing conferences use a combination of topic preferences from reviewers and topics of papers to automatically match reviewers to papers, supplemented by bidding on specific papers.

One class of bidding strategies requires reviewers to submit a minimum number of positive bids (papers a reviewer would like to review) and a maximum number of negative bids (papers a reviewer does not want to review); more positive/fewer negative bids gives the matching algorithm more discretion in choosing papers to assign to a reviewer [Jecmen et al, 2022]. Another approach allows a reviewer to bid only on a randomly chosen subset of the papers.

Assigning Reviewers

Good matching of reviewers to papers is important, but the number of reviews is important too, not least because more reviews dilute the influence of any single malicious reviewer. Three reviews should be the minimum, but many conferences devote four or five reviews to papers in the final round before decisions are made.

Randomness can be used in paper assignments to decrease the chances that a malicious reviewer is assigned a specific paper [Jecmen et al, 2020]. In the special case where reviewers are also authors, one should disallow cycles where a pair (or more) of reviewers are assigned to review each other's papers [Guo et al, 2018, Boehmer et al, 2021].

Geographical and organizational diversity introduce constraints into the paper-assignment process to ensure that not all reviewers and paper authors come from the same geographic region or same institution. The rationale is that geographically and organizationally diverse assignments make it less likely that reviewers will know each other, reducing the potential for collusion.

Because the assignment of papers is imperfect and based on imperfect information (particularly knowledge of all conflicts of interest), we cannot completely rely on the paper-reviewer matching process to result in a sufficient number of accurate and trustworthy reviews. Analysis of the reviews can be used to identify outliers and mitigate the effect of outliers on the automated parts of the decision-making process. While clearly some auditing of reviews is worthwhile and even necessary to identify and correct problems, there is no consensus on best practices. We list some approaches that have been proposed or in some cases tried; we consider this an area where additional work is needed.

Paper scores far from the average can be given a lower weight, preventing papers from initially being ranked very high or low based on outlying scores. A reviewer's scores are sometimes weighted based on how often they are outliers across all reviews. Weighting scores based on distance to the average score works better with more scores. The approach can also be improved if more scores, for categories, are assigned per paper by the same reviewer. Another

enhancement is measuring when multiple reviewers correlate in their far-from average reviews. We note that some of these methods require more data and may only be meaningful for large conferences. Similarly, approaches based on aggregate statistics of reviews are gamed more easily if multiple malicious reviewers manage to be assigned to the same paper. Again, the goal cannot be to address all potential threats all the time, but rather to systematically discourage bad behavior at every stage of the process.

A complementary strategy is to assign a trust score or alternatively an uncertainty to each reviewer-paper matching. Such a score could vary based on the knowledge of the reviewer in the area of the paper as well as, say, a potential conflict that is not severe enough to count as a conflict of interest; for example, the reviewer and an author have never published together but have several papers with a common co-author. Such a trust score could be used in a variety of ways, again, for weighting scores, or alternatively to check whether the system has enough 'trust' in the overall score for a paper to make a decision or whether additional reviewing is needed.

Consequences for Violations

Professional societies and other publishers have standard procedures for enforcing rules of behavior, including for violations of research integrity. These procedures are discussed in detail in the section on grievance procedures. We recommend that authors and reviewers be reminded of the consequences of bad behavior by requiring agreement to a plain language summary of the rules of the governing body of the publication at submission time.

Abuse of Meeting Sponsorship

The working group heard evidence that the creation of substandard, but sanctioned, meetings is increasing. The specific current issue centers on “special topic” meetings. Typically a professional society approves an application to sponsor a one-time meeting on a research topic, only to discover later that the proceedings of the meeting, which bears the society’s imprimatur and is generally included in the organization’s on-line digital library, was assembled without even minimal standards for selecting papers.

This problem has a similar character to the abuses of peer review, targeting specific points in the publication process to enable papers that would otherwise be rejected to be published. The difference is one of scale, targeting entire proceedings instead of individual papers.

Unlike peer review, there is a single point of control, namely the approval process for meeting sponsorship. The professional societies are aware that efforts to game this process exist and we expect that some combination of more vetting of applications for meetings, auditing of processes, and better communication of sanctions for abuse (including potentially withdrawing an entire proceedings) will be sufficient.

Fabrication and Falsification of Results

While incidents of fabrication and falsification of results do happen, this problem does not seem to be increasing or to have significantly changed in other ways in recent years. The computing research community's efforts in encouraging reproducibility, while not primarily aimed at improving research integrity as defined here, discourage outright misrepresentation of results.

Reproducibility is a key part of the scientific method. Requiring experimental results to be reproducible by others (for example, at least requiring that code and datasets to reproduce experiments be made available) has obvious benefits for incentivizing research integrity and has become widespread in the computing research community.

Reproducibility is a focus for funding agencies. For example, the August 2022 OSTP Nelson memo noted that "Scientific data underlying peer-reviewed scholarly publications resulting from federally funded research should be made freely available and publicly accessible by default at the time of publication" and such scientific data is defined as "recorded factual material commonly accepted in the scientific community as of sufficient quality to validate and replicate research findings" [Nelson, 2022]. Some publishers have also taken steps to encourage the release of materials needed to reproduce published results by awarding reproducibility badges.

It should be noted that reproducibility is not always fully possible even when code and data are shared, for example, if technical hardware or software environments are no longer available to reproduce the original results. Nevertheless, reproducibility and the sharing of underlying code and data remain key principles for maximizing the integrity of the scholarly record.

Plagiarism

Plagiarism has always been an issue. The working group heard evidence that the frequency and nature of plagiarism cases has been stable over the past decade. Furthermore, the community has developed and adopted countermeasures, such as plagiarism detection software and training. These methods may need to be revisited in the context of generative AI, however, as fluent automatic rewording can now be accomplished at scale.

Manipulation of Citation Counts

Like plagiarism, inappropriate efforts to increase citation counts is a long-standing issue. However, unlike plagiarism, there are no widely accepted and deployed mechanisms to discourage or even measure practices such as self-citation of unrelated works and agreeing with others to cross-cite papers for no legitimate scholarly reason. As persistent identifiers for authors become more widespread, it should become easier to automatically and reliably detect abusive citation practices.

It is worth noting that few would expend time and effort to increase their citation counts unless there were tangible rewards for doing so. Using citation counts, rather than assessing the quality and impact of the work, for evaluation of a researcher is a good example of valuing proxy research metrics over the actual research. Unfortunately, citation counts are used for some

important decisions over which the research community has no control, such as whether to award certain kinds of visas to foreign nationals in the United States.

Ignoring Standards of Authorship

Reported problems of authorship are relatively rare. One way in which such issues come to light is when a paper is found to have serious flaws and an author claims that it is not their fault because they were not very involved. As standards of authorship clearly state, by agreeing to be an author one is taking responsibility for the entire contents of the publication.

One way to minimize potential issues of authorship, if paired with the use of persistent IDs for reviewers and authors as discussed earlier, is for submission management software to email all listed authors on paper submissions by default.

The message that authorship is not just about receiving credit, but also potentially accepting blame, needs to be communicated broadly within the research community. (See the recommendations on training below.)

Training for Authors and Reviewers

Currently, members of the research community learn what is and is not acceptable behavior from their peers and mentors. Given that the research community is extremely diverse and has grown very rapidly, there is no reason to assume that everyone is familiar with the shared standards of integrity at the start of their careers.

We recommend the development and dissemination of training materials on the expectations and standards for authors and reviewers. We believe there is an opportunity to foster a collective commitment to research integrity by providing formal training for new members of the community. Some examples of such training include discussion of the standards for authors and reviewers with relevant anonymized examples to emphasize that the issues are not just theoretical, the use of “shadow program committees” to educate people about reviewing before being program committee members themselves, and reviewer mentoring programs (an example is here). Attention should be given to the additional burdens involving research with human subjects (especially the imperatives not to do harm and to have the approval of independent review boards), as human subjects research is part of computing today and likely to become more common in the future. Finally, this training should be supplemented with consistent statements at conferences and other related opportunities about publication standards and expectations.

Ideally, this training would reach beyond new community members so that eventually all have exposure to a baseline of information about the responsible and ethical conduct of research (RECR), analogous to the largely standardized training that already exists around research with human subjects. Professional societies could require that all authors and reviewers complete a standard course on RECR as a prerequisite for publishing in or reviewing for venues sponsored by the society. Universities should include such training on publications and reviewing for new faculty and graduate students, as well as for those submitting or reviewing grant proposals.

As a concrete example of required training, consider that the America COMPETES act, passed in the U.S. in 2007, requires each institution that applies for funding from the National Science Foundation (NSF) to have a training and oversight plan for RECR for undergraduates, graduate students, and postdoctoral researchers. The CHIPS and Science Act in 2022 expanded the requirement to include faculty and to include research security training. The research integrity training we are recommending could potentially be combined with this mandated training.

Grievance Procedures

Grievance processes address failures to conduct ethical research and reviewing. Well-crafted enforcement policies help inform the community about acceptable research practices as well as deter further research misconduct. This section considers grievance processes for both professional societies and universities, aspects of research integrity covered by policies, information about sharing in cases of findings of research integrity violations, and suggestions for the best advice for avoiding research misconduct.

Professional societies typically have extensive policies governing misconduct of multiple types. Most relevant to this report is misconduct related to plagiarism, misrepresentation of authors, falsification of data, as well as multiple types of misconduct related to the peer review process such as conflict of interest, coercion, and failure to maintain confidentiality. Critically, such policies are constantly being updated due to findings of new forms of misconduct (such as reviewer collusion rings) and new forms of technology that have the potential for misuse (such as generative AI).

As an example of grievance processes, the following general steps are used by the Association for Computing Machinery (ACM) in suspected cases of research misconduct:

1. Suspected misconduct is reported to the society by one or more individuals who have encountered what they believe might be a violation.
2. There is the consideration of whether the claim relates to policy issues and whether there is sufficient evidence to begin an investigation.
3. If an investigation is undertaken, the process can involve interviews and the collection and review of evidence.
4. Following investigation, ACM's volunteer Ethics and Plagiarism Committee makes a decision as to whether there has been a policy violation. If so, a penalty or sanction is often given, such as bans from publication or peer review for a period of time. In cases of plagiarism, papers may be retracted.
5. Individuals receiving a penalty or sanction may appeal.

Six months to a year is typically needed to resolve a reported case of research misconduct.

Professional societies can raise community awareness of the potential harms and consequences of violations of research integrity by notifying the community that the rules are being applied. Towards the goal of ethical and educational practice, ACM, for example, has been disclosing information about the annual incidence of research misconduct reports and investigations as well as the type of violations found.

Disclosure of the identities of violators by global professional societies has been restricted due to privacy laws such as the European Union's *General Data Protection Regulation* (GPDR), which limit the release of such information. In addition, further disclosure is restricted due to criminal charges that could potentially result for volunteers who make a third-party disclosure that might be determined to be slanderous or defamatory. As a result, the research community must work within a framework where repeat offenses across different institutions will be difficult to detect.

The current best advice to avoid research misconduct is promoting efforts aimed at educating authors and those involved in peer review processes for the purpose of avoiding research misconduct. For societies, this could take the form of the training materials we recommend above.

Many universities have established policies and procedures on research misconduct, which, at a minimum, cover fabrication, falsification, and plagiarism. A typical process related to a research misconduct case at a university includes the following high-level steps:

1. Assessment

- Receive a report of an allegation related to research misconduct. Such a report can come from outside the institution.
- Determine if the allegation is sufficiently credible and specific to warrant next steps. If not, the process ends.

2. Inquiry

- Conduct an inquiry to determine if an investigation is warranted.
- Make any required disclosures to research sponsors.
- If an investigation is not warranted, the process ends.

3. Investigation and Decision

- Conduct an investigation resulting in a report of findings.
- Make a determination of whether research misconduct has occurred.
- Make a determination of corrective actions and/or sanctions that may be warranted.

4. Appeal and Decision

- Provide a channel for appeal of the investigation report and proposed actions.
- Make a decision on appeal.

These processes are typically described in a discipline-agnostic manner, which can make it challenging to see how the terms and issues relate specifically to computing research. As with professional societies, these processes can be lengthy, typically requiring a year.

Conclusions

To reiterate our specific recommendations:

- We recommend practices to mitigate the influence of malicious reviews by avoiding self-reported information, limiting the influence on reviewer preferences on paper assignments, ensuring there are enough reviews that outcomes do not overly rely on any single reviewer, and flagging papers that have weak evidence for additional reviews.
- We recommend that professional societies tighten the processes for sponsoring special topic meetings with proceedings.
- We recommend the development of training materials to educate new members of the research community on the standards and expectations for authors and reviewers and the consequences for violations of those standards.
- To the extent possible, publishers should regularly provide the community with general information on the quantity and nature of violations and the penalties levied against offenders.

There are some areas that we believe merit further investigation. Regular auditing and analysis of reviews to detect problems seems unarguably worthwhile, but how to do such analyses in a way that reliably identifies problems, is not overly time consuming, does not violate privacy protections, and does not discourage reviewers from forthrightly stating their views is unclear.

Generative AI is likely to have at least some influence on the publication process, but at this time the nature of that influence and whether it is benign, beneficial, or a serious problem is a matter of speculation. While some organizations are already taking first steps to provide guidance on the use of generative AI in reviewing and publication, some time and experience is likely needed for the community to gain perspective.

Another area for further study that falls outside the scope of this report is alternative publication models. The trend towards hybrid conference/journal publications in computing, for example, changes the review process. While all of our specific recommendations for reviewing still apply to hybrid conferences/journals, there could be additional considerations. In contrast, a more fundamental change such as a shift to open publication (e.g., on sites such as arXiv.org) without the initial gatekeeping function of reviewing radically alters incentives and would certainly significantly alter where issues of research integrity arise.

We recommend that the working group be reconstituted in three to five years if there is still a perception that threats to research integrity in the computing community are either not well understood or not sufficiently controlled.

Working Group Process

The CRA Working Group on Research Integrity met biweekly from February through May 2023. Most of these meetings were devoted to hearing from community stakeholders, including the computing professional societies, university officials charged with investigating violations of

research integrity, recent program chairs of large conferences, representatives of the major conference management systems, and the National Science Foundation. The membership of the working group was also chosen to broadly reflect these stakeholders as well.

Additional Materials

Integrity violations and citations:

- ACM Publication Policies and Procedures
- ACM site for reporting a potential publications violation
- ACM site listing publication penalties for publication violations
- IEEE site for reporting an ethics complaint
- IEEE Publication Services and Products Board Operations Manual with content related to IEEE Principles of Ethical Publishing and allegations/investigations of possible misconduct
- Committee on Publication Ethics (COPE) site on promoting integrity in research and its publication
- Ethics and Plagiarism update
- National Academy of Sciences

Paper bidding and assignment:

- Matching Papers and Reviewers at Large Conferences, arXiv paper written by AAAI program chairs (and others)
- A Dataset on Malicious Paper Bidding in Peer Review, Steven Jecmen, Minji Yoon, Vincent Conitzer, Nihar B. Shah, Fei Fang
- Tradeoffs in Preventing Manipulation in Paper Bidding for Reviewer Assignment, Steven Jecmen, Nihar B. Shah, Fei Fang, Vincent Conitzer, Workshop on ML Evaluation Standards at ICLR 2022
- Mitigating Manipulation in Peer Review via Randomized Reviewer Assignments, Steven Jecmen, Hanrui Zhang, Ryan Liu, Nihar B. Shah, Vincent Conitzer, and Fei Fang, NeurIPS 2020.

References

Niclas Boehmer, Robert Bredereck, and André Nichterlein. *Combating collusion rings is hard but possible*. arXiv.2112.08444, 2021.

Sourav Bhowmick. *CLOSET: Data-Driven COI detection and management in peer-review venues.* Communications of the ACM 66(7):70-71, 2023.

Guillaume Cabanac and Thomas Preuss. *Capitalizing on order effects in the bids of peer-reviewed conferences to secure reviews by expert referees*. Journal of the American Society for Information Science and Technology, 64(2):405–415, 2013.

Laurel Haak, Martin Fenner, Laura Paglione, Ed Pentz, and Howard Ratner. *ORCID: A system to uniquely identify researchers.* Learned Publishing 25(4): 259-264, 2012.

Longhua Guo, Jie Wu, Wei Chang, Jun Wu, and Jianhua Li. *K-loop free assignment in conference review systems*. Proceedings of the International Conference on Computing, Networking and Communications, pp. 542–547. IEEE, 2018.

ResearcherID, https://webofscience.help.clarivate.com/en-us/Content/wos-researcher-id.htm, August, 2023.

Steven Jecmen, Nihar Shah, Fei Fang, and Vincent Conitzer, *Tradeoffs in preventing manipulation in paper bidding for reviewer assignment*, Proceedings of the ML Evaluation Standards Workshop, 2022.

Steven Jecmen, Hanrui Zhang, Ryan Liu, Nihar B. Shah, Vincent Conitzer, and Fei Fang. *Mitigating manipulation in peer review via randomized reviewer assignments*. Advances in Nueral Information Processing Systems 33:12533-12545, 2020.

Kevin Leyton-Brown, Mausam, Yatin Nandwani, Hedayat Zarkoob, Chris Cameron, Neil Newman, and Dinesh Raghu. *Matching papers and reviewers at large conferences*, arXiv:2022.12273, 2022.

Michael L Littman. *Collusion rings threaten the integrity of computer science research*. Communications of the ACM, 64(6):43–44, 2021.

Alondra Nelson. *Ensuring free, immediate, and equitable access to federally funded research.* Office of Science and Technology Policy Memorandum, https://www.whitehouse.gov/wp-content/uploads/2022/08/08-2022-OSTP-Public-Access-Memo.pdf, 2022.

CRA

Computing Research Association

© 2023 Computing Research Association (CRA). All rights reserved.

1828 L Street, NW, Suite 800
Washington, DC 20036

P: 202 234 2111
F: 202 667 1066

www.cra.org
info@cra.org