Hiring and Promotion Guidance for Computing Researchers

September 2025

CRA
Computing Research Association

Authors

Nancy M. Amato, University of Illinois Urbana-Champaign
https://scholar.google.com/citations?user=AmaB9c4AAAAJ&hl=en

Ovidiu Daescu, University of Texas at Dallas
https://profiles.utdallas.edu/ovidiu.daescu

Maria Gini, University of Minnesota
https://scholar.google.com/citations?user=2ewb-10AAAAJ&hl=en

Gillian R. Hayes, University of California, Irvine
https://scholar.google.com/citations?user=5mIkNyYAAAAJ&hl=en

Kate Larson, University of Waterloo
https://scholar.google.com/citations?user=QK-c5esAAAAJ&hl=en

Wei Li, Texas Southern University
https://scholar.google.com/citations?user=iMTvFmEAAAAJ&hl=en

Ming Lin, University of Maryland
https://scholar.google.com/citations?user=ugFNit4AAAAJ&hl=en

Dan Lopresti, Lehigh University
https://scholar.google.com/citations?user=dFCu3gMAAAAJ&hl=en

Katie A. Siek, Indiana University
https://scholar.google.com/citations?user=-R9W7IkAAAAJ&hl=en

Gene Spafford, Purdue University
https://scholar.google.com/citations?user=MkC5kHcAAAAJ&hl=en

Benjamin Zorn, Microsoft
https://scholar.google.com/citations?user=mpar-iAAAAAJ&hl=en

Suggested Citation

Amato, Nancy M., Ovidiu Daescu, Maria Gini, Gillian R. Hayes, Kate Larson, Wei Li, Ming Lin, Dan Lopresti, Katie A. Siek, Eugene H. Spafford, and Benjamin Zorn. 2025. *Hiring and Promotion Guidance for Computing Researchers*. Washington, D.C.: Computing Research Association (CRA).
http://cra.org/wp-content/uploads/2025/09/Hiring-and-Promotion-Guidance-for-Computing-Researchers.pdf

Abstract

In 2024, the CRA Board convened a workgroup to examine challenges in hiring, promotion, and tenure for computing research professionals across academia, industry, and government/non-profit labs. This report surveys the current state of the field and offers guidance for promotion and evaluation committees, department chairs, and managers in developing effective promotion packages. It may also serve as a resource for junior scholars seeking to explain their contributions and advance their careers.

A companion best practices document is directed toward senior leaders — such as division directors, deans, provosts, chancellors, and presidents — supporting high-level decision making about computing scholars.

About the Computing Research Association (CRA)

The Computing Research Association (CRA) represents nearly 300 North American academic units, laboratories, centers, and companies engaged in computing research. Since its founding in 1972, CRA has brought together academia, industry, and government to strengthen the computing research community and its contributions to society.

CRA's mission is to catalyze computing research by leading the community, informing policymakers and the public, and promoting the development of an innovative and responsible computing research workforce. Through its programs and initiatives, CRA supports researchers across career stages and helps shape the future directions of the field.

TABLE OF CONTENTS

EXECUTIVE SUMMARY 4

EVALUATION AT DIFFERENT TYPES OF INSTITUTIONS 4

STATE OF THE CURRENT COMPUTING RESEARCH ENTERPRISE 5

- Building a Community 5
- Changing Dissemination Mechanisms 5
- Hiring and Appointment Evolution 5
- Teaching and Mentoring Expectations 6
- Increased Quantity of Publications 6
- Engagement with Industry, Government, and Non-Profit Organizations 6

ASSESSING QUALITY AND IMPACT OF RESEARCH 7

EVALUATION OF TEACHING AND MENTORING 8

EVALUATION OF SERVICE 9

GUIDANCE FOR UNIT ADMINISTRATORS REGARDING EXTERNAL EVALUATION 10

EVALUATING INTERDISCIPLINARY CONTRIBUTIONS 11

REFERENCES 12

EXECUTIVE SUMMARY

This report provides recommendations for evaluating computing research professionals across academia, industry, government, and non-profit labs. Unlike earlier guidance, which was primarily focused on R1 universities, this report broadens guidance to reflect the wide range of settings where high-quality computing research occurs.

Key themes include:

- **Adapting to Context:** Evaluation standards must reflect institutional norms, joint appointments, and evolving career paths.
- **Research Assessment:** Emphasis should be on the quality and impact of contributions over simple counts of output.
- **Teaching & Mentoring:** Expanding enrollments, new technologies, and biases in student evaluations require broad, reliable assessment methods.
- **Service & Community:** Contributions to conferences, open-source efforts, and industry collaborations are vital and should be recognized.
- **External Review:** Institutions should streamline letter requirements, include evaluators from all sectors, and provide context to ensure fair comparisons.

Overall, the report underscores that computing researchers should be evaluated on the significance of their ideas and contributions to the field, not volume of output.

EVALUATION AT DIFFERENT TYPES OF INSTITUTIONS

High-quality computing research happens at a wide variety of institutions. While the two prior versions of this document were primarily written with an R1 university appointment in mind, we seek in this document to broaden the applicability of this guidance. In particular, we note that similar assessment of research quality might happen at research universities, teaching-focused colleges and universities, and industry and government labs, with variable expectations for output. Some faculty may have joint appointments to visit industry and government labs (Brachman, Schooler, and Wright 2025). At the same time, some researchers from those organizations may visit academic institutions for short periods or hold adjunct appointments and advise students. Expectations of teaching may be high at some colleges and universities and non-existent in some non-academic positions or even in research-focused faculty positions. Thus, those tasked with hiring and promotion assessments should use the guidance below as appropriate to their institution, ensuring that instructions regarding the norms of output for their department, division, school, college, or other unit are clear to all those involved in the evaluation process. They should also recognize and adjust expectations for researchers with partial appointments or leaves that cross over into other types of institutions.

STATE OF THE CURRENT COMPUTING RESEARCH ENTERPRISE

Building a Community

As a consequence of cost, environmental, and physical health concerns, many computing review panels (e.g., conference program committees, grant review panels) have transitioned from in-person to video conferences, and in some cases, asynchronous online decision-making. This transition has increased the opportunity for global participation; however, it has decreased the ability of researchers to network and build community within their subdisciplines (Zorn et al., 2020).

At the same time, the interdisciplinary nature of many (if not most) parts of computing research has necessitated that computing researchers build connections with those in the biological and medical sciences, physical sciences, social sciences and humanities, and more. This has led to a space in which the norms of computing are not as explicit nor as solid as they once may have been.

Changing Dissemination Mechanisms

Workshops and conferences on a wide variety of topics that did not exist in past years now provide venues for publication, some conferences have developed an "alt" track with peer review done by a broader community more akin to crowdsourcing than traditional peer review, and some researchers have begun to share preprints on sites (e.g., arXiv) before peer review.

Although conference dissemination continues to be the preferred mechanism in many computing subdisciplines (Patterson et al., 1999), some computing subdisciplines have adopted a journal review model before presentation at conferences. This model allows for more participation from fields that require journal articles, as well as more consistent recognition by funding bodies and others. However, in some cases, the journal model increases the review time and decreases opportunities to iterate and resubmit within short timelines.

With the increase in submissions, conference organizers and journal editors have struggled to recruit review committees with appropriate expertise (Zorn et al, 2020), thus making peer review, publication, and dissemination more challenging.

Hiring and Appointment Evolution

In terms of experience, the candidates hired have changed over time. In 2024, one-third (33%) of new assistant professors had completed a postdoc, compared with 30% who moved directly from a PhD. Back in 2018, only 15% had a postdoc, while 38% were hired straight from their doctoral program (CRA Taulbee Survey 2025). In part this is due to the financial challenges of recent years, as faculty positions became more limited, and programs like the CRA Computing Innovation Fellows were developed to provide substantial post-doc opportunities in computing during difficult budgetary shortfalls.

Some computing faculty have split appointments — between academic units to facilitate collaborative, multidisciplinary work (Agapie et al, 2024, Pollack & Snir 2008) or with industry to apply their innovations at scale (Boules et al, 2016, Brachman, Schooler, and Wright 2025; Morrisett et al., 2019) – which can increase broader, societal impact, but also decrease time for other P&T oriented considerations.

Teaching and Mentoring Expectations

Computing has experienced surging enrollments (Computing Research Association, 2017), which have led to increased course sizes for introductory, core, and some sub-area elective courses. Additionally, faculty and instructional staff have been tasked with creating online and hybrid courses in addition to their existing course loads. The increase in use of AI by students, sharing of course materials online, and other changes to the technological and social spaces of computing education add pressure on faculty to develop new models of assessment that allow for learning goals to be achieved in an environment that protects academic integrity.

Increased Quantity of Publications

Despite the computing community emphasizing quality over quantity (Friedman & Schneider, 2015; Vardi, 2015), the volume of computing research dissemination has increased. The increasing breadth of the field coupled with large numbers of publications means that some peers may not be able to evaluate a published corpus fully and critically. The resulting increased use of counting publications may be a contributing factor to a notable increase in predatory "write-only" journals, collusion rings in reviewing, "ghost" conferences, plagiarism, incremental result publication, and submissions written (and often reviewed) largely by AI (Aiken et al., 2023).

Combined with preprints that are not peer-reviewed, it is possible to build a seemingly impressive record of publications based on dubious or fraudulent items. CRA provides considerations for ethical and responsible publication practices that should be considered by all computing researchers in making decisions about what, how often, and where to publish their work (Aiken et al., 2023, Cranor et al., 2024).

Engagement with Industry, Government, and Non-Profit Organizations

Engagement between computing researchers in academia and industry is vital for advancing both fundamental research and real-world innovation. The rapid pace of technological change means that close collaboration with industry provides academics with access to cutting-edge tools, large-scale datasets, and pressing practical challenges that can shape impactful research agendas. Joint appointments, consulting roles, and industry advisory positions enable faculty to integrate insights from commercial practice into their teaching and scholarship, ensuring that academic work remains relevant and forward-looking. Similarly, sabbaticals and other forms of leave from the university offer opportunities for deep, sustained collaboration that can lead to transformative technologies, while also fostering knowledge transfer in both directions.

From the industry side, supporting these engagements is equally important. Structured programs that welcome academic researchers, through sponsored sabbaticals, collaborative research projects, or shared infrastructure, help bridge the gap between theoretical advances and their deployment in real-world systems. Such partnerships not only accelerate innovation but also help to cultivate a workforce and research culture that are both scientifically rigorous and industry-ready. Industry researchers who engage with academia, support these kinds of activities, and further research translation should be recognized and rewarded for their efforts.

For more information on these challenges, please see the CRA report on working with industry and dual appointments.

ASSESSING QUALITY AND IMPACT OF RESEARCH

The fundamental basis for achievement in research is the impact of one's ideas and scholarship on the field. It is the expectation that scholarly work will make scholarly contributions while also considering broader impacts beyond computing. Scholarly contributions add new knowledge to the field. These contributions may be engineering oriented (accomplishing something with a computational artifact that could not be accomplished previously), empirical (demonstrating a phenomenon through observation, experimentation, or the analysis of data), methodological (creating new or synthesizing best practices to facilitate research), or theoretical (new insights, frameworks, models, or theories). For the purpose of evaluating a computing researcher for hiring or promotion, it is the expectation that they make substantial scholarly contributions and that their impact is appropriate for the level of promotion. Both aspects can be documented, but a proper analysis should go beyond counting archival publications. CRA strongly encourages quality of scholarly work and its impact over quantity of publications (Friedman & Schneider, 2015, Amato et al., 2025).

In computing disciplines, prestigious conferences involve intense peer review for which the program committee must be satisfied, as in a high-quality academic journal, that the claims of scholarly contribution are validated and that the significance is great enough to meet the standards of the venue. The measure of the impact of any individual paper is at least partially embodied in the quality of the publication, i.e., if the publication's standards are high, then the significance is presumed to be high. However, not all papers in high-quality publications are of great significance, and high-quality papers can appear in lower-quality venues. Thus, the indirect approach to assessing impact by assessing the quality of the venue is useful, but not definitive.

Similarly, given the speed of progress in computing, researchers in some subdisciplines may find it useful and/or necessary to share their ideas before peer review in another venue (e.g., arXiv). A challenge then emerges when, for example, a pre-print may have markers of impact (e.g., high citation and download counts) that are stronger than those of its eventual peer-reviewed version. In such cases, authors and reviewers alike should take care to cross-reference these documents, disseminate updates, and generally make the community aware of the final peer-reviewed product.

Scholars may also contribute to the development of significant artifacts (e.g., new computational frameworks or algorithms, databases that allow for other analyses, etc). While editorial boards and the leadership of scholarly societies have begun to establish mechanisms for citing such work, it is not yet standard practice. Thus, traditional publication-oriented mechanisms for measuring the impact of such artifacts may be limited. Instead, in appraising this work, reviewers might consider downloads of the artifact, number of users, scoring from open source systems, granting of patents and licenses, and so on. However, just as with traditional publications, there is possibly no way to measure quantitatively the impact of such artifacts. Rather, expert academic opinion should be used to judge quality in these cases. Whenever possible, such analysis should also consider the scholarly and intellectual merits that contributed to the development of the artifact. Specifically, it is possible to write a valuable, widely-used piece of software, induce a large number of downloads, but not make any significant intellectual contribution.

EVALUATION OF TEACHING AND MENTORING

Computing has expanded in scope by including more design, ethics, and social science teaching, as well as increased levels of architecture and abstraction. At the same time, interest in computing has grown substantially with traditional computer science majors ballooning alongside the rise in new majors in data science, artificial intelligence, informatics, human-computer interaction, game design and development, software engineering, and bioinformatics, just to name a subset. These booming enrollments have necessarily changed the size and style of specific computing courses at many institutions. Similarly, expectations within industry for supervising interns and within academia for research supervision of students and postdoctoral trainees are highly differentiated by organization and type of research. Variability across institutions as well as sub-disciplines should be considered when assessing teaching and mentoring contributions. The specific expectations for the candidate should be shared with all those making recommendations and decisions about the candidate.

Given the well-documented bias observed repeatedly in teaching evaluations (Deo, 2015; Kreitzer & Sweet-Cushman, 2021; Mengel et al., 2019) and the continued challenges faced by underrepresented groups in computing, in particular (Aragón, Pietri, & Powell, 2023; Gordon & Alam, 2021), we recommend using teaching evaluations sparingly. Certainly, teaching evaluations do have some merit and can provide useful feedback to the instructor (Boring & Ottoboni 2016). However, other mechanisms for evaluating teaching and mentoring should be considered, either alongside or instead of traditional teaching evaluations. The following are options, although not a complete list:

- **Teaching philosophy statements** can be used to demonstrate an instructor's overall understanding of pedagogical approaches and be forward-looking to how they plan to apply such an approach in their teaching. These are particularly useful for junior faculty or new hires who have not yet had an opportunity to teach or mentor.
- **Reflective teaching statements** consider specific experiences during the review period and how the instructor might assess and revise their own practices. They should not focus on general ideas about teaching or a philosophy, but instead reflect on how certain approaches and practices worked well and where improvements can be made.

- **Peer evaluation:** Observations and evaluations by other faculty or teaching experts on campus can provide valuable information on teaching quality and actionable feedback on teaching strategies.

- **Undergraduate and Graduate Student Mentoring:** Researchers in industry and governmental labs may demonstrate mentoring through interns or as serving as outside committee members for graduate student theses. Faculty frequently mentor students beyond their own PhD students, by serving as committee members and/or informal mentors to graduate students. Similarly, many faculty advise undergraduate students through independent study, class projects, or student employment. All of these types of mentoring should be evaluated through evidence of output (e.g., inclusion of students as co-authors on publications) as well as through self-report (e.g., evaluative commentary by the trainees themselves) when possible.

- **Other evidence** might include evidence of student learning outcomes, teaching awards, feedback from trainees (e.g., interns, post-docs), graduation rates or dissertation awards from PhD students, or evidence of broad dissemination and impact of teaching materials and instructional activities.

EVALUATION OF SERVICE

As previously noted, conferences, rather than journals, are the primary venues for disseminating high-impact research for many subfields, and even in those areas with heavy journal participation, conferences still play an outsized role in knowledge dissemination. Consequently, service on conference organizing committees, technical program committees, and editorial boards for conference proceedings is as important, if not more so, than traditional journal editing in other fields. This service requires extensive peer review, curation of cutting-edge work, and shaping the direction of the discipline.

Similarly, computing researchers often contribute to the broader community by managing open-source code and data repositories, which serve as critical infrastructure for research reproducibility and benchmarking. Such efforts parallel the stewardship of archives or scholarly databases in other fields and should be recognized as an important service to the discipline.

Another distinguishing feature of service in computing is its close relationship with industry and standards-setting organizations. Many computing researchers serve on industry advisory boards for academic programs, act as advisors to government and non-profit organizations, or lead initiatives that translate academic research into real-world technologies, ensuring that innovations developed in research labs are deployed responsibly and effectively. Participation in technical standards bodies (e.g., IEEE, ACM, or W3C) and in influential community-driven standards efforts — such as the Message Passing Interface (MPI), a widely used standard for parallel computing — represents another important form of service and venue for dissemination, as these groups set the protocols and guidelines that shape global computing practices and foster interoperability at scale.

Additionally, organizing hackathons, coding competitions, and other hands-on events bridges service and education by fostering skill-building, community engagement, and broadens participation in computing.

Taken together, these activities underscore the unique ecosystem of service in computing, which blends academic rigor, community infrastructure, and direct impact on industry, education, and society.

GUIDANCE FOR UNIT ADMINISTRATORS REGARDING EXTERNAL EVALUATION

Universities in particular but also some government and non-profit labs and companies rely on external evaluation not only for hiring but also in promotion decisions. In particular, one frequently relied upon means of assessing quality and impact in hiring and promotion is by letters of evaluation from scholars at peer institutions, including faculty at other universities and researchers at leading industry and government research centers.

Over the years, the number of letters required at some institutions has ballooned to an untenable point, creating a substantial burden across the field. We therefore **recommend soliciting the minimal number of high-quality letters, typically around three for hiring and four or five for promotion, with an emphasis on the fit and quality of the assessment over the volume of writers.** It is expected that evaluators will have access to and knowledge of all scholarly work, including but not limited to publications, artifacts, and databases/datasets.

In the case of academic promotion, including tenure, we strongly suggest that universities allow at least some letters from researchers in government labs and industry. Some of the field's best researchers work in these institutions, and they are more than qualified to assess research contributions. While they may not have the expertise to assess teaching, they also are often highly involved in traditional academic service activities and able to assess those as well. Furthermore, academic-industry-government collaborations occur regularly in our field and should be recognized as a strength in assessing a scholar's case (Morrisett et al., 2019). Similarly, collaborations outside of computing are to be applauded when appropriate to the research. We encourage committees to consult the CRA/CCC report on interdisciplinary scholarship in computing (Agapie, Wynewski, & Siek 2025) as part of their assessment of interdisciplinary computing research.

We encourage units to provide letter writers with some context on the resources and teaching loads of the candidate to calibrate their evaluation. Units provide different resources and teaching loads to researchers based on the type of institution and its goals (e.g., information unit start-up funding for faculty tends to be smaller than computer science; private schools typically provide more robust hiring packages than public, the subdiscipline they contribute to, and the type of research the candidate may conduct. In addition, units have different expenses (e.g., graduate student costs, overhead expenses, technology costs, etc.). For example, the current cost of supporting graduate students may impact how many graduate students can be funded on a grant or the amount of time the candidate can charge their own salary. Thus, we encourage units to consider providing information about this, particularly if it varies significantly from the average, to external evaluators.

Letter writers should also be provided with representative publications and artifacts prepared by the candidate. A small number of representative publications can contribute to a quality and impact focused evaluation of a large corpus of publications rather than the counting of publications that might otherwise take place. Artifacts encapsulate information that cannot be captured easily in a traditional publication. They are frequently shared with the community as open source contributions. Many artifacts, such as software and datasets of dynamic nature, may need evaluators to conduct alternative forms of evaluation. Further, some artifacts are so complex that it is difficult to explain all their characteristics; it is better to observe them. Data or other information needed to reproduce results should also be provided with the artifacts.

EVALUATING INTERDISCIPLINARY CONTRIBUTIONS

Computing research continues to grow in depth and breadth as technology is increasingly leveraged to tackle our most complex societal problems; as such, interdisciplinarity also continues to grow in importance across our field. Yet, interdisciplinary computing researchers are still faced with disciplinary challenges when being evaluated for tenure and promotion.

- **Implementation of Reviews.** If the interdisciplinary norms of the candidate's field differ significantly from the tenure and promotion guidelines of the university/department, we recommend that the unit have a process for transparently documenting how the candidate will be fairly evaluated based on the success criteria in their interdisciplinary field (e.g., in the Chair letter or department/school policies).

- **Broadening Dossier Expectations.** A variety of types of artifacts may be more commonly accepted as research excellence when working with people outside computing, and these should be considered as first-order contributions when appropriate. Additionally, team science is frequent in cross-cutting and interdisciplinary work, meaning that contribution statements may help explain the particular elements of importance for the candidate's dossier.

- **Interdisciplinary Service and Leadership.** Interdisciplinary computing researchers may not meet the criteria for national and international prominence in a singular discipline, which is often required for tenure and promotion. Instead, they may demonstrate leadership across multiple adjacent communities at a lower level than if they focused on a single community.

REFERENCES

Agapie, E., Wisniewski, P., Siek, K., & Griffin, H. (2024). *Catalyzing Interdisciplinary Computing Research: Best Practices for Researchers.* Computing Research Association.https://cra.org/wp-content/uploads/2024/08/Catalyzing-Interdisciplinary-Computing-Research-Best-Practices-for-Researchers_August-2024_FINAL.pdf

Aiken, A., Cranor, L., Lopresti, D., & Zorn, B. (2023). *Report of the CRA Working Group on Research Integrity.* Computing Research Association. https://cra.org/wp-content/uploads/2023/08/Report-of-the-CRA-Working-Group-on-Research-Integrity_August-2023.pdf

Amato, N., Daescu, O., Gini, M., Hayes, G., Larson, K., Li, W., Lin, M., Lopresti, D., Siek, K., Spafford, G., & Zorn, B. (2025). *Unique Considerations for Evaluating Computing Researchers.* Computing Research Association. https://cra.org/wp-content/uploads/2025/07/Unique-Considerations-for-Evaluating-Computing-Researchers.pdf

Aragón, O. R., Pietri, E. S., & Powell, B. A. (2023). Gender Bias in Teaching Evaluations: The Causal Role of Department Gender Composition. *Proceedings of the National Academy of Sciences, 120*(4), e2118466120. https://doi.org/10.1073/pnas.2118466120

Boring, A., & Ottoboni, K. (2016). Student Evaluations of Teaching (Mostly) Do Not Measure Teaching Effectiveness. *ScienceOpen Research.* https://doi.org/10.14293/S2199-1006.1.SOR-EDU.AETBZC.v1

Boules, N., Douglas, K., Feldman, S., Fix, L., Hager, G., Hailpern, B., Wright, H., & Others. (2016). *The Future of Computing Research: Industry–Academic Collaborations.* arXiv. https://arxiv.org/abs/1606.09236

Brachman, R., Schooler, E., & Wright, H. (2025). *Evolving Trends in Dual Appointments: Shaping the Future of Talent-Sharing Between Academia and Industry in Computer Science.* Computing Research Association. http://cra.org/wp-content/uploads/2025/07/Evolving-Trends-in-Dual-Appointments-Report.pdf

Computing Research Association. (2017). *Generation CS: Computer Science Undergraduate Enrollments Surge Since 2006.* Computing Research Association. https://cra.org/wp-content/uploads/2017/02/Generation-CS.pdf

Computing Research Association. (n.d.). *CRA Taulbee Survey.* Retrieved September 2025, from https://cra.org/resources/taulbee-survey/

Cranor, L., Hazelwood, K., Lopresti, D., & Stent, A. (2024). *Conference Submission and Review Policies to Foster Responsible Computing Research.* Computing Research Association. https://cra.org/wp-content/uploads/2024/07/Report-Conference-Submission-and-Review-Policies.pdf

Deo, M. E. (2015). A Better Tenure Battle: Fighting Bias in Teaching Evaluations. *Columbia Journal of Gender and Law, 31*(1), 7–43.

Friedman, B., & Schneider, F. B. (2015). *Incentivizing Quality and Impact: Evaluating Scholarship in Hiring, Tenure, and Promotion.* Computing Research Association. https://cra.org/resources/best-practice-memos/incentivizing-quality-and-impact-evaluating-scholarship-in-hiring-tenure-and-promotion/

Gordon, N., & Alam, O. (2021, March). The Role of Race and Gender in Teaching Evaluation of Computer Science Professors: A Large-Scale Analysis on RateMyProfessor Data. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education* (pp. 980–986). ACM. https://doi.org/10.1145/3408877.3432379

Mengel, F., Sauermann, J., & Zölitz, U. (2019). Gender Bias in Teaching Evaluations. *Journal of the European Economic Association, 17*(2), 535–566. https://doi.org/10.1093/jeea/jvx057

Kreitzer, R. J., & Sweet-Cushman, J. (2021). Evaluating Student Evaluations of Teaching: A Review of Measurement and Equity Bias in SETs and Recommendations for Ethical Reform. *Journal of Academic Ethics.* Advance online publication. https://www.researchgate.net/publication/349185345_Evaluating_Student_Evaluations_of_Teaching_a_Review_of_Measurement_and_Equity_Bias_in_SETs_and_Recommendations_for_Ethical_Reform

Morrisett, G., Patel, S., Rexford, J., & Zorn, B. (2019). *Evolving Academia/Industry Relations in Computing Research.* arXiv. https://arxiv.org/abs/1903.10375

Patterson, D., Snyder, L., & Ullman, J. (1999). Evaluating Computer Scientists and Engineers for Promotion and Tenure. *Computing Research News.* https://archive2.cra.org/uploads/documents/resources/bpmemos/tenure_review.pdf

Pollack, M. E., & Snir, M. (2008). Best Practices in Promotion and Tenure of Interdisciplinary Faculty. *Computing Research News.* https://cra.org/crn/2008/09/best_practices_in_promotion_and_tenure_of_interdisciplinary_faculty/

Vardi, M. Y. (2015). Incentivizing Quality and Impact in Computing Research. *Communications of the ACM, 58*(5), 5–5. https://doi.org/10.1145/2714568

Zorn, B., Conte, T., Marzullo, K., & Venkatasubramanian, S. (2020). *Evolving Methods for Evaluating and Disseminating Computing Research.* arXiv. https://arxiv.org/abs/2007.01242