

AI and TCS: The Next Six Months

Working Group Report

September 2026 Working Group*

September 28, 2026

Abstract

This report is the product of a two-day working group held during September 9–10, 2026 at the Simons Institute. The participants considered four questions around how the TCS community can adapt in light of ongoing advances in AI. Via a series of facilitated small and large group discussions, they arrived at the twelve concrete near-term recommended actions presented here, which received broad consensus from the group. The questions as well as the discussions were informed by 134 responses to a broadly circulated survey.

Contents

1	Introduction	2
1.1	Choice of Questions and Participants	3
1.2	Workshop Process	4
1.3	Overview of the Recommended Actions	4
2	Attribution and Disclosure	6
2.1	Action 1: Standardized AI Methodology Section	6
2.2	Action 2: Tracing the Lineage of AI-Generated Ideas	9
3	TCS Conferences	10
3.1	Action 3: Conference Rules and Norms	10
3.2	Recommended Experimental Measures	11
3.2.1	Action 4: arXiv Posting with Submission	12
3.2.2	Action 5: Video Talks with Submissions	12
3.3	Action 6: Unified Public Repository for Conference Rules and Norms	13
4	Hiring	14
4.1	Action 7: Best Works and Technical Talks in Interviews	14
4.2	Action 8: Additional Postdocs	14
4.3	Action 9: Workshops for Identifying New Theory Directions	15
5	Graduate Education and Training	16
5.1	Action 10: “TCS in the Age of AI” Series at the Simons Institute	16
5.2	Action 11: Opportunities for Training and Evaluation without AI Assistance	17
5.3	Action 12: Coverage of Reasonable LLM Expenses from Research Grants	18
6	Conclusion	18
A	Examples of AI Methodology Sections	19
B	Workshop Format Details	21

*Organizers Tony Feng (UC Berkeley, Google DeepMind), Zeph Landau (UC Berkeley), Amit Sahai (UCLA), and Nikhil Srivastava (UC Berkeley). See Section 1.1 for participants.

1 Introduction

Recent developments in the reasoning capabilities of AI models have made a remarkable impact on both the scientific field and on the research community of Theoretical Computer Science (TCS). Large Language Models (LLMs) have unlocked new research possibilities, while simultaneously putting severe stress on the field’s pre-existing infrastructure, practices, and culture. The pace of this change and the uncertainty surrounding it have led to both excitement and anxiety, especially among junior researchers who are the most affected.

These events have prompted many discussions and introspections—in various formats including essays, podcasts, panels, seminars, etc.—on the nature and goals of the human endeavor to understand mathematics and theoretical computer science, and the rapidly evolving role of AI in this endeavor. Such discussions have often resulted in eloquent affirmations of human understanding and its irreplaceability, but have often left open the question of *precisely what adaptations* the community should take in response to AI.

The purpose of this document is to take a first step towards addressing this question, by formulating a collection of *concrete actions* to undertake in the short term. Specifically, we will describe recommendations, some of experimental nature, to try implementing *in the next six months*. In order to properly contextualize these actions, we will additionally describe their motivations and envisioned implementation.

We do not claim that our recommendations are sufficient to address all problems presented by AI, even in the short term. However, our list at least offers a *shared baseline* for action. Many departments, conferences, and funding bodies are already grappling with these questions; a common, coordinated set of actions would provide clarity in a chaotic time.

We briefly discuss the provenance of this list of recommendations. It was assembled during an intense two-day workshop at the Simons Institute during September 9 and 10, 2026 consisting of 28 invited members of the TCS community, 4 organizers, and one professional facilitator, and informed by a widely distributed survey answered by 134 community members. While there was a deliberate attempt to choose a diverse group of participants, we acknowledge that we did not achieve a perfect representation of the TCS community. Notably, the recommendations in this paper *all received broad support* from the full group of 32 workshop participants, in a sense described precisely in Section 1.2.

We stress that our action items should be viewed only as working proposals, to be considered and adjusted by the larger community. The next step is *implementation*, and turning these recommendations into real actions will require a concerted effort. We are optimistic that our proposals will resonate with many members of the larger community, who will then find ways to help implement them, possibly with modifications.

We also emphasize that our recommendations are meant to be a first step, with a limited lens into the future. Some of them are suggested only as experiments, on the way towards finding a longer-term equilibrium. Indeed, the workshop was itself an experiment in mustering consensus around actions for adapting to AI, an issue which has been divisive and destabilizing. If successful, the workshop format could be a useful format for future iteration. We envision that a regular series of such events will be necessary for the community to coevolve with AI.

This report is centered around the following four questions:

1. How should attribution and disclosure be handled for AI-assisted research?
2. How should TCS conferences shift, in submission, reviewing, and dissemination?
3. How should early career researchers be evaluated in academic hiring?
4. How should PhD education and training change?

The main content of this report is a list of suggested actions for each question.

Details about the choice of questions, the choice of participants, and the process used to generate the recommendations are provided in the subsections below. The actions themselves appear in Sections 2–5.

1.1 Choice of Questions and Participants

On July 14, 2026, the organizers sent out a survey to all current and former participants in the Simons Institute’s programs during 2012–2026, which is about 3,000 people, with an instruction to share with others. The goal of the survey was to identify issues of interest to the TCS community, asking community members to respond to the prompt:

In response to the ongoing advances in AI, what practices will lead to better scientific and socio-logical outcomes for the TCS community?

The organizers received 134 responses to the survey. Guided by these responses, they identified the above four questions of interest. There were other prominent issues in the survey which were not part of the focus of the workshop, namely: equity with regard to AI access and how the TCS community presents itself to the public. These important topics were excluded because the organizers felt that it would be difficult to come up with concrete actions to address them that members of the community could execute by themselves.

The organizers invited 54 people to attend the workshop, informing them of its focus on these questions, and asking them to accept the following value as the starting point for the discussion:

TCS research is fundamentally grounded in the purpose of human understanding (which requires human thinking, and human struggling with mathematics).

The organizers did their best to gather a diverse sample of the TCS community. The invitations were sent iteratively, attempting to balance across seniority, research area, institution, and gender, and to include recent PC/steering committee members as well as people active in the public AI/Math/TCS discourse. Of the 54 invitees, 22 were people who filled in the survey. Invitees were asked to commit to spending two full in-person days at the workshop. The short timeline and in-person constraint naturally biased the workshop towards participants who could come to Berkeley on short notice. Due to the small group format detailed below, the organizers made a conscious effort to limit the number of participants to 28.

A total of 28 people accepted the invitation; 19 people declined, and 7 did not reply. The final list of participants is listed here:

Noah Bergam (Columbia University), Simina Brânzei (Purdue University, Google), Moses Charikar (Stanford University), Cynthia Dwork (Harvard University), Sumegha Garg (Rutgers University), Parikshit Gopalan (Apple), Isaac Hair (UCSB / UCLA), Pooya Hatami (Ohio State University), Thore Husfeldt (IT University of Copenhagen), Sandy Irani (University of California, Irvine), Sanjeev Khanna (New York University), Dakshita Khurana (University of Illinois Urbana-Champaign), Pravesh Kothari (Princeton University), Jiayu Li (Massachusetts Institute of Technology), Shachar Lovett (University of California, San Diego), Peter Manohar (Institute for Advanced Study), Kunal Marwaha (University of Chicago), Jelani Nelson (Anthropic / on leave from UC Berkeley), Shayan Oveis Gharan (University of Washington), Prasad Raghavendra (OpenAI / on leave from UC Berkeley), Amit Rajaraman (Massachusetts Institute of Technology), Lev Reyzin (University of Illinois Chicago), Tim Roughgarden (Institute for Advanced Study / a16z crypto), Rocco Servedio (Columbia University), Min Jae Song (University of Chicago), Thuy-Duong Vuong (University of California, San Diego), David Woodruff (Google / on leave from Carnegie Mellon University), Rachel Zhang (UC Berkeley).

This list is comprised of 5 graduate students, 3 postdocs, 16 faculty, and 4 people from industry.

1.2 Workshop Process

Formation of Small Groups. The 28 participants were asked to indicate which questions they would be interested in thinking about, and 24 responded. Seven participants and one organizer were assigned to each question. Each participant who expressed a preference was assigned to a group they said they were interested in.

Workshop Format. The structure of the workshop was designed by the organizers in conjunction with a professional facilitator, Paul Sullivan. The structure consisted of repeated units of small group work in set teams of 8 followed by sharing the work with the larger group and collecting facilitated feedback that was incorporated into the small group work. The full schedule, the small group structure, and consensus measurement procedure are described in Section B, and may be useful for others wishing to organize such events.

Level of Consensus. The actions produced by the workshop are the outcome of the process described above. Each action was rated on a scale of 0 to 5 by all 32 participants during the final large group discussion. Most actions received zero votes of 0 or 1 and every action received at most 5 votes of 0, 1, or 2 out of a group of 32.

Writing. Each section of this report corresponding to one of the four questions was written by its working group of 8 people during the week following the workshop, with a round of feedback from the full group of 32. These sections were then edited and combined into a full report by the organizers. In the interest of preserving the authors’ careful work, minimal editorial intervention has been applied. The resulting document retains a multiplicity of voices and perhaps occasional quirks on the level of prose.

1.3 Overview of the Recommended Actions

We collect here brief descriptions of the recommended actions, with full details appearing in Sections 2.1–5.3.

1. **Standardized AI Methodology Section** (Section 2.1). Every publicly posted research paper should include a standardized AI Methodology section describing what the authors and AI tools contributed.
2. **Tracing the Lineage of AI-Generated Ideas** (Section 2.2). Authors should make their best effort to trace the lineage of AI-generated ideas, and explain their connections to prior work.
3. **Conference Rules and Norms** (Section 3.1). Calls for Papers should contain clear rules and guidelines, including the following.
 - (a) **AI use is permitted.** AI use in the ideas or writing of the submission is permitted, and reviewers may not penalize a paper for AI use as long as it does not inhibit human understanding.
 - (b) **Author attestation.** Submission platforms should require an attestation from the authors certifying that they understand and have verified the work they are submitting, and accept responsibility for the correctness and presentation of the work.
 - (c) **Requirement for an overview of results and key technical ideas.** Submissions must contain clear exposition, motivation of work and ideas, lineage of ideas, and effective presentation of the key ideas at the beginning of the paper, in the first 10–12 pages.
 - (d) **Poor exposition is sufficient grounds for rejection.** Papers may be rejected with a one-sentence review, regardless of the technical merit of the paper, solely because of exposition which is sufficiently unclear.
 - (e) **Submission limit.** Impose a limit (recommendation: 5) on how many submissions an author can have to a conference.

- (f) **AI assistance in reviewing, with disclosure.** Reviewers may use AI assistance to help check technical claims (proofs, references, consistency) but the assessment of novelty, significance, and interest must be the reviewer’s own.
- 4. **arXiv Posting with Submission** (Section 3.2.1; experimental). Conferences should require that conference submissions be published on the arXiv.
- 5. **Video Talks with Submissions** (Section 3.2.2; experimental). Conference submissions should require the submission of a talk video in addition to the paper.
- 6. **Unified Public Repository for Conference Rules and Norms** (Section 3.3). Establish a co-ordinated process whereby conferences establish standardized CfPs, along with an up-to-date public repository for publicizing this work.
- 7. **Best Works and Technical Talks in Interviews** (Section 4.1).
 - (a) **Emphasize quality over quantity.** Evaluate candidates on no more than **three** works of their choosing, with this fact explicit in the instructions for candidates and their letter writers.
 - (b) **Test technical depth.** Evaluate the depth of the candidates’ understanding of the aforementioned works more than was done normally until now.
- 8. **Additional Postdocs** (Section 4.2). Establish additional postdoctoral positions for the TCS community, a temporary measure that will be phased out after a few years.
- 9. **Workshops for Identifying New Theory Directions** (Section 4.3). First, organize a workshop (or several) with the goal of identifying a list of modern topics related to AI; and second, run a series of workshops on each of these topics to identify fundamental directions and establish foundations and agendas for future work.
- 10. **“TCS in the Age of AI” Series at the Simons Institute** (Section 5.1). Establish a recurring and widely accessible “TCS in the Age of AI” series at the Simons Institute.
- 11. **Opportunities for Training and Evaluation without AI Assistance** (Section 5.2). To ensure development of necessary skills for a successful career in academia, a PhD in TCS should partially include *training* and *evaluation* components where the use of AI is forbidden.
- 12. **Coverage of Reasonable LLM Expenses from Research Grants** (Section 5.3). Relevant professional organizations should investigate the barriers researchers have faced in obtaining access to the state-of-the-art AI tools.

Many of the actions above solve a common underlying problem and reinforce each other. Some cross-cutting goals are:

- *Human understanding.* Researchers are incentivized to thoroughly understand their AI-assisted work by Action 5 (conference submissions require videos) and Action 7 (interviews for hiring involve longer technical talks). Action 11 (opportunities for training students without AI assistance) aims to prepare PhD students to reach this level of understanding.
- *Managing the volume of research.* Researchers are incentivized to produce a small number of high-quality works rather than a large amount of AI slop both by Actions 3 and 4 (a limit of 5 papers per conference, with additional friction from the video requirement) and by Action 7 (a limit of 3 papers considered in hiring), which they then have a better chance of understanding.
- *Maintaining trust and transparency.* Action 1 (standardized AI methodology section) aims to normalize honest disclosure of AI assistance in research, which preserves an atmosphere of transparency and trust in the community during this period of rapid change. Action 3 makes this a requirement in conference paper submissions; conferences further state that AI use must not be penalized during reviews and demonstrate this through paper acceptances.

- *Adapting to the changing research landscape.* Action 8 proposes additional postdoc positions in the near term to give junior researchers time and space to develop the deeper level of human understanding needed in response to current AI advances and results, as well as to navigate changing norms. Action 9 proposes organizing workshops to articulate TCS research directions related to AI. Action 10 suggests hosting regular community-wide gatherings to treat both the scientific and sociological aspects of this transition.

Timing. All of the actions recommended in this report are framed to be executed as soon as possible, ideally in the next six months.

Acknowledgments

We are extremely grateful to our facilitator Paul Sullivan, who helped us design the structure of the workshop and whose guidance was essential for us to have productive discussions on this contentious topic. We thank Amy Ambrose and the Simons Institute staff for planning and organizing the meeting on very short notice, and Kristin Kane for feedback on the presentation. This workshop was funded by the Simons Institute’s Director’s Opportunity Fund.

2 Attribution and Disclosure

2.1 Action 1: Standardized AI Methodology Section

Action: Every publicly posted research paper should include a standardized *AI Methodology* section describing what the authors and AI tools contributed. We provide a template in Section 2.1 and examples of informative and uninformative AI Methodology sections in Appendix A. We prefer the word methodology to disclosure, as it sounds less punitive and emphasizes transparency of the research process and actual communication of methods.

Why: The goal of our proposed standard is to uphold three core values, described below. The first two directly support shared human understanding; the third sustains incentives to contribute to it. How we evaluate research contributions is likely to change over time, and documenting a clear record lets those judgments evolve without losing the factual basis for them.

1. *Documenting the scientific narrative.* A coherent scientific narrative connects new results to earlier ideas, helping the community understand them in context and build upon them. Maintaining it requires researchers to earnestly try to record the lineage of ideas, explaining precisely how they relate to previous work and where the new insights lie. Current AI models often fail to trace the origins of their results, producing results without any context, and researchers relying on them may be unable to correct these omissions when models combine ideas from areas outside their expertise. These missing connections fragment the scientific narrative and can inhibit human understanding and further scientific progress.

An AI Methodology section preserves a snapshot of the research process by recording which parts of a paper were produced by AI and by humans, which is relevant since AI is at the moment much more likely to produce arguments without appropriate attribution¹. This record gives others a starting point for retrospectively tracing any missing connections and integrating the result into the existing body of ideas and techniques, helping restore continuity in the scientific narrative.

2. *Trust and transparency in the research process.* Research thrives in an environment of trust, where people can exchange ideas and collaborate without questioning one another’s credibility. Normalizing

¹Humans are of course not perfect at attribution either, but they have well-understood failure modes and more accountability.

candid discussion of how AI models contributed to the research helps sustain this trust and provide useful data as research practices and authorship and credit norms evolve. A standardized AI Methodology section clarifies the role of human authors in producing the paper, reducing suspicion and guesswork.

3. *Recognition of contributions beyond idea generation.* Recognition of intellectual contributions is both an incentive and a value for research: researchers invest in long-term projects expecting their work to be recognized by the community. Recognition also connects researchers to those whose work they build upon, helping ideas spread and develop into shared understanding. AI separates forms of contribution that traditional authorship bundled together, making authorship alone less informative about the human work behind a paper. AI Methodology sections make these contributions explicit, and provide shared data for analyzing shifting authorship norms and revising how contributions are evaluated. In particular, these sections clarify how authors contributed to digesting, simplifying, and presenting results, providing a basis for intellectual recognition even when AI tools supplied the initial argument.

Overlapping with (1), recognition of contributions helps to organize the set of technical ideas in a field as clusters grouped by human beings, who have other attributes such as scientific perspective and taste. This is useful since if one wants to understand the context of an idea/technique, one could look at the prior work of the contributor(s).

Although AI disclosure/methodology sections are already common, many omit relevant details and are too vague to satisfy the above goals. Standardization helps authors provide a more informative account.

How: The proposed AI Methodology section has two components: (a) a checklist, indicating the categories in which AI tools contributed to the work; (b) a free form response, describing which tools were used² and how, elaborating on the binary responses in the checklist. Suggested locations for this section are after the introduction or at the end of the main text, before the appendices.

The checklist should address the following six categories as they apply to at least one of the main results in the paper. Categories may overlap. We recommend that the authors indicate all that apply.

For the first four categories, a checked box indicates meaningful influence on the research or its exposition, in that a comparable contribution from a person would previously have been recognized through an offer of co-authorship. The emphasis is on the significance of the contribution, rather than the number of interactions with a tool or the amount of generated text. We recognize that thresholds vary, but given the unprecedented nature of AI models being able to make substantial contributions, the co-authorship comparison anchors it to judgments researchers already make in practice.

1. *Asking the question:* AI's influence on the formulation of the main questions studied in the paper. Examples include proposing a model or definition, identifying a relevant assumption, or suggesting a conjecture. This category concerns decisions about what to study. The threshold for checking this box is if a hypothetical human who had made similar contributions would have been offered co-authorship.
2. *Coming up with the approach:* AI's influence on the overall conceptual approach to proving the paper's main results. Examples include suggesting a reduction, choosing a relaxation, or identifying a sequence of intermediate claims. This category concerns the organizing plan of the proof. The technical ideas that make that plan work belong in the next category. The threshold for checking this box is if a hypothetical human who had made similar contributions would have been offered co-authorship.
3. *Proof development:* AI's influence on the important technical ideas underlying the paper's main results. Examples include a proof technique for an important technical lemma, or a crucial inequality. As a guide, these are the ideas the authors would choose to highlight in a 50-minute technical talk on the work. The category is intended to capture central technical contributions, rather than every technical choice made in completing a proof. Optimizing constants or proving a particular lemma belongs here

²This choice is motivated by the goals of transparency, good scholarship, and accurately tracing the sources of ideas.

only when that work is itself one of the paper’s main contributions. The threshold for checking this box is if a hypothetical human who had made similar contributions would have been offered co-authorship.

4. *Writing and exposition:* AI’s influence on the substantive explanations and presentation of the paper. Examples include drafting a proof overview, organizing an argument, or developing an example for exposition. Routine copy-editing belongs in the final category below. Contributions to research ideas should also be reported in the relevant categories above, regardless of who wrote the final text. The threshold for checking this box is if a hypothetical human who had made similar contributions would have been offered co-authorship.
5. *Checking for bugs:* AI’s contribution to identifying or correcting consequential errors in the paper. Examples include exposing a circular argument, finding a missing hypothesis, or correcting a computational error that affects a reported result. Typographical corrections belong in the final category below. The threshold for checking this box is if the AI made any substantial contribution along these lines.
6. *Other supporting tasks:* AI’s assistance with other aspects of the research or preparation of the paper, which could have been done without AI. Examples include copy-editing, making figures, conducting numerical experiments, proofs of minor lemmas, searching the literature, and understanding past literature. Where these activities also influence the questions, strategy, or important technical ideas, they should additionally be reported in the corresponding categories above. The threshold for checking this box is if the AI made any substantial contribution along these lines.

We note that any of the first four categories could be the main contribution of a typical paper. For instance, when a paper makes progress on a well-known open problem, there may not be a non-trivial contribution in the paper in the first category. On the other hand, when a paper proposes a new model or a new mathematical object, the main contributions may only be in formulating the question.

For each checked category, in the “Details” portion of the AI methodology section, authors should identify the relevant contribution and describe what they supplied to the AI tool, what the tool proposed, what was retained or revised, and how the resulting claims and attributions were checked. The level of detail should be proportionate to the contribution.

Rationale for a checklist. We decided to include a checklist at the beginning of the AI Methodology section because it is much easier to fill in (for people who are less inclined to write a full methodology section) and makes it much harder to omit key information, a common failure mode we observed in existing AI disclosures on arXiv. The checklist also primes authors to write in more nuance about each of the six included items.

The checklist was developed by reviewing existing AI Methodology sections from arXiv papers and assessing their strengths and weaknesses (see Appendix A). Our recommendation is also inspired by the CRediT taxonomy [1] used by the broader scientific community to standardize descriptions of research contributions.

Incentive Alignment. Program committees should incentivize people to include such disclosures by making them required in submissions, with a clear statement that they will not be penalized for using AI. This neutrality would be effectively demonstrated to the community if many papers which acknowledge AI use are accepted to conferences, reducing the incentive to lie or mislead.

Template for AI Methodology section

The following is a template for the AI Methodology section, which is also available as a downloadable L^AT_EX file on the Simons Institute website³, along with the operational descriptions of the categories above.

³<https://simons.berkeley.edu/media/31229/download>

Instructions for authors.

1. Complete the checklist below to indicate whether AI tools were used for each activity. Replace each [Yes/No] with **Yes** or **No**, answering **Yes** if the AI made a contribution. For the first four boxes, this means at the level for which you would consider a human for co-authorship, and for the last two any substantial contribution is enough.
2. Replace the bracketed guidance under *Details* with your own account.
3. Delete the *Instructions for authors* heading and this entire numbered list from the completed paper.

Activity	AI used
Asking the research question	[Yes/No]
Coming up with the approach	[Yes/No]
Proof development	[Yes/No]
Writing and exposition	[Yes/No]
Checking for bugs	[Yes/No]
Other supporting tasks	[Yes/No]

Details (elaborating on the points above). [Replace this bracketed guidance with a free-form story of how the work developed and how AI was involved. One might describe what was asked or tried, what was useful or unsuccessful, what the authors contributed or changed, and how the resulting arguments, text, or code were checked. Name the models or tools used, including versions where known. Focus on the parts relevant to the work; there is no need to discuss every checklist item separately.]

Who needs to implement this action: Authors of papers and program committees.

2.2 Action 2: Tracing the Lineage of AI-Generated Ideas

Action: Authors should make their best effort to trace the lineage of AI-generated ideas, explain their connections to prior work, and give appropriate credit prominently when introducing and motivating their work. This requires particular care, as current AI models often fail to attribute ideas correctly, and can even reproduce arguments from the literature without recognizing them as such (examples of this “subconscious plagiarism” are documented in the case study [5]).⁴

Why: This action is motivated by preserving the scientific narrative and accurately tracing the lineage of ideas, as discussed in Section 2.1. Having mathematical arguments which exist in a vacuum obscures the scientific narrative and is demoralizing to the researchers whose work it uses without proper citation. LLMs make it much easier to generate such arguments, especially when they draw on ideas outside the researcher’s expertise.

How: Authors should be aware of this particular pitfall of being overly reliant on AI. They should put in their best effort to discover the source of the ideas they use; this should of course include using LLMs to search through the literature. We encourage authors to consult the original papers they discover and ensure that their attributions are accurate. A related work section is not the appropriate place for lineage discussion; lineage should be explained when the relevant ideas or results are first discussed, ideally early in an Introduction.

⁴This is of course not meant to exclude tracing the lineage of human-generated ideas.

Program committees should make it clear in their Call for Papers that best-effort tracing of the lineage of ideas is an important aspect of quality of exposition. Section 3.1 collects the other Call for Papers rules we recommend.

Who needs to implement this action: Authors and program committees.

3 TCS Conferences

3.1 Action 3: Conference Rules and Norms

Problem: With the recent dramatic increase in AI capabilities there has been an accompanying dramatic increase in the number of papers posted to online repositories such as arXiv, ECCC, and the Cryptology ePrint Archive. Conferences have also seen dramatic increases in the number of submitted papers, and this trend seems to be accelerating. Many of these papers are poorly written and hard for humans to understand. It is increasingly possible that some authors do not understand the contents of their own papers.

Action: Calls for Papers should contain clear rules and guidelines. Below are some rules that we believe should be incorporated into future CfPs. We note that the STOC 2027 PC has decided to adopt almost all of our guidelines below.

1. **AI use is permitted.** AI use in the ideas or writing of the submission is permitted, and reviewers may not penalize a paper for AI use as long as it does not inhibit human understanding. Further suggestions regarding how authors should transparently describe AI usage can be found in Section 2.1.

Why: The reality is that the way our community produces papers has changed rapidly, and AI use is now widespread as part of the research and writing process. This should not be penalized as long as the resulting papers are aligned with the goal of deepening human understanding.

2. **Author attestation.** Submission platforms should require an attestation from the authors certifying that they understand and have verified the work they are submitting, and accept responsibility for the correctness and presentation of the work. This could take the form of a signed checkbox.

Why: For the continued health of the community it is important that authors understand and explicitly take responsibility for their papers.

3. **Requirement for an overview of results and key technical ideas.** We expect submissions to contain clear exposition, motivation of work and ideas, lineage of ideas, and effective presentation of the key ideas at the beginning of the paper, in the first 10–12 pages.

Why: This has been the community standard for many years. We wish to make this expectation explicit in the rules, even for AI-assisted papers.

4. **Poor exposition is sufficient grounds for rejection.** Theoretical computer science CfPs should contain clear and unambiguous language stating that papers may be rejected with a one-sentence review, regardless of the technical merit of the paper, solely because of exposition which is sufficiently unclear so as to significantly hinder an interested and informed reader from understanding the context, motivation, and underlying ideas.

Why: Theoretical computer science has a long and cherished tradition of valuing clear communication and presentation of ideas. The purpose of a TCS conference paper is not just to stamp a theorem as having been proved, but also to explain the work and its context and significance to the community of present and future human TCS researchers.

5. **Submission limit.** Impose a limit (recommendation: 5) on how many submissions an author can have to a conference. This recommendation is not meant to subvert or undermine the jurisdiction of

the PC chair. There could be exceptional cases, and those could be handled by the PC chair as deemed appropriate.

Why: This action helps with several problems: (1) reduces PC load, to allow informed decision making; (2) keeps conference size manageable so that the most important results get visibility; (3) reduces the mental load on students and junior authors by signaling that there is no expectation to submit an unreasonable number of papers to top conferences.⁵ The recommended limit of 5 was chosen after considering top percentile author submission counts in recent years: we do not expect this to be a constraining limit for most authors, and we certainly do not intend for it to set an expectation of five submissions. The specific limit will likely need to be altered as future conditions warrant.

6. **AI assistance in reviewing, with disclosure.** Reviewers may use AI assistance to help check technical claims (proofs, references, consistency) but the assessment of novelty, significance, and interest must be the reviewer’s own. A reviewer is responsible for everything in the submitted review. The review should include a disclosure report, as a confidential comment to the PC, that specifies the tool used, and the parts of the review that are based on the tool use. To ensure that the confidentiality of submissions is not compromised, the implementation of this idea could use one of the following approaches: (a) arXiv posting with submission requirement (see Section 3.2.1), or (b) an explicit permission from author(s) for their paper to be reviewed using AI tools, or (c) by setting up a dedicated LLM server that retains no submission data for training or otherwise. We do not recommend allowing for authors to opt-out, as this creates a two-tiered reviewing process.

Why: Submission volume is outpacing expert review capacity. Conferences are receiving record numbers of submissions, AI-assisted to varying degrees, with proofs that are potentially longer and more intricate than before. Verifying them takes expert time that has not grown, and PC members cannot cover this volume within the short reviewing timeframe. This negatively impacts overall quality of decision-making. AI use can help verify complex technical details and spot potentially fatal issues. It also allows human reviewers to focus their energy on novelty, significance, and interest aspects of the results.

Concern: A concern is that multiple reviews of a paper may be shaped by a single LLM tool. This proposal strongly requires that AI assistance is limited only to checking technical details, and the judgement/recommendation should be the reviewer’s own. Conferences should collect data on which tools were used by reviewers, and on whether they in fact produced overly similar voices across the reviews of a paper, to refine implementation of this idea if needed.

Who needs to implement this action: Steering committees, PC chairs, and workshop participants with administrative connections to conferences.

How: This group will contact relevant steering committees and PC chairs to obtain buy-in for these policies, motivating the policies by describing the “Problem” and the “Why” from above. PC chairs will add clear language stating these policies (or, if disagreement, clear policies addressing each of the issues about exposition, author responsibility, and reviews) to CfPs for their conferences.

Where: The CfPs of TCS conferences. These rules should be collated and stated in the public repository for conference rules and norms. Such rules may link to this white paper for justification.

3.2 Recommended Experimental Measures

In this section, we give a couple more experimental measures we believe would be fruitful for conferences to try in the coming months. As part of implementation, conferences should evaluate the impact of each experimental measure, and evaluate whether the measure should be made into a more permanent rule.

⁵As we mention in Section 4.1, we recommend that hiring decisions be made on the basis of a candidate’s best three works. In this sense, the limit of 5 already exceeds the number of papers on which a candidate will be evaluated.

3.2.1 Action 4: arXiv Posting with Submission

Problems: The usage of LLMs by reviewers compromises the privacy of the submission. AI makes it easier for authors to submit many low-quality papers.

Action: On an experimental basis, require that submissions be published on the arXiv. The arXiv submission timestamp should be no later than the conference submission deadline. Authors should be required to affirm that the arXiv version is the same as their submitted version. There should be some mechanism for PC members and reviewers to access an anonymized version if they choose to do so, to help concerned reviewers protect against bias (and this method of access can even be the default).

Why: Having a publicly available version would allow reviewers to use LLMs without concerns about privacy. Requiring authors to publicly commit to their submission will encourage higher quality submissions.

How: Authors are required to submit their arXiv submission confirmation email with their submission form. There should be some process for authors to provide an arXiv link after conference submission if the link is not available before the deadline. Only the arXiv version that is submitted will be evaluated, not any future revisions. PC members may check the links for the papers that are assigned to them.

Who needs to implement this action: This idea will require broad buy-in on the part of conference steering committees. Aside from QIP, which has required arXiv posting for the past couple years, this would be a new custom for the TCS community.

Why (Experimental): Many people in our community have come to count on an additional period after submission to improve their results or presentation before committing to a public version. There is also some concern that LLMs make it easier for other researchers to behave unethically by scooping and improving results after they are posted, even while the original research group is continuing work (in such situations, improvements should be offered to the original authors, who should accept them with an added acknowledgement to the group that suggested the improvement). Nonetheless, we feel that it's better for authors to timestamp their results through the arXiv since there is a much greater risk of inadvertently leaking technical ideas through the use of LLMs in reviewing. Because of the concerns related to this idea, we suggest this item be implemented on an experimental basis.

3.2.2 Action 5: Video Talks with Submissions

Problem: AI-generated papers mean that people can submit papers that they do not understand and can also submit many papers with minimal human input.

Action: Conference submissions will require the submission of a talk video in addition to the paper.

How: The talk submission deadline should be 1–2 weeks after the paper submission deadline. The talk video length should be about the length of a conference talk, although we recommend giving significant leeway with regard to talk duration, as the primary motivation is to demonstrate human understanding. We recommend that the talk be required to be a board talk in which the author can be clearly seen and heard. Authors must certify that the talk is a camera-recorded talk of actual author(s) explaining their result, with no computer-generated content. The point of the talk is to provide evidence of human understanding, through a clear presentation of the merits of the paper, including a discussion of the paper's importance within the context of prior work and a description of the key technical and conceptual ideas used to achieve its main claims. Reviewers are not required to watch the talk, but are permitted to reject a paper with a poor talk *solely* due to the quality of the talk with no consideration of the technical merits of the paper.

Why: First, the talk provides some nontrivial evidence that the authors know something about the paper. Sections 4.1 and 5.2 recommend evaluating human understanding through live presentation in the hiring and PhD settings respectively, and so preparation of such talks for conference submissions helps prepare students for essential aspects of their research careers. In particular, it acts as a “proof of work” that the authors have invested some minimal amount of human effort into the paper, and this minimal amount of friction will provide a small barrier against flooding the submission system with papers generated via AI with near-zero human input. Second, we believe that conferences are a venue for human communication of ideas and learning, and as such should take into account presentation quality as a criterion for acceptance. We recommend that the talk is a board talk (or otherwise low-tech and handwritten) to emphasize the human component, and to prevent fully AI-generated talks from being submitted.

Experiment and Evaluation: The requirement of a talk video with paper submission is a large shift in norms for our community, and as such ought to be treated as an experiment. Conferences should require reviewers (when they submit reviews) to rate how much the talk video influenced the review.

Who needs to implement this action: Steering committees and PC chairs.

Concerns: Talk videos de-anonymize at least one author of the submission. The talk also introduces a greater potential for bias in the reviewing process based on characteristics of the speaker, as opposed to the quality of the technical results or exposition. However, we believe that having papers be accepted without the authors understanding the paper is far more destructive to our field.

3.3 Action 6: Unified Public Repository for Conference Rules and Norms

Problem: Rules and norms for publication venues may change rapidly in reaction to Artificial Intelligence, including in response to recommendations such as the present document. This leads to confusion for authors (“do we need to post our paper on arXiv first?”) as well as Steering Committees, PC chairs, PC members, and reviewers (“what’s SODA 2027 planning to do?”, “am I expected/encouraged/forbidden to use an LLM for writing my review?”)

Action: Establish a coordinated process whereby conferences establish standardized CfPs, along with an up-to-date public repository for publicizing this work. The relevant organizers and governance of each conference should work together to create and update a consistent set of norms and rules for conferences. To do this, we suggest establishing a public repository of clear formulations of policies, including relevant sections from current and past Calls for Papers. This public repository will be accompanied with a governance and coordination structure for curating this repository. We give specific policy recommendations in Section 3.1, along with experimental recommendations in Section 3.2. In their Calls for Papers, TCS conferences are encouraged to use the formulations from this repository verbatim.

Why: Reliable Source, Mandate for Curation. Coordinated rules and guidelines among thematically related venues with overlapping communities can reduce confusion, frustration, and workload for submission authors, steering committee members, and PC chairs and members.

Who needs to implement this action: The Steering Committees of rank A conferences and the current PC chairs of those conferences.

Where: A GitHub repository for curated and accessible hosting; coordination via GitHub Discussions (or Slack).

How: Contact relevant steering committees to appoint representatives (say, the SC chair). These representatives add and remove current PC chairs as terms rotate. Set up a GitHub organization (say, `simons-berkeley` or `tcs-cfp-policies`) to host the repository.

4 Hiring

4.1 Action 7: Best Works and Technical Talks in Interviews

Action: A comprehensive review of the hiring process is needed. We recommend the following immediate changes to committees for hiring junior faculty:

1. **Emphasize quality over quantity.** Evaluate candidates on no more than **three** works of their choosing, with this fact explicit in the instructions for candidates and their letter writers. The expectation is that letters of reference speak in detail on these works, and that job talks will be based on these works. This is standard practice in other fields like economics.
2. **Test technical depth.** Evaluate the depth of the candidates’ understanding of the aforementioned works more than was done normally until now. This evaluation could take the form of a “chalk talk,” probing questions during meetings, and/or other methods of the department’s choosing. Section 5.2 recommends related forms of AI-free evaluation during the PhD.

In addition, we recommend that hiring committees carefully review their process in order to prioritize enduring traits. Examples of such traits might include clarity of thought, the ability to pose novel questions, effective exposition, insightfulness, and creativity.

Why: The changing landscape of today has exacerbated the difficulties of the hiring process on two fronts:

1. it has become less clear which traits to look for, and
2. their identification based on accomplishments may become harder in the presence of AI.

This changing landscape is also obfuscating what early career researchers view as success and how their mentors should guide them. In particular, AI risks accelerating an “arms race” in the number of research publications.

How: We propose that theory faculty advocate to their department chairs and hiring committees and make other local changes depending on their circumstances. Obstacles may include overcoming local rules, coordination with the rest of the department, and accommodating the transition period associated with the change in the process. To help overcome these obstacles, the Simons Institute will set up a moderated Discord channel for faculty to share their experiences with this task.

We recommend that these changes be made starting with the 2026–2027 hiring season.

Who needs to implement this action: Theory faculty and hiring committees at universities.

4.2 Action 8: Additional Postdocs

Action: We recommend that additional postdoctoral positions be established for the TCS community, a temporary measure that will be phased out after a few years.

Why: With the rapid pace of current advances and results, it is reasonable to anticipate a much deeper level of human understanding across many subfields of TCS, which will lead to many new questions and research areas. Indeed, we believe that we are in the midst of a TCS renaissance, with the introduction of new techniques for well-established areas, as well as new problems arising in the context of AI. Both of

these developments, together with uncertainties around hiring arising from changing norms, not only pose transitional challenges for junior researchers, but provide generational opportunities to shape the future of theory. The creation of these postdoctoral positions will provide the time and space needed to pursue these opportunities.

How: The creation of postdoctoral positions will require the support of existing and new funding organizations. We imagine that this will require a coordinated effort between leadership at the various institutions and non-profits that have an interest in supporting this effort. The situation is complicated by the reduced federal funding for research in recent years—we thus suggest that universities and other institutes consider temporarily redirecting resources if necessary.

Similar measures have historically been taken in times of upheaval, such as the CI Fellows program during the Great Recession, and the Simons “bridge” positions during the COVID pandemic.

Who needs to implement this action: Various entities could support this action, including institutes such as the Institute for Advanced Study, the Simons Institute, IDEAL, CMSA, IFML, and IPAM; industry collaborations; non-profits such as the Mathematical AI Safety Institute (MAISI) and the Institute for Responsible Superintelligence (RESI); PIs and senior theorists; and departments at universities.

4.3 Action 9: Workshops for Identifying New Theory Directions

Action: In addition to the deeper human understanding AI will permit in well-established subfields of theoretical computer science, the study of AI itself poses many fascinating new domains for theoretical research. See, for example, the workshops during the 2024–2025 Special Year on Large Language Models and Transformers at the Simons Institute.

We recommend that the community provide catalytic support by providing a platform that permits theorists to develop new fields of research that arise in the context of AI. We propose two steps.

1. First, that a workshop (or several) be organized soon with the goal of identifying a list of modern topics related to AI. Examples include alignment, interpretability, harness design, and the existence and nature of “recursive self-improvement.”
2. Second, that a series of workshops be run on each of these topics to identify fundamental directions and establish foundations and agendas for future work.

Why: Theorists play a unique role in understanding the structure and limits of computation. Such questions, especially in a world with extremely powerful computational tools like LLMs, are becoming increasingly relevant. For example, theory already has a tradition of studying what can be achieved with the help of a powerful but untrusted oracle. Prioritizing human understanding will permit building and controlling better AI systems.

How: Towards laying the preliminary groundwork for future research, we expect that numerous entities, including institutes, departments, and conferences, are already interested in making such efforts. We propose that the Simons Institute provide an online venue for organizing and coordinating these efforts.

Who needs to implement this action: As outlined above, there is a role to play for various entities, including institutes such as the Institute for Advanced Study, the Simons Institute, IDEAL, IPAM, CMSA, and IFML; industry collaborations; non-profits such as the Mathematical AI Safety Institute (MAISI) and the Institute for Responsible Superintelligence (RESI); and departments at universities.

5 Graduate Education and Training

5.1 Action 10: “TCS in the Age of AI” Series at the Simons Institute

Action: A recurring and widely accessible “TCS in the Age of AI” series is held at the Simons Institute.

Why: The rapid development of AI tools is leading to new opportunities and significant uncertainty in the TCS community. This gathering would serve several purposes:

1. A large fraction of the TCS community is learning and trying to adapt to using the rapidly evolving AI tools in their research, teaching, and other aspects of their job. Our proposed program can be used to share useful practices and lessons to better prepare the members of our community as this technology develops.
2. There is already significant anxiety and stress among the PhD students and postdocs working in theory of computing. This includes natural concerns about what a future academic or industrial career in theoretical computer science will look like in the coming years. Students could use more information and perspective on some of these concerns.
3. Faculty are likewise facing new challenges in advising and encouraging their students in the age of AI: what skills should students develop, how should AI tools be incorporated in research training, how should graduate students be evaluated on their training progress and independence? This program can address such concerns and offer new practices on students’ training and evaluation.
4. The Math community is already organizing and managing AI and Math events. Within TCS, in comparison, there have been very few such events, including an event at STOC 2026 and one that is planned for the upcoming FOCS. We believe that such events need to be organized regularly because of rapid advances in AI tools and significant need from the community.

How: We propose regular online or hybrid gatherings on the interaction between AI and TCS. We propose the following possible components:

1. Talks by relevant community members on the changing education and research practices (e.g. sharing successful or unsuccessful stories on teaching or research, how to use AI for research).
2. Panels on future of TCS (e.g., research funding, new research topics) and career opportunities for TCS students.
3. Social gatherings in which graduate students, postdocs and faculty share their experiences around the use of LLMs for TCS, and help identify other topics related to AI and TCS that they would like to see addressed (e.g. work-life balance, learning with AI assistance).

The format and topics should stay flexible to be able to respond to the rapidly changing technology as well as new concerns that may arise as time passes.

We recommend that these talks and panels are run as a part of a regular monthly series; the frequency of the gathering may be adjusted based on outcome and the needs of the community.

Who needs to implement this action: We recommend that the hybrid talks and panels be primarily sponsored and organized at the Simons Institute. The institute can also consider in-person events in which graduate students participate and train on how to conduct successful research using AI tools.

5.2 Action 11: Opportunities for Training and Evaluation without AI Assistance

Action: To ensure development of necessary skills for a successful career in academia, we recommend that a PhD in TCS partially include *training* and *evaluation* components where the use of AI is forbidden. Specifically, we recommend explicitly “AI-free” intervals and tasks throughout the PhD program, especially in the first two years of the PhD. AI-free evaluation can occur at appropriate milestones throughout the program.

Why: We recognize that AI integration in education and research brings new opportunities and is inevitable. The purpose of our recommendation is *not* to discourage or push back on the use of AI in TCS education or research. In particular, we are *not* recommending, e.g., a full year of studying without the use of AI or that students conduct research entirely without AI. Our intent is to preserve opportunities for students to develop and demonstrate important research skills whose development may otherwise be hampered by the free availability of powerful AI tools.

1. We believe that even as AI excels at specific research tasks, other skills will remain essential for conducting high-quality research and deepening human understanding. Without deliberate training and corresponding changes to evaluation, however, graduate students may have fewer opportunities or incentives to develop these skills.⁶
 - (a) These skills could include, among others, mathematical problem solving, internalizing the intuition behind mathematical proofs, and the ability to communicate mathematical results and ideas clearly.
 - (b) As an example, it can be detrimental to a student’s development to repeatedly prompt AI on a research problem over and over and collect counterexamples without spending substantial time thinking about and understanding the underlying barrier that these examples reveal.⁷
 - (c) Similarly, we believe that thorough AI-free evaluation, such as interactive in-person presentations, can help students build and demonstrate the ability to understand, present, and defend their ideas when challenged; skills that remain crucial for a successful career as a scholar.
2. In principle, some of these skills could also be developed alongside appropriate and thoughtful use of AI. We also encourage advisors and graduate students themselves to explore other ways to ensure that they are gaining the necessary skills that over-AI-usage would hinder. Nevertheless, we believe that reserving some training and evaluation activities for work without AI assistance would be broadly helpful across various research settings, skill-development goals, and levels of AI proficiency.

How:

1. For *training*: We recommend that advisors or PhD Program Committees consider:
 - (a) Assigning small or incremental research projects (possibly with an already known solution or approach) or textbook exercises for students, as part of the PhD program, while disallowing the use of AI for brainstorming or developing solutions to these problems. The primary purpose of these exercises and problems would be skill-development rather than producing new research results.
 - (b) Organizing reading groups or courses where PhD students learn to solve problems, review literature and/or communicate and present their/others’ work without the help of AI.

⁶There is a growing body of empirical evidence that using AI can lead to a decline in learning [2, 4, 3].

⁷The study [2] showed that in a preregistered randomized trial, students given unrestricted access to an AI assistant during practice performed substantially better on practice problems, but performed *worse* than the unassisted control group on subsequent exams taken without AI aid.

2. For *evaluation*, we recommend that PhD Program Committees consider evaluating the progress of students using means that make the use of AI impossible: for example, in-person qualification exams (written or oral), and an in-person thesis defense. Section 4.1 recommends a comparable evaluation of technical depth in faculty hiring. The PhD Program committees should insist that the candidates rigorously understand and present their research contributions.

Who needs to implement this action: Individual advisors and PhD Program Committees.

5.3 Action 12: Coverage of Reasonable LLM Expenses from Research Grants

Action: We recommend that relevant professional organizations investigate the barriers researchers have faced in obtaining access to the state-of-the-art AI tools. If deemed necessary, they may issue a public statement recognizing access to LLMs as a valuable research resource, and recommend to relevant funding agencies and entities such as universities or school systems that their policies explicitly allow reasonable LLM-use expenditures.⁸

Why: Some faculty at U.S. institutions encountered barriers to paying for LLM access using federal research grants. At the same time, LLMs are playing an increasingly important role in theoretical computer science research, including in the work of postdocs and graduate students. A clear statement from a professional scientific organization of stature could provide valuable support for researchers facing such barriers to advocate for updated policies at their institutions. Background research is first needed to determine the exact nature and source of the barriers present at different institutions.

How: We recommend conducting a targeted research survey during the “TCS in the Age of AI” program at the Simons Institute to better understand the barriers that researchers face and the root cause of these barriers; whether they arise from federal or state funding rules or university-level policies.

Following the review, the Simons Institute may connect with SIGACT, ACM, or a broader scientific organization to request issuance of a statement recognizing that access to state-of-the-art LLMs is a legitimate and important research expense. In addition, the organization may communicate with relevant funding agencies and recommend they clarify and (when necessary) update their policies to permit reasonable expenditures for LLM access for research. The statement and recommendations should make it clear that the support for AI tools is not intended to displace support for students, postdocs, and researchers. For example, guidance may recommend that expenditures on LLM access from a grant remain a modest fraction of a research budget.

The statement may also aim to identify other means to enable LLM access to the broader research community.

Who needs to implement this action: SIGACT, ACM, and other governing bodies; Simons Institute.

6 Conclusion

We hope that our recommendations can provide a common denominator for coordinated change. But by itself, a document with a list of recommended actions is worth very little. So what now? The power comes from what happens next, and that will involve individuals, departments, and institutions considering these recommendations, modifying them as they see fit, and tackling the challenge of implementation.

We stress that this is a moment when *every person in our community has a role to play*. To enact versions of each of the 12 actions across the community will undoubtedly take real efforts, both big and small, from many people. We tried our best to articulate the path forward on implementation and did so with the firm

⁸We also acknowledge the lack of LLM access for the TCS community outside of the U.S. as a broader issue that is beyond the scope of the workshop.

belief in the capabilities of our community to collectively solve the challenges that arise. That belief extends to you — a fellow community member, with likely the same concerns and questions that led us to initiate this work.

In a time of such uncertainty, we see the call to implement these actions not as yet another chore, but as an opportunity to rally towards a future where we can continue to explore and share beautiful and interesting ideas with each other.

A Examples of AI Methodology Sections

Consider the following fictitious, uninformative methodology statement:

“We used AI in developing and verifying our proofs, code, and write-up. We take full responsibility for the content of our paper.”

This statement lists broad activities but leaves unclear what the authors supplied, what AI contributed, and what the checking involved. The scientific narrative behind the paper remains opaque, and the claim of responsibility offers little assurance.

The most informative methodology sections we have seen explain **how AI shaped the research**, rather than merely naming a model or asserting that its output was checked. They make human and AI contributions explicit by describing:

- Which AI models were used, what authors had already established, and what information (e.g., research notes on specific topics, pseudocode) was supplied to the models.
- Which key ideas came from AI models and how the authors built upon them.
- What errors or limitations the authors found in AI outputs and how they addressed them.

The following five arXiv papers provide concrete examples of such informative accounts.

Semi-Streaming Matching in a Single Pass II: Greedy is Optimal

Sepehr Assadi, Max Jiang, and Mars Xiang

[arXiv:2607.14656v2](#) · July 16, 2026 · Primary cs.DS · PDF p. 18

Clarifying the roles of humans and AI in the research process.

“Specifically, the authors provided an optimal solution to a *relaxation* of blueprints ... and a combination of Claude Fable and GPT 5.6 Sol gave the idea for using random walks to take this solution from the relaxations to blueprints. These random walks formed the crux of our blueprints.”

The authors specify the inputs provided to AI models and clearly state the nature of the contributions: the authors guided the research and wrote the exposition, while AI supplied a key technical idea. They also describe which AI models were helpful and which were not during their research process.

An iterative rounding 2-approximation for Feedback Vertex Set via AI-assisted proof of an extreme point property

Karthekeyan Chandrasekaran, Chandra Chekuri, and Shubhang Kulkarni

[arXiv:2609.04414v1](#) · September 3, 2026 · Primary cs.DS · PDF pp. 5–6

Tracing how human and AI contributions built on one another.

“We used AI tools Claude (Opus 4.7) and Gemini (3.1 Pro) (via their desktop versions) that were fed our research notes consisting of our almost-laminar basis proof and other observations. Over multiple rounds of interaction, the tools helped unravel the proof. Two key ideas that the tools proposed were the following ...”

The authors specify the input provided to the AI models, namely their research notes on specific topics, identify two key ideas the models proposed, and explain how they incorporated and extended them. This preserves the scientific narrative by connecting the human groundwork, the AI-proposed insights, and their subsequent adaptation by the authors.

The Query Complexity of Local Search in Rounds on General Graphs

Simina Brânzei, Ioannis Panageas, and Dimitris Paparas

[arXiv:2601.13266v2](#) · January 19, 2026 · Cross-listed from cs.CC · PDF pp. 22–23

Describing errors in AI outputs and how authors resolved them.

“... the authors identified a circular argument in the model’s initial proof for the bijection between histories and candidate sets ... and provided the model with a hint for resolving it.”

By explicitly describing the issue in the AI-generated proof and how they helped resolve it, the authors make clear what their checking involved. They also report that the model produced a parallel steepest-descent algorithm ruling out a lower bound they had expected to hold, illustrating its role as a source of counterexamples as well as a proof assistant.

Three-edge-coloring apex cubic graphs

Yuta Inoue, Ken-ichi Kawarabayashi, Rintaro Matsuo, Atsuyuki Miyashita, Bojan Mohar, and Tomohiro Sonobe

[arXiv:2608.22870v1](#) · August 24, 2026 · Cross-listed from math.CO · PDF pp. 1, 3, 41

Specifying inputs to AI and withheld information to support a reproducibility claim.

“... independent implementations were reconstructed by ChatGPT Pro and Claude Code, directly from the pseudocode ... and the public input data, without access to the authors’ source code. Both independently reproduced the computational results required in this paper”

The authors clearly specify that AI models received the pseudocode and public inputs, but not the original source code. This helps readers interpret the reconstruction as a reproducibility test of the method described in their pseudocode.

k -Coloring is Faster than Computing the Chromatic Number

Or Zamir

[arXiv:2607.25973v1](#) · July 28, 2026 · Primary cs.DS · PDF p. 40, “AI Storytime”

A caveat about inputs to AI: model context may extend beyond the authors’ intended inputs.

The author notes that ChatGPT-5.6-Sol (memory enabled) used constructions from earlier conversations without clear attribution:

“I even ‘caught’ it using constructions I had described in earlier chats without attribution: a reference to the previous chat appeared in the animated ‘thinking’ glimpse but not in the final answer. Thus, my account of the prompts I supplied might not reveal the full context I inadvertently provided the model with.”

The author also reports the model reproducing an existing proof framework without attribution, which a less experienced reader could mistake for a novel contribution. This highlights current limitations of AI models in attribution and how detailed AI Methodology sections can help readers retrospectively connect AI-generated proofs to prior work.

B Workshop Format Details

Schedule. Each day began with a meeting of all participants and included an outline and description of the goals and processes for the day. The structure consisted of a basic repeated unit of separate small group work in set teams of 8 (see Formation of Small Groups) followed by a sharing of the work to the larger group and getting facilitated feedback that was then incorporated into the small group work. The schedule was as follows:

Day 1 Morning

- Small Group: Identify a small number of problems that the small group subject faced.
- Large Group: Small groups share out and get feedback

Day 1 Afternoon

- Small Group: Incorporate large group feedback.
- Small Group: Identify a small number of recommended actions to address the identified problems
- Large Group: Small groups share out and get feedback

Day 2 Morning

- Small Group: Incorporate large group feedback; work out details of the recommended actions.

Day 2 Afternoon

- Large Group: Small Groups share out followed by consensus measurement and feedback.
- Small Group: Incorporate large group feedback and writing.

Day 1: Small group structure. The morning and afternoon sessions were structured. They began with 3–5 minutes of silent brainstorming where group members would write one idea related to the session’s prompt on a sticky note. Then each person in turn would explain their ideas and the representative sticky notes were put on a board. In preparation for the workshop, the organizers had mined the survey for content that could contribute to this brainstorming phase and included relevant ideas during this process. Once all the ideas for everyone were incorporated the small group collectively gathered the sticky notes into themes and consolidated overlapping ideas.

After discussion there was a distilling process to identify the ideas that had consensus. Some groups found this naturally came out of the discussions, other groups used an additional tool called *dot voting* to narrow things down. Dot voting consisted of each participant apportioning five stickers among the ideas they supported. Each group assigned one or more people to share out their findings to the larger group along with someone to carefully note feedback to be incorporated later. The incorporation of feedback followed an open discussion format within the small groups. It should be noted that this feedback often led to important changes.

Day 2: Small group structure. The morning session on day two was focused on taking the ideas from day 1, incorporating the feedback and building out the details. In many groups it featured smaller groups working on different actions and then sharing out within their group of 8. The afternoon session featured incorporating the feedback from the big share out session and writing. No group produced a finished draft by the end of the workshop; separate Discord channels were set up to facilitate completing the writing up of the actions.

Day 2: Large group feedback and consensus measurement. Between the morning and afternoon small group sessions was a facilitated sharing out session. We used a group rating tool that worked as follows. Participants are asked to rate their agreement (from zero to five) in response to a statement. 5 meant full support, 4 meant mostly support, 3 meant neutral, 2 meant a mild objection, 1 meant a strong objection, and 0 meant that if you had veto power you would use it. Loosely speaking, the process was that groups would share out individual actions, clarifying questions would be answered, and then the facilitator called for a rating from everyone. Those with 0's, 1's, and 2's were asked for feedback and the small group was entrusted to take this feedback in their subsequent work. Had any action received a tepid response, it would not have been included in the final list, but that did not happen in this case.

Possible Improvements. Participants generally felt that the format worked well. We received the feedback that a three-day format would have made the writing easier. The graduate students and postdocs in the group made many of the most interesting suggestions, and since they are also the most impacted, we plan to invite a larger fraction of young researchers to future such events, closer to 1/2 rather than 1/4. Of course, we hope that with a longer lead time we can have more participation in the future from researchers in other countries.

References

- [1] L. Allen, A. Brand, J. Scott, M. Altman, and M. Hlava. Publishing: Credit where credit is due. *Nature*, 508(7496):312–313, 2014.
- [2] H. Bastani, O. Bastani, A. Sungu, H. Ge, Ö. Kabakcı, and R. Mariman. Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26):e2422633122, 2025.
- [3] A. Darvishi, H. Khosravi, S. Sadiq, D. Gašević, and G. Siemens. Impact of AI assistance on student agency. *Computers & Education*, 210:104967, 2024.
- [4] Y. Fan, L. Tang, H. Le, K. Shen, S. Tan, Y. Zhao, Y. Shen, X. Li, and D. Gašević. Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, 56(2):489–530, 2025.
- [5] T. Feng, T. Trinh, G. Bingham, J. Kang, S. Zhang, S.-h. Kim, K. Barreto, C. Schildkraut, J. Jung, J. Seo, C. Pagano, Y. Chervonyi, D. Hwang, K. Hou, S. Gukov, C.-C. Tsai, H. Choi, Y. Jin, W.-Y. Li, H.-A. Wu, R.-A. Shiu, Y.-S. Shih, Q. V. Le, and T. Luong. Semi-autonomous mathematics discovery with Gemini: A case study on the Erdős problems. [arXiv:2601.22401](https://arxiv.org/abs/2601.22401), 2026.