

STOC 2027: Explaining the New Policies

STOC 2027 is implementing several policy changes in response to the rise of gen-AI. The goal of this document is to explain the reasoning behind them. These changes followed intense debates within the STOC PC, which consists of 99 members of the TCS community, and are approved by SIGACT EC which is the steering committee of STOC.

What is the problem we are trying to solve?

TCS conferences serve a dual role:

1. Select top papers
2. Provide a meeting place for the community to learn about new developments

AI tools have made it significantly easier to generate papers that make progress on or solve important open problems without the need for the authors to understand the problem, the area, or the motivations for studying the problem in the first place. This means that such authors would most likely not be able to explain their result well or answer questions about why they study the problem or what the natural next directions are.

There are three concrete problems we want to tackle:

1. **Volume** - if STOC gets flooded with submissions, there is no way we can have meaningful PC discussions about the papers. To give some data points:
 - STOC'25 and STOC'26 had ~800 submissions
 - SODA'27 got 50% more submissions than SODA'26
 - ITCS'27 got 66% more submissions than ITCS'26
 - ICLR'27 has 60k submissions, compared to 20k for ICLR'26 and 12k for ICLR'25.
 - It is reasonable to expect STOC'27 to receive 1,200-1,600 submissions, or perhaps many more if we do not adopt policies to curtail this growth
2. **Quality** - as it becomes easier to generate papers by querying AI models, we expect to get many more mediocre papers. Typically, the TCS community self-moderates STOC submissions, as STOC has a high bar. But if we get submissions from authors who are unfamiliar with the standards of the TCS community, then we should not expect the same self-moderation.
3. **Presentations** - as we already discussed, one of the main goals of STOC is to be a useful meeting place for TCS researchers. It does not help to have talks where the authors cannot answer basic questions about their paper. We are concerned that this might happen when authors do not understand the papers they generate with AI models.

How do we address these challenges?

These are all new problems for our community, and so the solutions below should be thought of as experiments. We will test which policies work well and which ones don't, and use what we learn to improve future iterations of STOC (and other TCS conferences). The policies below were adopted by a vote of the PC after lengthy discussions in which we considered many other policies. We chose them because we think they are both useful in solving the problems described above and practical to administer at the scale of STOC.

Policy changes:

1. **Limit the number of submissions per author.** We decided on a limit of 5. This is meant to help control submission volume so that the PC can have meaningful discussions about the papers.
2. **Authors are not anonymous.** The main reasons are: (1) authors should be willing to stand behind their work; (2) it makes it much easier for the PC to find reviewers while avoiding conflicts of interest; (3) it meshes well with our other policies below, which would be hard to implement with anonymous authors.
3. **Submissions are public.** Authors will submit their papers through HotCRP in the usual way. The organizers will post all submitted papers publicly, and authors must explicitly consent to this when submitting. The main reasons are: (1) authors should submit complete papers rather than unfinished drafts rushed to meet the deadline; (2) making submissions public allows the PC and reviewers to use LLMs to help with review and evaluation without violating the confidentiality of submissions.
4. **Video requirement.** As a major goal of STOC is to have authors present results that they understand, we require each submission to include a video, presented by one or more of its authors, explaining the paper and its contributions. The main goal is to make sure that the authors understand the result and can present it properly. Note that the main artifact being evaluated is still the paper, not the video. The video does not need to be polished and will not be made publicly available. PC members and reviewers may use the video to better understand the result, particularly in cases where there is suspicion that the paper is AI-generated and the authors do not understand it well enough to present it properly at the conference.
5. **Everyone can use LLMs.** This applies to authors, PC members and reviewers. LLMs are powerful tools and, when used responsibly, can help authors with research and exposition. They can also help PC members and reviewers validate correctness, check previous work, and understand technical aspects of the work. We expect PC members and reviewers to evaluate the work themselves and use LLMs to support, not replace, their judgment. We expect everyone to fully disclose how they used AI tools in their work. We suggest adding an "AI methodology" paragraph to all papers, even if it says that no AI was used. Reviews will include a section describing how the reviewers used AI.