Episode Book

EPISODE #6466

Conferences Rule Machine Learning, Not Journals

PUBLISHED	5 November 2026 at 20:02 UTC
DURATION	27:32
LISTEN	myweirdprompts.com/episode/machine-learning-peer-review-conferences-87441/
HOSTS	Corn, Herman, Hilbert
SCRIPT MODEL	z-ai/glm-5.3
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.26
SOURCES	25 cited by the research (listed at the end)
BOOK GENERATED	6 November 2026 at 11:12 UTC

SYNOPSIS

In machine learning, flagship conferences serve as archival records rather than stepping stones to academic journals. From the massive scale of NeurIPS, ICML, and ICLR to specialized venues like CVPR and ACL, acceptance rates have converged around twenty to twenty-seven percent. This episode explores how open-access journals like JMLR and JAIR fit into the ecosystem, the astronomical fees of legacy-brand publishers, and why arXiv recently began demanding peer-reviewed acceptance stamps on survey papers.

TRANSCRIPT



Corn

Twenty-one thousand five hundred and seventy-five. That's how many papers landed at one machine learning conference last year, and they found almost twenty-two thousand reviewers to read them. Not a typo, and not a typo the second time either.



Herman

And that's just one conference, Corn. One of dozens.



Corn

Daniel's latest is about exactly that. We've done the arXiv firehose to death — he says so himself. What he wants now is the tier above it, and he's careful to say he doesn't mean snobbish. He means: where do the serious machine learning researchers actually present, what are the peer-reviewed journals, and how does the whole thing hang together with the preprint server underneath it. His phrase is that it's a broad ask for one episode, so we'll pick a few conferences and a few journals and get to know them properly rather than surveying everything.



Herman

Which is the right instinct, because if you try to list them all you get a phone book.



Corn

And here's the thing you need before any of the names make sense. In this field, the conference comes first. Not the journal.



Herman

That's the part that breaks people from medicine, physics, biology. The Computing Research Association put it in their career guide almost word for word: newcomers from other disciplines are often surprised that in most of computer science, the flagship venue is a peer-reviewed conference, and its published proceedings are the archival record, not a stepping stone to a later journal version. The conference paper is the finished thing.



Corn

You spent decades in medicine. When you published, the journal was the prize and the conference was the trailer for the journal.



Herman

You stood up at a conference with your slides, you got the feedback, and then the real work of getting it into a journal started — months of revision, another round of reviewers, proofs. In machine learning, that whole structure is inverted. The conference is the prize. NeurIPS, ICML, ICLR — those are where the field's best work appears first, and the proceedings are what people cite.



Corn

Stanford tells its undergrads the same thing, apparently. Computer science is a conference-oriented field.



Herman

So the answer to Daniel's "second tier" question has a twist in it. The tier above arXiv isn't mostly journals at all. It's the big three general conferences, then the domain conferences, and the journals are a smaller, more specialized part of the picture — still respected, sometimes the right home, but not the center of gravity.



Corn

Then let's do the big three. NeurIPS first, since we've already got its number on the table.



Herman

Twenty-one thousand five hundred seventy-five valid submissions at NeurIPS 2025, five thousand two hundred ninety accepted. Call it twenty-four and a half percent. To run that they needed twenty-one thousand nine hundred reviewers, nearly two thousand area chairs, and a couple hundred senior area chairs above them. It's a small city that exists for one week and then dissolves.



Corn

Twenty-two thousand reviewers for twenty-one thousand papers. The math is one reviewer per paper, and each paper gets three reviews.



Herman

So most reviewers are carrying three, four, five papers. Keep that number in your pocket — it comes back later, and it's the reason this whole system is creaking. But yes: NeurIPS tripled in four years. Ten and a half thousand submissions in 2022, twenty-one and a half thousand in 2025, and this year's edition is reported north of thirty thousand.



Corn

Thirty thousand. At that point "conference" is a misnomer. It's a mid-sized industry with an annual meeting attached.



Herman

Then ICML, the International Conference on Machine Learning. Twelve thousand one hundred seven submissions in 2025, three thousand two hundred sixty accepted, about twenty-seven percent. Same family, slightly smaller, held in Vancouver that year. And ICLR, the International Conference on Learning Representations — that one's the youngest and the fastest-growing. Nineteen and a half thousand submissions for the 2026 edition, five thousand three hundred fifty-five accepted, twenty-seven point four percent, which the organizers themselves called the largest and most competitive edition in its history. Two years earlier it was seven thousand submissions.



Corn

So ICLR nearly tripled in two years and its acceptance rate barely moved. That means the bar didn't rise. The queue did.



Herman

That's a sharp read, and it's right. The acceptance rate is a ratio, and when both sides of it explode together, the ratio tells you almost nothing about quality. It just tells you the field got huge.



Corn

Which is why I'd argue acceptance rate is the most-quoted and least-informative statistic in this entire landscape.



Herman

Look at the spread. NeurIPS twenty-four and a half, ICML twenty-seven, ICLR twenty-seven and a half, CVPR twenty-two, AAAI twenty-three, ACL twenty. They all live in the same band between roughly twenty and twenty-seven percent. If you line up the flagship venues and ask which is most selective, the honest answer is: by that measure, none of them. They've converged.



Corn

It's a club where every member claims the door is equally heavy.



Herman

What actually differs is scale — NeurIPS at twenty-one thousand submissions against IJCAI at around five thousand — and what differs inside. That's the part outsiders miss. A NeurIPS acceptance isn't one thing. Roughly four and a half thousand of those accepted papers are posters. A few hundred are spotlights. And the orals — the ones they put on stage — are under half a percent of everything submitted.



Corn

Half a percent. The conference accepts a quarter of everything and parades half a percent.



Herman

CVPR does the same thing with its own numbers. Thirteen thousand submissions, about twenty-two percent accepted, and ninety-six papers — ninety-six — selected for oral presentation. The poster hall is where the conference actually lives; the orals are the showcase window.



Corn

And the domain conferences. Say a word about those before we do journals, because Daniel asked to get to know a few, and NeurIPS-ICML-ICLR is really only the machine learning slice.



Herman

Right — the general AI story continues with AAAI and IJCAI, the two older, broader artificial intelligence conferences. AAAI took nearly thirteen thousand submissions in 2025 and accepted twenty-three percent. IJCAI is smaller and tighter, around nineteen percent. Then you drop into subfields. CVPR is computer vision — it's enormous, nearly ten thousand people show up in person. ICCV is its sibling, held on the odd years. For language, the ACL family: the Association for Computational Linguistics is a proper professional society, going back to 1962, and ACL and EMNLP are its two big annual conferences, both sitting right around twenty to twenty-two percent for the main track.



Corn

One thing worth flagging for Daniel, since he does ASR work — NLP has a second landing zone called Findings. Papers that don't make the main conference but the reviewers thought were worth existing.



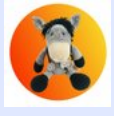
Herman

Yes, that's an important structural detail. ACL 2025 accepted about seventeen hundred papers to the main conference and another fourteen hundred to Findings, which is a separate, slightly less prestigious section of the same proceedings. EMNLP does the same split. It's a soft rejection — and it's also a pressure valve, because it lets the venue say no to the main track without saying the work is worthless.



Corn

A tier inside the tier. Like being told you didn't make the team but you can stand near the bench.



Herman

Corn, that's basically the Findings logo.



Corn

Now the journals. Because they do exist, and Daniel asked.



Herman

The one I'd put first is the Journal of Machine Learning Research, JMLR. Founded in 2000, fully open access, no author fees, and — this is the tell — its stated commitment is to rigorous yet rapid reviewing. In a field where the conference is the archive, a journal survives by being better at review, not by being more prestigious. Its impact factor is around seven, which sounds modest next to the giant we'll get to in a minute, but nobody in machine learning thinks JMLR is a seven.



Corn

Which is the impact-factor paradox in this field, isn't it. The most competitive venues are conferences, and conference papers don't feed journal citation metrics at all. So the journals' numbers are systematically understated relative to how the field actually ranks things.



Herman

Impact factor is a weak signal in ML for that structural reason — the league table it measures isn't the league the field plays in. TPAMI is the other heavyweight — the IEEE Transactions on Pattern Analysis and Machine Intelligence, running since 1979, impact factor up around eighteen, and the tradeoff is a review cycle of seven to nine months. Vision people who want the archival version of a conference paper send it there. Then JAIR, the Journal of Artificial Intelligence Research — since 1993, one of the first journals that existed properly on the web, completely open access, no fees, review in about three months. It's the scrappy one. It's fast because it has to be.



Corn

Because a journal that takes nine months is competing with a conference that decides in three.



Herman

And then there's the one at the other end of the money spectrum. Nature Machine Intelligence. Impact factor reported anywhere from about eighteen to nearly thirty depending on which aggregator you ask — I can't pin it down, the sources disagree — and



Herman (continued)

an article processing charge of twelve thousand six hundred ninety dollars.



Corn

Twelve thousand seven hundred dollars. To publish. Per paper. JMLR charges nothing and JAIR charges nothing, and the one with the nature-brand letterhead wants the price of a used car. For a PDF.



Herman

To be fair to them, you're buying the letterhead. That's the product. Whether it's worth it is a question every lab answers differently.



Corn

Now the relationship with arXiv, because that's the piece Daniel is really poking at. The firehose underneath the tier.



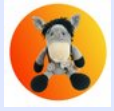
Herman

Preprinting is standard practice. Nearly everything that ends up at these conferences goes on arXiv first — sometimes a year before the conference decides, because the submission cycles are that slow relative to how fast the field moves. Most venues permit it, with rules about anonymity and not submitting the same work to two places at once. ACM, for instance, doesn't normally allow simultaneous submission. So the picture is: arXiv is where the work appears, the conference is where it gets judged.



Corn

Or it was, until last October. Something changed there.



Herman

Something changed. As of the thirty-first of October last year, arXiv's computer science category requires that review articles — surveys, position papers — be accepted at a journal or conference with documented peer review before they can be posted at all. And



Herman (continued)

arXiv said the quiet part out loud: they were flooded, they don't have the resources to do quality control in-house for content types they don't officially cover, and reasonable and trusted outside refereed venues already exist, so they'll rely on those. They said they were receiving hundreds of review articles every month and the majority were little more than annotated bibliographies.



Corn

They blamed it on us, more or less. Language models making papers fast and easy to write.



Herman

They did say that, yes. And there's a second line in that policy worth pulling out — they explicitly said workshop review generally does not meet the same standard of rigor as traditional peer review. Which is arXiv, the most permissive institution in scientific publishing, drawing a line between kinds of peer review.



Corn

So the layering flipped. It used to be: arXiv takes everything, the conferences sort it. Now arXiv is saying, for some content, show me the conference's stamp first. The firehose has started checking papers at the door.



Herman

The venue that defined the no-gatekeeper model now needs gatekeepers. That's the most important structural fact in this whole story, and it sets up the hard question, which is whether the gatekeepers are holding.



Corn

Are they?



Herman

Depends what you mean by holding. Let me give you the numbers and you can judge. Nature reported last year that twenty-one percent of the manuscript reviews at one AI conference were found to be entirely AI-generated. Not AI-assisted. Generated.



Corn

A fifth of the reviewing at a conference, done by the same class of model the conference exists to evaluate.



Herman

Then a Cornell group, in January, did the first large careful study — a hundred twenty-five thousand paper-review pairs across ICLR, NeurIPS and ICML. At NeurIPS 2025 they classified sixteen point nine percent of reviews as LLM-aided, against about one point seven percent of the papers. So roughly ten times as many reviews as papers were substantially machine-touched. Ten to one, Corn.



Corn

The reviewers are using the models ten times more than the authors. The referees are the ones outsourcing.



Herman

And the effect on the judgments is what you'd expect. The LLM-assisted reviews came back more lenient toward weaker papers, and fully generated reviews showed what the study calls severe rating compression — everything clusters toward the middle, the strong paper gets a three, the weak paper gets a three. Peer review stops discriminating, which is the entire job of peer review.



Corn

So what do you do about it? You can't have twenty-two thousand humans read twenty-one thousand papers. The arithmetic was already brutal before the models.



Herman

Corn, before I answer that, I want to know what you'd actually do if you were a program chair. You've heard the constraint — you can't just hire more reviewers, the reviewer pool is the author pool.



Corn

You do what the real ones are already doing, I think. You stop pretending everyone reads everything and you police the edges hard instead. You make reviewing a condition of publishing — that part NeurIPS already runs, reviewers who don't deliver can have their own papers pulled — and you desk-reject, publicly, the cases that are beyond doubt. You accept that the middle of the distribution is mush and you defend the ends.



Herman

That's strikingly close to what's happening. NeurIPS 2025 ran what they called Responsible Reviewing — program chairs could withhold reviews from authors who hadn't met their own reviewing obligations, and papers co-authored by grossly negligent reviewers could be desk-rejected outright. Eleven confirmed cases. Eleven papers pulled not because of what was in them, but because of who wrote them.



Corn

Collective punishment as an incentive scheme.



Herman

A very old one, yes. CVPR did the same dance — desk-rejected nineteen papers by reviewers they called highly irresponsible, some of whom had submitted LLM-generated reviews of other people's work. And ICLR 2026's retrospective shows the scale of the churn: out of nineteen and a half thousand valid submissions, seven hundred seventy-nine desk rejections and five thousand forty-two withdrawals. A quarter of everything submitted never reached a decision on the merits.



Corn

Five thousand withdrawals. People yanking their own papers mid-review. Some of that is the standard poaching between deadlines — you withdraw from ICLR because NeurIPS came through. But some of it has to be authors reading their reviews, seeing the compression, and concluding the dice are loaded either way.



Herman

And here's where I think the ranking question Daniel didn't quite ask gets its real answer. People want to know which venues are respected, and there are formal systems — the China Computer Federation publishes an A, B, C catalogue, seven A-class AI conferences on it; the Australians run one with an A-star tier on top. The problem is they disagree with each other. IJCAI is CCF A on the federation's own site and CCF B on at least one ranking site that tracks the same catalogue. There's no single authoritative ranking, and the discrepancies aren't typos — they're different snapshots of a list that different institutions weight differently.



Corn

So the leagues are administrative artifacts. A committee's opinion with a filing number.



Herman

The CCF catalogue's seventh edition runs six hundred eighty-one entries. It's enormous, it's useful for hiring paperwork in China, and it is not physics.



Corn

Then what's the real sorting mechanism? If acceptance rates cluster at twenty-five percent everywhere and the rank lists contradict each other, what is a skeptical outsider like Daniel actually supposed to check?



Herman

Three things, I'd say. Scale — because twenty-one thousand submissions with a functioning process is a different object from five thousand with the same rate. Internal stratification — whether the venue has a real oral tier that means something, that half-a-percent cut, versus one where everything accepted is the same. And now, increasingly, whether it polices its own reviews. NeurIPS naming eleven negligent co-authors, CVPR pulling nineteen papers — that's the new prestige signal. Not the selectivity, which everyone has. The willingness to burn some of your own submissions to keep the referee pool honest.



Corn

The venue equivalent of a doctor who'll tell a colleague they're wrong.



Herman

I'd take it, honestly. Twenty-two thousand reviewers, none of whom you chose, none of whom are paid, and a nontrivial fraction of whom may not have read the paper. Against that, holding a twenty-five percent rate at all is a minor miracle.



Corn

Or a machine no one is steering.



Herman

Both, possibly. You'd know about machines no one steers.



Corn

I do my best work that way. Is that the whole picture, then? Because there's a piece missing for me — the rebuttal. The authors get to answer the reviews before the decision, yes?



Herman

Yes, most of these venues run a rebuttal phase — authors see the reviews, get a window, usually about a week, to respond, and the discussion between reviewers and area chairs happens partly in public, on the OpenReview platform at ICLR especially. The whole exchange is visible. Which is how we know so much about what the reviews look like — a lot of it is just sitting there readable.



Corn

Wait — the reviews are public? So anyone can read twenty thousand strangers' anonymous opinions about papers?



Herman

At ICLR, largely yes. Signed or anonymized depending on the venue, but posted with the papers.



Corn

So a fifth of them being machine-written was discoverable by anyone with a free afternoon. The scandal was sitting in the open with a URL.



Herman

That's how the studies got done. And there was even a security incident at ICLR 2026 — an API problem reset review scores to their pre-rebuttal state, and they had to work through what that meant for decisions already influenced by the later scores. The machinery is visible enough to break in public.



Corn

Which brings me back to the arXiv flip. If arXiv is now outsourcing its quality control to these conferences, and the conferences are watching a fifth of their reviews come off a model — what exactly is arXiv buying with that policy?



Herman

They're buying the stamp, not the guarantee. The conference stamp says: three or four named humans agreed to review this, an area chair supervised, there was a rebuttal. It doesn't certify the reviews were good. But it does certify there's a process, a paper trail, and someone accountable if it fails. For a minimal-gatekeeper institution, that's a lot. It's the difference between content moderation and content judgment — arXiv is hiring out judgment and keeping moderation.



Herman

And for what it's worth, the venues are starting to push on the other side of the pipeline too. NeurIPS told authors for this year's position paper track that papers must be substantially written by human authors, AI for copy-editing only, and violations get a standard desk rejection that isn't subject to appeal. No appeal, Corn. That's how tired they are.



Corn

I have a question about the reviews themselves. You said rating compression. What did the fully generated reviews actually look like?



Herman

Structured, grammatical, generic. The studies detect them partly through lexical fingerprints — the models have favorite words and turns of phrase that show up at abnormal rates, the same polite hedging constructions, paragraph after paragraph. And the rating compression means the distinguishing content — the specific criticism — thins out. Everything sounds like a thoughtful review and almost nothing behaves like one.



Corn

The well-formatted nothing. Reviewers producing that would have failed kindergarten show-and-tell.



Herman

There was even a GPTZero pass over NeurIPS submissions that turned up hallucinated citations in papers — references to works that don't exist. Both sides of the desk are contaminated, is the honest summary.



Hilbert

The handwriting on the reviews was fine. That was the give-away, in the end.



Herman

Hilbert, we were just talking about the Cornell numbers —



Hilbert

I heard. Your one-point-seven percent of papers and your sixteen-point nine of reviews. I'm telling you it was going on before there was a model to blame. Late eighties, I had a paper in at an instrumentation conference — signal conditioning for load cells, nothing clever. Three reviews came back. Two of them had the identical opening sentence, word



Hilbert (continued)

for word, and the third said, in its entirety, "This is fine." Not fine how. Not fine apart from the gain drift. Fine.



Corn

A two-word review.



Hilbert

I photocopied all three, mailed them to the program chair with a note asking whether the reviewer had been a large language model or just a large idiot. This was some years before that was a natural thing to ask a man, but I've never let timing govern my correspondence. A colleague of ours, Keith, could produce a review in forty minutes flat. I timed him over a season, average forty-three. He was proud of it. That was the part I couldn't stomach.



Herman

What did the chair say? Did he write back?



Hilbert

He wrote back six weeks later to say the review process was "robust as constituted." I've kept the letter. The reviews are in a folder marked Exhibit A, which I've been carrying to some purpose for a while now. The proceedings volume's at my mother's, under the stepladder — the thing's four inches thick, no shelf in the flat takes it, and I wasn't going to pay for storage on a book I'm in.



Corn

You kept the folder for forty years waiting for someone to ask.



Hilbert

You asked about review quality for twenty minutes. I didn't interrupt.



Herman

The identical opening sentence in two reviews — that's the fingerprint they're detecting now, with the lexical markers. You found it by eye.



Hilbert

Anybody would. The point you two keep circling — that it got worse when the models arrived — I'd say it got cheaper. The forty-minute Keith became a four-minute Keith, and they gave the four-minute one to everybody. The conferences were already too big for anyone to read anything; the machine just made the not-reading affordable. The paper got in, mind. Fine was a pass.



Corn

Forty-three minutes. He knew his average.



Hilbert

To the minute. Gave it as a defence.



Corn

There's the cutting-room-floor item, and it's a good one. The CRA career guide — the same document with the conference-first line — also notes that the peer-reviewed publication is "the coin of the realm" in computing research. It certifies new knowledge, measures careers, and sits behind nearly every funding, hiring, and promotion decision. So this entire structure we've been describing, the tripling conferences and the swamped reviewers, is load-bearing for the careers of everyone in the field. It isn't a side channel.



Herman

And it's under more strain every cycle it grows.



Corn

Which leaves the open question I'd want Daniel to sit with. arXiv has outsourced its quality control to the conferences. The conferences are running on a reviewer pool that is a fifth machine and shrinking under load. So what's the next tier — who quality-controls the quality controllers? Whether that's reputational systems layered above the venues, or something that looks nothing like peer review today, that's where I'd watch.



Herman

And I'd watch the venues that survive being the ones that solve the integrity problem, not the selectivity one. Everyone's already selective.



Corn

Thanks as always to Hilbert Flumingtop for producing the show. For more along these lines, there's episode twenty, Architectural AI; episode twelve, The AI Breakthrough; and episode nineteen, AI Images. This has been My Weird Prompts, with me and the walking encyclopedia. If you've enjoyed this one, a review on Spotify helps other people find us — my weird prompts dot com has everything else. See you tomorrow, Herman.



Herman

Tomorrow, Corn. Watch the withdrawal numbers.

SOURCES

What the research for this episode read before the script was written. Primary sources first.

1. **arXiv blog, Attention Authors: Updated Practice for Review Articles and Position Papers in arXiv CS Category, 2025-10-31, retrieved 2026-11-05. [primary]**
   https://blog.arxiv.org/2025/10/31/attention-authors-updated-practice-for-review-articles-and-position-papers-in-arxiv-cs-category/

2. **NeurIPS 2025 Fact Sheet, retrieved 2026-11-05. [primary]**
   https://media.neurips.cc/Conferences/NeurIPS2025/press/NeurIPS2025-Fact_Sheet.pdf

3. **NeurIPS Blog, Reflections on the 2025 Review Process from the Program Committee Chairs, 2025-09-30, retrieved 2026-11-05. [primary]**
https://blog.neurips.cc/2025/09/30/reflections-on-the-2025-review-process-from-the-program-committee-chairs/

4. **NeurIPS Blog, AI-Generated Papers in the NeurIPS 2026 Position Paper Track, 2026-06-02, retrieved 2026-11-05. [primary]**
https://blog.neurips.cc/2026/06/02/ai-generated-papers-in-the-neurips-2026-position-paper-track/

5. **NeurIPS 2025 LLM Policy, retrieved 2026-11-05. [primary]**
https://neurips.cc/Conferences/2025/LLM

6. **ICLR Blog, A Retrospective on the ICLR 2026 Review Process, 2026-03-31, retrieved 2026-11-05. [primary]**
https://blog.iclr.cc/2026/03/31/a-retrospective-on-the-iclr-2026-review-process/

7. **IEEE/CVF, New Research Presented at CVPR 2025, 2025-06-10, retrieved 2026-11-05. [primary]**
https://cvpr.thecvf.com/Conferences/2025/News/Technical_Program

8. **IEEE/CVF, CVPR 2025 wrap, 2025-06-19, retrieved 2026-11-05. [primary]**
https://cvpr.thecvf.com/Conferences/2025/News/Wrap

9. **CVPR 2025 Program Guide (main conference), retrieved 2026-11-05. [primary]**
https://media.eventhosts.cc/Conferences/CVPR2025/CVPR_main_conf_2025.pdf

10. **ACL 2025 Findings front matter, retrieved 2026-11-05. [primary]**
https://aclanthology.org/2025.findings-acl.0.pdf

11. **EMNLP 2025 Findings front matter, retrieved 2026-11-05. [primary]**
https://aclanthology.org/2025.findings-emnlp.0.pdf

12. **EMNLP 2025 Main front matter (1,811 main at 22.16%, 1,417 Findings at 17.34%), retrieved 2026-11-05. [primary]**
https://doi.org/10.18653/v1/2025.emnlp-main.0

13. **ACL Admin Wiki, 2025Q3 General Chair report, retrieved 2026-11-05. [primary]**
https://www.aclweb.org/adminwiki/index.php/2025Q3_Reports:_General_Chair

14. **Journal of Machine Learning Research homepage, retrieved 2026-11-05. [primary]**
https://jmlr.org/

15. **JAIR, About the Journal, retrieved 2026-11-05. [primary]**
https://www.jair.org/index.php/jair/about

16. **JAIR current issue (Vol. 87, 2026), retrieved 2026-11-05. [primary]**
https://www.jair.org/index.php/jair/

17. **China Computer Federation, AI venue catalogue (A-class list), retrieved 2026-11-05. [primary]**
https://www.ccf.org.cn/Academic_Evaluation/AI/

18. **Sharma, Joachims & Dean, Do LLMs Favor LLMs? Quantifying Interaction Effects in Peer Review, Cornell, 2026-01-30, retrieved 2026-11-05. [primary]**
http://arxiv.org/pdf/2601.20920v1

19. **Liang et al., Monitoring AI-Modified Content at Scale, ICML 2024 proceedings (PMLR v235), retrieved 2026-11-05. [primary]**
https://proceedings.mlr.press/v235/liang24b

20. **Computing Research Association Career Center, Academic Publishing and Peer Review in Computing Research, published 2026-07-06, retrieved 2026-11-05.**
https://careercenter.cra.org/career-resources/on-the-job-3/academic-publishing-and-peer-review-for-computing-researcher-2026-107

21. **OpenAccept, NeurIPS acceptance rates and submission stats, retrieved 2026-11-05.**
https://openaccept.org/c/ai/neurips/

22. **OpenAccept, ICML stats, retrieved 2026-11-05.**
https://openaccept.org/c/ai/icml/

23. **OpenAccept, AAAI 2025 stats, retrieved 2026-11-05.**
https://openaccept.org/c/ai/aaai/2025/

24. **OpenAccept, IJCAI stats, retrieved 2026-11-05.**
https://openaccept.org/c/ai/ijcai/

25. **OpenAccept, ACL 2025 stats, retrieved 2026-11-05.**
https://openaccept.org/c/ai/acl/2025/