Episode Book

EPISODE #6333

Filtering arXiv in the Age of Generative AI

PUBLISHED	28 October 2026 at 17:25 UTC
DURATION	23:07
LISTEN	myweirdprompts.com/episode/filtering-arxiv-submission-caps-67319/
HOSTS	Corn, Herman, Desmond, Neville
SCRIPT MODEL	openai/gpt-6.1-sol
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.25
SOURCES	15 cited by the research (listed at the end)
BOOK GENERATED	29 October 2026 at 08:41 UTC

SYNOPSIS

arXiv's monthly submissions doubled over two years to surpass forty thousand, driven by computer science and artificial intelligence. To manage the strain, the platform introduced a monthly submission cap of two papers per person. Filtering this volume requires navigating strict API limits, avoiding silent truncation in broad RSS feeds, and building digests that distinguish between institutional gates, generated text, and actual scientific claims without relying on superficial quality scores or unvetted author affiliations.

TRANSCRIPT



Corn

You can make the reading pile shorter without making the reading easier, Herman. Daniel is considering a weekly arXiv synopsis tool, and before building it he wants to know what he would actually be filtering. Roughly thirteen hundred new papers a day, a record September approaching forty thousand... that is a substantial commitment to having tabs open.



Herman

The daily average also hides the announcement schedule. Announcements go out Sunday through Thursday, so an actual mailing can carry roughly eighteen hundred to two thousand new papers. A weekly tool has to handle those batches without treating every announcement as a separate discovery.



Corn

Daniel wants the groundwork underneath the synopsis. What kind of research makes up that volume, who gets through the submission gate, what access permits, and whether existing tools already solve the useful part. I suspect generating a pleasant paragraph is going to be the least interesting engineering problem.



Herman

Especially if it is a pleasant paragraph about a paper you would never have wanted to read.



Corn

The abstract has already applied for that job.



Herman

Abstracts at least come attached to the author's claims. A generated synopsis can quietly make those claims sound independently verified. That matters before we even get to ranking.



Corn

Your instinct will be to build the ingestion system first. I can see you arranging the queues already.



Herman

I was arranging the identifiers, actually.



Corn

Forgive me. Entirely different furniture.



Herman

The identifiers matter because the underlying stream has just changed its rules. On October first, arXiv imposed a cap of two submissions per person per calendar month, with no more than three active submissions at once. That is a major intervention in a service built for rapid distribution.



Corn

A builder starts measuring demand and discovers the supplier is rationing it.



Herman

arXiv describes the cap as a stopgap while it develops better moderation tools and norms for AI-era submissions. Its scientific director, Steinn Sigurðsson, told Nature the steep rise was driven mainly by slop generated by artificial intelligence tools. That is his characterization of the growth, not a measurement assigning every paper a quality score.



Corn

Important distinction. Otherwise a filter built to rescue serious research starts by assuming everything newly arriving is rubbish. You could produce a remarkably compact digest that way.



Herman

Empty, punctual, cheap.



Corn

Your most reliable installation.



Herman

My installations are reliable once you stop changing things.



Corn

I close windows, Herman.



Herman

Some of those windows are services. Anyway... the cap means historical growth is a poor basis for forecasting the next mailing. A catch-up tool should cope with a changed stream, rather than depend on submissions continuing to accelerate at the same rate.



Corn

It also has to separate importance from volume. A category generating more papers should not automatically occupy more of my week.



Herman

Would you give each field an equal allocation?



Corn

No. I would give my attention an allocation. If everything relevant expands, the digest still has to stop.



Herman

Then the product needs an explicit attention budget. Otherwise improved discovery makes the user's problem worse. Every successful match becomes another thing they are supposed to read.



Herman

September brought forty thousand three hundred and sixty-three submissions. The comparable month two years earlier was about twenty thousand six hundred, and a decade earlier about ninety-nine hundred. Roughly doubled over two years, roughly quadrupled over a decade.



Corn

A weekly synopsis built around the older volume would now need its own weekly synopsis.



Herman

The operational burden is concrete. September generated almost nine thousand support tickets for staff and moderators. Submission volume means questions, disputes, corrections, moderation decisions. It does not simply mean more files sitting cheaply on storage.



Corn

Can we actually say most of those papers barely get read? It sounds plausible, but I would distrust a tool that treated obscurity as evidence of worthlessness. A paper on some narrow measurement might matter enormously to the people making that measurement.



Herman

I cannot give you a reliable fraction of arXiv papers that go unread or uncited. Attention is uneven, certainly. Citations also arrive late, differ across disciplines, and tell you little about a paper posted yesterday. For a weekly digest, immediate citation counts would reward established visibility rather than identify emerging value.



Corn

My guess is the useful long tail and the useless long tail look rather alike at posting time. Both might have unfamiliar authors and a title that attracts six people. The difference could be buried in whether the experiment actually answers its question.



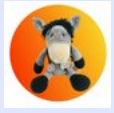
Herman

That pushes the tool toward explaining why a paper was selected. Shared subject, a method the reader follows, a result contradicting something in their library. Those are inspectable reasons. A single importance score hides the assumptions.



Corn

Especially if the score mostly means the author already has an audience.



Herman

Computer science now contributes nearly half the submissions. Its share was forty-five point three percent last year, compared with thirty-four point eight percent in twenty twenty-two. That changes what an archive-wide digest looks like even before personalization.



Corn

The default general-interest digest slowly becomes an AI newsletter wearing a university scarf.



Herman

The AI subject class is growing particularly quickly. Submissions from January through September rose one hundred and fifty percent against the same period last year. Over two years, that class has grown more than sixfold. That concentration makes the overall growth less representative of every field.



Corn

Although AI is also a active research area. A category can have valuable new work and a dreadful signal problem simultaneously. I would not want an anti-slop filter that simply penalized whichever vocabulary was fashionable.



Herman

Nor would I. Detecting AI-written prose is a poor substitute for evaluating a contribution. Writing assistance, generated filler, and an unsupported scientific claim are different things. A useful digest should preserve distinctions between a proposed method, a demonstrated result, and an author's speculation.



Corn

You have spent years reading pediatric trials. You are going to click every summary that says breakthrough, then complain about the controls.



Herman

I will complain before clicking. It saves time.



Corn

There is our personalization model.



Herman

What would make you trust an unfamiliar contributor, Corn? Suppose the affiliation field is empty.



Corn

A clear claim, enough detail to examine it, and some account of what would prove it wrong. Affiliation could help me locate the author. It would not settle whether the paper worked. I would resist making university membership a ranking feature.



Herman

arXiv's endorsement system does create an institutional fast path. Institutional email plus claimed authorship can provide automatic endorsement. Otherwise an established arXiv author can personally endorse someone. First submissions, including entry into a new category, need that gate.



Corn

Permission to enter the conversation, rather than a certificate that the conversation is correct.



Herman

arXiv explicitly says endorsement is not peer review. The purpose is to establish membership in the scientific community at a scale the service can sustain. It does not validate the paper's argument.



Corn

What we cannot infer is how many contributors are independent. Institutional access being easier tells us something about the gate, not the population behind it. Someone could also have an affiliation when submitting and do the actual work independently.



Herman

I do not have a documented independent-versus-affiliated split. Nor would I infer one from email domains. A builder can record affiliations when supplied without turning missing affiliation into a claim about someone's competence.



Corn

The new quota adds another distinction. The person pressing submit is not necessarily the person who did most of the research.



Herman

The cap applies across categories, and rejected submissions still count. Coauthors are unaffected by somebody else's submission quota. Collaborating groups could change who submits without changing the author list or the volume of work.



Corn

The submit button acquires a rota.



Herman

That is a possible response, not something we can assume has happened. Daniel's October count through the twenty-eighth is already about thirty-seven thousand seven hundred. That alone does not establish what the cap achieved. We need comparable periods and the usual submission rhythms.



Corn

It does establish that nobody should design around the flood having conveniently disappeared.



Herman

Or train a suspicious-author detector on submission frequency without understanding that policy change. The administrative record may shift while the scientific collaboration stays much the same.



Corn

Access is where I would prefer boring infrastructure. Download once, remember what arrived, update what changed. Your version will require a dashboard before breakfast.



Herman

A small dashboard.



Corn

An observatory for missed emails.



Herman

The public interfaces give you several routes. RSS and Atom are useful for current category announcements. The legacy API supports queries. OAI-PMH supports metadata harvesting. Bulk downloads handle large file collections. The terms also cover SWORD, although depositing papers is not the central job of Daniel's reader.



Corn

A weekly service should not need to hammer the website to discover that nothing happened.



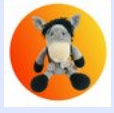
Herman

The API terms specify one request every three seconds, one connection at a time, across all machines under your control. Spreading requests across extra servers to evade that limit is explicitly forbidden.



Corn

Distributed computing meets distributed manners.



Herman

Batching makes routine metadata collection quite manageable. The limit is per request, not per paper. I should narrow that... a full archive metadata harvest is a different workload from downloading every full-text file. Those need different access routes.



Corn

Otherwise somebody hears three million papers and starts calculating three million individual requests.



Herman

Precisely the design to avoid. RSS feeds update daily at midnight Eastern time and can combine categories, but a feed request has a two-thousand-result limit. A broad feed near that ceiling should not be treated as proof you collected the whole archive.



Corn

Quiet truncation would be particularly nasty here. The tool could confidently explain the week's research after silently losing the end of the week.



Herman

What would you want it to do when collection is incomplete?



Corn

Tell me. Keep the previous material, mark the affected interval, retry politely. Never fill the missing space with the model's general impression of recent science.



Herman

That also means separating collection from generation. A failed download should not become a fluent answer. Stable identifiers, version tracking and retrieval timestamps give the synopsis something concrete to stand on.



Corn

And now, as ever, here is Desmond with Desmond's Money Tip.

D

Desmond

This is Desmond's Money Tip. Reuse one sheet of aluminium foil for a full calendar year. Washed, dried, folded flat in the drawer, it saves you maybe four pounds over twelve months, which over a working lifetime is about one hundred and sixty pounds, or roughly the price of a modest fridge nobody will thank you for.

D

Desmond

The embarrassing detail is that I have a drawer of foil sheets labelled by year, and I know which one is two thousand and nineteen by feel.

D

Desmond

Money is just the slow counting of things you did not throw away. Mind the drawer.



Corn

I could label a drawer... Herman, stop looking appalled. Bulk files still leave the question of what Daniel may actually serve.



Herman

Metadata is the generous part. Titles, abstracts, authors, identifiers and classification terms are available under CC0. You can retrieve, store, transform and share that descriptive metadata. A search interface, notifications, topic histories and links back to abstract pages fit comfortably into that model.



Corn

The abstract being included matters. A useful first version could rank metadata and produce an explicitly abstract-based synopsis without ingesting the entire PDF collection.



Herman

Full text has different rights. arXiv is not the copyright holder. Most papers use its non-exclusive distribution license, which lets arXiv distribute them without granting everyone else permission to rehost the PDF or source.



Corn

Open to read does not mean available to republish.



Herman

Personal use and research storage are permitted within the terms. Serving files from your own servers requires a license that allows it. Some papers have Creative Commons licenses with different permissions and conditions. The tool should retain the license for the particular version it uses.



Corn

A summary is not automatically the same legal act as republishing a PDF, though. We should avoid implying that every generated explanation is prohibited.



Herman

Yes. The clear prohibition here is unauthorized full-text redistribution. Whether a particular commercial summarization workflow is permitted needs its own assessment, including what material it copies or exposes. A summary generator does not acquire blanket rights merely because its output is shorter.



Corn

The download button has not hired a solicitor.



Herman

For full-text bulk access, S3 is requester-pays. The complete file set was about nine point two terabytes in April last year, growing by roughly a hundred gigabytes a month. Those are dated capacity figures, not a claim about its exact size today.



Corn

Requester-pays is unusually clear product naming.



Herman

Files arrive in archive chunks with XML manifests. That is useful for substantial research ingestion, but unnecessary for a weekly metadata-first digest. Buying the whole collection before proving relevance would be an expensive way to postpone deciding what the tool is for.



Corn

Existing tools suggest people are quite good at postponing that decision with features.



Herman

SummarizePaper is the substantial example in this small open-source digest ecosystem. It offers summaries, full-text chat, collections, saved searches, daily digests and weekly letters. It also has citation alerts and journal-club slides.



Corn

The paper now has a more organized social life than I do.



Herman

Its repository has roughly three hundred stars and current activity. Several smaller projects have only a handful. The recurring design is category retrieval, keyword filtering, an LLM summary, then email delivery.



Corn

Stars measure attention to a repository, not reader retention. A small project might quietly serve its author perfectly. An inactive one might even be finished, although dependencies have strong opinions about that.



Herman

We cannot diagnose abandonment from one last-push date. Still, the landscape is heavily personal-scale and category-specific. Materials-science digests, AI digests, language-processing digests. There is even a meme-style summary project.



Corn

Peer review by reaction image.



Herman

The straightforward loop is clearly reproducible. Sustained relevance is harder to demonstrate. A command-line interface can improve delivery for one reader, and a recommender can learn preferences, but neither automatically solves unfamiliar terminology or changing interests.



Corn

A weekly overview also differs from a pile of paper summaries. If several papers test the same idea, I want the relationship. Did they agree? Were they measuring different things? Is a claimed improvement just a different comparison? Compressing each abstract separately misses that work.



Herman

Cross-category discovery is another opening. Something useful may be classified outside the field you habitually browse. The system needs to find that without broadening your feed until everything is relevant.



Corn

You are about to recommend learning from every click.



Herman

I was going to ask whether you would tolerate that.



Corn

Not alone. I click things because they look wrong. Give me explicit reasons to save or reject a recommendation, and let a new interest expire. Otherwise one afternoon of curiosity becomes a lifelong subscription to an obscure apparatus.



Herman

Negative feedback needs care too. Not useful this week differs from never show this subject. The preference model should preserve that distinction.



Corn

Audio makes selection even less forgiving. I can skim past an irrelevant paragraph. I cannot glance past eight minutes while listening.



Herman

The small open-source digest landscape does not establish an audio market gap. I cannot name a confirmed dedicated paper-podcast product from these examples. It does suggest that format alone is not much protection. Email, terminal or audio, the user still needs a reason to trust the selection.

N

Neville

...which is covered in my expert briefing. Well, the title is covered. The page says Weekly Paper Triage and then there is quite a lot of space. I was going to fill that in after finishing my intern shift.



Corn

Neville Muddledorf, our specialist in the remaining space.

N

Neville

I did try the idea for a course project, Corn. My flatmate wanted papers about robots handling soft things. I added cloth, then fruit, then hands. Actually, hands was a mistake. The Friday email was enormous. I printed it at the library for two pounds and put a little square beside everything he might like. He ticked the square beside the unsubscribe instructions.



Herman

Your keywords widened the selection without telling you which matches mattered. A relevance check against papers your flatmate actually wanted would help more than another keyword.

N

Neville

Oh, that is much better, actually. I had read half the brief and thought the endorsement meant someone had checked the findings. That would have saved me reading them.



Corn

Endorsement checks entry into the scientific community. It does not peer-review the findings. A summary still needs to keep the author's claim separate from what has been established.

N

Neville

Right. I shall move that sentence. Uncle Hilbert keeps everything in boxes, labelled and dated, with boxes inside boxes, so I thought I ought to give every paper a folder. He is technically my great uncle twice removed, then sideways. Anyway, I made the folders before asking my flatmate what he wanted. Now he wants the weekly email back, but only the bits I marked in pencil. I have not worked out how to automate my pencil.



Herman

Those pencil choices are useful examples. Keep the reasons alongside them, including why you rejected plausible matches. Otherwise you will have a collection of answers without the decisions that produced them.



Corn

His flatmate has given a better specification than unlimited personalization. Show me the choices someone can defend. That could become a small evaluation set before Daniel spends money on full-text processing.



Herman

Include wanted papers and plausible distractions. Test whether a changed filter retrieves the wanted material without making the weekly selection unmanageable. Not a universal benchmark... a working check against the reader's actual needs.



Corn

Plus a route to the things excluded. I want a short weekly selection, but I do not want the machine quietly deciding the rest of science no longer exists.



Herman

A compact main digest with inspectable search results underneath would do that. Selection can be strict while discovery remains open. Maintenance then includes checking collection completeness, model changes, licenses and whether recommendations still match the reader.



Corn

The common misconception is that once fetching and summarizing work, the tool is finished.



Herman

That produces readable output. A useful catch-up tool also has to maintain selection quality, preserve provenance and respect content permissions as both the service and the reader change.



Corn

I would build the narrow version first: metadata, explainable selection, links to originals. The open question worth watching is whether the October cap changes the incoming stream enough to alter what a good weekly filter must reject.



Herman

Our thanks to Hilbert Flumingtop for producing the show, and to Neville for covering the desk.



Corn

If you want more of this, try episode twenty-five twenty-six, How Peer Review Actually Works (and Fails); episode eleven ten, The arXiv Effect; and episode seven, Building Custom ASR Tools. This has been My Weird Prompts, the human-AI collaboration podcast. No question too obscure, no rabbit hole too deep.



Herman

You can find the show at my weird prompts dot com.



Corn

See you tomorrow.

SOURCES

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

1. **arXiv API Access [primary]**
https://info.arxiv.org/help/api/index.html

2. **Terms of Use for arXiv APIs [primary]**
https://info.arxiv.org/help/api/tou.html

3. **Bulk Data Access via S3, updated 2023-03-23 [primary]**
https://info.arxiv.org/help/bulk_data_s3.html

4. **RSS news feeds [primary]**
https://info.arxiv.org/help/rss.html

5. **arXiv License Information [primary]**
https://info.arxiv.org/help/license/index.html

6. **The arXiv endorsement system [primary]**
https://info.arxiv.org/help/endorsement.html

7. **Monthly Submissions, as of 2026-10-28 [primary]**
https://arxiv.org/stats/monthly_submissions

8. **Updated rate limit policy, 2026-10-01 (via search snippet; direct fetch rate-limited) [primary]**
https://blog.arxiv.org/2026/10/01/updated-rate-limit-policy/

9. **Nature, ArXiv versus AI slop (via search snippet)**
https://www.nature.com/articles/d41586-026-03194-w

10. **New Scientist (via search snippet)**
https://www.newscientist.com/article/2591679-arxiv-might-not-survive-ai-slop-onslaught-warn-mathematicians/

11. **The Register, 2026-10-02**
https://www.theregister.com/ai-and-ml/2026/10/02/arxiv-imposes-rate-limit-on-paper-submissions-to-stem-the-ai-slop-tide/5300899

12. **Daniel Lemire, 2026-10-04**
https://lemire.me/blog/2026/10/04/arxiv-capped-submissions-at-two-a-month/

13. **SummarizePaper statistics guide**
https://www.summarizepaper.com/en/guides/arxiv-statistics/

14. **312, MIT, last push 2026-10-28**
https://github.com/summarizepaper/summarizepaper

15. **arXiv blog, 2026-07-09**
https://blog.arxiv.org/2026/07/09/arxiv-now-hosts-over-3-million-articles/