

32 TOOLS ACROSS 14 CATEGORIES. 16 OF THEM COMPLETELY FREE.

The Research Stack (2026)

EVERY PRICE VERIFIED AUGUST 17 AND AUGUST 20, 2026

- We opened every one of these and checked the price on the vendor's own page, not a roundup. The date above is real.
- No affiliate links, no sponsored placements, no paid spots. Nobody is here because they asked.
- Every tool gets a "watch out". If the free tier has a catch that bites you at submission, that is the part we lead with.
- You will not find a journal-matching tool in here. We build one, so we are not the right people to tell you whose is best. Everything else on this list is a category we have no product in.

Most tool roundups are affiliate pages wearing a lab coat. This one is not. Every price below was read off the vendor's own page. Where a vendor blocks that, we say so in the entry and quote no figure rather than repeat a number from a roundup, and two tools were dropped entirely for it.

The useful finding: a researcher with no budget can still cover almost the whole list. Sixteen of these thirty-two tools are free with no paid tier at all, and the one category where free genuinely struggles is figures, for a licensing reason most people discover far too late.

Quick reference

The whole stack in one place. Save this page and skip the rest until you need it.

FINDING PAPERS	<u>Semantic Scholar</u>	FREE
FINDING PAPERS	<u>Connected Papers</u>	
FINDING PAPERS	<u>Elicit</u>	
FINDING PAPERS	<u>Research Rabbit</u>	
FINDING PAPERS	<u>Litmaps</u>	
FINDING PAPERS	<u>Consensus</u>	
SYSTEMATIC REVIEW SCREENING	<u>Rayyan</u>	
SYSTEMATIC REVIEW SCREENING	<u>Covidence</u>	
SYSTEMATIC REVIEW SCREENING	<u>ASReview</u>	FREE
REFERENCE MANAGEMENT	<u>Zotero</u>	FREE
REFERENCE MANAGEMENT	<u>Mendeley</u>	
READING PAPERS	<u>Gemini Notebook (formerly NotebookLM)</u>	
EQUATION AND TABLE OCR	<u>Mathpix Snip</u>	
TRANSCRIPTION	<u>Otter.ai</u>	
TRANSCRIPTION	<u>Whisper</u>	FREE
WRITING AND LANGUAGE	<u>Overleaf</u>	
WRITING AND LANGUAGE	<u>LanguageTool</u>	
FIGURES	<u>BioRender</u>	
FIGURES	<u>Inkscape</u>	FREE
STATISTICS	<u>JASP</u>	FREE
STATISTICS	<u>jamovi</u>	FREE
STATISTICS	<u>G*Power</u>	FREE
LAB NOTEBOOKS	<u>eLabFTW</u>	
SURVEY AND DATA CAPTURE	<u>REDCap</u>	FREE
CHECKING A JOURNAL IS REAL	<u>Think. Check. Submit.</u>	FREE

PREPRINTS AND SHARING	<u>bioRxiv and medRxiv</u>	FREE
PREPRINTS AND SHARING	<u>OSF</u>	FREE
PREPRINTS AND SHARING	<u>Zenodo</u>	FREE
PREPRINTS AND SHARING	<u>Dryad</u>	
PREPRINTS AND SHARING	<u>arXiv</u>	FREE
INTEGRITY CHECKS	<u>Retraction Watch Database</u>	FREE
INTEGRITY CHECKS	<u>PubPeer</u>	FREE

HOW TO READ THIS

Start with the two free ones that everything else hangs off: Zotero for your library and Semantic Scholar for search. Get those right and half the other tools stop being necessary.

FINDING PAPERS



Semantic Scholar

Free search across 237 million papers, with one-line summaries

[semanticscholar.org](https://www.semanticscholar.org) (<https://www.semanticscholar.org>)

Free. There is no paid tier.

WHY IT MATTERS

This is the corpus the paid tools are built on. Elicit, Consensus and most of the "AI research assistants" sit on top of it, so you can search the same 237 million papers for nothing. Every result carries a one-line summary, and citations are ranked by how influential they actually were instead of by raw count, which is how you find the three papers that matter rather than the thirty that cite politely. It is run by the Allen Institute for AI, a nonprofit, so there is no upsell waiting at the end.

WATCH OUT

The one-line summaries are machine-written and will turn a hedged finding into a confident one. Read the abstract before you cite it. Coverage outside STEM, and outside English, is noticeably thinner.

Connected Papers

Turns one paper into a visual map of its whole field

[connectedpapers.com](https://www.connectedpapers.com) (<https://www.connectedpapers.com>)

Free tier is capped at a small number of new graphs each month. A low-cost academic tier lifts the cap.

WHY IT MATTERS

Give it one paper and it builds a graph of the surrounding literature using co-citation, not just direct citations. That matters because it surfaces relevant work that never cited your seed paper, which is exactly the work you miss and a reviewer finds. The Prior Works and Derivative Works views separate the foundational papers from the follow-ups, so you can see at a glance whether you are standing on a 2003 method or a 2025 one.

WATCH OUT

A graph is only as good as the seed paper you feed it, so a niche or badly-chosen starting paper gives you a thin map. The paid tier only lifts the monthly graph cap, it does not unlock different features. We could not verify their exact price first-hand because the pricing page will not render for a fetcher, so check it yourself before paying.

Elicit

Pulls sample sizes and outcomes into one comparison table

elicit.com (<https://elicit.com>)

Freemium. Free tier includes unlimited search across 138 million papers and Zotero import. Pro is \$49/month billed annually.

WHY IT MATTERS

The reason to use this over a chatbot is the table. Name a variable, say sample size or outcome measure or effect size, and it extracts that field across dozens of papers into one grid you can actually scan. For a systematic review or a first pass at a literature landscape, that collapses a week into an afternoon. The free tier gives you unlimited search over 138 million papers, which is genuinely usable without paying anything.

WATCH OUT

It reads real papers rather than inventing them, so it hallucinates far less than a raw chatbot, but 2026 studies still find retrieval-based systems fabricate a small share of citation links, and DOIs are the least reliable field it extracts. Never cite anything it surfaces without opening the actual paper.

Research Rabbit

Maps how papers connect, and learns what you care about

[researchrabbit.ai](https://www.researchrabbit.ai) (<https://www.researchrabbit.ai>)

Free to sign up and use.

WHY IT MATTERS

Keyword search finds papers that use your words. It misses the paper that solved your problem in a different vocabulary, which is the one that gets you scooped or desk-rejected for missing prior work. Research Rabbit works from citations and co-

authorship across more than 310 million papers, and it keeps adjusting to what you actually open. Point it at four or five papers you already trust and it will surface the neighbours you did not know about.

WATCH OUT

It has a pricing page, so treat "free" as true today rather than guaranteed forever. Citation-based discovery also inherits citation bias: it finds what is well cited, not what is right.

Litmaps

Citation maps with alerts when new work cites your seeds

[litmaps.com](https://www.litmaps.com) (<https://www.litmaps.com>)

Free: 2 maps, 100 articles per map, basic search up to 20 inputs. Pro is \$10 per month, or \$120 a year.

WHY IT MATTERS

The feature worth paying attention to is the alert. You build a map around your seed papers and Litmaps tells you when something new cites into that neighbourhood, which is how you avoid discovering a competing paper at review. The free tier is two maps of 100 articles, enough to test whether the visual approach suits how you think before spending anything.

WATCH OUT

Two maps fills up quickly if you work across several projects, and the free basic search caps at 20 inputs. Budget for Pro if you run more than one review at a time.

Consensus

Ask a research question, get answers with the papers attached

consensus.app (<https://consensus.app>)

Has a free tier and paid tiers. We could not read their pricing page first-hand, so we are not quoting a figure.

WHY IT MATTERS

You ask a question in plain language and it answers from the literature with the underlying papers cited, which is a different motion from searching keywords and reading twenty abstracts to find out whether anyone agrees. It is most useful early, when you are working out whether a claim is contested or settled before committing a paragraph of your introduction to it.

WATCH OUT

It summarises, and a summary is not a reading. Open the papers before you cite them, because an AI restatement of a finding is exactly the kind of thing that becomes a misattributed citation in your reference list. Their pricing page blocks automated readers, so verify the current tiers yourself.

SYSTEMATIC REVIEW SCREENING

Rayyan

Blind title and abstract screening with a co-reviewer

[rayyan.ai](https://www.rayyan.ai) (<https://www.rayyan.ai>)

Free forever: 3 active reviews, 2 invited reviewers, 1 sample. Essential is \$4.99 per seat per month billed annually.

WHY IT MATTERS

Screening thousands of abstracts in a spreadsheet is where systematic reviews quietly go wrong, because two reviewers can see each other's calls and drift into agreement. Rayyan blinds you to your co-reviewer until you both decide, then surfaces the conflicts to resolve deliberately. The free tier carries 3 active reviews and 2 invited reviewers, which covers most single-project theses outright. Paid tiers add PRISMA flow diagrams, which you will need for the methods section anyway.

WATCH OUT

Duplicate auto-resolution and PICO extraction are paid-only, so on the free tier you are still deduplicating by hand. That is the job that eats an afternoon.

Covidence

Cochrane-aligned screening and extraction for review teams

[covidence.org](https://www.covidence.org) (<https://www.covidence.org>)

\$339 per year for a single review, \$907 per year for up to three. Free trial covers 500 records.

WHY IT MATTERS

This is the tool most Cochrane-style reviews are actually run on, and its structure is the reason: screening, full-text review, risk-of-bias assessment and extraction all live in one audited trail, so your PRISMA numbers come out of the software rather than out of your memory. Unlimited collaborators on every plan means the price does not scale with your team. Check first whether your university library already holds a licence, because many do and it is the single biggest saving here.

WATCH OUT

The licence is per review and expires after twelve months, so a review that overruns its year needs renewing. Authors in low and lower-middle income countries can contact them about pricing.

ASReview

Open-source AI screening that reorders your unread queue

asreview.nl (<https://asreview.nl>)

Free and open source, no paid tier. Runs locally on your own machine.

WHY IT MATTERS

Coordinated at Utrecht University, it keeps re-sorting your unscreened records as you screen, so the likely-relevant ones surface first. In practice you find most of what matters early and can make an evidence-based decision about when to stop, instead of grinding through every abstract to be safe. It runs entirely on your own machine, which matters when your records are not yours to upload to a commercial service.

WATCH OUT

It reorders your queue, it does not decide for you, and stopping early is a methodological choice you have to defend in your methods section. Reviewers will ask.

REFERENCE MANAGEMENT

Zotero

Free open-source reference manager with 9,000+ citation styles

[zotero.org](https://www.zotero.org) (<https://www.zotero.org>)

Free and open source, forever. Only cloud sync storage costs money: 300 MB free, then \$20/year for 2 GB.

WHY IT MATTERS

This is the one tool on the list we would tell every researcher to install today. It handles 9,000+ citation styles, captures a reference from your browser in one click, and shared group libraries cost nothing no matter which storage tier you are on. Because it is open source there is no vendor lock-in and no publisher owns your reading history. If you outgrow the free storage you can point sync at your own WebDAV server and pay nothing at all.

WATCH OUT

The free 300 MB fills up fast once you sync PDFs, which is roughly 100 to 150 papers. What you eventually pay for is attachment storage, never the software itself.

Mendeley

Reference manager with 2 GB free storage, owned by Elsevier

[mendeley.com](https://www.mendeley.com) (<https://www.mendeley.com>)

Freemium. 2 GB free storage, which is about 6x Zotero's free tier. Paid upgrades above that.

WHY IT MATTERS

The 2 GB free allowance is the most generous no-cost PDF storage among the mainstream managers, and the Word and LibreOffice citation plugins work without configuration. If your lab already standardises on it, that shared habit is worth more than switching to something marginally better on your own.

WATCH OUT

It is owned by Elsevier, a publisher with a direct commercial interest in your reading data, which some researchers treat as disqualifying on its own.

Mendeley Desktop is also in limbo: Elsevier announced its retirement, reversed that in July 2025, and has said sign-ins will eventually stop with notice. Build on Mendeley Reference Manager, not Desktop.

READING PAPERS

Gemini Notebook (formerly NotebookLM)

Answers only from the papers you upload, with passage citations

notebook.google (<https://notebook.google>)

Freemium. Free tier: 100 notebooks, 50 sources per notebook, 50 chat queries a day.

WHY IT MATTERS

The difference from a normal chatbot is that it is grounded. It answers only from the sources you upload and links each claim back to the exact passage, so you can check it in one click instead of trusting it. Fifty sources per notebook on the free tier is enough for a real literature review. The audio overview feature turns a stack of papers into something you can listen to on a commute, which is the most underrated part.

WATCH OUT

It was renamed from NotebookLM in July 2026, so older bookmarks and guides point at a redirect. The per-source caps are the same on every tier, meaning a very large file fails to import no matter what you pay. And uploading an unpublished manuscript to a Google product is a data-governance decision your institution may already have rules about.

EQUATION AND TABLE OCR

Mathpix Snip

Screenshots equations and tables straight into LaTeX

mathpix.com (<https://mathpix.com>)

Free tier available. Pro starts at \$4.99 per month.

WHY IT MATTERS

Retyping an equation out of a PDF is both slow and the easiest place in a manuscript to introduce a silent error, because a wrong subscript looks exactly like a right one. Snip screenshots the equation and returns the LaTeX. It does the same for tables, which is the other thing nobody wants to retype. If you write in Overleaf this removes a genuinely tedious step.

WATCH OUT

We could not confirm the free tier's monthly snip allowance on the vendor's own pricing page, so we are not quoting one. Check it before you build a workflow around the free tier.

TRANSCRIPTION

Otter.ai

Transcribes interviews and meetings with speaker labels

otter.ai (<https://otter.ai>)

Free: 300 minutes per month, 30 minutes per conversation, 3 lifetime file imports. Pro is \$8.33 per user per month billed annually.

WHY IT MATTERS

If you run interviews, focus groups or oral histories, transcription is the step that silently eats weeks. The free tier gives you 300 minutes a month with speaker identification, which is roughly ten half-hour interviews. The 30-minute cap per conversation is the real constraint, so long sessions need splitting before you record. Paid moves you to 1,200 minutes and a 90-minute cap.

WATCH OUT

Only 3 file imports on the free tier, ever, so existing recordings are effectively excluded. Check your ethics approval before uploading identifiable participant audio to any cloud service.

WH Whisper

Runs transcription on your own machine, nothing uploaded

github.com/openai/whisper (<https://github.com/openai/whisper>)

Free and open source under the MIT licence. No account, no upload, no per-minute cost.

WHY IT MATTERS

The answer when your ethics approval does not permit sending participant audio to a third party, which is more often than people assume. Whisper runs locally, handles dozens of languages, and has no minute cap because there is no service behind it. For a lab with a decent GPU it is the cheapest transcription you will ever do, and the recordings never leave your control.

WATCH OUT

You have to install it, and accuracy depends on the model size you can afford to run. It does not label speakers, so multi-person interviews need a

separate diarisation step.

WRITING AND LANGUAGE

Overleaf

Browser LaTeX with journal templates and live collaboration

[overleaf.com](https://www.overleaf.com) (<https://www.overleaf.com>)

Freemium. Free tier: unlimited projects but only 1 collaborator per project. Student is \$8.25/month billed annually.

WHY IT MATTERS

The headline on the free tier is unlimited projects, so you are limited on collaborators rather than on output. It carries submission templates for most major journals, which saves you the formatting pass entirely. Check with your library before you pay: a lot of universities hold a site licence that hands you the top tier free, and researchers routinely pay for something they already have.

WATCH OUT

The one-collaborator limit on the free tier bites the moment you co-author, which in academia is nearly always. Long documents with heavy figures can also hit the free compile timeout partway through a thesis, and the paid tiers exist largely to lift that.

LanguageTool

Grammar and style in 30+ languages, and you can self-host it

languagetool.org (<https://languagetool.org>)

Freemium, open-source core. Free tier checks 2,000 characters at a time across 30+ languages.

WHY IT MATTERS

If English is not your first language this beats the obvious alternative, because it covers 30+ languages properly rather than treating everything as English. The part that matters for unpublished work: the core engine is open source and can be self-hosted, so a manuscript you have not submitted yet never has to leave your own machine. That is a real answer to the confidentiality objection every other writing tool has.

WATCH OUT

The free 2,000-character limit is about one paragraph of a manuscript, so checking a whole paper free means tedious chunking. Free-tier style suggestions are also noticeably shallower than the paid rule set.

FIGURES

BioRender

Publication-quality life-science figures, but read the license first

[biorender.com](https://www.biorender.com) (<https://www.biorender.com>)

Free tier: 3 figures, and no publication rights. Individual is \$35/month billed annually.

WHY IT MATTERS

Nothing else gets you a clean mechanism schematic in ten minutes. The icon library covers cells, organisms, pathways and lab equipment, all drawn to a consistent style, so you stop rebuilding a mitochondrion by hand. Just be clear about what the money buys: the paid plan is what grants permission to publish the figure in a journal, plus watermark-free high-resolution export. The drawing is not the thing you are paying for. The license is.

WATCH OUT

The free plan says it in writing: no commercial use or publication, and 3 figures total. Researchers find this out at submission, with the figure already in the manuscript. At \$35/month it is also the most expensive tool on this list by a wide margin. Inkscape is free with no publication restriction. It just makes you draw it yourself.

Inkscape

Free vector editor that exports the EPS and PDF journals demand

inkscape.org (<https://inkscape.org>)

Free and open source. No tiers, no caps, no publication restrictions.

WHY IT MATTERS

This is the free answer to "the journal rejected my figure because it is not vector". It works natively in SVG and exports to EPS, PDF, PostScript and AI, which is exactly the format list a production editor asks for. It is actively maintained rather than abandoned, and because it is open source there is no license question hanging over anything you publish with it.

WATCH OUT

The learning curve is real, and there is no scientific icon library, so every cell and arrow is yours to draw. Budget an afternoon the first time. That is the honest trade against BioRender: you pay in effort instead of dollars.

STATISTICS

JASP

Free SPSS-style stats with Bayesian versions of every test

jasp-stats.org (<https://jasp-stats.org>)

Free and open source, supported by the University of Amsterdam. No paid tier.

WHY IT MATTERS

Every standard analysis appears in both frequentist and Bayesian form side by side, which is rare in free software and increasingly what a reviewer means when they ask for a Bayes factor. Output comes out APA-formatted by default, so the table you paste into the manuscript needs almost no reworking, and it connects to OSF directly for sharing your analysis file.

WATCH OUT

It is less flexible than R once your model stops being standard, and because it is GUI-first your analysis is less reproducible than a script unless you deliberately save and share the project file.

jamovi

Point-and-click stats that shows you the R code it wrote

[jamovi.org](https://www.jamovi.org) (<https://www.jamovi.org>)

Free and open source.

WHY IT MATTERS

It is a graphical front end to R, so you can click your way through an ANOVA or a mixed model and then reveal the R syntax it generated underneath. That makes it the best on-ramp we know of for someone leaving SPSS and heading toward reproducible code, because you learn the script by seeing it appear next to the thing you just clicked.

WATCH OUT

The module ecosystem is smaller than R proper, and you will outgrow it once your models get complicated. Treat it as a bridge rather than a destination.

G* **G*Power**

Free power analysis that justifies your sample size to reviewers

[psychologie.hhu.de/arbeitsgruppen/allgemeine-psychologie-und-arbeitspsychologie/gpower](https://www.psychologie.hhu.de/arbeitsgruppen/allgemeine-psychologie-und-arbeitspsychologie/gpower) (<https://www.psychologie.hhu.de/arbeitsgruppen/allgemeine-psychologie-und-arbeitspsychologie/gpower>)

Free for everyone, including commercial use.

WHY IT MATTERS

This is the tool that produces the a priori sample-size justification that reviewers and ethics boards now ask for as standard. It covers t, F, chi-square and z tests plus ANOVA, repeated measures, correlation, logistic and Poisson regression, and it draws the power curves you paste straight into the methods section. Running it before you collect data is the cheapest way to avoid an unanswerable reviewer question later.

WATCH OUT

It is effectively frozen. The current builds date from 2020, so macOS compatibility gets shakier with each OS release and there is no support for modern mixed or multilevel designs. For those, use the R packages instead.

eLabFTW

Open-source electronic lab notebook you can self-host

elabftw.net (<https://www.elabftw.net>)

Self-hosted: free and open source under AGPLv3, with no features behind a paywall. A hosted option is sold separately.

WHY IT MATTERS

Paper lab notebooks fail the moment somebody asks you to produce the raw record for a paper published three years ago. eLabFTW gives you timestamped, searchable entries with a real audit trail, and because the self-hosted version is AGPLv3 with nothing paywalled, your institution can run it without a per-seat licence. When a reviewer or an integrity office asks what you did and when, this is the artefact that answers.

WATCH OUT

Self-hosting means somebody has to run and back up the server. If that person is you and you leave, the lab has a problem. Their hosted pricing is not listed on the main site.

SURVEY AND DATA CAPTURE

REDCap

Secure survey and study database, free to non-profits

project-redcap.org (<https://project-redcap.org>)

Free to non-profit organisations that join the consortium. Self-hosted by your institution.

WHY IT MATTERS

If you are collecting data from human participants, the question at review is not whether your spreadsheet was tidy, it is whether the capture was auditable. REDCap gives you validated entry, an audit trail and exports that drop straight into your stats package, and it is free to non-profits through the consortium: 8,392 institutional partners across 166 countries and over 2.8 million projects. Your institution very likely already runs it.

WATCH OUT

You cannot just sign up as an individual. Access comes through your institution being a consortium member and hosting it, so the first step is finding out whether yours already does.

CHECKING A JOURNAL IS REAL

Think. Check. Submit.

The checklist that tells you if a journal is real

thinkchecksubmit.org (<https://thinkchecksubmit.org>)

Free, CC BY 4.0 licensed.

WHY IT MATTERS

Two structured checklists, one for journals and one for book publishers, backed by a coalition that includes COPE, DOAJ, ISSN, OASPA and STM. It teaches a repeatable process instead of maintaining a blacklist, which is why it does not go stale and why it avoids the defamation problem that killed earlier predatory-journal lists. The materials are translated into many languages, which matters if English is not your first.

WATCH OUT

It is a process, not a verdict. It will never tell you that a specific journal is predatory, and a well-built fake journal can pass a superficial pass of the checklist. You still have to actually do the checking.

PREPRINTS AND SHARING

bioRxiv and medRxiv

Free preprint posting that timestamps your finding in 48 hours

[biorxiv.org](https://www.biorxiv.org) (<https://www.biorxiv.org>)

Free to post and free to read.

WHY IT MATTERS

Papers appear within 48 to 72 hours of submission, which establishes your priority long before a journal decision. Here is the number that makes the argument: roughly 80% of preprints are eventually published in a peer-reviewed journal, at an average lag of around 250 days. That is eight months of visibility, citation and feedback you simply give away by waiting. In 2025 the two servers took in over 64,000 new preprints between them.

WATCH OUT

Governance changed in March 2025, moving from Cold Spring Harbor Laboratory to a new nonprofit called openRxiv. More practically, a minority of journals still treat a preprint as prior publication, so check your target journal's policy first. Screening is for scope and safety, not peer review.

OSF

Free preregistration and project hosting, 50 GB if you go public

osf.io (<https://osf.io>)

Free and open source, run by the nonprofit Center for Open Science. 5 GB per private project, 50 GB per public one.

WHY IT MATTERS

Preregistration is the cheapest credibility upgrade available to an early-career researcher: a timestamped, frozen analysis plan that answers the p-hacking question before anyone asks it. OSF is where most of it happens. Note the storage ratio, 5 GB private against 50 GB public, which is a deliberate nudge toward openness, and it connects free to Dropbox, Google Drive and GitHub if you go past the caps.

WATCH OUT

The caps are per project, so once a private project hits 5 GB you cannot add files until you make it public or attach outside storage. Preregistrations are also permanent and cannot be deleted, which is the entire point but still surprises people the first time.

Zenodo

Free CERN-run repository that gives any output a DOI

zenodo.org (<https://zenodo.org>)

Free, unlimited records.

WHY IT MATTERS

Built and run by CERN with OpenAIRE, so its preservation backing is about as durable as research infrastructure gets. Every upload gets a DOI within seconds, which makes a dataset, a slide deck, a script or a negative result independently citable. The GitHub integration is the quiet win: connect a repository and every release mints its own DOI automatically, so your code becomes a citable artifact without any extra work.

WATCH OUT

Zenodo does no curation, so a DOI there gives you citability, not credibility. Where a proper discipline-specific repository exists, use that instead. Per-record storage has a generous but real ceiling, so check it before uploading imaging or sequencing data.

Dryad

Curated data repository with a citable DOI per dataset

datadryad.org (<https://datadryad.org>)

Sponsored free at partner institutions up to 10 GB. Otherwise \$150 for up to 5 GB. Fee waivers exist for unfunded authors.

WHY IT MATTERS

Journals increasingly require a data availability statement, and the ones that link to a real repository tend to be cited more than the ones that say "available on request". Dryad curates every submission before release, so you get a citable DOI and a human check rather than a bare file dump. Check your institution first: partner

members have submissions up to 10 GB sponsored, which turns a \$150 charge into nothing.

WATCH OUT

Charges scale hard with size, from \$150 at 5 GB to \$808 at 100 GB, and waivers are not approved above 10 GB. Size your deposit before you promise it in a cover letter.

arXiv

Free preprint server for physics, maths, CS and more

arxiv.org (<https://arxiv.org>)

Free to submit and free to read. No author charges.

WHY IT MATTERS

In several fields arXiv is where the work actually appears first and the journal version is the formality that follows. Posting there timestamps your claim publicly, which is the cheapest protection you have against being scooped while a journal sits on your manuscript. Submission costs nothing and there are no author charges at any stage.

WATCH OUT

New authors, or anyone posting to an unfamiliar category, may need an endorsement from an established author before they can submit. Everything is moderated, so allow time. Check your target journal permits preprints first, though most now do.

INTEGRITY CHECKS

Retraction Watch Database

Free retraction database, updated every working day

[crossref.org/documentation/retrieve-metadata/retraction-watch](https://www.crossref.org/documentation/retrieve-metadata/retraction-watch)

(<https://www.crossref.org/documentation/retrieve-metadata/retraction-watch/>)

Free and public.

WHY IT MATTERS

Crossref bought this database in September 2023 specifically to make it permanently public, so what used to be proprietary is now a free CSV and API updated once per working day. Running your reference list against it before you submit takes minutes and catches the genuinely damaging mistake of citing a retracted paper. Editors do check, and a retracted citation in your reference list is the kind of thing that follows a paper around.

WATCH OUT

Retractions are logged with real lag, so a very recent one may not be in there yet. The database tells you a paper was retracted, not whether the retraction was the author's own correction or a misconduct finding, so read the notice itself before you draw a conclusion.

PubPeer

Browser extension that warns you when a paper is disputed

pubpeer.com (<https://pubpeer.com>)

Free, including the browser extension.

WHY IT MATTERS

Install the extension and it quietly flags any paper you are reading that has post-publication comments on it, which catches problems that never made it as far as a formal retraction. Image-manipulation reports are the most common category by a distance. Threads here have triggered real retractions and fraud investigations, so it is the closest thing to an early-warning system for the literature you are about to build on.

WATCH OUT

Comments are anonymous and unrefereed, which cuts both ways: some are rigorous forensic work and some are unfounded sniping. Treat a thread as a prompt to look harder yourself, never as a verdict on a paper or on the people who wrote it.

ONE MORE THING

A tool stack does not fix a paper

These get the mechanics out of your way: your library, your figures, your stats, your sample-size justification. What they cannot tell you is whether an editor will send your manuscript out for review, or desk-reject it in the first ten minutes.

That judgment is the thing worth borrowing before you submit.

[Run the free readiness check](#)