

NC RAA ANNUAL CONFERENCE · CHAPEL HILL · JULY 2026

The Verification Gap

Preparing research administration for AI-enabled research

When research gets faster, more interdisciplinary, and harder to verify

George Koomullil

An email arrives in your office

The attachment

A complete R01 draft: aims, biosketches, budget justification. Polished and well cited.

The note

"AI helped me pull this together, but I checked it. Can we submit?"

The clock

Due Monday. Your office signs before it ships.

Polished. Plausible. Due Monday.

THE QUESTION

The question is not whether AI wrote it.

Can your office stand behind it?

The bottleneck in research is moving from producing the work to trusting it.

40+

grant applications filed by a single PI in
one round, many with AI assistance



6

the cap NIH set per PI per year, effective
September 25, 2025

The same notice, [NOT-OD-25-132](#), says applications substantially developed by AI are not the applicant's own original ideas.

This is not a future problem. The funders are already reacting.

Where I come from

RESEARCHER

PhD in computational engineering.
Postdoc at Duke. NRC fellowship
at NASA Ames.

BUILDER

Two decades of building AI tools
for research and IP: finding
expertise, protecting inventions,
moving discovery to market.

FOUNDER

Three companies. The current
one is a PBC: legally bound to
mission and stakeholders, not
just returns.

I have sat on both sides of this system.


WHERE I STAND

As research gets faster and cheaper, the scarce resource becomes trust.

This hour: what is changing upstream, why it reaches your desk first, four working principles, and a 90-day plan that costs discipline, not money.

THE ROUTE

Where we are headed this hour



Read the river

What is changing upstream, before it reaches your desk?

- upstream
- downstream
- discipline
- 90-day plan

Which of these reached your office this year?

- A** An AI-drafted proposal or aims page crossed your desk
- B** A faculty member asked what AI use they must disclose
- C** A citation in a document did not check out
- D** A worry about confidential data entering an open AI tool
- E** A capability or team claim you could not verify


The deepest change is not in your inbox yet. It is inside research itself.

Two public signals from the last eighteen months. Both directional, and I will flag that as we go.

57% to 84% in one year

Researchers using AI tools in their work, 2024 to 2025. For research and publication tasks specifically, 45% to 62%.


Wiley ExplanAltions researcher surveys, 2024 and 2025. Self-reported; directional.



+57% applications

Grant applications to major funders rose 17% by 2024 and 57% by 2025 against the 2022 baseline. Some schemes saw up to +142%. Reviewer capacity did not move.

Research on Research Institute;
Rees and Wilsdon, *Nature*, 2026.



Why research gets faster and cheaper

1 Reading the field

Literature that took a semester to survey is summarized, translated, and cross-linked in an afternoon.

2 Designing the work

Hypotheses, methods, and analysis plans drafted on demand, in any field, at near-zero marginal cost.

3 Running the analysis

Code gets written conversationally; pipelines that needed a programmer now need a prompt.

4 Writing it up

A polished draft in hours, not weeks: papers, proposals, and reports, each fluent and fully cited.

Faster, cheaper, wider, deeper

AlphaFold predicted structures for about 200 million proteins, nearly every protein known to science. Each once took years in the lab.

DeepMind and EMBL-EBI

GNoME proposed 2.2 million new crystal structures, roughly 800 years of materials discovery at the prior pace.

GNoME, Nature, 2023

An autonomous AI Scientist drafts complete papers, experiments included, for about \$15 in compute each.

Sakana AI, 2024

Frontier models now answer PhD-level exam questions in biology, physics, and chemistry at expert level, in every field at once.

GPQA benchmark, 2024-25

Cheaper by orders of magnitude, capable in every field, and layered too deep to check with a skim.

Four changes coming to research

1

Velocity and volume

More research produced faster, everywhere, at once.

2

Interdisciplinary by default

A capable collaborator in every field, on demand.

3

Expertise gets hard to read


When anyone can sound expert, a CV tells you less.

4

Trust becomes the bottleneck

Confident output arrives faster than anyone can check it.

What expertise can you trust in AI-enabled research?



When every project is interdisciplinary: who is the expert here, and expert in what?

The early studies point one direction

AI-assisted writing clusters toward what already exists; distinctiveness narrows.

Qian [et. al.](#), 2026

Across 41 million papers: output and citations up, breadth of focus down.

41M-paper study, 2026

Fluent answers create illusions of understanding; they feel right either way.

Messeri and Crockett, 2024

A growing share of scholarly text shows signs of LLM modification.

Liang et al., 2025

This year, the gap became measurable

A 2025 audit flagged roughly 147,000 suspect citations across 2.5 million papers, about twelve times the 2023 rate.

Citation audit, 2025

About one hundred fabricated citations passed peer review at a top AI conference.

NeurIPS, 2025

A major preprint server paused some submission types; volume had outrun its moderators.

arXiv, 2025

Authors of an autonomous AI Scientist warn their own system could overwhelm review.

Nature coverage, 2026

The reviewers who caught these were doing verification work. So will you.

THE MOMENT WE ARE IN

**Everything can sound expert now.
And the acceleration is real.**

Both are true. The faster the work gets, the wider the distance between what is produced and what is verified.

NAME THE PROBLEM


The verification gap

The widening distance between how fast research output appears and how fast an institution can confirm that it is accurate, original, compliant, and real.

When text becomes cheap, trust becomes the work.

THE SHAPE OF IT

Output is outrunning verification



Illustrative. The shaded region is where unverified work accumulates: the institution's risk.


The gap reaches you first

Where does it show up, and whose desk does it land on?

- upstream
- downstream
- discipline
- 90-day plan


DOWNSTREAM

You sit downstream



When the current speeds up, it reaches you first.

The same gap, in your daily workflows



Each one is a place where polished output meets your signature.

Volume and capacity

Triage becomes the bottleneck; the work outruns the people.

- More proposals, more subawards, more complexity, on the same headcount
- AI pours volume into a pipeline built for scarcity
- This is the burnout conversation, with the pressure turned up




Compliance and integrity

Brand-new obligations, and no playbook yet.

- AI-use and disclosure rules, written differently by every funder
- Fabricated citations and generated text to catch before submission
- Confidential material that must never enter an open tool



The funders moved first



Each rule is the same admission: output has outrun verification.

Expertise and team formation

When everyone can sound expert, who can actually do the work?

- Teams form across fields faster than anyone can vet them
- When every proposal reads well, fit beats polish, and fit must be proven
- The signal moves from what is claimed to what has been demonstrated



Where the funders go next

1 Claims become data

Free-text narratives shrink; structured, checkable claims grow. SciENCv, ORCID, and persistent IDs are the plumbing, already mandated.

2 Evidence outweighs polish

When every proposal reads well, prose stops carrying signal. Funded work, published papers, and preliminary data decide.

3 Research stays open; claims get policed

The same agencies keep funding AI-enabled science at scale. What tightens is the application: disclosure, originality, review.

4 The people mission gains weight

The NIH cap exempts training and education awards. Discovery can be accelerated; trained researchers cannot be generated.

5 Verification lands on your signature

Funders will lean on institutional certification. The authorized signature stops being a formality and becomes the audit point.

THE STAKES

The risk is that something untrue ships with your institution's name on it.

A fabricated citation. An overstated capability. An undisclosed use of AI.

THE TURN

**The offices that close the gap
stop being overhead.
They become assurance.**

Research administration does not get automated away. It gets promoted.

What it asks of the profession

What does closing the gap actually take?

- upstream
- downstream
- discipline
- 90-day plan

Evidence over assertion

Trust what has been demonstrated, not what is asserted.


Funded grants, published papers, and issued patents keep their meaning when prose gets cheap. Claims do not.

IN PRACTICE

Before asking how good it sounds, ask what this person, team, or tool has actually produced.



Keep a trust boundary



Automate below the line. Decide above it. Never let fluency blur it.

Human judgment stays the point

AI proposes. People decide, and stay accountable.

A tool can inform your judgment. It cannot carry your accountability.

IN PRACTICE

Every consequential output gets a named human owner before it leaves the office.



Label honestly

Say plainly what is verified, and what is not.

In a flood of confident, unlabeled output, candor is rare, and it compounds.

IN PRACTICE

Label what you circulate: verified, estimated, or unverified. The field guide defines each.



THE SHIFT

From processors of paperwork to stewards of trust.

EVIDENCE · BOUNDARY · JUDGMENT · HONESTY

Not AI skills. Administration skills. This profession was built for this moment.

The 90-day plan

What can you start Monday morning, without a budget line?

- upstream
- downstream
- discipline
- 90-day plan

Try the habit before the plan

"Dr. Smith has led three NSF awards in AI-enabled coastal resilience and is the natural PI for this proposal."

What would you check before you rely on that sentence?

15 min

to verify one claim by hand, when the sources cooperate



100+

citations in a single proposal, before you reach the capability claims

Multiply by every proposal, biosketch, subaward, and report your office signs in a year, while volume rises and headcount stays flat.

The two-minute drill works. By hand, it does not scale.

Fight scale with scale

Hand machines the mechanical checks

- Does every citation resolve, open, and say what is claimed
- Do the numbers reconcile: budget, effort, dates, totals
- Is AI use disclosed the way this funder requires
- Pull the public evidence: awards, papers, patents, in seconds

Keep people on the judgment

- What a flagged mismatch means, and what to do about it
- Whether the evidence actually supports the claim
- The decision to sign, with a named owner
- Relationships, strategy, and the genuinely hard cases

The technology that created the flood can also run the checks. Automate below the trust boundary; decide above it.

Four questions for any tool

PROVENANCE

Where did this answer come from? Can I see the source?

CONTROL

Whose hands does our data end up in?

AUDIT

Can someone check the work after the fact?

ACCOUNTABILITY

Is a named human answerable for the decision?

Ask them of every vendor.

Six moves that form a loop

1 See the field

Funding and opportunity
intelligence from public evidence.

2 Map capability

Expertise read from grants,
papers, and patents

3 Govern AI-use

A one-page policy: where AI may
work, what gets disclosed

4 Verify before it ships

A release checklist and honest
labels on every output

5 Invest in people

Shift time toward judgment,
relationships, verification

6 Measure what matters

Three metrics beyond volume and
win rate

One loop. Run it on one live opportunity, then run it again.

Open these this week

FUNDING	PEOPLE AND PAPERS	IP AND INTEGRITY
NIH RePORTER and Matchmaker NSF Award Search USAspending Grants.gov and forecasts	OpenAlex ORCID Semantic Scholar Lens.org scholarly search	USPTO full-text search Google Patents Retraction Watch doi.org resolver

All free. No vendor required. The cost is discipline, not money.

A one-page policy beats a perfect one

- 1 Where AI may be used, and where it may not
- 2 How AI use gets disclosed, matched to each funder's rules
- 3 Who reviews AI-assisted work before it leaves the office
- 4 How confidential material stays out of open tools

Anchor it to the guidance from your funders: [NOT-OD-23-149](#), [NOT-OD-25-132](#). Revisit quarterly; the rules are moving.

The last check before your name goes on it

- 1 Every citation opens, and says what the document claims it says
- 2 Capability claims match demonstrated evidence: awards, papers, patents
- 3 AI use is disclosed the way this funder requires
- 4 A named person has signed before it leaves the office

Then label it: verified, estimated, or unverified.

People first, then the scoreboard

Move people up the value chain

- Shift time from formatting to judgment and relationships
- Train verification as a skill: provenance, disclosure, checking
- The answer to burnout is value, not speed

Measure what matters

- Quality and originality of what goes out
- Equity of reach across departments and career stages
- How much ships verified, not just how much ships
- What the institution learned this quarter

PUT IT TOGETHER

The first ninety days

WEEKS 1-4 · SEE CLEARLY

- Open the free sources; give each a named owner
- Draft the one-page AI policy

WEEKS 5-8 · PILOT SMALL

- Evidence map for one live opportunity
- Run the release checklist on one real proposal

WEEKS 9-12 · MAKE IT STICK

- Pick three metrics beyond volume
- Labels on everything you circulate

Day 91: one page to leadership. What changed, what it caught, what is next.

READ THE RIVER AHEAD

When text becomes cheap, trust becomes the work.

Demand evidence. Keep the boundary. Keep judgment human. Label honestly.

Thank you.

The field guide can be downloaded from the QR code on the right: the 90-day plan, the policy starter, the release checklist, the sources.

One ask: pick one live opportunity Monday and run the loop on it.

Questions welcome.



field guide and sources

George Koomullil · Founder and CEO, Ascendr, a public benefit corporation

george@ascendr.ai

Sources you can check

POLICY, ADOPTION, AND THE SCIENCE

NIH NOT-OD-25-132 - six-app cap; AI-developed work not original
grants.nih.gov/NOT-OD-25-132

NIH NOT-OD-23-149 - generative AI barred from peer review
grants.nih.gov/NOT-OD-23-149

NSF - generative AI in the merit-review process
nsf.gov/notice-to-the-research-community-on-ai

Wiley ExplanAltions 2025 - AI use 57->84%, 45->62%
wiley.com/ai-study/for-researchers

Rees & Wilsdon, Nature 2026 - volume +57%, MSCA +142%
researchonresearch.org/could-agentic-ai-topple-grant-funding-systems

AlphaFold - ~200M structures; 2024 Nobel in Chemistry
embl.org/news/science-technology/alphafold-wins-nobel-prize-chemistry-2024

GNoME, Nature 2023 - 2.2M crystal structures
nature.com/s41586-023-06735-9

Sakana AI - The AI Scientist (~\$15 per paper)
sakana.ai/ai-scientist

GPQA - graduate-level Google-proof Q&A benchmark
arxiv.org/abs/2311.12022

WHAT THE EARLY STUDIES FIND

Qian et.al., 2026 - LLMs narrow proposal distinctiveness
arxiv.org/ab/2601.15485

Hao et al., arXiv 2024 (accepted at Nature) - 41.3M papers; focus narrows
arxiv.org/abs/2412.07727

Messeri & Crockett, Nature 2024 - illusions of understanding
nature.com/s41586-024-07146-0

Liang et al., Nature Human Behaviour 2025 - LLM-modified text
nature.com/s41562-025-02273-8

Hallucinated citations at scale - Nature analysis, 2026
nature.com/d41586-026-00969-z

~100 fabricated citations at NeurIPS 2025
arxiv.org/abs/2602.05930

arXiv pauses CS reviews/position papers - Nature, 2025
nature.com/d41586-025-03664-7

The AI Scientist could overwhelm review - Nature, 2026
nature.com/s41586-026-10265-5

Free tools and funder policy

THE FREE TOOLKIT (MOVES 1 AND 2)

NIH RePORTER & Matchmaker
reporter.nih.gov

NSF Award Search
nsf.gov/awardsearch

USAspending
usaspending.gov

Grants.gov
grants.gov

OpenAlex
openalex.org

ORCID
orcid.org

Semantic Scholar
semanticscholar.org

Lens.org
lens.org

USPTO Patent Public Search
uspto.gov/patents/search

Google Patents
patents.google.com

THE FREE TOOLKIT, CONTINUED

Retraction Watch
retractionwatch.com

DOI resolver
doi.org

FUNDER POLICY PLUMBING

SciENCv - NSF-mandatory since Oct 2023 (biosketch, C&P)
[nsf.gov / faq-using-sciency](https://nsf.gov/faq-using-sciency)

NIH Grants Policy Statement
[grants.nih.gov / nihgps](https://grants.nih.gov/nihgps)