

ARCHON: The Scientific Swarm

Autonomous Agentic Workflows for Verifiable, Publication-Ready Manuscripts

Ian Scott Kinney, Marcos Guzmán, and Tomáš Zvolenský

Abstract—Publishing is where research slows down. A study can be finished, the data clean and the findings solid, and it will still take months to become a manuscript a Q1 journal will accept. Writing, formatting, citing, and defending every claim consume the time of the people least able to spare it. ARCHON gives that time back. It is a swarm of 55 specialized agents that carries a project from raw input to a peer-review-hardened manuscript, engineered for submission to Q1 and Q2 journals indexed in Web of Science and Scopus. ARCHON accepts six starting points: a finished thesis, a researcher’s own dataset, a topic, a rejected manuscript, a half-finished draft, or open data ARCHON gathers itself. The destination is always the same. Everything in the finished paper is verifiable. Before any writing begins, ARCHON reads the sources and locks their contents into a frozen, fully-traced evidence ledger. Every sentence draws from that ledger, and every number carries a trail back to the exact place it came from. Every reference is checked against Crossref, PubMed, OpenAlex, and Semantic Scholar, and trusted only when at least two of them agree, because a resolving DOI proves nothing on its own. When a citation cannot be confirmed, the pipeline halts. The finished draft then faces an adversarial review panel that reasons against the frozen evidence, under a rule that no manuscript with an unresolved critical finding can be accepted, however well it reads. None of this is a theory about how publishing works. ARCHON encodes a method that has taken more than 3,250 manuscripts to publication for over 100 universities across fourteen years. The researcher’s name goes on the paper. ARCHON makes sure they can defend every line of it.

Index Terms—Agentic workflows, scientific publishing, evidence grounding, citation verification, provenance, adversarial review, retrieval-augmented generation, Q1–Q2 journals.

I. FOURTEEN YEARS OF GETTING PAPERS PUBLISHED

IT started with a question from a mid-sized university in the south of Chile: *can you help us publish 80 papers this year?*

That was fourteen years ago. The answer we chose was person-first. Rather than writing papers on the institution’s behalf, we taught its researchers to publish, on the premise that an institution whose people can publish grows on its own. To teach it, we had to master it. We learned publishing as a craft, one manuscript, one rejection, one revision at a time. We were journalologists before the term was in wide circulation.

The first finding surprised us. Publishing has far less to do with subject matter than anyone assumes. A chemical engineering paper and a history paper reach acceptance through the same sequence of decisions. Author guidelines differ, word counts differ, formats differ. Underneath sits one constant: a communicational strategy that gets a manuscript taken seriously

enough to reach substantive peer review. Field expertise belongs to the researcher and cannot be outsourced. Getting the work read is a separate skill, and it transfers across every discipline.

We have tested that claim about as widely as it can be tested. Animal science and astronomy through zoonotics and zythology, the science of brewing beer. More than 3,250 manuscripts have gone through the method, at more than 100 universities. Several have climbed into the major international ranking systems on the strength of what their own researchers learned to produce.

ARCHON is that method, transferred to agents.

We did not prompt a language model into a plausible workflow and hope it held. We built the swarm around steps we already knew worked, and held it to the standard we have used for over a decade: find the viable and practical niche, establish the novelty and the original contribution, communicate clearly and directly, and verify everything before a researcher’s name goes near it. There are no tricks in the method. There were none before the agents existed, and there are none now.

We do not ask the AI how to publish a paper. We spent fourteen years learning that. The agents do the labour.

II. THE VERIFICATION GAP

A language model drafts a methods section in seconds, summarizes a literature in a paragraph, and never tires. For work that consumes weeks of a scientist’s year, that speed looks like a gift.

The trap is just as real. A model produces text that reads as authoritative whether or not it is true. It writes a perfectly formatted citation for a paper that was never published. It reports a coefficient to three decimals that appears nowhere in the dataset. It states a conclusion the cited study does not support, in prose indistinguishable from a grounded claim. Fluency tells the reader nothing about which is which.

That is the verification gap. A model generates language. It does not confirm that the language is true. Three failure modes follow.

Fabricated references. A plausible source with real-sounding authors, a real journal, and a well-formed DOI that resolves to nothing or to something unrelated. It survives a quick glance and fails a careful one.

Invented or misreported numbers. A statistic generated rather than measured, or a real value copied into the wrong place. The prose reads precisely, and the precision is false.

Claims that outrun their source. The paper says more than the study behind it actually found. A careful “may be associated with” becomes “causes.” One small study becomes “the literature shows.” The reference is genuine, and the sentence attached to it is not what the study reported.

In fourteen years of reading reviewer reports, we have learned which errors an editor forgives and which end a submission on sight. These belong to the second group. Peer review is adversarial by design, and a reviewer who finds one fabricated reference assumes there are others. Editors now screen for exactly these signatures. A fabricated citation that reaches print invites correction or retraction, and that cost lands on the researcher whose name is on the paper.

The quieter cost is verification itself. A researcher who cannot trust the output must check all of it: every reference chased to its source, every number traced back to the data, every claim checked against the study it cites. Done properly, that takes as long as writing the section by hand. The time the model appeared to save is spent again on catching what it got wrong.

Speed was never the bottleneck for a careful researcher. What they lack is a manuscript that verifies itself before a human has to.

These three failure modes set the standard for the rest of this paper. Section V returns to each one.

III. THREE LEVELS OF GROUNDING

Every system that writes a manuscript answers one question: where does the text come from? There are three answers in practice, and they are not equivalent.

A. Level 1 · The prompt window

Most researchers meet AI writing through a chat window. They describe the study, paste in what fits, and ask for a draft. The model writes from what it absorbed in training plus whatever was pasted. It has no access to the sources themselves, so it reconstructs them from memory.

Every failure mode in Section II is available at this level, and nothing in the system is positioned to catch any of them. The model cannot distinguish a reference it retained accurately from one it is inventing, because both are produced by the same process. Verification falls entirely on the researcher, after the fact, with no record of where any claim came from.

B. Level 2 · Retrieval

The obvious improvement is to give the model the sources. Most automated writing tools do this through retrieval-augmented generation, or RAG. Documents are split into fragments, each fragment is indexed by meaning, and when the model writes about a topic the system fetches the fragments closest to the query and places them in the model's working context.

This is a real improvement over Level 1, and for search or question answering it is the right tool. For a manuscript it leaves the gap open, for reasons that are structural rather than incidental.

Retrieval returns a subset. Similarity search surfaces the fragments that look closest to the query. If the passage containing the needed value ranks just outside the cutoff, it is never retrieved, and the model completes the sentence from memory instead. That substitution is where a hallucinated

number comes from. Coverage is probabilistic, and the system cannot tell the difference.

Retrieval grounds the input, not the output. It improves what the model reads before it writes. It checks nothing afterward. A model handed the correct source can still misattribute it, overstate it, or cite a different paper, and the pipeline emits that sentence unexamined.

A resolving DOI is treated as proof. A citation passes if it is well-formed and its identifier resolves. That test misses the common failure where a genuine DOI carries the wrong title or an altered author list.

The finished claim keeps no trail. Once a sentence is written, the link between the claim and the fragment that supported it is gone. A reviewer who questions a number has no way to trace it, and neither does the system.

Nothing is empowered to stop. A pipeline built to produce a draft treats completion as success. Where a review pass exists, it checks internal consistency rather than hunting the fabricated reference that should block submission.

Levels 1 and 2 differ in quality and share one property. The burden of trust ends on the researcher, who becomes the verifier of last resort for work the machine never confirmed.

C. Level 3 · Complete, frozen evidence

ARCHON does not retrieve while it writes. It reads the sources in full first, extracts everything into one complete ledger, records where each item came from, and freezes the ledger before a single sentence is drafted. The writing agents draw only from that record. There is no moment when the right passage sits outside a retrieval window and gets replaced by a guess, because nothing is fetched on demand and nothing is missing.

Verification then runs against that same frozen record: every reference confirmed across Crossref, PubMed, OpenAlex, and Semantic Scholar, every claim checked against the source it rests on, every number bound to its exact location. When something cannot be confirmed, the pipeline halts rather than writing around it. Table I sets the three levels side by side.

IV. THE ARCHON ARCHITECTURE

Five commitments define the system. Each answers a failure named in Section III, and each is built into the architecture rather than added to a language model afterward.

- 1) **Complete evidence rather than sampled evidence.** The whole record of the sources, held in one place.
- 2) **Verification at every stage rather than at the end.** Checking is part of writing.
- 3) **The researcher stays in control.** The agents do the labour. The scientist decides, at checkpoints that cannot be skipped.
- 4) **The authority to stop.** When something cannot be confirmed, ARCHON halts and says so.
- 5) **Best over Fast.** Given a quicker answer and a defensible one, ARCHON chooses defensible.

TABLE I
THREE LEVELS OF GROUNDING COMPARED

	Prompt window	Retrieval pipeline	ARCHON
Source of the text	Training memory	Fragments fetched per query	A complete, frozen evidence ledger
Coverage	None of the actual sources	An approximate subset	The whole record
Missing evidence	Filled from memory	Filled from memory	Recorded and surfaced
Verification	None	Grounds the input only	Checks every output claim against its source
Citations	Unchecked	Trusted when the DOI resolves	Trusted only when at least two of Crossref, PubMed, OpenAlex, and Semantic Scholar agree
Provenance	None	Lost once written	Travels with every claim
On a fabricated reference	Emits it	Emits it	Halts
Review	None	Internal consistency, cannot block	Adversarial, evidence-grounded, can block
The researcher's role	Verifier of last resort	Verifier of last resort	In control at the checkpoints

A. How the work flows

Whatever a project starts from, it travels one sequence.

Ingest. ARCHON reads the source material in full: the thesis, the dataset, the prior manuscript, or the open data it gathered. Nothing is skimmed.

Ground. ARCHON extracts the evidence into a single complete ledger, recording for every item which document, which page, which value. When the ledger is complete, ARCHON freezes it. Nothing can then add to it or alter it. This forensic integrity architecture is what makes everything downstream auditable. A claim with no support in the record cannot be written as though it had one.

Write. Specialized agents draft the manuscript section by section, each with a narrow job, each working only from the frozen ledger. Every number carries its trail back to the source, so the draft arrives already traceable.

Verify. Every reference is confirmed across Crossref, PubMed, OpenAlex, and Semantic Scholar. Every factual claim is checked against the source it rests on. Anything the sources do not support is recorded as a gap and surfaced.

Review and revise. An adversarial review panel reasons against the frozen evidence rather than the prose. The panel can block. Revision addresses the findings and only the findings, and the manuscript is checked again.

Deliver. A formatted, submission-ready manuscript with a clean, verified reference list.

B. A swarm rather than a single model

ARCHON runs as 55 specialized agents, each responsible for one task and orchestrated in sequence. A single model asked to do everything at once is hard to inspect and easy to lead astray. A swarm of narrow agents, each with a defined job and a defined output, can be checked at every handoff. When one agent's only responsibility is to confirm references and another's is to trace numbers, integrity becomes a property of the architecture rather than a hope about a model's behaviour.

The division of labour mirrors how our editorial teams have always worked. No one person ever checked a manuscript's references, statistics, structure, and argument in a single pass. The swarm is that team, at scale.

C. The researcher in command

Automation is not delegation of judgment. At defined points ARCHON pauses and hands the decision back: approve the



Fig. 1. The ARCHON pipeline. Six kinds of input converge into one sequence. The evidence ledger is frozen before any writing begins, and an adversarial review with the authority to stop stands between the draft and delivery. The researcher approves at checkpoints that cannot be skipped.

evidence before writing begins, approve the direction before the draft is finalized, approve the manuscript before it is called done. Top scientists are not ready to hand their name to an

One sentence in the manuscript

schematic

Treatment reduced recurrence by **18.4 %** over 24 months [7].




Fig. 2. Anatomy of a verified sentence. The value resolves to an entry in the frozen ledger, and the reference clears the agreement threshold across independent databases. Every claim in the manuscript carries this apparatus behind it, so when a reviewer asks where a figure came from, the answer already exists. The example is schematic.




Fig. 3. Reference verification. A reference enters the manuscript only when at least two of the four databases agree on the title and the authors. A conventional pipeline emits an unconfirmed reference and continues.

Fig. 4. Value placement. A number may enter a sentence only with an entry in the frozen ledger behind it, and the locator travels with the claim from that point on.

agent, and they should not have to. The human stays the author and the decision-maker, with the swarm as the instrument.

V. WHAT THE ARCHITECTURE GUARANTEES

Section II named three failure modes. Each has a mechanism against it, and one review stage stands behind all three. The figures in this section show how ARCHON behaves. They are not its implementation.

A. Against fabricated references: consensus verification

ARCHON treats every citation as a claim to be proven. Before a reference enters the manuscript, ARCHON confirms it against Crossref, PubMed, OpenAlex, and Semantic Scholar, and accepts it only when at least two of them agree on both the title and the author list. A resolving DOI is insufficient alone, because a genuine identifier can be attached to the wrong paper.

A reference that cannot clear the bar is never quietly dropped or written around. ARCHON stops and flags it.

A reference in an ARCHON manuscript is one that at least two of those four authorities confirm, or it does not appear. Fig. 3 shows the decision and its two outcomes.

B. Against invented numbers: provenance on every value

Sample sizes, coefficients, p-values, effect sizes. Each is bound to the exact place in the source it came from. A writing agent cannot state a value without pointing to the item in the frozen ledger that supports it, and that pointer travels with the claim through the rest of the pipeline. When a reviewer asks where a figure came from, the answer already exists. Fig. 2 follows a single sentence through both checks, and Fig. 4 gives the rule that governs the value.




Fig. 5. The review decision. One unresolved critical finding blocks acceptance, however well the draft reads, and revision returns it to the panel.

C. Against claims that outrun their source: knowing the limit in advance

A sentence can quietly say more than the study behind it found, and nothing in the prose gives it away. ARCHON already knows what the sources do and do not establish, because the evidence was gathered in full and frozen before the writing began. Verification reads each finished sentence back against that record and asks one question: does the source actually say this? A sentence that goes further than its source is caught by that comparison, rather than by someone judging how confident the writing sounds.

D. Behind all three: an adversary between the draft and the journal

A finished draft is not treated as done. An adversarial review panel attacks it, looking for the overstated claim, the unsupported number, the citation that does not say what the sentence claims. The panel reasons against the frozen evidence rather than an impression of the prose, and it has teeth. One unresolved critical finding blocks acceptance regardless of how well the draft reads. Revision then addresses those findings and only those findings, and the draft returns for another check.

A manuscript earns its way past review by surviving an attack on its evidence. Fig. 5 states the rule.

VI. SIX ENTRY POINTS, ONE STANDARD

Section I explained why one method serves every discipline. The same logic applies to starting points.

From a finished thesis. A doctoral thesis usually contains a publishable paper, buried in three hundred pages no journal will read. ARCHON isolates the contribution, reframes the results for a Q1 or Q2 audience, and carries every figure and citation across intact and verified.

From your own data. You have the dataset and the analysis. ARCHON ingests your data files directly, extracts the results into the frozen ledger, and writes the manuscript around them. Every number is bound to your data from the start, so the Results section is traceable the moment it exists.

From a topic. Give ARCHON a topic and the variables of interest, and it runs a systematic review to the PRISMA 2020 standard: searching the literature, screening studies against a

TABLE II
ENTRY POINTS AND DELIVERABLES

You start with	ARCHON delivers
A finished thesis	A focused journal article drawn from it
Your own dataset	A manuscript built around your verified results
A topic and variables	A PRISMA 2020 systematic review
A rejected manuscript and its reviews	A revision that answers each reviewer, point by point
A half-finished draft	A completed paper in your own voice
Only a research domain	Gathered open data carried through to a finished paper

protocol, appraising them, and synthesizing what they show. The search can be reproduced.

From a rejected manuscript. ARCHON reads the manuscript and the reviews together, separates each reviewer concern into a discrete item, and works through them one at a time against the evidence. The response is traceable point by point, which is what an editor expects on resubmission.

From a half-finished draft. ARCHON grounds what is already written, finds the gaps, completes the missing sections in the author's own voice, and runs the whole through the same verification as any other manuscript.

From open data ARCHON gathers. Name a domain, and ARCHON discovers and gathers relevant open data, assembles it into a structured dossier, and carries it through the full pipeline.

Six doors into the same building. The input changes. The standard does not.

VII. BEST OVER FAST

Every guarantee in Section V costs time. Confirming references across four independent databases is slower than accepting a DOI. Tracing every number is slower than writing a plausible one. Assembling complete evidence before writing is slower than fetching fragments on demand. Surviving an adversarial review is slower than calling a draft finished.

ARCHON pays that cost deliberately. That is why its governing principle is Best over Fast. In research, the expensive error is the fast paper that has to be corrected, retracted, or defended after it is in print. A slow paper costs hours. A wrong one can cost a career. ARCHON is built to be slow in exactly the places where speed would cost a researcher their credibility.

VIII. WORKING WITH US

Language models can already write a competent paragraph. What they cannot do alone is make it trustworthy: references that exist, numbers with a source, claims that stay inside the evidence. ARCHON closes that gap in the architecture rather than moving it onto the human, using a method that has already taken more than 3,250 manuscripts to publication.

The researcher stays in command throughout. The agents assemble, source, and check. The scientist decides, writes with the swarm, and keeps the authorship. The heavy lifting moves to the agents. The paper stays the researcher's.

Early-access cohort. ARCHON is opening to a limited number of research groups and institutions. We are onboarding

principal investigators who want to run their own work through the pipeline, whether that is a finished thesis, a dataset ready to become a paper, a draft that stalled, or a manuscript returned with reviews. Places are limited, and each group works directly with us during onboarding.

Tomáš Zvolenský · tomas@journalrevisions.com

Our experience, our agents. Save time, better papers.

APPENDIX A GLOSSARY

Adversarial review. A review stage in which the system tries to break its own draft, searching for overstated claims, unsupported numbers, and citations that do not match the source. It reasons against the recorded evidence rather than the prose.

Agent swarm. ARCHON’s design of many specialized single-purpose agents, each responsible for one task and orchestrated in sequence, so the work can be inspected at every handoff.

Best over Fast. ARCHON’s governing principle. Given a quicker result and a more defensible one, the system chooses defensible and accepts the added time as the cost of trust.

Claim drift. A sentence that says more than the study behind it found, for example turning “may be linked to” into “causes,” or one small study into “the literature.” The reference is genuine. What the sentence claims is not what the source reported.

Consensus verification. ARCHON’s rule for citations. A reference is trusted only when at least two of Crossref, PubMed, OpenAlex, and Semantic Scholar agree on its details. A single lookup, or a DOI that merely resolves, is not enough.

Forensic integrity architecture. The design by which every claim in an ARCHON manuscript remains traceable to a fixed record of the sources, from first draft to delivery.

Frozen evidence ledger. The complete, fixed record of what the sources say, assembled before any writing begins and closed to further change. Every sentence is written from this record, and every claim can be traced back into it. The method ARCHON uses to build and freeze the ledger is proprietary.

Hard-stop rule. A single unresolved critical finding blocks a manuscript from acceptance, regardless of how well it reads.

Journalology. The study and practice of scholarly publishing itself: how manuscripts are prepared, evaluated, and accepted, independent of subject matter.

Phantom citation. A reference that looks legitimate but cannot be confirmed, such as plausible authors and a well-formed DOI attached to a paper that does not exist or does not match. ARCHON halts when it encounters one.

PRISMA 2020. A widely used international standard for reporting systematic reviews and meta-analyses. ARCHON’s systematic-review pipeline follows it, including a search that can be reproduced.

Provenance. The trail tying a claim in the manuscript to the exact place in the source it came from: which document, which page, which value.

Q1–Q2 journals (WoS / Scopus). Journals ranked in the top two quartiles of their field in Web of Science or Scopus. ARCHON targets manuscripts suitable for submission to them.

Retrieval-augmented generation (RAG). A technique that grounds a language model by fetching the stored text fragments most similar to a query and placing them in the model’s context. Effective for search and question answering. ARCHON does not use it to ground manuscripts, because it returns an approximate subset rather than the complete record.

© 2026 Ian Scott Kinney, Marcos Guzmán, and Tomáš Zvolenský. All rights reserved.