

Secondbrain

Architecture for Cumulative Expertise

Maintained knowledge, an explicit ingestion skill,
and evidence-grounded connections between topics

A technical companion to *From Daily Reading to Cumulative Expertise*. This paper explains how a personal reading archive becomes a maintained, navigable body of knowledge, and where human learning still requires its own work and evidence.

Francesco Di Costanzo

8 October 2026 / Version 1.0

Implementation snapshot: b6a550a18ddc

Descriptive architecture and proposed evaluation. No measured claim of improved human expertise.

1 Scope, evidence and the design claim

Secondbrain implements an approach suggested by Andrej Karpathy’s *LLM Wiki*: preserve incoming material and ask an agent to maintain an interlinked interpretation of it over time [1]. The system’s useful unit is an explanation with a history and a route back to evidence. It supports repeated reading across related subjects, including developments whose meaning changes as later reporting arrives.

The design claim is conditional. A maintained explanation can reduce the work of reconstructing context, expose differences among sources and suggest useful comparisons. A reader develops expertise only if these opportunities support better explanation, application and correction. Research on organised knowledge, retrieval practice and the distinction between learning and performance motivates that hypothesis; it does not demonstrate an effect for this implementation [2, 3, 4].

This paper examines the author’s local Secondbrain repository on 8 October 2026 at commit `b6a550a18ddc9609f63e736db3bfcfce7b5e3d54`. Its source archive is not supplied as a publicly accessible research dataset. The implementation locator in Appendix A identifies the files inspected. Schemas, diagrams and pseudocode below make the central mechanisms understandable without access to that checkout.

Three evidence levels are kept separate. **Implemented** denotes behaviour inspected in executable code. **Contract** denotes instructions in the wiki skill and its references; compliance still depends on execution and review. **Proposed** denotes evaluation or extensions that are not claimed to exist in the app.

Snapshot inventory	Files	Interpretation
News raw briefs	11,451	Captured current-reading material
Library raw briefs	92	Saved material from the Library channel
Compiled source pages	11,543	Individual source interpretations
Concept pages	1,197	Explanations of recurring subjects
Entity pages	1,149	Accounts of named subjects

These are Markdown file counts, not audited measures of completeness or quality. Equality between the total raw count and source-page count does not prove a one-to-one mapping. The paper’s preparation included 18 passing parser, search and Pathfinder tests. It did not run a whole-corpus editorial review or a human learning trial.

The main architectural consequence is a separation of responsibilities: generative work maintains the archive; deterministic application code exposes its contents; the reader evaluates and uses the explanation. Each boundary introduces its own possible errors.

2 Architecture and responsibilities

The repository contains two intake collections, one shared compiled wiki, an instruction package governing compilation, and a Next.js reading application. Markdown holds the durable content. Generated JSON indexes make that content practical to browse and search. The graph is derived from authored relationships rather than maintained in a separate graph database. [L1–L4]




Figure 1. The content path. Capturing a source does not itself authorise compilation. Rebuilding the index makes saved editorial changes available to the indexed application.

The News collection and Library collection differ by intake, not by evidential rank. A podcast saved through Readwise remains an attributed argument; a PDF can contain forecasts or marketing. Library does not mean peer reviewed, timeless or true. The preserved raw may be a source brief instead of a complete article, so preservation cannot recover details absent from the capture.

The wiki skill is a package of operating instructions, templates and review references. It is not a trained model, an autonomous background process or a guarantee enforced entirely by code. An agent executes it when the relevant operation is requested. The skill recognises INGEST, QUERY and LINT, with distinct scopes and constraints.

The application’s Ask endpoint is also distinct from an agent running the skill’s QUERY mode. The endpoint assembles stored passages without a model call. Confusing those two surfaces would misstate the architecture and make evaluation difficult.

3 Data model, identity and provenance

The model separates *what was captured*, *what one source says* and *what a topic currently means*. A raw brief belongs to News or Library. A compiled source page represents exactly one raw under the skill's canonical contract. Concepts and entities combine material from multiple sources. [L2,L4,L5]

Stable identity should survive editorial title changes. News source IDs derive from the raw filename stem with underscores normalised to hyphens. Library source IDs use the captured Readwise identifier. Concept and entity IDs use their respective type and canonical slug. A slug must not identify both a concept and an entity: wikilinks resolve by their final path segment.

```
# Illustrative Library source metadata, not a live record
id: "source:library:<readwise-id>"
title: "A source-specific explanatory title"
collection: "library"
origin: "readwise"
readwise_id: "<readwise-id>"
format: "pdf"
sources:
  - "library/raw/<canonical-file>.md"
source_url: "https://publisher.example/document"
published_at: "2026-08-11"
captured_at: "2026-08-16"
last_updated: "2026-09-20"
provenance_status: "resolved"
related:
  - "[[ai-infrastructure]]"
```

A resolved source has exactly one repository-relative raw path. An unresolved record can have an empty source list. Resolved describes the provenance mapping; it does not certify the truth of every claim, the completeness of a capture or the independence of supporting reports. Unknown metadata must remain unknown when the evidence cannot supply it.

Aliases preserve former source-page paths when a duplicate is absorbed or a page is renamed. They are compatibility metadata, not new evidence and not alternative canonical identities. Treating aliases as separate documents would inflate apparent coverage.

Source pages use six stable sections: Summary, Key claims, Concepts and entities, Useful details, Contradictions or caveats, and Links. Their common structure helps the parser and reader locate different kinds of content, while the writing contract requires source-specific explanations.

W3C PROV provides a broader conceptual vocabulary for derivation and attribution; SKOS provides one for concepts and semantic relationships [5, 6]. Secondbrain does not implement these standards. Its smaller design retains useful distinctions while accepting weaker formal semantics: a generic related link does not encode a typed causal predicate or a logically validated proposition.

4 The ingestion skill as an editorial protocol

The skill defines a repeatable mutation process around preserved evidence. Capture and compilation remain separate actions. The first creates or receives raw briefs; the second is an explicit request to incorporate eligible material into the wiki. This makes a queue visible without treating every captured item as already understood. [L2,L3]

An INGEST run begins by auditing coverage and reading the wiki index and log. Coverage is established from parsed source paths, not guessed from similar filenames. The agent reads each selected raw in full, including metadata and caveats. Reconciling a disagreement with older coverage requires reading the relevant compiled/raw pair as well.

For each raw, the agent creates or updates the source page and then the directly affected hubs. It must determine what changed in the explanation, whether an existing concept already fits, and which relationships are supported. Adding a new page for every new phrase would fragment the account; linking every shared keyword would dilute it.

The skill prescribes bounded topical batches for larger queues, commonly eight to fifteen raw items. The complete manifest must partition the selected raw set without overlaps or omissions. Where workers are used under that workflow, their ownership is bounded to assigned source and hub files. Shared outputs such as INDEX, LOG and the final briefing are integrated centrally. This paper describes that contract; no ingestion or worker execution was performed to prepare it.

```
# Contract-level pseudocode, not executable application code
selected = audit_unrepresented_raws(scope)
manifest = partition_into_topical_batches(selected)
assert disjoint_union(manifest) == selected

for batch in manifest:
    read_complete_raws_and_relevant_prior_pairs(batch)
    write_one_source_per_raw(batch)
    revise_directly_affected_hubs(batch)
    preserve_identity_dates_and_disagreement(batch)
    review_complete_saved_changed_pages(batch)

integrate_index_log_and_reader_outputs()
validate_structure()
require_empty_news_and_library_queue_for_completion()
rebuild_research_index()
```

This is an editorial protocol with executable checks around it. The checks can reject malformed identities or missing paths, but they cannot determine whether a synthesis has chosen the right causal explanation. Human-readable instructions carry part of the system's reliability burden, which is why documenting them matters as much as documenting the UI.

5 Evidence fidelity and correction

The evidence-fidelity reference supplies rules that a structural schema cannot express. It distinguishes author, publisher, quoted speaker and organisation; separates a proposal from execution; and requires numbers to retain their unit, period, population and comparison base. A forecast does not become an observed result when its target date passes. [L3]

This discipline applies independently to titles and captions. A heading can turn negotiations into completion even when the paragraph below remains qualified. The contract therefore requires checking what a reader would believe from the heading alone. It also requires corrections to reach the directly affected representations that repeat the mistake.

A bounded correction may involve a source page, the relevant hub passage, Current understanding, an INDEX label and a current draft briefing. It should not spread into pages that never made the claim or rewrite a historical briefing with hindsight. Wider restructuring and changes to approved or posted publishing outputs are separate work.

The distinction between repeated reporting and independent corroboration is especially important. Several captured articles may trace back to one announcement. Counting those records gives an archive size, not the number of independent observations. Likewise, two different estimates may concern different time horizons and therefore need no reconciliation. Where figures conflict on the same basis, an attributed unresolved disagreement is more accurate than an invented explanation.

Limited captures receive honest treatment. Title-only or unrelated material can justify a provenance-only record, but it cannot support a substantive inference. This restriction is principally an editorial rule. A resolved raw mapping alone does not automatically identify every weak capture to downstream retrieval code.

After the final edit, the skill requires a full read of saved changed artifacts, including headings and caveats, and a record of reviewed versions or hashes. Subsequent edits invalidate the earlier review for the changed material. Completion cannot be inferred from an agent's assertion that it reviewed everything.

Research on citation quality and factual evaluation supports separating these checks. ALCE evaluates citation support as a distinct property; FActScore evaluates supported atomic facts rather than a single impression of correctness [7, 8]. Those are useful evaluation concepts, not validators currently embedded in this wiki.

Failure boundary. The protocol reduces opportunities for ungrounded synthesis, but no measured corpus-wide faithfulness rate is reported here. A model can follow the file format while misunderstanding the evidence.

6 Time, current understanding and history

A durable knowledge system needs to preserve the difference between an earlier account and a current one. Secondbrain's contract does this through dated hub updates alongside a separate Current understanding section. The former preserves the development of the account; the latter explains the latest supported interpretation and states its evidence horizon. [L2,L3]

Several clocks can coexist. An event can occur before publication. A reader can save it later, a capture can be created later still, and ingestion can occur on another date. The date in a canonical update heading is the ingestion date. The prose must provide the event or publication context when that controls the meaning.

```
## Current understanding

As of 2026-09-27, the saved evidence supports ...
[Explain the present account and remaining uncertainty.]

## Development

### 2026-08-16 update [library] (batch-id): Caption
[Explain what this source added, with its own dates,
attribution, qualifications and source wikilinks.]
```

This is an illustrative structure. The exact enclosing section names can vary across hubs; the dated update grammar is the important parser contract. A historical prediction remains historical even after a later update changes the view. Correcting a demonstrable editorial error differs from replacing a faithful earlier forecast with hindsight.

The same distinction affects Pathfinder. It derives a passage's date and collection from a matching update heading and resets that context at a level-two heading. Its date filter therefore concerns the authored update context. It is not a general temporal database answering which relationships were true on an event date, and it does not reconstruct every prior version of a page.

Freshness checks are advisory. The health code flags hubs with at least five dated updates when Current understanding is missing or its date trails the latest update by more than thirty days. It can also suggest possible overlaps between similarly titled hubs with shared sources. These signals identify review candidates; they do not prove obsolescence or justify an automatic merge. [L8]

There is a further coverage limit. An evidence horizon should describe the evidence the account uses, not imply exhaustive awareness of the outside world. A page can be current relative to saved sources while missing an important publication. The intake policy and the selection of reading remain part of the system's epistemic boundary.

7 Derived indexes and deterministic Ask

The research index is a serving representation of Markdown. In the inspected v3 design, metadata lives in a generated JSON index, bodies live in separate buckets, and search text is stored in shards. Concept and entity bodies share a hubs bucket; other kinds are distributed using the first hexadecimal digit of a hash of the item ID. Body filenames include a content hash. [L4,L6]

The builder writes body artifacts before replacing index metadata by atomic rename. This reduces exposure to metadata pointing at not-yet-written bodies. It is not a database transaction over the whole repository or deployment. Source edits remain invisible to a stale index until a rebuild. Generated data can be reconstructed from the durable Markdown.

Search uses deterministic word matching. It normalises text with Unicode NFKC, lowercases it, removes possessive endings and filters a stopword set from the query. For a multiword query, a candidate needs at least $\lceil 0.65|Q| \rceil$ matching query words in its search text. A weighted score rewards coverage, title matches, an exact title phrase or named subject, and an exact slug phrase. Date and title break ties. It is neither embedding retrieval nor BM25. [L7]

Ask searches compiled non-raw items and excludes unresolved provenance. An optional focus item can be mapped from a raw to its compiled source. The response retains at most five candidate citations, deduplicated by title, and loads passages from the first three contributors. Title deduplication can suppress distinct records with identical titles.

Passage candidates come from introductions, Current understanding, storylines, facts and why-it-matters content. Token overlap and modest section, recency and focus boosts determine selection. Up to four distinct passages longer than forty characters are returned if they have positive scores. Otherwise the endpoint reports insufficient matching material. [L4,L7]

No language model is called by this endpoint, and its response is not automatically persisted as new wiki knowledge. Attribution identifies the page contributing a passage, but it does not independently verify the passage's underlying claims. A lexical search can miss a synonymous expression or rank a recent but incomplete account highly.

RAG and GraphRAG offer different designs: generation conditioned on retrieved material, or graph/community summaries used to answer broader questions [9, 10]. Secondbrain's current Ask should be evaluated as attributed extractive retrieval. Calling it conversational RAG would obscure both its lower query-time complexity and its limitations.

8 Building the topic graph

Pathfinder's nodes are compiled concept and entity records whose provenance status is not unresolved. Source pages contribute evidence but are not intermediate topic nodes. Raw captures are not graph nodes in this view. Each eligible hub has a stable ID and a slug used to resolve authored references. [L5,L9]

An edge exists when a hub's related metadata or body wikilink names another eligible hub. The pair is stored in canonical order, and the code remembers which endpoint authored the link. Traversal is undirected. A reciprocal link can improve its ranking, but two editors or two directions do not create two independent pieces of evidence.




Figure 2. An illustrative evidence distinction, not an exported live route. A shared mention can accompany an existing edge, but cannot create one.

For each source page, the builder collects body-linked hub IDs. Intersecting the source sets for two hubs produces shared supporting mentions. The user-facing edge object keeps these separate from `citedEvidence`, which comes from source links in the displayed explanatory passages. Shared mention is context, not proof that a source supports the relationship being asserted.

Passage extraction is deliberately simple and has a consequential boundary. It scans physical body lines containing a wikilink to the target hub, skips headings, rejects very short text and generic relationship-list openings, and collects source links on that same line. A citation on the following line will not be attached by this extraction step.

Thus line layout can affect what evidence appears on an edge. That is a reason to test and document the representation, not to infer that absent extracted evidence means no evidence exists anywhere in the note. The architecture offers readable, inspectable relationships without supplying a formal entailment engine.

9 Path ranking, filtering and limits

Pathfinder searches for short, useful routes through authored relationships. It uses a minimum-cost priority queue over simple paths, excluding revisited nodes. Initial exploration allows four hops; an explicitly branched route can use six. Search stops after sixteen completed candidate paths or twenty thousand expansions, then retains up to three sufficiently different alternatives. A breadth-first shortest-route fallback preserves a feasible shortest route if the weighted search misses it. [L9]

For an edge, let s be the number of shared source mentions capped at five, $r = 3$ for a reciprocal relationship and zero otherwise. Let $e = -2$ if any explanatory passage has an explicit source citation, $e = -1$ if a passage exists without one, and $e = 3$ otherwise. If the destination has degree d , define

$$h(d) = \begin{cases} 0, & d \leq 24, \\ \min(4, 2 \log_2(d/24)), & d > 24. \end{cases}$$

The step cost is

$$c = \max(2, 10 - s - r + e) + h(d).$$

The final destination receives no degree penalty. Path cost is the sum of its step costs; the exposed score is $1000/(1+C)$, rounded to two decimals. It is a ranking heuristic, not a calibrated probability, confidence level or measure of causal strength.

Alternative routes are suppressed when overlap among their internal nodes exceeds 75%, using the larger internal-node count as the denominator. A shortest mode prioritises hop count. Branch suggestions must still reach the destination within the six-hop bound, avoid the existing path and pass prefix validation; at most eight suggestions are retained.

Each public edge displays up to three passages, ordered primarily by date and then explicit citation count. Its qualitative strength label is based on shared-mention count: at least five yields *very-strong*, at least two *strong*, otherwise *linked*. Those names should be interpreted only in that implementation-specific sense. They do not assess source independence.

Collection and date filters require at least one matching authored passage. Shared mentions alone cannot keep a filtered edge eligible. Undated passages cannot satisfy a date-bound filter, even if relevant evidence exists elsewhere on the page.

The code also warns when adjoining explanations share little vocabulary, or when a broad intermediate hub joins weakly overlapping contexts. This is a useful prompt to read both passages, but lexical overlap is an imperfect proxy for semantic continuity. Neither the warning nor a high path score determines whether a multi-step argument is valid.

10 Worked trace: AI infrastructure and credit

The inspected archive contains a Library brief titled *How AI Debt Is Reshaping Credit Markets*, attributed to Goldman Sachs Exchanges, and another titled *2026 Midyear Credit Outlook*, attributed to Apollo. Their captured metadata records publication dates of 5 and 11 August 2026 respectively, saving on 14 August and capture on 16 August. This section describes saved records, not an independent re-verification of the underlying market data. [L11]

Their compiled source pages use the titles *Goldman Locates AI Debt Risk in Absorption and Duration* and *Apollo Separates AI Credit Quality From Market Capacity*. The source identities remain the captured Readwise IDs. Retitling therefore changes the reader's entry point without breaking the raw-to-source mapping.

The Goldman account distinguishes borrowers' financial capacity from investors' appetite for repeated, long-duration issuance. The Apollo account additionally examines exposure to one AI spending cycle across different investment categories. Both discuss alternative financing channels, but their estimates use different periods and scopes. Neither account licenses substituting an annual borrowing estimate for a multi-year investment forecast.

The AI infrastructure concept page combines these accounts in a 16 August Library update. It explicitly preserves the differing bases of their estimates and relates financing channels to common underlying demand. It also identifies the institutions' participation in the markets discussed. This is an example of a maintained explanation linking source-specific claims while retaining attribution and caveats. [L12]

The next reading task can now be more precise. Given a new financing announcement, a reader can distinguish the identity of the capital provider, contractual loss protection, the time horizon and dependence on future computing demand. Changing one of these does not necessarily change the others. This question is an analytical use of the archive, rather than a market recommendation.

The trace also shows why graph semantics matter. Both source pages link to AI infrastructure and private credit, so they can contribute shared mentions for an already-authored hub relationship. To count as explicitly cited edge evidence, however, a hub passage must link the destination topic and the source on the same physical line under the current extractor.

The historical update examined here is not itself asserted to be a complete Pathfinder route. It demonstrates source reconciliation. A live route would require checking the actual eligible hub links, extracted passages, filters and ranking at the selected snapshot.

Learning use, proposed. Before reading the synthesis, ask the reader to explain how new funding can expand capacity while leaving a common demand exposure. Then compare the explanation with the two source accounts and apply it to a different financing structure. The app does not currently score this exercise.

11 Quality gates and known failure modes

Structural validation and editorial review answer different questions. The maintenance script checks canonical metadata and paths, duplicate source mappings, aliases, Library IDs, update-heading rules and wikilink resolution. Normal validation allows queued raw briefs. Complete validation additionally requires the entire News and Library queue to be represented by resolved source pages. [L2,L10]

An empty queue is therefore a completion condition, not a quality score. A weak capture can have a valid provenance record, and a syntactically correct hub can still combine incompatible claims. The evidence-fidelity review must examine the material the structure cannot judge.

Failure	Existing response	Remaining limit
Duplicate identity	Canonical IDs, parsed paths, validator	Equivalent claims can remain in distinct sources
Missing capture	Unresolved provenance state	Resolution does not establish capture quality
Stale synthesis	Dated Current understanding; health hints	Dates cannot detect omitted outside evidence
Misleading connection	Authored-edge requirement; passages	Authorship is not entailment
Repeated reporting	Editorial independence rule	No automatic evidence-family clustering
Query mismatch	Attributed passages; insufficient-match response	Lexical synonyms and scope can still fail
Broken correction	Propagation and final-read contract	Review completeness still needs evidence

The completion gate has an operational trade-off. Because new raws arriving during an ingest count towards the full queue, continuous intake can keep the gate open. A deliberately frozen intake window could make a bounded run easier to complete, but that is an operating choice rather than a behaviour this paper claims the application enforces.

Briefings and Brainwaves are downstream publishing artifacts. The skill retains review boundaries and distinguishes draft work from completed or approved outputs. Book recommendations are not part of the wiki's evidence graph. These boundaries prevent suggested reading or a draft publishing angle from silently becoming supporting evidence for a topic.

The read-only tests executed for this paper covered parser, search and Pathfinder behaviour: eighteen tests passed with no failures. They include authored-edge requirements, evidence distinctions, ranking, route bounds, branching, fallback and filtering. They do not establish corpus-wide accuracy, security, user learning, retrieval relevance on real questions or production performance.

The ordinary validation command writes a generated validation report. It was not invoked against the reference repository during this read-only architecture study. Consequently, this paper reports neither a newly verified empty queue nor a fresh full-corpus structural acceptance result.

12 Evaluation: maintained knowledge and human expertise

Evaluation should separate the quality of the archive from the capability of its user. Otherwise a faster answer, a longer reading streak or a larger graph can be mistaken for learning. The distinction follows research on learning versus performance and on the gap between access to explanations and internal understanding [4, 11].

Proposed archive evaluation. Freeze a corpus and its index, then sample sources, hubs and relationships across both collections. Have reviewers compare individual claims with preserved raw material and, where authorised and available, original publications. Record source fidelity, unsupported inference, temporal qualification, attribution, missing caveats and whether citations support the precise relationship. Reviewers should not know which system produced a candidate explanation.

For retrieval, prepare questions with answerable evidence and genuine no-answer cases. Measure relevant passage recall, misleading inclusions, source attribution and abstention. Compare lexical retrieval with an embedding retriever or RAG baseline using the same permitted evidence. For Pathfinder, judge each edge independently before judging the route. Measure whether the route helps answer a question, rather than rewarding the existence of a path.

Proposed human evaluation. Compare ordinary reading and search, maintained-wiki access, and wiki access paired with explanation/retrieval exercises. Equalise source material and time budgets; measure prior knowledge. Use random allocation where feasible, a prespecified delayed assessment and unfamiliar transfer cases. Include an unaided phase and a supported-work phase so benefits from external access remain distinguishable from retained understanding.

Score factual accuracy, mechanism quality, calibration, recognition of misleading analogies and correction after contrary evidence. Analyse task and participant variation rather than reporting only one average. Select sample size from the chosen design and expected variance; this paper supplies no invented power calculation or effect size.

The learning intervention should make the reader generate explanations and compare cases. Work on retrieval practice, self-explanation and analogical encoding motivates those actions [3, 12, 13]. An automatically generated map and a reader constructing a relationship are different interventions.

AI-learning studies do not justify a universal forecast. Bastani and colleagues found harm from unguarded assistance in one mathematics setting; Kestin and colleagues found benefits from a designed physics tutor; Tutor CoPilot supported human tutors [14, 15, 16]. Their results support evaluating the instructional design and task. None supplies an effect estimate for Secondbrain.

Maintenance cost belongs in the outcome set: capture effort, agent compute, editorial time, correction burden and time spent finding the relevant account. A system that improves fidelity but takes disproportionate upkeep may still be the wrong choice for a narrow, stable subject.

13 Operating trade-offs and extension points

Markdown keeps the intellectual record inspectable through ordinary tools and version control. A generated index lets the reading interface evolve independently from the source format. Deterministic Ask and Pathfinder make many retrieval decisions reproducible at a fixed snapshot. These properties are useful for an individual maintained archive; they do not imply unlimited scale or zero maintenance.

The architecture invests effort before a question arrives. That can pay off for recurring topics with substantial overlap, but it also creates work on pages that may never be read again. Query-time RAG makes a different trade-off by assembling material when demanded. GraphRAG adds another option through entity graphs and community summaries [9, 10]. The choice should follow the reader's questions and correction requirements.

A small archive might need only a clear index and disciplined reading. At this snapshot's scale, broad hubs can accumulate lengthy history, and generic relationships can become less discriminating. Health hints and path penalties make this visible but cannot decide how a topic should be split. Editorial boundaries remain necessary.

Several extensions are plausible but **not implemented claims**. A typed relationship layer could distinguish causation, financing, ownership and analogy. Evidence families could group articles derived from one primary announcement. A paragraph-aware passage parser could reduce dependence on physical line breaks. A temporal model could separate ingestion, publication and event intervals more formally. Each would need migration and regression checks against existing content.

A learning layer could store a reader's explanation before revealing the wiki, schedule delayed retrieval and preserve responses to unfamiliar cases. Those responses should remain distinguishable from source-grounded wiki claims. Personal reflections or practice errors should not enter the evidence graph as factual support.

NIST's generative-AI profile recommends verifying source grounding and avoiding broad extrapolation from narrow evaluations [17]. That is an appropriate standard for extending this system. A new feature should have a stated failure mode, an observable acceptance test and a way to inspect the evidence behind its output.

The practical next step is a bounded comparison on a recurring subject. Freeze the available material, record the reader's initial explanation, expose the maintained account and assess a new case after a delay. Keep both answers and the source snapshot. That would begin to test whether the maintenance architecture supports cumulative expertise, rather than only a better organised archive.

A Implementation locator and reproducibility

The following paths are relative to the inspected Secondbrain checkout at the full commit recorded in Section 1. They are locators for the author’s implementation evidence, not public URLs or a claim that the repository is an open research artifact.

ID	File	Evidence used
L1	README.md	Collection boundaries, application and serving architecture
L2	knowledge/skill/SKILL.md	Canonical schemas, modes, ingestion and completion contract
L3	knowledge/skill/references/evidence-fidelity.md; human-writing.md in the same directory	Evidence fidelity, chronology, correction and saved-state review
L4	lib/research.ts	Index creation/loading, content derivation and Ask assembly
L5	lib/research-types.ts	Item types, passage and graph response structures
L6	scripts/build-research-index.ts	Generated index entry point
L7	lib/search.ts; app/api/research/ask/route.ts	Token ranking and HTTP Ask behaviour
L8	lib/knowledge-health.ts	Advisory freshness and overlap heuristics
L9	lib/pathfinder.ts	Graph construction, evidence extraction, route scoring and filtering
L10	scripts/knowledge-maintenance.ts	Structural and complete-validation distinction; report writes

The focused verification command was:

```
node --import tsx --test lib/pathfinder.test.ts \
  lib/search.test.ts lib/parser.test.ts
```

Result: 18 tests passed, 0 failed. The command exercises fixtures and code behaviour, not the complete live knowledge corpus.

The inventory counts Markdown files directly under `news/raw`, `library/raw`, and the wiki’s `sources`, `concepts` and `entities` directories. No ingestion, correction, index rebuild or validation-report write was performed in the reference repository for this paper.

A reader reproducing the architecture in another system should retain stable identities, explicit raw mappings and the distinction between authored relationships and supporting mentions. The particular score weights and hop limits are implementation choices, not requirements derived from the learning literature.

B Worked-example locators and interpretation

L11: raw and source pairs. The following exact paths identify the evidence examined for the financing example. The source titles differ from the captured titles by design; their stable IDs and raw mappings preserve continuity.

```
library/raw/
  2026-08-14_how-ai-debt-is-reshaping-credit-markets_p5rbqzjphp.md
knowledge/wiki/sources/
  2026-08-14-how-ai-debt-is-reshaping-credit-markets-p5rbqzjphp.md

library/raw/
  2026-08-14_2026-midyear-credit-outlook_a4raqg81tr.md
knowledge/wiki/sources/
  2026-08-14-2026-midyear-credit-outlook-a4raqg81tr.md
```

The stable Library IDs are:

```
source:library:01m006pvwenx71b1p5rbqzjphp
source:library:01m00e7xsfmsfz2ja4raqg81tr
```

L12: maintained topic account. The relevant hub is `knowledge/wiki/concepts/ai-infrastructure.md`. The inspected dated block begins:

```
### 2026-08-16 update [library] (b24-library):
Private finance expands capacity without breaking
the common AI factor
```

The heading is line-wrapped here for print. In the actual Markdown it is a single heading line. The block preserves the differing bases of the two accounts and explicitly links their compiled source pages.

The paper does not redistribute the raw collection or claim to verify the original publications' numerical forecasts. Its use of the pair is narrower: demonstrate how a maintained account can retain source identity, compare different scopes, expose commercial interests and formulate a question that transfers to another case.

The implementation's graph should be read with the same restraint. Evidence can name a source attached to a relationship, a shared mention or the contents of a captured brief. These are not interchangeable levels of proof. A technically reproducible path can still be an intellectually weak explanation.

Publication and regeneration

The companion article supplies the PDF through the site's existing download metadata. The editable source is `papers/2026/secondbrain-cumulative-expertise/main.tex`; the public artifact is `public/downloads/secondbrain/secondbrain-cumulative-expertise.pdf`. Build the source with Tectonic, using a separate temporary output directory. The document uses standard LaTeX packages and Latin Modern fonts.

References

- [1] Andrej Karpathy. *LLM Wiki*. 2026. <https://gist.github.com/karpathy/442a6bf555914893e9891c11519de94f>
- [2] National Academies of Sciences, Engineering, and Medicine. *How People Learn II: Learners, Contexts, and Cultures*, Chapter 5. 2018. <https://www.nationalacademies.org/read/24783/chapter/7>
- [3] Jeffrey D. Karpicke and Janell R. Blunt. *Retrieval practice produces more learning than elaborative studying with concept mapping*. Science, 2011. <https://pubmed.ncbi.nlm.nih.gov/21252317/>
- [4] Nicholas C. Soderstrom and Robert A. Bjork. *Learning Versus Performance: An Integrative Review*. Perspectives on Psychological Science, 2015. <https://doi.org/10.1177/1745691615569000>
- [5] W3C. *PROV-DM: The PROV Data Model*. Recommendation, 2013. <https://www.w3.org/TR/prov-dm/>
- [6] W3C. *SKOS Simple Knowledge Organization System Reference*. Recommendation, 2009. <https://www.w3.org/TR/skos-reference/>
- [7] Tianyu Gao, Howard Yen, Jiatong Yu and Danqi Chen. *Enabling Large Language Models to Generate Text with Citations*. EMNLP, 2023. <https://aclanthology.org/2023.emnlp-main.398/>
- [8] Sewon Min et al. *FActScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation*. EMNLP, 2023. <https://aclanthology.org/2023.emnlp-main.741/>
- [9] Patrick Lewis et al. *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks*. 2020. <https://arxiv.org/abs/2005.11401>
- [10] Darren Edge et al. *From Local to Global: A Graph RAG Approach to Query-Focused Summarization*. 2024; revised 2025. <https://arxiv.org/abs/2404.16130>
- [11] Matthew Fisher, Mariel K. Goddu and Frank C. Keil. *Searching for Explanations: How the Internet Inflates Estimates of Internal Knowledge*. Journal of Experimental Psychology: General, 2015. <https://doi.org/10.1037/xge0000070>
- [12] Michelene T. H. Chi and Ruth Wylie. *The ICAP Framework: Linking Cognitive Engagement to Active Learning Outcomes*. Educational Psychologist, 2014. <https://doi.org/10.1080/00461520.2014.965823>
- [13] Dedre Gentner, Jeffrey Loewenstein and Leigh Thompson. *Learning and Transfer: A General Role for Analogical Encoding*. Journal of Educational Psychology, 2003. <https://groups.psych.northwestern.edu/gentner/papers/GentnerLoewensteinThompson03.pdf>
- [14] Hamsa Bastani et al. *Generative AI Without Guardrails Can Harm Learning: Evidence from High School Mathematics*. PNAS, 2025. Author manuscript: https://hamsabastani.github.io/education_lm.pdf
- [15] Greg Kestin et al. *AI tutoring outperforms in-class active learning: an RCT introducing a novel research-based design in an authentic educational setting*. Scientific Reports, 2025. <https://www.nature.com/articles/s41598-025-97652-6>
- [16] Rose E. Wang et al. *Tutor CoPilot: A Human-AI Approach for Scaling Real-Time Expertise*. 2024; revised 2025. <https://arxiv.org/abs/2410.03017>
- [17] NIST. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST AI 600-1, 2024. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>