

# ARCHORA

*archive + agora*

## Product Design Document

Version 1.2 · Wiki-Native Knowledge Architecture

### Vision

Replace editorial gatekeeping with peer-driven natural selection, amplified by RL and wiki-compiled knowledge.

### Core Engine

Three-layer knowledge architecture: Raw Sources → Personal Wiki → Archora Space, with dual RL loops.

### Target Users

Academic researchers, independent scientists, R&D labs — globally, starting with a focused seed community.

### Key Innovation (v1.2)

Wiki structural similarity as a novel SNS recommendation primitive — topology-based, not engagement-based.

Archora is a research-native social platform that replaces the arbitrary selection of traditional academic publishing with the natural selection pressure of expert peer communities, continuously sharpened by reinforcement learning.

The name combines two foundational concepts:

- Archive — the permanent accumulation of research knowledge, sources, and evolving hypotheses
- Agora — the open marketplace of ideas where research competes, connects, and propagates on merit

#### Core Belief

The next great scientific ideas will not be chosen by editors optimizing for journal prestige. They will emerge from the honest friction of expert peers, the compounding intelligence of reinforcement learning, and the radical openness of a global research commons — mediated through each researcher's personal knowledge wiki.

## 1.1 What Archora Is Not

Archora is not a preprint server (though it integrates with them). It is not a reference manager (though it includes archiving). It is not a social network for general academic networking. It is specifically engineered for the lifecycle of research ideas: from raw collection, through LLM-compiled synthesis wiki, to peer-evaluated public knowledge.

## 1.2 The Central Tension Archora Resolves

Traditional publishing creates a structural mismatch: the people best positioned to evaluate a paper have the least institutional power to surface it, while editors with gatekeeping power are often generalists. Archora inverts this: the weight of endorsement scales with demonstrated domain expertise, not institutional title. The personal wiki layer ensures this endorsement is grounded in a researcher's full, compounded knowledge — not just a single reaction.

## 2 Three-Layer Knowledge Architecture

Archora's architectural core is directly inspired by the LLM Wiki pattern: rather than leaving researchers with a passive raw archive and re-deriving synthesis on every query, Archora has the LLM incrementally build and maintain a persistent personal knowledge wiki between the raw sources and the social surface. This is the foundational change from v1.1.

### Architecture Principle (Karpathy LLM Wiki Pattern)

Raw Sources (immutable) → Personal Wiki (LLM-maintained, persistent markdown) → Schema File (ARCHORA.md, governing structure) → Archora Space (social surface). The wiki is a compounding artifact — cross-references pre-built, contradictions pre-flagged, synthesis already reflecting everything the researcher has read.

### 2.1 Layer 1 — Raw Sources (Research Archive)

The private, immutable source-of-truth layer. Papers, notes, datasets, annotations, web clippings, conference talks, preprints — archived with full provenance. The LLM reads from this layer but never modifies it.

<b>Supported Types</b>	Academic papers (DOI-linked, auto-metadata), personal notes & margin commentary, raw datasets, web clippings, preprints, hypothesis drafts, conference talks.
<b>Privacy</b>	Permanently private by default. Zero content used without explicit researcher action. Immutability is enforced — no LLM writes to this layer.

### 2.2 Layer 2 — Personal Wiki (LLM-Maintained) **NEW**

The compounding intelligence layer. The LLM incrementally builds and maintains a structured, interlinked collection of markdown files that sits between the raw sources and the public surface. Every time a researcher archives new sources, the LLM updates the wiki — not just indexing for later retrieval, but integrating into the existing structure.

#### Wiki Operations

<b>Ingest</b>	Triggered by: archiving 3+ new items, significant new Archora Space content, tracked hypothesis receives citations, or manual request. The LLM reads new sources, writes/updates entity and concept pages, updates index.md, flags contradictions with existing pages, and logs the operation in log.md. A single source may touch 10–15 wiki pages.
<b>Query</b>	Researcher asks questions against the wiki. LLM searches index.md, reads relevant pages, synthesizes with citations. Good answers are filed back into the wiki as new pages — explorations compound just like ingested sources. This is the AI Synthesis channel output.
<b>Lint</b>	Periodic wiki health check: orphaned pages (no inbound links), contradiction density, stale claims superseded by newer sources, missing cross-references, hub density analysis. Lint signals feed Loop 3 (structural health regularization).

## Wiki File Structure

<b>index.md</b>	Content-oriented catalog: every page listed with link, one-line summary, source count. LLM reads index first on every query. Updated on every ingest.
<b>log.md</b>	Append-only chronological record: ingests, queries, lint passes. Parseable with consistent prefix format. Enables wiki evolution tracking as a reputation signal.
<b>Entity/Concept pages</b>	Per-concept and per-entity markdown files with YAML frontmatter (tags, source count, last updated, confidence level). Cross-linked with <code>[[wiki-link]]</code> syntax. Source of the structural similarity signal.

## 2.3 Layer 3 — Schema File (ARCHORA.md) **NEW**

A governance document that tells the LLM how each researcher's wiki is structured, what conventions to follow, and what workflows to apply during ingest, query, and lint. This is the researcher's semi-public epistemology declaration — the document that makes the LLM a disciplined wiki maintainer rather than a generic chatbot.

<b>Private layer</b>	Full schema file: page conventions, domain-specific workflows, researcher-specific ontology choices. Never shared.
<b>Semi-public layer</b>	Researcher can opt to expose schema structure (not content) to Archora's recommendation engine as an epistemological alignment signal. Used in wiki-similarity matching.

## 2.4 Layer 4 — Archora Space (Public)

The global research commons — the social surface of Archora. Publishing to Archora Space is equivalent to exposing a wiki node publicly: the researcher selects a wiki page, review, and promotes it. Once public, the recommendation engine routes it to relevant researchers based on the multi-dimensional recommendation engine (now including wiki-similarity).

- AI analyzes content on publication and routes to relevant researchers via weighted recommendation
- Community engagement signals feed the reputation system
- Content can be cited by researchers and AI systems — both tracked separately
- Versioning supported: full diff history between published wiki node versions

## 3 Reinforcement Learning Architecture

Three nested RL loops operate at different timescales and with different reward structures. Loop 1 and Loop 2 are carried over from v1.1, now with wiki-structure signals incorporated. Loop 3 (structural health) is formalized in v1.2.

### 3.1 Loop 1 — Personal Wiki Optimization

This loop runs at the individual researcher level. Its goal is to make AI-generated wiki content — synthesis pages, hypothesis drafts, gap analyses — increasingly aligned with each researcher's intellectual trajectory, quality standards, and preferred structure.

#### Signal Sources

- Explicit disposition (Publish / Save / Discard) — primary signal, high weight
- Time-to-action on AI outputs — fast action = stronger signal
- Edit behavior before publishing — heavy editing suggests format or framing mismatch
- Wiki adoption: does the synthesis become a stable wiki page, or get discarded after initial save?
- Cross-link density: synthesis that generates many outbound links to existing wiki pages = high relevance
- Passive engagement — opened but not acted on within 48h = mild negative

#### What the Loop Learns

- Topic depth and adjacency preferences
- Preferred output format (hypothesis vs. gap analysis vs. synthesis report)
- Level of speculation acceptable to the researcher
- Which wiki page types (entity / concept / comparison / timeline) researcher finds valuable
- Skill mix: which AI reasoning skills produce accepted wiki content

#### Skill Randomization

AI applies a portfolio of research skills (analogical reasoning, systematic gap finding, cross-domain synthesis, Socratic questioning, etc.). Skills are randomly sampled initially, then weighted toward those producing accepted wiki pages per researcher. New skills injected periodically to prevent local optima.

### 3.2 Loop 2 — Community Natural Selection

This loop runs at the community level. It uses collective researcher behavior as a natural selection mechanism, elevating wiki nodes that survive peer scrutiny and demoting those that fail to propagate.

#### Signal Sources

- Peer citations — weighted by citer reputation in relevant domain
- Downstream hypotheses — how many new ideas are generated citing this work
- Expert endorsements — explicit positive signals from high-reputation peers
- Archive adoption — how many researchers save a public post to their own archives and wikis
- AI citation count — how often AI synthesis engines cite this work (starts at 5% weight)
- Cross-domain propagation — work bridging domain boundaries receives novelty bonus

### AI Citation Weighting Strategy

AI citation count starts at low weight (5%) to prevent a closed AI-reinforces-AI loop. Weight increases gradually as the signal proves predictive of long-term human citation impact. Quarterly validation check: if AI citation count and human citation impact diverge significantly, AI weight is automatically reduced.

## 3.3 Loop 3 — Structural Wiki Health (Regularization) **NEW**

This loop acts as a regularizer on Loop 1, preventing the personal wiki from degrading into an incoherent or orphaned structure. Health signals are computed during lint passes and fed back as context to the synthesis engine.

Health Signal	Role in Regularization
Orphan page ratio	Pages with no inbound links indicate under-synthesized knowledge islands. High orphan ratio triggers synthesis suggestions to reconnect.
Hub density	Pages with many inbound links are intellectual hubs. Hub coverage over time is a proxy for synthesis quality and topic coherence.
Contradiction density	Flagged contradictions between pages that have not been resolved. Acts as a notification primitive — surfaces open epistemic questions.
Revision rate	Pages with many revisions indicate active live topics. Signals Loop 1 to prioritize synthesis in those areas.
Source-to-page ratio	How many sources have been compiled into wiki pages vs. left as raw. Low ratio indicates Archive is richer than the wiki — triggers ingest pressure.

**Design Hypothesis**

SNS recommendation in research contexts should be driven by structural similarity between compiled knowledge wikis, not by content consumption patterns. Two researchers whose wikis are topologically adjacent — overlapping concepts, complementary gaps, shared methodological hubs — are better collaborators and content partners than two researchers who happened to read the same papers.

This section defines the experimental recommendation primitive introduced in v1.2. It replaces the pure semantic similarity approach of v1.1 with a multi-dimensional wiki topology comparison. This is the primary technical experiment for Phase 1 validation.

**4.1 Why Wiki-Similarity Differs From Prior Approaches**

Approach	Signal Used	Limitation
Traditional SNS	Likes, follows, time-on-page	Optimizes for dopamine, not epistemic value
Archora v1.1	Semantic similarity of published content	Misses private knowledge topology; only surfaces public posture
<b>Archora v1.2 (Experiment)</b>	<b>Wiki topology: concept overlap, gap complementarity, schema alignment, evolution trajectory</b>	Requires user opt-in to expose anonymized wiki structure; privacy-first by

## 4.2 Wiki Similarity Signals

Wiki similarity is computed from anonymized structural features, never raw content. Researchers must opt in to wiki-similarity matching. Five distinct signals compose the similarity score:

Signal	Definition & Recommendation Implication
<b>Concept Overlap</b>	How many concept pages exist in both wikis (matched by embedding similarity, not exact name). High overlap → researchers are working in the same intellectual territory. Strong signal for content recommendation.
<b>Gap Complementarity</b>	User A has dense coverage where User B has orphan/missing pages, and vice versa. High complementarity → each researcher can fill the other's blind spots. Strong signal for collaboration matching.
<b>Schema Alignment</b>	Similarity between the semi-public portions of each researcher's ARCHORA.md schema file — how they categorize knowledge, what ontologies they use, what workflows they prefer. High alignment → epistemological compatibility.
<b>Evolution Trajectory</b>	Which concept clusters are growing in each researcher's wiki over time (tracked via log.md revision rates). Parallel growth trajectories → researchers are converging on the same frontier. Strong signal for early-stage collaboration.
<b>Contradiction Adjacency</b>	Contradiction pages in User A's wiki that overlap with confident claim pages in User B's wiki. Indicates productive epistemic friction. Low similarity, but high value — surfaces potential for critical peer review.

## 4.3 Recommendation Modes

<b>Content Feed</b>	Archora Space posts are routed based on weighted combination: concept overlap (50%) + reputation in domain (30%) + novelty factor for bridging unseen cluster combinations (20%). Wiki-similarity replaces the v1.1 semantic similarity of published content alone.
<b>Collaboration Matching</b>	When researcher opens collaboration status, gap complementarity score becomes primary matching signal. This produces meaningfully different matches than concept overlap alone — finding researchers who know what you don't, not just researchers who know what you know.
<b>Anti-Filter-Bubble</b>	A fixed 10% of recommendations are drawn from wikis with low concept overlap but high contradiction adjacency — deliberately surfacing researchers whose compiled knowledge challenges the researcher's existing wiki structure.

## 4.4 Privacy Architecture for Wiki Similarity

### Privacy Principle

Wiki similarity matching is always opt-in. The system never accesses raw wiki content — only anonymized structural features (page count per cluster, edge density, schema category labels). Researchers control which structural features are exposed. Full opt-out returns to v1.1 semantic-similarity-only recommendation.

- Default: off. No wiki structure exposed to recommendation engine.
- Level 1 opt-in: expose topic cluster distribution only (no page names, no content)
- Level 2 opt-in: expose schema file structure (category taxonomy, not content)
- Level 3 opt-in: expose gap map (which topic clusters have low coverage) — enables complementarity matching

## 5 Reputation & Discovery System

Archora replaces the blunt instrument of journal impact factor with a multi-dimensional, dynamically weighted reputation graph. Reputation is domain-specific — a leading virologist does not have inherently high reputation weight in materials science. In v1.2, wiki evolution is added as a mature-stage signal.

### 5.1 Reputation Signal Matrix

Signal	Weight (Early)	Weight (Mature)	Description
Peer Citation Weight	High (35%)	<b>High (30%)</b>	Citations from high-rep researchers multiply in value proportional to citer domain expertise score
Expert Endorsement	High (30%)	<b>Medium (20%)</b>	Explicit endorsements from verified researchers with strong track records in the relevant field
Track Record Accuracy	Medium (15%)	<b>High (20%)</b>	Historical ratio of adopted hypotheses, accurate predictions, and community-validated insights
Community Influence	Medium (10%)	<b>Medium (15%)</b>	Downstream impact: how often your work shapes others' research directions, archives, and wikis
AI Citation Count	Low (5%)	<b>High (25%)</b>	How frequently AI synthesis engines cite your work when generating new hypotheses — rises as data accumulates
Wiki Evolution	None	<b>Medium (5%)</b>	Rate and quality of wiki growth over time — tracked via log.md, signals sustained intellectual engagement with a domain
Institutional Prestige	Low (5%)	<b>Low (5%)</b>	Affiliation and venue signals. One factor among many; never dominant. Decays if not reinforced by peer signals.

### 5.2 Domain Specificity

Reputation scores are computed per domain cluster, not globally. Domain clusters are learned from the knowledge graph — not manually defined taxonomies. A researcher can have high reputation in one cluster and low in another.

### 5.3 Visibility Rules

Reputation scores are never shown as raw numbers. Researchers see: Domain standing (Emerging / Established / Leading / Foundational), influence indicators on individual posts, trend direction over 6-month and 24-month windows. Raw score manipulation avoided by design.

## 6 Cold Start & Seed Community Strategy

The RL loops and wiki-similarity recommendation both require sufficient data to function. The cold start strategy is unchanged from v1.1 in structure, with an additional consideration: seed researchers should be willing to opt into wiki-similarity matching early to bootstrap the recommendation graph.

<b>1. Domain Focus</b>	Launch within a single dense research domain. Target: computational neuroscience or climate modeling. Goal: 50–150 researchers in Year 0.
<b>2. Invitation Model</b>	Closed beta with invitation chains from 10–15 seed researchers. Each invites 5–10 peers. Ensures social density from day one.
<b>3. Reputation Bootstrap</b>	Seed researchers pre-assigned reputation scores from verifiable external signals (Google Scholar h-index, existing citation networks) to ensure Loop 2 has a working baseline.
<b>4. Wiki Seeding</b>	NEW in v1.2: Seed researchers optionally pre-populate their wiki structure during onboarding (guided import from existing Obsidian vaults, Zotero libraries, Notion workspaces). Faster convergence of wiki-similarity recommendation.
<b>5. Expansion Trigger</b>	Open next domain only when Loop 1 achieves >60% accept rate, Loop 2 produces stable reputation rankings, AND wiki-similarity recommendation CTR exceeds semantic-only baseline in A/B test.

## 7.1 Knowledge Graph & Research Map

A visual representation of connections between researchers, hypotheses, citations, and institutions. In v1.2, wiki topology (hub pages, cluster density, cross-domain bridges) is also visualized — making structural similarity relationships visible.

- Researcher nodes scaled by domain reputation
- Edge weight represents citation, collaboration, and wiki-similarity
- Cluster detection reveals emergent research communities
- NEW: Wiki topology overlay — shows which researchers' compiled knowledge structures are topologically adjacent

*Phase: Beta launch, expanding globally in Phase 2.*

## 7.2 Hypothesis Versioning & Evolution Tracking

Hypotheses are versioned like code. Every iteration preserved with timestamps, commentary, and community signals. In v1.2, versioning also tracks the underlying wiki page that generated the hypothesis.

- Full diff view between hypothesis versions
- Citation propagation tracking: which version is being cited
- NEW: Wiki lineage — which source pages in the raw archive contributed to this hypothesis

*Phase: Core feature, Day 1.*

## 7.3 AI Confidence Scoring

Each AI-generated synthesis item includes explicit confidence across Literature Consistency, Testability, and Novelty. In v1.2, a fourth dimension is added: Wiki Integration — how well the synthesis connects into the researcher's existing wiki structure.

*Phase: Beta Phase 2.*

## 7.4 Preprint Fast-Track

- One-click submission to arXiv, bioRxiv, SSRN
- Archora post and preprint bi-directionally linked
- Preprint citation counts imported back into reputation

*Phase: Phase 2.*

## 7.5 Multilingual Content Architecture

- Full content translation on demand for any Archora Space post
- Recommendation engine accounts for language — non-English posts not penalized
- Citation and reputation signals language-agnostic

*Phase: Phase 2.*

## 8 Development Roadmap

Phase	Timeframe	Focus	Key Milestones
<b>Phase 0</b>	Months 1–6	<b>Seed Community</b>	50–150 researchers in one domain; core channel architecture; basic RL loop instrumentation; wiki layer MVP (ingest + lint); human-augmented AI synthesis
<b>Phase 1</b>	Months 7–12	<b>RL + Wiki Activation</b>	Loop 1 >60% accept rate; wiki-similarity recommendation A/B test launched; reputation graph operational; knowledge graph visualization; hypothesis versioning live
<b>Phase 2</b>	Months 13–24	<b>Expansion</b>	Second domain launched; wiki-similarity CTR validation; preprint fast-track; multilingual support; AI confidence scoring; replication registry
<b>Phase 3</b>	Year 3+	<b>Platform Maturity</b>	Global open access; AI citation weight increases as signal validates; institutional partnerships; API for external integrations; wiki topology as public epistemology layer

Risk	Mitigation
Cold Start / RL data scarcity	Deliberate seed community; reputation bootstrapping; human editorial augmentation in Phase 0; optional wiki pre-seeding from existing tools.
<b>Wiki similarity privacy violation</b>	NEW RISK. Structural features could leak intellectual roadmaps. Mitigated by: strict opt-in, anonymized cluster-level aggregation only, no raw page names or content exposed, differential privacy on cluster distributions.
AI citation loop closure	AI citation weight starts at 5%; quarterly validation checks; automatic weight reduction if divergence detected.
<b>Wiki quality degradation</b>	NEW RISK. Wikis become incoherent over time as LLM writes poorly connected content. Mitigated by Loop 3 structural health regularization; lint operations; researcher retains ability to lock or delete wiki pages.
Reputation gaming	Scores hidden; domain-specific scoring; citation velocity anomaly detection; wiki-similarity not gameable because structural topology is harder to fake than explicit engagement.
Filter bubble / insularity	10% off-domain recommendation (non-disableable); contradiction adjacency matching surfaces epistemic friction by design.
Sycophantic AI outputs	Negative signal from discarded outputs; skill randomization injection; novelty scoring penalizes redundant synthesis; wiki integration score penalizes synthesis that simply echoes existing pages.

***"Knowledge should be shaped by those who know it best."***

Archora is built on the conviction that scientific progress is impeded not by lack of intelligence, but by structural misalignment: the wrong people with the wrong incentives deciding what matters. The platform does not replace peer review. It rebuilds the selection environment so that peer pressure — from the most knowledgeable, most engaged researchers in a domain — is the dominant force shaping what rises.

The wiki layer, inspired by the LLM Wiki pattern, ensures that this selection pressure is grounded in compounded knowledge — not just reactions to individual papers. Two researchers whose compiled intellectual structures are topologically adjacent are not just reading the same things. They are building the same kind of understanding, from different vantage points. Archora finds them for each other.

---

**Version History** v1.0 — Initial concept design. | v1.1 — Event-driven AI generation, graduated AI citation weighting, multilingual architecture, opt-in peer matching, seed community cold start strategy, AI confidence scoring. | v1.2 — Three-layer wiki architecture (LLM Wiki pattern), personal wiki layer, ARCHORA.md schema file, Loop 3 structural health, wiki-similarity recommendation engine experiment, wiki privacy architecture.