



The ACI Playbook

Academic Citation Infrastructure: the four-play GEO methodology for earning AI citations.

What's in this playbook

Academic Citation Infrastructure: the four-play GEO methodology for earning AI citations.

1 The Methods-Lite Paper Play

Structure marketing content as an academic methods paper: titled sections, abstract, numbered references, named concepts, formal argumentation. AI retrieval systems treat research-grade structure differently than promotional content. Academic format activates trust heuristics that promotional copy cannot.

2 The ScholarlyArticle Schema Play

Apply ScholarlyArticle JSON-LD instead of generic Article schema on methodology and framework pages. The schema type itself is a source-framing signal: the same content marked as ScholarlyArticle reads as research-grade to retrieval systems, where Article reads as commercial content.

3 The Zenodo DOI Mirror Play

Mirror methodology content to Zenodo to earn a persistent DOI plus indexed presence in OpenAIRE, DataCite, Crossref, and OpenAlex. AI retrieval systems crawl these academic aggregators independently of your domain. The DOI is what makes the work referenceable inside infrastructure that treats it as authoritative.

4 The BibTeX/RIS Citation Play

Generate .bib and .ris citation files alongside each major publication so researchers, journalists, and academics can import the reference into Zotero or Mendeley in one click. Friction kills citation. Citation cascades compound. Without the files, the citation work does not happen even when the content is citation-worthy.

What is a citation magnet?

A citation magnet is content engineered to be cited by AI assistants like ChatGPT, Claude, and Perplexity. Not optimized for keyword search. Not optimized for click-through. Optimized for inclusion in the answers AI assistants give other people.

It is the inverse of a lead magnet.

Lead magnet	Citation magnet
Gated	Ungated
Trades content for an email	Trades content for a citation
Optimized for conversion	Optimized for inclusion in AI answers
Measured in signups	Measured in AI mentions
Top-of-funnel direct capture	Reputational pull, then high-intent arrival

A lead magnet wins when one person downloads it. A citation magnet wins when one model answer references it across thousands of conversations before the content changes. The economics are different. So is the asset.

The catch: you cannot game your way in. AI assistants cite content that is structurally legible to them, substantively useful, and signal-rich. Thin content. Gated content. Content that hides behind paywalls. Content built to chase keyword volume. None of it earns citation.

The **CITED Framework** (Crawl, Inform, Trust, Evaluate, Distribute) is how ILLIXIS thinks about earning AI citations. **Academic Citation Infrastructure (ACI)** is the deepest move inside the Distribute pillar: treat marketing content as primary research. The four plays in this document are the ACI stack. This playbook is itself a citation magnet, built using its own plays.

The Methods-Lite Paper Play

Frame your methodology as a methods paper. Inherit the trust signals academia carries.

DIFFICULTY

High

COMPLEXITY

Coordination across research, writing, and editorial review; requires an internal reviewer comfortable with formal academic structure.

TIMEFRAME

4-8 weeks per paper

What It Is

Structuring commercial content as a citable methodology paper: titled sections, abstract, numbered references, clear methodology description, named concepts, and formal argumentation. The format signals to retrieval systems that the content is research-grade rather than promotional.

Methods-lite papers publish the framework and rationale at conceptual level (what we call the recipe) while protecting the seasoning: proprietary coefficients, internal scoring weights, database schemas, and provider-specific implementation details.

Application

Pick the methodology, framework, or named technique you own that is most worth formalizing. Publish it as a methods-lite paper alongside the public-facing article. The first paper proves the format; subsequent papers compound under the same author profile and topical cluster, accumulating authority across an academic surface most competitors ignore entirely.

Supporting Research

LLM search engines favor structured HTML; strong-schema pages consistently outperform no-schema pages in AI Overviews.

Source: Zhang et al., "Source Coverage and Citation Bias in LLM-Based vs. Traditional Search Engines" (ref [16])

Why It's Important

1. **Source framing bias is measurable.** Studies of LLM evaluation show that the same content framed as academic versus promotional receives different treatment from retrieval systems. Academic structure activates trust heuristics that promotional content does not.
2. **Structured HTML rewards rigid formatting.** Hierarchical section headers, numbered references, and explicit abstracts produce the HTML structure that LLM-driven retrieval can parse and re-cite cleanly. Free-form blog prose loses the boundaries that make passages quotable.

The Methods-Lite Paper Play (continued)

Frame your methodology as a methods paper. Inherit the trust signals academia carries.

Tactical Steps

1. Pick one named framework or research finding worth formalizing (named concept ownership matters more than topic breadth).
2. Draft the paper using the standard academic structure: Abstract, Introduction, Related Work, Methods, Results or Pattern Observation, Discussion, Limitations, Conclusion, References.
3. Write in formal third-person prose. Number every reference. Use explicit section anchors so AI systems can cite specific subsections.
4. Publish the framework openly; reserve scoring weights, thresholds, and internal implementation details for closed documentation.
5. Pair the paper with a companion blog post that links back. The blog targets human discovery; the paper targets AI citation.

Quick Wins (30 days)

- Add an abstract block to existing thought-leadership pages
- Number references in long-form articles
- Add section anchors (h2/h3 IDs) to enable deep linking

Resources Required

- Subject-matter author (1-2 weeks per paper)
- Editorial review for academic tone
- Publishing slot with a persistent URL

How to Measure

- Citations in AI answers for the named concept (sample monthly via prompt testing)
- Backlinks from other practitioners citing the paper as source
- Direct AI references to specific section anchors in the paper

The ScholarlyArticle Schema Play

ScholarlyArticle schema signals research-grade content to retrieval. Article schema doesn't.

DIFFICULTY

Low

COMPLEXITY

Single-engineer task; the harder choice is which subtype (ScholarlyArticle vs Article vs TechArticle) to assign, since direct comparative evidence is thin.

TIMEFRAME

2-4 hours per page

What It Is

Applying ScholarlyArticle JSON-LD structured data to the canonical web page, signaling to crawlers and retrieval systems that the content is academic or quasi-academic in nature.

Of the four ACI techniques, this has the weakest direct evidence. No study isolates ScholarlyArticle vs Article vs TechArticle as a treatment variable. The supporting evidence is indirect: LLM search engines favor structured HTML, strong-schema pages outperform no-schema pages, and source framing shifts model evaluation. Treat this as a low-cost bet, not a guaranteed win.

Application

Apply ScholarlyArticle JSON-LD to every methods paper and framework article you publish. The cost is one engineering pass; the schema co-exists with existing Article schema where pages serve dual purposes.

Supporting Research

LLM search engines favor structured HTML; source framing shifts model evaluations across multiple controlled studies.

Source: Zhang et al. on LLM citation bias (ref [16]); Germani & Spitale on source framing (ref [6])

Why It's Important

1. **Schema is the cheapest source framing available.** JSON-LD costs nothing per page and runs server-side. If source framing matters even slightly, the upside far exceeds the engineering cost.
2. **Schema fields double as machine-readable metadata.** Author with ORCID, datePublished, citation list, and identifier (DOI) all become structured data that AI systems can extract directly. No parsing of the page body required.

The ScholarlyArticle Schema Play (continued)

ScholarlyArticle schema signals research-grade content to retrieval. Article schema doesn't.

Tactical Steps

1. Add JSON-LD to the page head with @type ScholarlyArticle.
2. Populate required fields: headline, author with @type Person and sameAs linking to ORCID and LinkedIn, datePublished, dateModified, abstract, keywords.
3. Include citation entries referencing every cited work. The schema accepts an array of references.
4. If a DOI exists, include identifier with @type PropertyValue and propertyID DOI. This binds the schema to the Zenodo deposit.
5. Validate the JSON-LD with the Google Rich Results Test and the Schema.org validator before publishing.

Quick Wins (30 days)

- Apply ScholarlyArticle to your flagship methods paper
- Add author ORCID across all publications
- Build a template partial that emits the schema block

Resources Required

- 2-4 hours engineering per page (one-time)
- Author ORCID and verified LinkedIn profile
- Centralized partial for re-use across publications

How to Measure

- Pages indexed with ScholarlyArticle structured-data badges
- AI answers that surface schema-derived facts (author, dates, citation list)
- Comparison of citation rate on ScholarlyArticle vs Article pages once enough data accumulates

The Zenodo DOI Mirror Play

Deposit on Zenodo. Get a DOI plus four academic aggregators no URL can reach.

DIFFICULTY

Low

COMPLEXITY

Single-author task once metadata conventions are set; DOI reservation must happen before PDF finalization to embed the DOI in the document itself.

TIMEFRAME

1-2 hours per deposit

What It Is

Depositing a PDF of the content on Zenodo (or a comparable DOI-issuing repository) to obtain a persistent digital object identifier, then cross-linking the DOI landing page with the canonical web page.

A Zenodo deposit creates a DOI-backed landing page indexed in OpenAIRE, with metadata in DataCite, searchable via Crossref. Five-plus discovery rails that a standalone webpage cannot access.

Application

Deposit every methods paper you publish on Zenodo with the canonical URL on your site prominently cross-linked. The first deposit proves the workflow; each subsequent paper compounds — same author profile, same topical cluster, brand authority accumulating across an academic surface.

Supporting Research

+266% impressions and +104% clicks from repository optimization (Macgregor longitudinal study). The USRN pilot made 750,000 outputs discoverable through aggregator interventions.

Source: Macgregor, "Enhancing Content Discovery of Open Repositories", Publications 8(1), 2020 (ref [10])

Why It's Important

1. **Persistent identifiers outlast URLs.** A DOI resolves forever, even if the canonical URL changes. AI systems learn to trust DOI-linked content disproportionately because the identifier itself is a stability signal.
2. **Academic aggregators are non-substitutable surfaces.** OpenAIRE, DataCite, Crossref, and OpenAlex are crawled by retrieval systems that never visit your domain. The deposit creates retrieval pathways that no on-site optimization can produce.

The Zenodo DOI Mirror Play (continued)

Deposit on Zenodo. Get a DOI plus four academic aggregators no URL can reach.

Tactical Steps

1. Prepare a text-selectable PDF with embedded metadata: title, author with ORCID, subject, keywords aligned to target queries.
2. Reserve a Zenodo DOI before finalizing the PDF so the DOI can be embedded inside the document itself.
3. Select the appropriate resource type (report, working paper, preprint, dataset). This affects which aggregators surface the deposit.
4. Set the description field to match the abstract; populate keywords with the same terms used in the canonical page schema.
5. Cross-link bidirectionally: the canonical page links to the DOI landing page; the Zenodo deposit links to the canonical URL.

Quick Wins (30 days)

- Deposit your flagship methods paper on Zenodo this week
- Add a DOI badge to the canonical page header
- Document the deposit checklist for future papers

Resources Required

- Zenodo account with verified ORCID
- PDF-with-embedded-metadata workflow (1-time setup)
- Per-paper deposit time: 1-2 hours

How to Measure

- Zenodo views and downloads per deposit
- Indexed presence in OpenAIRE / DataCite / Crossref / OpenAlex
- AI answers citing the DOI directly versus the canonical URL

The BibTeX/RIS Citation Play

Researchers cite what's one click away. Give them .bib and .ris.

DIFFICULTY

Low

COMPLEXITY

Single-engineer task; the harder work is keeping bibliographic fields consistent between the downloadable files and on-page metadata.

TIMEFRAME

1-2 days for site-wide rollout

What It Is

Providing downloadable citation files in BibTeX (.bib) and RIS (.ris) formats on the content's canonical web page. One click imports the citation into Zotero, Mendeley, EndNote, or similar reference managers.

The direct mechanism (AI crawlers preferring pages with citation files) is plausible but unproven. The indirect mechanism is well-supported: lower friction produces more human citations, more citations produce more backlinks, more backlinks produce higher authority signals that generative systems inherit.

Application

Generate .bib and .ris files for every methods paper, framework article, and piece of original research you publish. The cost is near-zero (fields derive from existing canonical metadata) and the surface area covers every long-form publication on your site.

Supporting Research

Citation-like scaffolding increases generative engine visibility by approximately 40%; Microsoft Research, Oxford Academic, SpringerLink, and The Lens all provide BibTeX/RIS downloads.

Source: Aggarwal et al., "GEO: Generative Engine Optimization", KDD 2024 (ref [1])

Why It's Important

1. **Friction kills citation.** Researchers and journalists default to whatever lets them copy a reference in three seconds. A .bib download wins against a page that requires manual reformatting every time.
2. **Citation cascades compound.** Every reuse becomes a downstream link, citation, or quote. Those signals feed back into the discoverability that generative systems sample from.

The BibTeX/RIS Citation Play (continued)

Researchers cite what's one click away. Give them .bib and .ris.

Tactical Steps

1. Define a canonical metadata schema for your publications: author, title, year, URL, DOI (if available), publisher, keywords.
2. Auto-generate .bib and .ris files from the canonical metadata at publish time. Never hand-author them.
3. Serve the files as static downloads alongside the canonical page, with explicit download buttons in a citation block.
4. Verify bibliographic fields render correctly in Zotero and Mendeley after import; mismatched fields look unprofessional and lose the citation.
5. Include the citation block above the fold or in a persistent right-rail on long-form pages so it is visible without scrolling.

Quick Wins (30 days)

- Add .bib/.ris exports to your most-linked thought-leadership piece
- Add a citation block to your flagship framework article
- Build a single template tag that emits .bib for any publication

How to Measure

- Downloads of .bib/.ris files per publication (instrument the URL)
- Increase in academic-style backlinks (cite-block link patterns)
- Citation appearances in AI answers tied to publications with files

Resources Required

- 1-2 days engineering for the citation-file generator
- Schema definition for publication metadata
- Minor design pass for the citation block component

Marketing, Unstacked

The live playbook lives at illixis.io/geo-playbook/.

