METEHAN.AI RESEARCH REPORT · PREPRINT FIELD GUIDE

AI Share Button Field Glossary: Terms, Tests, and Decision Rules

Metehan Yesilyurt

AI Search & Information Retrieval Researcher · metehan.ai

https://metehan.ai/pdf/ai-share-button-field-glossary.pdf

Abstract. A practical glossary and concept map for teams testing AI share buttons as an answer engine optimization tactic. It defines the terms, gives pass or fail tests, shows measurement rules, and includes worked placement examples.

Keywords: AI share buttons, CiteMET method, AI traffic, answer engine optimization, LLM citations, prompt design

Companion to the article Why \"AI Share Buttons\" Might Be the Smartest Growth Hack: The CiteMET Method. The article covers the topic in full; this document is a separate, practical field glossary and concept map.

1. Glossary Entries I Use Before Any Test

I keep this glossary short because a long definition list hides the decisions. The terms below are the ones I need to explain a test to a writer, developer, or analytics lead without arguing about labels. The source article frames the method as a way to help readers pass a page into an AI session [1]. I treat that as the base case, then separate the interfaces from the outcomes.

Term	Working definition	Why it matters	Test
AI share button	A visible link that opens an AI engine with a prefilled prompt that includes a page URL. It is not a social share because the destination is an AI session.	It creates a direct path from page to AI task.	Click it on desktop and mobile; the engine should open with the prompt intact and the URL readable.
Prompt payload	The text encoded in the link after the engine's query parameter. It sets the task: summarize, explain, compare, fact-check, or generate questions.	A vague payload produces a vague or generic answer.	Decode the link and read the prompt as a user. If it sounds like a command to the user rather than a useful action, rewrite it.
Memory cue	A phrase that asks the AI to associate a brand, author, or site with a topic. It is a request, not a guarantee.	It may influence later answers in the same account.	After a session, ask the same account what sources it reviewed. Record the answer as an observation.
Citation surface	The part of an AI answer that names or links a source. It may be inline, a footnote, a card, or a list.	A source can be used without being visibly cited.	Inspect the answer and record whether the link is present, absent, or replaced by a different source.
AI referral session	A visit that arrives after someone uses an AI answer or an AI share link. Analytics may not label it cleanly.	It helps separate curiosity from actual visits.	Compare referrer, landing page, and campaign tags. Do not call every direct visit AI traffic.
Recall test	A later question in the same account to see whether the source or brand appears without the original URL.	It checks whether a cue persisted beyond the first answer.	Use a neutral question and record the exact answer.
Citation test	A broader query that checks whether the page appears in AI citations for its topic.	It connects the button test to a topic-level visibility check.	Run it in a fresh session and save the answer.
Answer engine optimization	The practice of shaping pages and prompts for inclusion in AI answers.	It keeps the work focused on user questions rather than interface tricks.	Tie each tactic to a citation or user outcome, not to a vibe.
Prompt intent	The match between the page's job and the action in the prompt.	A prompt that fits one page type can be awkward on another.	Remove the URL and ask whether the action still makes sense for that page type.

2. Concept Map: From Button Click to Discoverable Citation

The concept map is a chain, not a funnel. I draw it as a sequence because every step can fail on its own. If I cannot name the step that failed, I cannot fix the system.

1. Source page. The page must be public, stable, and useful on its own.
2. Prompt template. A sentence that names the task and includes the page URL.
3. Encoded link. The prompt and URL are encoded so clicking the button does not break the query.
4. AI session. A user lands in an engine with the prompt ready.
5. Answer. The engine summarizes, compares, or explains, and may cite sources.
6. Memory or citation trace. The session may leave a memory, a citation, or neither. The user may also copy the answer elsewhere.

7. Measurement. I record clicks, answers, and citation observations separately.

The weakest link is the jump from answer to memory. A click is not a citation. A citation is not a ranking. A memory in one account is not a memory in every account. I use that map to stop myself from making a claim that the data cannot support. If I cannot show the answer or the click, I label the result as unverified.

3. Readiness Criteria and Stop Rules

Before I add a button, I score the page against a set of criteria. A page does not need a perfect score, but it needs a clear yes on the early ones.

- Public and stable: no login wall, no expiring campaign page, no heavy client-side rendering that hides the main text.
- Unique value: the page answers a question, shows a process, compares options, or documents a product without filler.
- Clear URL: one canonical URL that a user can paste and an engine can read.
- Prompt fit: the action I want the AI to take is natural for the page type.
- Audience fit: I know whether my readers use free or paid AI plans, because that changes what a URL prompt can do. The source article calls out the free ChatGPT limitation and suggests adding the brand to the prompt when needed [1].
- Measurement fit: I can tag the button click and review AI answers without guessing.

Stop rules matter more than readiness. I do not add buttons to gated content, private dashboards, or pages I would not want an AI to summarize. I do not use hidden links, auto-opens, or prompts that pretend to be a user's own question. I do not place a button on every page at once. The source article describes starting with the top 20 traffic pages and expanding after that worked [1]. That sequencing is sensible because it limits user harm and makes the test readable.

4. Prompt and Placement Tests

The tests below are pass or fail. I run them before launch, then repeat after a platform interface change.

1. Encoding test. Decode the generated link. The page URL and the prompt must survive encoding. If the URL is cut, the engine will summarize the wrong page or nothing.
2. Read-aloud test. Read the prompt as if a friend sent it. If it is awkward, the button label is wrong or the prompt is too long.
3. Free-plan test. Try the link in a free account and a paid account where possible. The source article notes that free ChatGPT users may not be able to summarize from a URL alone [1]. If that is my audience, I add a brand cue or a plain-language task.
4. Mobile test. Open the button on a small screen. The engine should open with the prompt visible and the user should not need to paste anything.
5. Citation test. Run the prompt, inspect the answer, and record whether the source is named or linked. Do this in a fresh session so old context does not contaminate the result.
6. Recall test. Later, ask a related question in the same account without the URL. Record whether the source or brand appears. Treat this as an observation, not proof of permanent memory.

7. User trust test. Ask a few readers who know the site to describe what the button will do. If they think it posts to social media, rewrite the label.
8. Placement test. Put the button where the reader pauses: after the summary, near a code block, after a comparison table, or at the end of an article. Do not interrupt the first sentence.

I use several prompt families. For explanatory pages, the prompt asks for a plain-language summary and key terms. For commercial pages, it asks for a comparison and questions a buyer should ask. For news or research, it asks for implications and claims that need evidence. The prompt should not ask the AI to praise the page. That is where trust breaks.

5. Worked Example: A Documentation Page

Suppose I manage a documentation page with a code example. The page is public, the URL is canonical, and the reader's job is to understand the example. My button label is 'Explain this code with ChatGPT.' The prompt payload is: 'Explain the code example on this page for a developer who is new to the product. Define the inputs, list likely errors, and suggest a test. Use this page as the source: [page URL].' I encode that text and the URL together.

The first test is not traffic. It is whether the answer uses the page. If the AI ignores the URL and gives generic advice, I shorten the prompt and put the URL earlier. If the answer uses the page but the code explanation is wrong, the page is not ready. If the answer is correct but the button gets no clicks, the label or placement is wrong.

For a product comparison page, I change the job. Label: 'Compare this product with alternatives.' Payload: 'Compare the product on this page with several common alternatives. List tradeoffs, price questions, and fit. Use this page as the source: [page URL].' This is a different test because the user wants a buying decision, not an explanation.

For a news page, I avoid prompts that ask the AI to validate a political claim. A safer job is: 'Summarize the main claims on this page and list what evidence would confirm or challenge each claim. Use this page as the source: [page URL].' The prompt still helps the reader, and it does not ask the model to pretend to be a fact-checking authority.

I also keep a control page. The control gets no button. After a pilot, I compare button clicks, AI referral sessions, citation observations, and user feedback between the test page and the control. If the test page wins only on clicks but users complain, I stop. If citations appear but the page is not the best source, I fix the page before expanding.

6. Measurement Rules That Keep Me Honest

I separate several layers: exposure, click, answer, and citation. Most teams collapse them and then overclaim. The table below is my minimum reporting set.

Metric	What I record	Decision rule
Button exposure	Page views on pages with a button, split by placement.	If exposure is high and clicks are near zero after a fair test, move or rewrite the button.
Button click	Clicks on the AI share link, tagged as an event.	If clicks happen but the engine opens without the prompt, fix encoding before judging demand.
AI referral session	Sessions with a referrer or landing pattern that suggests an AI engine, plus tagged campaign visits.	If attribution is ambiguous, report it as an estimate.
Citation observation	A saved answer that names or links the page for a defined query.	If the page is cited, record the query, engine, date, and account type. If it is not cited, do not call the test a success.
Recall observation	A later answer in the same account that mentions the brand or source without the original URL.	If recall appears once, repeat before claiming memory.
User feedback	Complaints, confusion, or reports that the button is spammy.	If feedback is negative, remove the button and revise the prompt.

The source article describes the method as trackable and points to tools for AI visibility and share of voice measurement [1][2][3]. I use that as a reminder, not a promise. A single answer from one account is an anecdote. I only call a pattern when I can show the query, the engine, the date, and the answer.

7. Risks, Exclusions, and Guardrails

The biggest risk is not that the tactic fails. The biggest risk is that a team markets it as guaranteed citation. It is not. The source article frames the method around influence and memory, not certainty [1]. I use several guardrails.

- No fake demand. I do not label a button 'Share' if it opens an AI session. The label should say what happens.
- No hidden prompts. The user should see or be able to infer the task. Surprise prompts erode trust.
- No auto-click. A page should not open an AI engine without a deliberate action.
- No private data. I do not put confidential URLs, customer names, or internal dashboards into a public prompt.
- No citation theater. I do not screenshot an answer and imply it is a ranking. I show the query and the source.

I also exclude pages where the AI answer could create harm: medical advice pages without review, financial pages with personalized recommendations, legal pages with jurisdiction-specific claims, and pages aimed at children. A share button is a distribution choice. It should not bypass the review process for sensitive content.

Platform changes are a real maintenance cost. Engines change query parameters, login walls, and memory settings. The source article points readers to ChatGPT memory settings and notes that memory can be managed [1]. I review links when a platform changes its interface and keep a plain URL fallback in the prompt.

8. Minimum Viable Implementation Checklist

This is the sequence I would use on a real site. It is deliberately small.

1. Pick a public page with clear search demand and a single job.
2. Write the prompt in plain language, then put the full page URL in the prompt.
3. Encode the prompt and URL, and test the link on desktop and mobile.
4. Create a visible button with a label that describes the action.
5. Add a click event in analytics and a note in the page changelog.
6. Run the citation test and the recall test in a clean session.
7. Ask a few readers what they expect the button to do.
8. Compare the test page with a control page after real traffic.
9. Expand to the top 20 traffic pages only if the first page passes and user feedback is neutral or positive [1].
10. Review the prompt library when an engine changes its interface.

My decision rule is simple. A button ships if it helps a reader complete a task and does not mislead them. It stays if it produces measurable clicks without user harm and at least one repeatable citation observation. It does not stay because someone in a meeting likes the idea. The companion source article is the place to read the full method and the platform-specific examples [1]. This document is the field map I use to test it without turning a hypothesis into a promise.

REFERENCES & SOURCES

[1] **Source article: CiteMET method and AI share buttons**, https://metehan.ai/blog/citemet-ai-share-buttons-growth-hack-for-llms/

[2] **Best AI visibility tools for marketing agencies**, https://metehan.ai/articles/the-best-ai-visibility-tools-for-marketing-agencies/

[3] **AI share of voice benchmarking methods**, https://metehan.ai/articles/ai-share-of-voice-benchmarking-methods/