

Answer Engine: A New Open-Source Tool for AI and Semantic Search That Cites Its Sources — or Says "I Don't Know"

Luke F. Walton released Answer Engine, an open-source AI search reference tool that cites sources, refuses when unsure, and keeps judgment outside the model.

SAN DIEGO, CA, UNITED STATES, June 29, 2026 /EINPresswire.com/ -- Luke F. Walton, an independent researcher on [AI ethics and answerability](#) and the [founder of Surmado, Inc.](#), has released [Answer Engine](#), an open-source reference implementation for AI-powered search that is built to cite its sources or honestly say "I don't know."

The project is available now on GitHub

under the permissive Apache-2.0 license, with a live deployment, "Ask the Archive," running on the author's own published writing and podcast transcriptions.



Luke F. Walton in Studio

As AI answer engines like Google's AI Overviews, Perplexity, ChatGPT, Copilot, and Gemini replace the list of links with a single composed answer, a new failure has come with them: these tools can return confident judgments in their own voice, sometimes citing sources that do not support the claim, sometimes describing work that does not exist, and often without a clear author a reader can question. Walton's research calls this the answerability gap: the point at which the account for what a system says comes loose from anyone who can be asked to answer for it.

“

The point isn't a smarter assistant. It's search that's helpful, guarded against hallucination, and that will accurately and reliably tell you when it doesn't know.”

Luke F. Walton

The deeper concern is not only that a model can be wrong. As answer-engine optimization grows into a paid discipline, the composed answer itself becomes a target. A party with a stake in the outcome can shape what an engine says while the surface still reads as neutral synthesis, and

whoever can spend the most on that shaping is the one rewarded. The public acting on the answer sees neither the price nor the author. Walton's writing argues that the remedy is not more disclosure of sources or models, but a visible owner standing behind the frame an engine voices.

Answer Engine is a deliberately small, readable demonstration of a different contract. You point it at a body of work: essays, documentation, lyrics, letters, research notes. It answers one question at a time. There is no chatbot, no conversation memory, no persona improvising on the author's behalf. By design, it is not a hosted app, a framework, or a search-optimization product. It is only the smallest useful version of an answer contract. Each answer is a single transaction: a cited answer drawn from the corpus, a pointer to where relevant private material lives, or an explicit refusal. As the project's documentation puts it, "a chatbot that's right most of the time speaks for you; an answer engine that cites or declines speaks from you."

Three mechanisms make that contract checkable rather than promised:

A structural privacy boundary. The corpus splits into public records, which are quotable and cited, and private notes, which are searchable but never shown to the model. The internal path that assembles evidence for the model carries a data type with no field for private text, so the leak is structurally unavailable on the supported prompt-building path, not a rule someone has to remember to follow.

Honest citation modes. Every answer declares one of four modes (supported, partial, related-material, or not-found), and the mode is re-derived from the evidence actually retrieved, not taken from the model's word for it. An invented source is treated as an error, not a footnote. Refusal as a tested behavior. A fixed suite of "gold" questions, including ones the system is required to refuse, runs like regression tests, so the engine's willingness to say "I don't know" is verified rather than assumed, and held constant across quantization types.

Asked about a subject the corpus covers, it returns the claim together with the source it came from. Asked about something that lives only in a private note, it points to where that moment can be found without quoting it. Asked about something the corpus never addresses, it returns nothing and says so, rather than assembling a plausible-sounding answer out of nearby material.

That private-but-not-quotable layer is the capability Walton expects institutions to care about most. A library, museum, lab, or private collection can let restricted material such as sealed transcripts, donor notes, or embargoed findings improve what the system finds and where it points researchers, without that material's text being quoted or reproduced. The private layer shapes results and can route a reader to where a source lives; its contents are never restated in the answer.

Answer Engine is a reference implementation, not turnkey software. It is meant to be cloned,

read in a single sitting, and adapted by the developers who serve archives, nonprofits, scientists, and small businesses. Walton is candid about where it stops. The design holds soundness: nothing enters an answer that isn't grounded in retrieved evidence or honestly refused. But it does not, and cannot, certify that everything that should have been retrieved was. "This is the bounded case on purpose," the documentation notes, "not a proof about the unbounded one." Proving the same boundaries hold at public, contested scale, Walton writes, belongs to the builders of public-scale systems, not to a teaching repository.

The repository runs locally on Node.js with a single API key, and ships with a synthetic example corpus, the evaluation suite, and citation metadata. To adopt it, a developer replaces the example content with their own records, adjusts how retrieval scores them, and rewrites the gold questions to define what their collection should and should not answer. The same suite that checks today's answers is what an adopter tightens as the archive grows. Contributions are welcome.

The tool is the engineering companion to Walton's broader work on answerability in AI, including papers on machine-mediated responsibility, answer-engine optimization, and citation-grounded AI systems.

Answer Engine (v2.0.0) is available at github.com/lukefwalton/answer-engine under the Apache-2.0 license and is archived on Zenodo (concept DOI: 10.5281/zenodo.20676773). The live deployment, Ask the Archive, and full project documentation are available at lukefwalton.com. The accompanying research is available at <https://philpeople.org/profiles/luke-f-walton>.

Answer Engine is a personal open-source project by Luke F. Walton, released under the Apache-2.0 license; it is not a Surmado product.

Luke F. Walton
Luke F. Walton
luke@lukefwalton.com

This press release can be viewed online at: <https://www.einpresswire.com/article/923084645>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.