Vickery Machine Logic, LLC
Greenville, South Carolina · vickerymachinelogic.com

PRESS RELEASE

FOR IMMEDIATE RELEASE

Whose Values Is an AI Agent Following? Greenville Developer John Vickery Shares an Early Release of Wisdom Core, an Open-Source Framework That Makes the Answer Explicit and Inspectable

> *Wisdom Core is a project in its infancy: an alpha (0.8.1), shared early for testing and feedback. Its mark, a glowing brain with a heart at its center, reflects the aim of the work: an agent whose values are written down, checked against a root principle, and kept on the record.*

GREENVILLE, S.C., Nov. 3, 2026 — Greenville developer John Vickery has shared an early release of Wisdom Core, an open-source Python framework (MIT, alpha 0.8.1) that gives AI agents a persistent, auditable value system. It is built around a question many people now ask about AI: whose values is an agent following? Wisdom Core makes the answer explicit, inspectable, and owned. The project is in its infancy, still in active development, and Vickery is inviting developers who build agents to try it, report what they find, and help shape where it goes.

Vickery designed and wrote the framework himself and publishes it through his company, Vickery Machine Logic, LLC. It is available on GitHub and the Python Package Index (PyPI), and installs with one command: pip install wisdom-core.

How Wisdom Core works

- **Core Axiom.** A root principle, written in plain language, is injected into every decision. Every principle the agent learns is evaluated against it.
- **Learned principles.** After each exchange, a synthesis pass extracts what the interaction taught and records it as a principle. Principles live in a local SQLite database and are embedded for semantic retrieval, so the most relevant ones (12 by default) shape each decision.
- **Distillation.** Near-duplicate principles are merged by embedding similarity, with thresholds calibrated to the embedding model in use. Principles that conflict with the axiom or with each other are removed, with a stated reason.
- **Feedback-weighted review.** The framework tracks which principles shaped each answer and how it was received. A principle behind answers that draw disproportionate pushback is flagged, then revised, removed, or confirmed. Nothing is removed on a count alone.
- **Full audit trail.** Every principle, merge, removal, and axiom change is logged as an event with its reason and timestamp, in a local file under the owner's control.
- **Model-agnostic.** Runs on Ollama, any OpenAI-compatible server, or Anthropic, with only the Python standard library.

Developers can use the framework's ready-made agent, or add two calls to an existing agent loop: `core.inject(task)` before a decision and `core.synthesize(task, result)` after it.

Where the project hopes it can help

In practice, an agent built on Wisdom Core runs a simple loop. Before each decision, the Core Axiom and the most relevant learned principles are placed into the prompt exactly as written. After each exchange, a synthesis pass records what the interaction taught as a principle that must fit the axiom, and skips the step when the lesson is already known. Over time, distillation is designed to keep that set small, consistent and readable.

Vickery's hope is that Wisdom Core can become a useful foundation wherever people need to see, and shape, the values an AI agent works from. Possible uses include customer support agents that keep staff corrections as traceable principles, personal assistants that accumulate standing preferences, review-heavy settings such as finance, healthcare administration and legal operations, tutoring tools anchored to a stated teaching principle, internal company copilots with an organization's own axiom, and research into how an agent's values change over time. These are possibilities the project is working toward, not proven deployments.

What it does not claim

Wisdom Core records the principles an agent has written down and how they change over time. It does not show what happens inside the model itself, and it is not a safety system. It is one practical, early step toward agents whose working values can be read and checked.

"The question I kept hearing was, whose values is this AI following? I think the answer should be: yours, written down where you can read it," said Vickery. "Wisdom Core is very new. I'm sharing it early because the people who try it will shape what it becomes. If it speaks to a worry you have, I'd love for you to try it and tell me what you find."

A project in its infancy, open to the people who try it

Wisdom Core is a new project released as an alpha. The architecture is in place, and testing and refinement are underway. Vickery plans to publish testing and research results as they become available. Questions, bug reports and ideas are welcome through the project's GitHub Issues page, and anyone is free to use, modify and build on the code under the MIT license. Code contributions are not being accepted yet.

About the creator

Vickery spent more than ten years running high-stakes U.S. payroll, tax and benefits operations for thousands of employees across many states, work that had to be right under audit. That work built a habit of automating complex workflows that demand precision, using Excel, Power Query and VBA to cut hundreds of hours of manual work. He now brings that same discipline to Python, SQL and AI through Vickery Machine Logic. He also drew the Wisdom Core mark himself: a glowing brain with a heart at its center.

Availability

- Project page: vickerymachinelogic.com/resources/wisdom-core/
- Source code and documentation: github.com/VickeryMachineLogic/wisdom-core
- Package: pypi.org/project/wisdom-core/ (`pip install wisdom-core`)

About John Vickery

John Vickery is the founder of Vickery Machine Logic, LLC, and the creator and sole maintainer of Wisdom Core. He is based in Greenville, South Carolina.

About Vickery Machine Logic, LLC

Vickery Machine Logic, LLC is a Greenville, South Carolina company that builds custom AI agents and apps for businesses, remote or local. Learn more at vickerymachinelogic.com.

Media contact

John Vickery, Vickery Machine Logic, LLC
Available for phone and video interviews.
john@vickerymachinelogic.com
864-436-8757

# # #