SNM05 AI-Powered Applications Meetup · University of Nizwa · 15 Nov 2026 Proposal · 10-minute talk with a live demo

The AI Suggests, the Code Decides

Agatha: a local-first AI research assistant that checks every quote against the real source

AUTHORS
Surena Karimpour Ghannadi, Behrad Azimi

PRESENTER
Behrad Azimi

ORGANISATION
Immortal Horizon Technologies

CONTACT
Behradazimi1@gmail.com

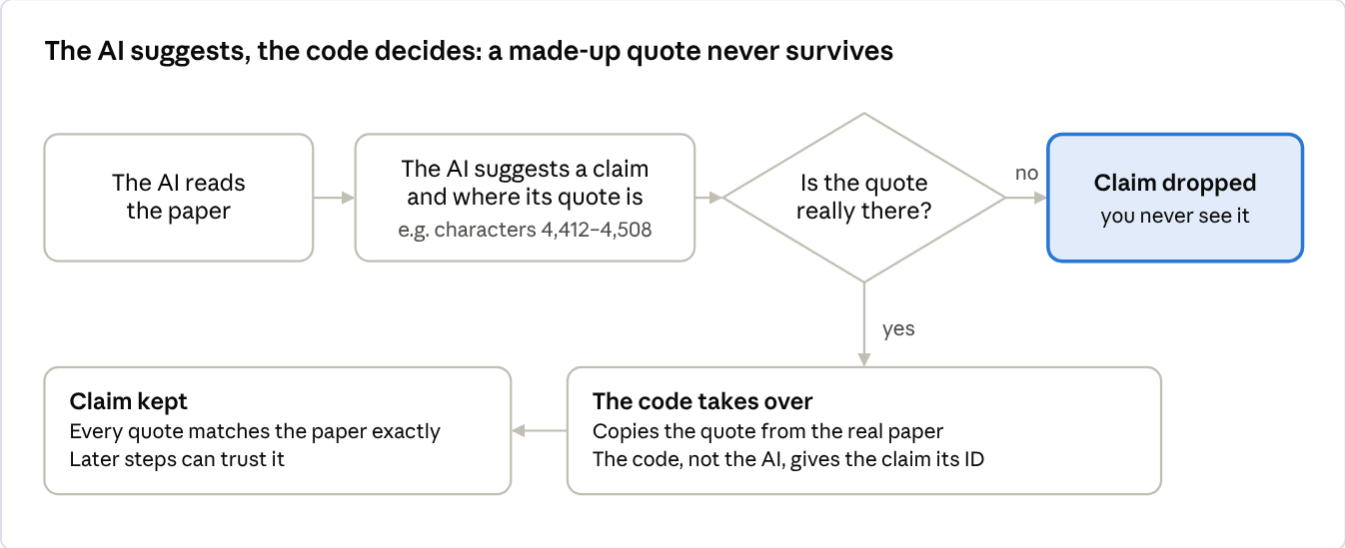
Abstract

AI assistants are now used for literature reviews, but they still invent sources and quotes, and AI agents can take actions nobody approved. Better prompts, human review and bigger models make these errors rarer, but they do not stop them.

We present a different approach: the language model only proposes, and plain, deterministic code decides. In Agatha, our research workbench, the model suggests a claim and where its supporting quote sits in a paper. The code checks that position against the real article and copies the quote from the source itself. A claim whose quote cannot be found is dropped, so a fabrication becomes a missing record rather than a convincing false one.

The same principle governs actions: each of Agatha's 112 agent tools declares what it reads, writes and trusts, and a missing declaration defaults to the safest case. Search, storage and checking run on the user's own computer, and Agatha sends no usage data to us. We will show a short live demo and discuss the open problem of extending these checks from reading to writing.

How a claim is checked



The model never supplies the quote text or the claim's ID; both come from the code and the source document.

Why it fits this meetup

- **Practical AI with a built-in check.** When Agatha reads papers, a made-up quote cannot pass as real, because every quote is copied from the source by code.
- **Responsible and compliant AI.** Documents stay on the user's device, three of its eleven AI provider options run fully offline, and tool permissions fail safe. This speaks directly to data protection, cybersecurity and ethics.
- **University support and the Oman market.** Agatha is built for researchers. We want to pilot it with universities in Oman and test it on their real research work.

What exists today

PART	WHAT IT IS	STATUS
Agatha	Desktop research workbench: searches 12 scholarly sources, extracts claims with exact quotes, and drafts in LaTeX.	Shipping, version 9.0.0
Horizon core	The shared engine for AI agents: tools, permissions, 11 AI providers (3 fully local) and multi-agent teamwork.	Shipping
AshaSwarm	A background service that runs coding tasks with one agent or a team of agents.	Pre-release

152,032
lines of code

4,031
automated tests

112
agent tools, all declared

Windows & Linux
tested automatically

10-minute talk plan

TIME	PART
0:00–1:30	The problem: AI that invents sources, and agents that act without permission
1:30–3:30	The idea: the model proposes, the code decides
3:30–6:30	Live demo of Agatha
6:30–7:30	What exists today
7:30–9:00	Why it fits Oman: data stays on the device, local AI models, built for universities
9:00–10:00	What we are looking for, and questions

What we hope to gain from the meetup

- **A university pilot in Oman:** one or two departments using Agatha on real research, with results we report back.
- **Research partners:** faculty interested in extending the checks from reading to writing, aiming for a peer-reviewed paper. We would welcome consideration for the Springer Nature publication opportunity.
- **Industry feedback:** other fields where AI output must be checked against a source, such as legal, compliance and technical documents.

Scope note. Agatha is an early-stage product: it is shipping, but it has no commercial customers yet. Extending the checks from reading to writing is open research, and we will present it as such.

Surena Karimpour Ghannadi & Behrad Azimi · Immortal Horizon Technologies

Final slides will be sent by 10 November 2026.