

Knowledge you can follow. Science you can build on.

Peel is a symbolic runtime that makes AI answers traceable to their source — built for research, medicine, and the tools people cannot afford to have wrong.

THE PROBLEM

AI is fluent. It is not grounded.

Language models write confident answers that can be wrong — and cannot reliably say why they said what they said. In research, clinical work, and tools for vulnerable users, an eloquent guess is a liability. Bigger models and bolt-on retrieval make the guesses smoother, not the grounding real.

THE INVERSION

Let the symbols author the facts.

Peel flips the usual architecture. A typed symbolic store is the **only author of facts**. The neural model is demoted to two jobs — propose (it must pass symbolic extraction) and render (it cannot add content). Hallucination is prevented structurally, not statistically.

THE PRODUCT

One console. Every answer sourced.

- ◆ **Peel** — a typed knowledge graph where every fact carries its provenance.
- ◆ **Lois** — a reasoning core that answers supported, conflicting, or unknown.
- ◆ **The Perslis runtime** — 44 source-pinned research tools on one console, driven in plain language.

SYSTEM CONNECTION

Get your lab hooked up. No special equipment.

One console connects the systems a lab already uses — 19 public science services (PubMed, UniProt, AlphaFold, ClinVar, KEGG and more), a compute bench, and WetHands lab automation reaching LIMS, electronic notebooks, and instrument APIs. We connect the instruments that were never meant to connect. Nothing to rip out.

THE MOAT

A contract nothing gets past.

- ◆ **Source contract** — a result that cannot name its source is refused, by construction.
- ◆ **Air-gap ready** — the lab record is local files that keep working with the cable pulled.
- ◆ **Metal to wet bench** — pure C++ on the metal and a wet-lab lane on the same console.
- ◆ **A symbolic admission layer** — not trained guardrails that can be talked around.

AIRTIGHT BY DEFAULT

Your research never leaves your computer.

- ◆ **Local models by default.** Ollama runs the reasoning on your own machine — no prompt, file, or graph is sent to a cloud model to answer a question.
- ◆ **We never train on your data.** Nothing leaves your devices to train any model, ever. Your research is not someone else's training set.
- ◆ **External models are opt-in.** Want a hosted model for a task? Choose from many providers — only if and when you decide. Off by default, always your call.
- ◆ **The record is yours.** Graph, ledger, notebook, and files live on your disk and keep working with the cable pulled.

PROOF IT IS REAL

Built. Tested. Live to click.

44

tools on one console

117/117

adversarial scenarios · 0 findings

6

public research papers

~372k

structured records in the prototype

A working demo runs today at perslis.com/peel/demo. Open datasets and an on-device model are public.

WHY NOW

Adoption is here. Trust is the bottleneck.

AI is entering regulated and high-stakes work faster than it can be trusted there. Agentic tools now **act**, not just chat — so access has to be governed and answers have to be checkable. On-prem and air-gapped demand is real. Grounding is the gate everyone hits, and it is the gate Peel is built to be.

WHERE IT LANDS

Research runs on tools like these.

- ◆ **Bioinformatics & research tooling** — the teams that live inside these databases every day.
- ◆ **Assistive & AAC** — where Peel began: context and dependable expression for vulnerable users.
- ◆ **Regulated & enterprise research** — pharma, clinical, and government, where provenance is non-negotiable.

HOW WE MAKE MONEY

Land a pilot. Expand the lab.

We stand up a pilot lab with a research partner around one bounded question and a clear definition of success. Pilots convert to per-server and per-seat licensing; enterprise deploys on-prem. Integration and support compound the relationship. Land narrow, expand across the lab.

THE PATH

Prototype to enterprise.

TODAY

Research prototype and a live, clickable demo.



NEXT

Pilot deployments with named research partners.



THEN

Domain validation against expert reference answers.



SCALE

Governed, on-prem enterprise deployments.

THE ASK

Fund the runtime research can trust.

We are raising to productize the runtime for on-prem and enterprise, validate it with research partners against expert references, and build go-to-market. The vision: a research environment where every conclusion carries its evidence — and no answer outruns what can be shown.

perslis.com · [research collaborations](#) →