



10 August 2026

## PROFESSIONAL ASSESSMENT OF PR4XIS

### To Whom It May Concern:

I am the President of tbay.tk LLC, a software company in Helena, Montana. We build kôdex, a platform that orchestrates and supervises AI coding agents. Over the past several months we conducted a detailed technical evaluation of the reasoning engine **pr4xis**, to determine whether it belonged in our own product. I am writing to describe what we found it can do. I offer this assessment freely and without compensation.

pr4xis answers questions and *shows its work*. Where a conventional AI system gives you an answer you must take on faith, pr4xis gives you the answer together with the complete trail of established facts it rests on. When it is right, you can audit exactly why. When it is wrong, it names the specific assumption that failed. That proved to be the difference that matters; an engineer can act on a named fault immediately, where a vague warning only starts an investigation.

Four capabilities stood out as genuinely uncommon:

- **It gives the same answer every time.** Identical questions produce identical answers and identical reasoning, run after run. That makes its output usable as a record; it can be filed, re-checked months later, and relied upon in an audit or a regulatory submission.
- **It declines to guess.** When a question falls outside what it can establish, pr4xis says so and states what it could not resolve, rather than producing a confident-sounding answer. For safety-critical work this is the most valuable property a system can have, and the one most commonly absent.
- **It finds the gaps in its own knowledge.** pr4xis can compare two bodies of knowledge and identify, on its own, distinctions that one of them cannot express; this is not error-checking, but the discovery of blind spots nobody knew to look for. The project documents a case in biological research where this surfaced a meaningful distinction the existing model had collapsed and obscured.
- **It grows without breaking.** New subject areas can be added and connected to what pr4xis already knows, and each connection is checked for soundness rather than assumed. Carefully assembled knowledge keeps its value as the system expands, instead of degrading as more is piled on.

The breadth already built is substantial. At the time of our review pr4xis carried more than one hundred and eighty distinct subject areas, each traced to the published literature it is drawn from, spanning navigation and sensor systems, aerospace and orbital mechanics, autonomous vehicles, industrial process control, biomedical science, United States law, and the English language itself. All of it is verified continuously by an automated suite of several thousand checks, and it runs on ordinary hardware (including entirely inside a web browser).

My professional judgment is that pr4xis is serious, carefully built work of a kind I do not often encounter. It addresses a real and widely acknowledged weakness in current artificial intelligence (that these systems cannot reliably tell you when they are wrong), and it does so with unusual rigour and intellectual honesty. It is well suited to any setting where an answer must be defensible rather than merely plausible: aerospace and vehicle autonomy, medical decision support, industrial safety, regulatory compliance, and the checking of work produced by other AI systems. My own interest is the last of these, and I regard the approach as sound.

I would be glad to expand on any of the above; please feel free to contact me directly.

With best regards,

**Alexander William Forrest Milton**

President, tbay.tk LLC  
Helena, Montana · alex@tbay.tk