Zain Dana Harper

Seattle, WA (Local / Hybrid / Remote Open) | zaindharper@gmail.com | harperz9.github.io
github.com/HarperZ9 | linkedin.com/in/zaindanaharper | orcid.org/0009-0001-7175-5393

AI Evaluation, Research Synthesis & Systems Direction

November 26, 2026 | Application: AI Evaluation, Research Synthesis & Systems Direction

Dear Hiring Lead,

My best work happens where architecture, empirical verification, and cross-domain synthesis meet. Rather than treating AI systems as black-box demos, I specify tight invariants, direct multi-agent and human workflows against deterministic oracles, and publish both what breaks and what holds. Recently I trained and evaluated JFC3P-7B-Oracle-v1 (4.68 GB Q4_K_M), lifting adversarial trap rejection from 16.67% to 95.83% (+33.33 pp, Fisher exact p = 4e-6) under Arm C, while my Three-Plate verification suite caught 115/115 planted defects with 0 false accepts where single-plate checks caught 6.09%.

Across five consolidated open-source platforms and 159 public repositories (including bio-research and crypto-margins), I connect deep technical infrastructure with real-world grounding: the 12-target BuildLang 1.15.0 compiler (Universal Hourglass lifters, V5 e-graph superoptimizer at 1.88x to 2.59x -O3 with 0.0 error, 1.18M-record corpus engine in 156 ms / 86 MB RSS, and BuildDB qdb.exe at 2.89 MB / 0/500 fuzz); Rowan unified local/subscription/API routing (46.73% AVTR-4R token savings); RAW DX12 Wave 72 (522/522 runtime tests, 14/14 DXC SM 6.10 shaders) and Elder ENB (1,004,693 downloads); CMMNS (1,735 tests); fail-closed 1-command installers; and a 14.45M-file / 6.01M-word FTS5 (3.66 ms) multimodal engine. Across 30 empirical studies (13 Zenodo DOIs), I published 12 null results openly to guard against metric overfitting.

I also bring practical team reliability from outside tech: 11 years leading commercial tree-care crews under high-hazard conditions taught me how to communicate clearly, verify load paths before cutting, and stay calm under pressure. In open source, external maintainers across 34 repositories (DeepEval, Pydantic AI, AgentFence, python-grimp/grimp, Datasette, TOMLKit, Biome) have merged 41 of my pull requests.

I do my strongest work on real architecture, evaluation, and synthesis problems rather than timed trivia quizzes. Every system, benchmark, and cryptographic receipt is public at harperz9.github.io and github.com/HarperZ9. I would welcome a 15-minute live walkthrough of a work sample or a scoped, practical evaluation task directly relevant to your roadmap.

30-Second Work-Sample Proof Strip

- **Verification & Model Proof:** JFC3P-7B-Oracle-v1 (95.83% trap reject, p = 4e-6), NOH-4 (78.0x recovery win), Three-Plate (115/115, 0 FP).
- **Systems & External Proof:** 12-target BuildLang + V5 e-graph (1.88x-2.59x -O3), 1.18M-record engine (156 ms / 86 MB), Rowan router, 159 public repos, 30 studies (12 nulls), and 41 merged external PRs.

Best regards,

Zain Dana Harper
zaindharper@gmail.com | harperz9.github.io | github.com/HarperZ9