

TrojanSpec-Bench


Adversarial Specification Elicitation
in AI-Assisted Formal Verification

Mohammad Zeeshan

Eotvos Lorand University, Budapest

Apart Research | Spec Pipelines Hackathon | May 2026

AI is now writing the specifications that verifiers trust



Verifier proves I satisfies S. But what if S does not capture what the user actually asked?

Real-world example: Kobeissi's libcrux audits (IACR ePrint 2026/192 and 2026/670) document **13 vulnerabilities** inside verified post-quantum cryptography. Every bug was found by hand. All passed the verifier.

What if the spec elicitor is the attacker?

Dolev-Yao threat model applied to the spec elicitor:

the strongest reasonable adversary sits at the boundary between human intent and machine-readable contract.

An adversarial elicitor produces a specification that:



Passes every verifier check



Passes mutation testing



Passes single-shot LLM-as-judge alignment review



Is satisfied by a backdoored implementation

The bug lives in the gap $R \setminus S$, where no verifier can see.

Four attack patterns model real spec-elicitation failures

Vacuity

Postcondition is logically equivalent to true

Trivially satisfiable

Implementation Leak

Postcondition delegates to an unverified axiom

Hides logic behind opaque assumption

Domain Restriction

Precondition silently excludes the dangerous input

Anchor: libcrux ML-KEM $d \leq 1$

Predicate Swap

An operator or constant flipped while paraphrase still reads correct

Anchor: Ed25519 double-clamping

All four attack classes are anchored to bugs documented in real cryptographic disclosures.

We built a verifier-admitted trojan benchmark across three languages

1,024

verifier-admitted
trojan triples

Every triple passes a dual-property
verifier admission test:

- (a) verifier accepts witness against trojan spec
- (b) verifier rejects witness against honest spec

Coverage


Dafny	319 admitted
Lean 4	417 admitted
Verus	288 admitted

Cryptographic anchors

Thirteen primitive families modeled on libcrux disclosure classes: ML-KEM, ML-DSA, SLH-DSA, Ed25519, X25519, AES-GCM, ChaCha20-Poly1305, SHA-3.

Generators: Bedrock Sonnet 4.6 (primary); Haiku 4.5 and Llama-3.3 70B as cross-family ablation.


The defender decomposes the LLM judgment into four atomic criteria



Why $K \geq 2$: a real trojan typically fails more than one criterion; a single noisy criterion does not flip a faithful spec to trojan.

Architecture inspired by FormalJudge (arXiv:2602.11136) and Epistemic Ensemble (arXiv:2506.10903) decomposition.

Atomic decomposition lifts detection without hurting recall



What changes

Recall held at full coverage

Every trojan in the test set is caught. The detector adds no missed-detection cost.

False positives drop sharply

The $K \geq 2$ vote suppresses single-criterion noise that the coarse baseline let through.

Cost stays modest

Four LLM calls per spec, deployable in a pre-commit hook.

SSC baseline: three-judge consensus voting, the published Specification Self-Correction protocol.

Static mutation testing structurally misses three of four attack classes

Vacuity

Implementation leak

Domain restriction

Predicate swap

MutDafny (static mutation)	Atomic $K \geq 2$ (ours)
 Caught	 Caught
 Missed	 Caught
 Missed	 Caught
 Missed	 Caught

Why this is structural

Mutation testing asks: can any mutant of the spec still be proved? A vacuous postcondition fails this test, so vacuity is caught.


But an axiom leak, a tighter precondition, or a flipped predicate each remain trivially provable under mutation. The verifier has no honest oracle to compare against.

Static mutation testing of specifications is structurally a vacuity detector.

Head-to-head on 319 admitted Dafny trojans. Baseline: MutDafny (Amaral et al., ICSE 2026).


The detector stays tight on real formal mathematics

Atomic $K \geq 2$



3 of 100 honest Mathlib lemmas flagged

Monitor consensus baseline



23 of 100 honest Mathlib lemmas flagged

Why this matters

A common worry about LLM-as-judge detectors: do they fire on legitimate proofs?

Atomic decomposition tightens precision without losing any true positive on the benchmark side.

Strict subset:

every Mathlib lemma the atomic detector flags is also flagged by the coarse baseline. Atomic narrows, never widens.

Honest limitations and the next experiment

Same-family attacker and judge

Sonnet 4.6 generated most of the benchmark and also serves as the atomic judge. Cross-vendor partition is the top follow-up.

Atomic detector not yet stressed adaptively

We measured static-detector fragility under adaptive attack. Running the same protocol against the atomic $K \geq 2$ detector is queued.

Mathlib evaluation is preliminary

100 honest lemmas give a useful first signal. A 1,000-lemma scale-up is on deck.

Generator-shaped distribution

Per-class proportions reflect the elicitor's failure modes, not an empirical bug-frequency survey of real verified code.

Treat the spec elicitor as adversarial

1

A new benchmark for a new threat

1,024 verifier-admitted trojan triples across Dafny, Lean 4, and Verus, anchored to real cryptographic disclosures.

2

Atomic decomposition is the deployable defense

Four Yes/No criteria with a $K \geq 2$ vote beat a coarse LLM judge on F1, drop FPR sharply, and validate on real formal math.

3

Static mutation testing is not enough

It catches one attack class out of four. Pair it with an atomic-criteria semantic check.

Mohammad Zeeshan hdglit@inf.elte.hu

Code: github.com/m-zest/TrojanSpec-Bench-v1.0
demo

Dataset: huggingface.co/datasets/m-zest/trojanspec-bench

Demo: huggingface.co/spaces/m-zest/specguard-

References

- [1] Kobeissi, N. Verification Theatre: Auditing formally verified post-quantum cryptography in libcrux. IACR ePrint 2026/192.
- [2] Kobeissi, N. Verification Facade: The SampleNTT axiom gap and other libcrux failures. IACR ePrint 2026/670.
- [3] Amaral, I., Mendes, A., Campos, J. MutDafny: A mutation-based approach to assess Dafny specifications. arXiv:2511.15403, ICSE 2026.
- [4] Zhou, J., Sheng, Y., Lou, H., Yang, Y., Fu, J. FormalJudge: A neuro-symbolic paradigm for agentic oversight. arXiv:2602.11136, 2026.
- [5] Zhang, L., Valentino, M., Freitas, A. Beyond gold standards: Epistemic ensemble of LLM judges for formal mathematical reasoning. arXiv:2506.10903, 2025.
- [6] Galle, A. et al. Specification self-correction: Mitigating in-context reward hacking through test-time refinement. arXiv:2507.04500, 2025.
- [7] Greenblatt, R. et al. AI Control: Improving safety despite intentional subversion. arXiv:2312.06942, 2023.
- [8] Dolev, D., Yao, A. C. On the security of public key protocols. IEEE Trans. Info. Theory 29(2), 1983.
- [9] Beer, I., Ben-David, S., Eisner, C., Rodeh, Y. Efficient detection of vacuity in temporal model checking. Formal Methods in System Design 18(2), 2001.