

THD-AI-001 v1.0 counterclaim

Rick Gaur | Provable Industries Inc. | admin@provableintelligence.com
Prepared September 11, 2026. Candidate for technical review.

Result and scope

The completed registered schedule-allocation experiment selects P48 in PRIMARY and confirms it against P23, P24 and P25 under 24 fresh REPLICATION seeds. All 216 PRIMARY and 96 REPLICATION conditions completed. The exact statistical falsifier is satisfied on the disclosed Apple MPS/shared-prefix implementation. Acceptance of that implementation and sponsor-side independent reproduction remain unresolved.

P48 minus	PRIMARY gain (pp)	REPLICATION gain (pp)	Replication p
P23	0.93125	0.80167	1.7180e-16
P24	0.81917	0.70708	7.0579e-15
P25	0.82667	0.70500	6.8511e-12

PRIMARY: P32 and P48 pass the registered intersection-union/Holm procedure across six challengers. The frozen selection rule chooses P48. P48 PRIMARY intersection-union $p = 1.8703578394e-13$. REPLICATION: maximum one-sided paired $p = 6.8511398984e-12$, below 0.05. Gains are percentage points, not relative percentages.

Exact claim and contradiction

THD-AI-001 v1.0 predicts that none of the six registered outside schedules reproducibly outperforms all three target schedules in the frozen CIFAR-10/ResNet-20 system. P48 is registered outside the target window, qualifies in PRIMARY, is frozen before REPLICATION, and again beats all three targets under the same decision rule. Conditional on the execution satisfying the registered premises, this contradicts that specific prediction. No claim is made about all THD propositions or every possible schedule.

Claim identification and timing

The official claim page identifies THD-AI-001 version 1.0 as ACTIVE on September 11, 2026. The exact original publication date is not displayed on the inspected claim page and needs sponsor confirmation; the supplied registered specification was preserved before completion. The result completed September 11 at 14:52:48 UTC. This document does not assert that preparation constitutes receipt or freezes the applicable rules.

Official claim: <https://creationunified.com/thd-ai-001-v1-0>
Bounty rules: <https://creationunified.com/bounty>

Premises and implementation

Frozen experimental system

System boundary: supervised CIFAR-10 image classification with CIFAR-ResNet-20. Inputs are training images, labels, registered paired seeds and the schedule P. State consists of model and BatchNorm parameters, optimizer momentum and random states. Outputs are final confirmatory top-1 accuracy and cross-entropy. Training feedback is SGD on the Optimization Set; the confirmatory endpoint is excluded from parameter optimization.

Data: official CIFAR-10 training images, deterministically ranked within class by SHA-256 with salt THD-AI-001-ISSP-PROBE-v1. Per class, 2,000 images form the 20,000-image Optimization Set, the next 500 form the excluded Diagnostic Set, and the next 1,000 form the 10,000-image Confirmatory Holdout. Remaining training images and the official test split are not used for adjudication. Split hashes, index audit and integer correct counts are supplied.

Model and training: CIFAR-ResNet-20, stage widths 16/32/64, three BasicBlocks per stage, adaptive pooling and 10 outputs. FP32 SGD, momentum 0.9, weight decay 0.0005, batch size 256, and 102 logical epochs per condition without early stopping. Six warm-up epochs increase learning rate from 0.001 to 0.1; P plateau epochs follow, then D=96-P epochs decay to 0.001. Paired seeds determine initialization, order, crops and flips independently of P.

Mapping the general bounty criteria to this claim

The registered performance metric is final top-1 accuracy, rather than an additional invented coherence metric. The tested intervention changes plateau/decay allocation under a fixed training budget; it acts through the learning-rate update during actual optimization. The activation and persistence requirements are the registered PRIMARY selection and fresh-seed REPLICATION rules, with 102 epochs per condition. This claim supplies a statistical schedule prediction, not extra independently measurable Integration/release/transition pathways. The resolution audit therefore compares P48 to each of P23/P24/P25 under the four explicit falsifier conditions, without substituting generic THD criteria.

Material execution differences

Execution used Python 3.12.14, PyTorch 2.8.0, torchvision 0.23.0, FP32 Apple MPS on arm64 macOS 26.5.2, deterministic algorithms enabled, CPU fallback disabled and MPS fast math disabled. The supplied reference training code constructs default-CPU tensors. The inspected claim specifies FP32 but does not expressly prohibit MPS. The saved checks do not establish CPU/MPS trajectory equivalence for all seeds.

Schedules share identical pre-divergence prefixes. Branching retains model/BatchNorm state, optimizer momentum and random-state handling. Frozen source, pilot/prefix checks and recovery receipts are retained. This is not 312 independently executed, untouched reference-code jobs. The ledger records 24,288 nominal physical epochs, 24,303 completed attempts including 15 recovery repeats, and two interrupted attempts with unknown work; there are 31,824 logical condition-epochs.

Evidence and reproduction

Review package

The companion ZIP contains all 312 per-condition logical result records, PRIMARY and REPLICATION decisions, split/runtime/source records, frozen execution source, sponsor package code/configuration, and the three completed audit reports. SHA256.json binds every included payload. Unmodified source files remain distinguishable from this presentation document.

Logical/statistical audit: all conditions, fixed partitions, integer endpoints, complete-stage selection and fresh-seed replication checked. The collection project reran the frozen logical checker successfully on copied completed records. This rerun trains no model. Provenance audit: 48 seed trees and physical-work ledger reconciled. Checkpoint audit: 576 retained payloads checked for structure, finite tensors, optimizer/random-state representations and internal-log consistency without executing a model.

Recompute the saved logical/statistical result

In a separate Python environment, install `mpmath==1.3.0`. From the extracted evidence directory run:

```
python verify_thd_registered_logical_results.py --results thd-shared-registered-results
```

The expected status is `COMPLETE_REGISTERED_LOGICAL_EVIDENCE_VERIFIED` and the replication decision is `SCIENTIFIC_FALSIFIER_SATISFIED`. Python 3.12 was used for this check.

Independent training reproduction and large evidence

For a fresh reference-code run, use `thd-package/README.md` and `requirements.txt`, validate the package, then run PRIMARY and freeze the challenger chosen by its decision before running REPLICATION. This is a separate reproduction and need not have the same trajectories as the MPS implementation. The exact executed shared-prefix runner and dependency locks are supplied for review, but its full execution gates require the retained pilot/prefix data and additional large artifacts.

The compact ZIP intentionally omits the 576 checkpoint binaries, the full physical-ledger tree and CIFAR-10 dataset. These remain preserved in the original completed result tree. The supplied audit reports bind relevant hashes; these summaries are not a substitute for sponsor access to the underlying artifacts. Please provide the preferred transfer method so the large evidence and gate artifacts can be delivered for independent review. No external download link has yet been published.

Requested adjudication

Please acknowledge receipt against THD-AI-001 v1.0 and identify the applicable bounty-rule version. Please assess the disclosed MPS/shared-prefix implementation against the frozen premises and specify any further reproduction or evidence required. Sponsor verification and acceptance are not established. The advertised USD 10,000 is requested only if the counterclaim is verified as qualifying; payment arrangements can follow review.

Primary sources checked September 11, 2026: the official claim and bounty pages above, and <https://creationunified.com/contact-us> (Bounty submissions option). The public ledger lists two older closed/unclaimed submissions and no verified falsifications; private reports are unknown. The current registry page could not be refreshed separately.