Evaluating Governed LLM-Driven ML Experimentation: Lifecycle, Provenance, and Failure Metrics*

Dowling Wong
dowling.wong@kit.edu
Karlsruhe Institute of Technology
Karlsruhe, Germany

arXiv:2611.23678v1 [cs.AI] 18 Nov 2026

Abstract

Autonomous agents that run ML experiments are usually judged by the best score they reach. That score says nothing about how it was reached: whether the agent stayed within the files it was allowed to edit, whether failed trials were kept or silently dropped, and whether each keep or discard decision can be traced to its evidence. We propose evaluating such systems on their experimental process as well as their outcome, and we build a governed harness that makes this possible. A control plane, not the LLM that proposes changes, owns each trial's lifecycle, checks its edit scope, decides keep or discard, and appends one record per trial to an append-only ledger. Governance metrics such as the invalid-action rate, the repeated-bad rate and provenance completeness are computed from that ledger. Across six ML tasks, from synthetic logistic regression to language-model training, and 1,445 trials, the same contract applied unchanged. The ledger exposed failures that a task score hides: a no-memory arm in which all 120 trials failed, a benchmark adapter whose bounded proposals ran out, and a reset failure that reproducibility hashes revealed, after which the affected campaigns were rerun. In a matched counterfactual, a loop that logs only successful trials left no record of 43 of 50 trials, while the governed ledger recorded all 50. Memory ablations show that memory's effect on repeated bad proposals depends on the task. We release the ledgers of all 1,445 trials, so that every number reported for them can be recomputed.

CCS Concepts

• **Computing methodologies** → **Artificial intelligence**; • **Software and its engineering** → *Software testing and debugging*; • **General and reference** → *Evaluation.*

Keywords

agentic AI evaluation, autonomous ML experimentation, governed harness, control plane, append-only ledger, memory ablation, failure taxonomy, reproducibility, bounded execution

1 Introduction

Suppose an LLM agent that tunes a training script reports that validation AUC rose by 0.05. The score does not say how the gain was obtained. The agent may have edited a file it was not supposed to touch, a run that crashed may have been dropped from the record, or the same bad idea may have been tried many times before one variant worked. For scientific ML, where an experiment is only as useful as the evidence behind it, these questions matter as much as the score: an autonomous experiment needs *auditability*, a traceable and reproducible record of every decision and its evidence.

Current practice does not provide this. ML-agent benchmarks such as MLAgentBench and MLE-bench ask whether the agent improved the task [5, 11]; agent systems such as AIDE keep their search state inside the same loop that makes the decisions [14]; and experiment trackers record runs after they exist, without governing which runs are attempted or how failures are classified [2, 29]. What is missing is a way to measure the experimental process itself, and a record of that process that the agent does not control.

We address both with a governed harness. Following the view of an agent as a model plus a harness [24], which together form a compound AI system [30], we separate three roles. A manager, typically an LLM, proposes one bounded change per trial; a worker applies it and runs training; and a control plane owns everything else: the trial's lifecycle, the check that the patch touches only editable files, metric parsing, the keep or discard decision, and one append-only ledger record per trial, written whatever the outcome. None of these is delegated to the agent, so the manager can be replaced without changing what is recorded. From the ledger we compute six governance metrics alongside the task metric: acceptance rate, invalid-action rate, repeated-bad rate, provenance completeness, artifact-reference completeness, and a failure taxonomy (Section 4.4).

We evaluate this contract on six ML tasks over 1,445 trials: synthetic logistic regression, two public OpenML tabular tasks, an MLAgentBench vectorization task, a ResNet classifier for a scientific trigger task, and GPT pretraining. Four findings stand out (Figure 1 gives the per-node overview). First, the same contract applied unchanged across tasks, worker types and failure regimes. Second, the ledger exposed failures that a task score hides: a no-memory arm in which all 120 trials failed, a benchmark adapter whose bounded proposals ran out, and a reset failure that node-state hashes revealed, after which the affected campaigns were rerun. Third, a matched counterfactual shows what is lost without the control plane: a loop that logs only successful trials lost 43 of 50 trials, and a campaign in which every trial fails leaves an empty ledger, indistinguishable from one that never ran (Section 5.3). Complete provenance, which the design guarantees (Section 3.4), held in every trial. Fourth, memory ablations used as diagnostic probes show that memory's effect on repeated bad proposals depends on the task: summary memory gave the lowest repeated-bad rate in two of three ResNet replicates, but not on the synthetic task (Section 5.4).

*Presented at the KDD 2026 Workshop on Evaluation and Trustworthiness of Agentic AI, Jeju, Korea, 9 August 2026. The workshop is non-archival: its papers are posted on the workshop website and are not part of the KDD proceedings. This version corrects the workshop version: it rewords several claims, reports the starting point of the L40S replication, corrects and adds references, updates the artifact-availability statement, and includes the supplement as appendices.

(a) Per-node scorecard (%)

Node	Camp.	Trials	Kept	Disc.	Fail	Prov.
ResNet-trigger	16	245	44	51	5	100%
LR-synthetic	15	450	8	45	46	100%
OpenML credit-g	9	230	13	60	26	100%
OpenML bank	9	230	13	55	32	100%
MLAgentBench	3	50	28	24	48	100%
Autoresearch Linux	8	240	6	24	70	100%

(b) Decision breakdown per node

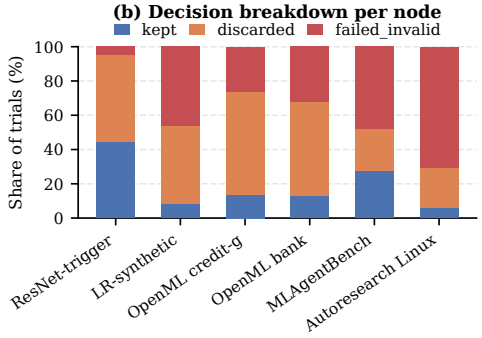


Figure 1: Governance overview across all six benchmark nodes (1,445 trials). (a) Per-node scorecard: number of campaigns (Camp.), trials, and the shares of kept, discarded and failed-invalid trials, with provenance completeness (Prov.); the Fail share is shaded amber below 20% and red at or above 20%. (b) The same decision shares as stacked bars. Provenance completeness is 100% on every node regardless of task type, worker class, or failure rate—including the 120-trial total-failure arm on `autoresearch_linux`.

Our contributions are:

- an evaluation protocol that reports governance metrics next to task metrics, with each metric assigned by the control plane rather than reported by the agent (Section 4);
- a control-plane design that separates proposal, execution, validation and keep/discard authority, so that the LLM manager is a replaceable component (Section 3);
- an empirical test across six tasks and 1,445 trials, with the ledgers of all 1,445 trials released so that every number reported for them can be recomputed (Section 5 and Appendix A.1).

Section 2 reviews related work, Section 3 describes the control plane, Section 4 gives the evaluation protocol, Section 5 reports the results, and Section 6 discusses limitations.

2 Background and Related Work

The key distinction of this paper is that we evaluate the *harness* that governs the agent loop—lifecycle states, scope enforcement, provenance, and failure classification—not only the model that proposes edits.

ML-agent benchmarks and code-space exploration. AIDE [14] is the closest academic system: both use LLMs to propose atomic code changes evaluated by a hard-coded metric, but they differ fundamentally in where lifecycle authority lives. AIDE is a coupled optimisation loop — search policy, evaluator, and solution summarisation share state inside a single process — and keeps its solution tree in memory. In the released code (commit `ad73d22`, checked in November 2026) [26], the main loop saves a snapshot of the whole journal after every step, so a crash loses only the step in progress; but the snapshot is overwritten rather than appended, there is no pending-trial guard, and a failed node carries a buggy flag (set from an LLM review, an exception, or a missing metric) and the exception type rather than a control-plane failure category. AIDE generates whole solution scripts, so it has no editable-path whitelist or scope enforcement. The present work inverts this design: the control-plane contract is fixed and externally visible; the manager and worker are replaceable without altering governance; every outcome—whether kept, discarded, or crashed—produces an immutable typed audit record before the next trial begins. The editable-path whitelist is a precondition guard, not a post-hoc filter.

MLAgentBench [11], MLE-bench [5], FML-bench [31], and MLGym [20] address a different evaluation target: they ask *did the agent improve performance on this task?* This paper asks *did the harness guarantee that every trial was bounded, classified, and provenance-tracked, regardless of outcome?* The two questions are complementary rather than competing. This paper wraps an MLAgentBench-style task (NumPy vectorization speed) as the `mlagentbench_vec.` node, demonstrating that any task expressible as a NodeSpec receives the same governance coverage, but the governance evaluation is independent of task-metric outcome. SHARP [3] governs reproduction of *existing* experiments; this paper governs autonomous *generation* of new ones.

Experiment tracking, AutoML, and NAS. Trackers such as MLflow and Weights & Biases [2, 29] record outcomes after a run exists but do not govern the agent loop or classify invalid edit attempts before training. AutoML and neural architecture search systems [8, 13] optimise within predefined search spaces; they are not designed to audit self-directed exploration, explain why a proposed change was kept or discarded, or expose lifecycle diversity. This paper makes the opposite tradeoff: a narrower set of controlled nodes with complete per-trial provenance.

Harness engineering and governance. This framing is consistent with practitioner writing on harness engineering [4, 9, 18, 24, 28]: Trivedy frames the harness as everything around the model, including tools, sandboxes, memory, middleware, and execution hooks. Böckeler argues that trust in agent output requires constraining the solution space with computational guides and feedback sensors. This paper is not another task benchmark; it is a governance layer into which any benchmark task can be wrapped via a NodeSpec contract. The node specification is a computational guide; the pending-trial guard and state machine are computational sensors; memory injection is an inferential guide; the repeated-bad detector is an inferential sensor. Agent memory and scaffolding frameworks [7, 23] inform the design of context-resident memory

injection and reflective self-improvement, but do not provide an evaluation protocol for governed autonomous experimentation.

3 Governed Control Plane

The system is a governed harness for autonomous ML experimentation. The manager proposes an experimental change; the worker materialises a patch and run artifacts. The control plane validates scope, executes the lifecycle, parses metrics, assigns a kept, discarded, or failed-invalid decision, updates memory, and appends the authoritative ledger record. Managers and workers are replaceable components; the control plane contract is fixed.

A worked trial. Before the full abstraction stack, we trace one governed trial end to end, drawn from the ResNet-trigger case study of Section 5.1. (i) *Proposal:* the manager emits one bounded change — summary, rationale, target parameter, intended effect — and the control plane assigns it a proposal provenance ID. (ii) *Scope validation:* the worker materialises a patch touching only `train.py`, the single whitelisted file; a patch touching frozen data or split logic would terminate here as `invalid_edit_scope`, before any training runs. (iii) *Execution under guard:* the control plane writes a pending-guard file, then the worker runs training; a crash here would leave a stale guard that forces a `failed_invalid` record at the next campaign start rather than a silent gap. (iv) *Metric parse and decision:* the parser extracts `val_auc` and the control plane — not the manager — compares it against the current best (this campaign improves 0.7747 → 0.7827 over 20 trials); an improvement is `kept`, a valid non-improvement `discarded`. (v) *Ledger append:* exactly one immutable JSONL record is appended — proposal text, patch reference and hash, raw-log reference, parsed metric, decision, failure category if any, and five provenance IDs (proposal, patch, run, metric, and decision IDs, each derived by hashing the trial ID) — the guard is deleted, and the next budget slot may open. All 1,445 trials in this paper, including all failures, follow exactly this path; the only variation is which guard or comparison produces the terminal state.

To make auditability concrete, we operationalise it through the five desiderata in Table 1. These are not claimed to exhaust every possible failure mode of agentic experimentation, nor to be statistically independent axes. They define the minimum operational conditions needed for the claims in this paper to be inspectable. In particular, complete attempt accounting and provenance linkage are related but distinct: the former asks whether every attempted budget slot appears in the ledger, while the latter asks whether each recorded attempt carries enough evidence to inspect it.

Table 1: Five auditability desiderata for governed agentic experimentation.

Desideratum	Operational meaning
D1. Bounded edit scope	The agent cannot change protected benchmark, data, or evaluation files.
D2. Complete attempt accounting	Every proposed attempt is recorded, including invalid proposals and crashes.
D3. Independent adjudication	Acceptance or rejection is decided by fixed control-plane rules, not by the agent.
D4. Provenance linkage	Each recorded attempt links to patch, configuration, run, metric, and decision evidence.
D5. Recomputability	Aggregate governance claims can be recomputed from the ledger without re-running training.

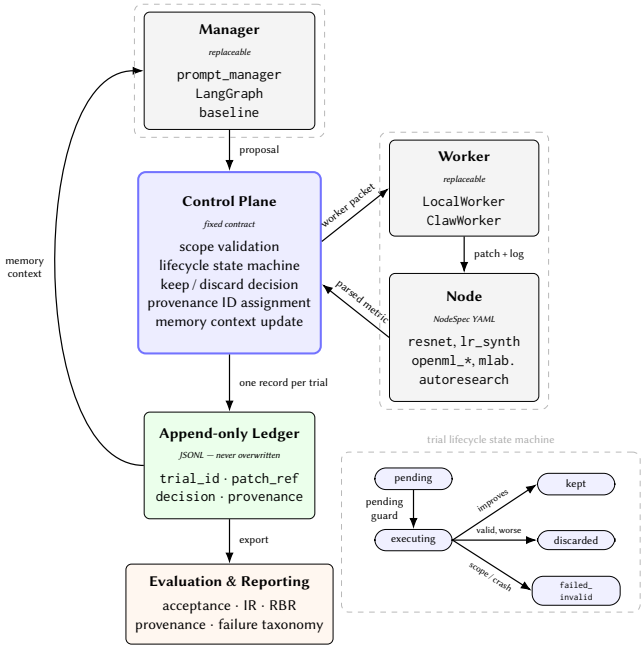


Figure 2: System architecture and trial lifecycle. *Replaceable components* (manager, worker, node) are interchangeable without changing the governance contract; the control plane contract is fixed across all six nodes and every manager backend tested in this paper. The control plane owns scope validation, lifecycle transitions, keep/discard decisions, provenance ID assignment, and memory context updates. Inset (bottom right): the trial lifecycle state machine (Section 3.4).

3.1 Alignment with Agent Evaluation Criteria

Table 2 maps the Evaluation and Trustworthiness of Agentic AI workshop concerns to the harness mechanisms and metrics reported in this paper.

3.2 Manager, Worker, and Replacement Principle

The manager generates proposals (summary, rationale, target, intended effect). The worker executes bounded actions (patch + run artifacts). The control plane owns authority: validity, state transitions, budget, memory updates, provenance IDs, and final decision. This separation is the *replacement principle*: any manager backend — baseline heuristic, `prompt_manager`, `LangGraphManager`, or an AIDE-style tree-search — can be swapped without changing the governance contract. Two worker implementations span the reproducibility spectrum: `LocalWorker` applies "Change PARAM from X to Y" directives with no external dependencies; `ClawWorker` invokes an external coding-agent path for free-form proposals, routing structured proposals through a deterministic constant-patch bridge that requires only Ollama [21].

Table 2: Workshop evaluation concerns mapped to harness mechanisms.

Evaluation concern	Harness mechanism	Reported metric
Repeated poor behaviour	Append-only memory + repeated-bad detector	Repeated-bad rate
No ground truth for intermediate choices	Decision authority separated from manager; provenance IDs per trial	Provenance completeness
Deployment and scope safety	NodeSpec whitelist, no-op guard, pending-trial guard	Invalid-action rate
Long-horizon failure accumulation	State machine; terminal `kept/disc./failed_inv.`	Failure taxonomy
Benchmark reproducibility	NodeSpec YAML, deterministic patch bridge, seeded replicates	Artifact manifest

3.3 Benchmark Node Inventory

Table 3 lists the six nodes reported in this paper. The control-plane contract is identical across all nodes: the only per-node configuration is the NodeSpec YAML (editable paths, frozen paths, run command, metric parser, and budget range).

3.4 Trial Lifecycle

Each trial follows the fixed lifecycle shown in Figure 2 (state machine inset). Every opened budget slot must produce a `kept`, `discarded`, or `failed_invalid` audit object.

The control plane writes a pending guard before worker execution, making mid-trial crashes auditable. A scope validator checks the patch against the NodeSpec whitelist before training begins: patches touching frozen files are rejected, recorded, and prevented from corrupting node state. The append-only JSONL ledger stores, per trial: proposal summary and rationale, patch and log references, parsed metrics, validity status, decision, failure category, provenance IDs, and reproducibility hashes.

Table 4 formalises the lifecycle.

Lifecycle invariant. A trial begins in PENDING, then enters EXECUTING. It ends in exactly one terminal state: KEPT, DISCARDED, or FAILED_INVALID. The control plane maintains the following invariant:

> *For every budget slot that enters EXECUTING, the control plane appends exactly one terminal record—*`kept`*,* `discarded`*, or* `failed_invalid`*—to the append-only ledger before the next slot may open. No slot is silently skipped; no slot produces more than one record; no record is overwritten or deleted.*

Enforcement has two paths. *Normal path:* the pending guard is written before the worker starts and deleted on the terminal transition; any deviation raises an error before the ledger append. *Crash-recovery path:* if the process crashes mid-trial, the guard file persists; at the next campaign start `PendingTrialError` is raised, the stale guard is inspected via `scripts/recover_pending.py`, and a `failed_invalid` record is appended before any new trial begins. Neither path allows a budget slot to close without a ledger record.

The invariant is structural rather than empirical: it holds by construction of the control-plane code, not because nothing crashed during the reported experiments. Complete provenance is accordingly observed in all 1,445 trials, as the invariant requires.

3.5 Failure Taxonomy and Metric Validation

Failure labels are assigned by the control plane at the terminal transition that writes the ledger record, not post hoc by the manager. The reported evidence uses five observed `failed_invalid` categories: `invalid_edit_scope`, `no_op_patch`, `proposal_precondition_failed`, `runtime_error`, and node-level `invalid_config`. The schema also supports `metric_missing` and `edit_failed`; those labels do not appear in the primary 1,445-trial evidence set. Valid non-improvements receive `discarded`, not `failed_invalid`; the label `degraded_metric` on their records marks a discarded trial, not a failure. The taxonomy is exhaustive for the current NodeSpec contract because every terminal transition is produced by one of those guards or by a parsed metric comparison. New nodes may add subcategories only if they map to the same three terminal outcomes before the ledger append. Structural validation uses forced stress campaigns, paired deterministic reruns, kept-artifact re-execution, and rule-based relabeling from observable ledger fields. Readers who want the taxonomy grounded before the aggregate results should consult Appendix E, which gives one concrete ledger example per category with its detection point and guard type.

3.6 Memory Modes and Decision Authority

The memory interface has three modes: no prior-trial context, append-only summary, and rationale-augmented summary. The rationale arm adds per-trial failure rationales and can truncate them with `rationale_max_tokens`. The ablation is a diagnostic probe: it tests whether the manager changes behavior when exposed to prior governed outcomes, and whether repeated-bad rate can detect that behavioral change. It is not treated as evidence that any memory format is generally superior.

The keep/discard decision lives in the control plane, not the manager. A valid candidate is kept only if the configured metric improves over the current best; a valid non-improving candidate is `discarded`; an invalid candidate is `failed_invalid` with a specific failure category. This deterministic, externally-owned decision rule follows practitioner guidance for long-running agents: harnesses should provide review loops and external checks rather than letting the model own success criteria [22].

4 Evaluation Protocol

The evaluation asks whether an autonomous ML experimentation loop is bounded, auditable, failure-aware, reproducible, and behaviorally affected by governance memory. Governance metrics are primary; task-metric AUC is secondary evidence that the loop runs end-to-end.

4.1 Benchmark Nodes

The primary scientific benchmark is `resnet_trigger`, a near-threshold detector waveform classification task with frozen H5 corpora, split logic, preparation code, and node metadata; workers may edit only `nodes/ResNet_trigger/train.py`. The synthetic transfer benchmark, `lr_synthetic`, is a pure-NumPy logistic-regression classifier run via `LocalWorker` for dependency-free protocol validation; its data-generation constants are frozen.

Table 3: The six benchmark nodes. All nodes share the same control-plane contract; only the NodeSpec YAML differs. Direction: max = maximise metric, min = minimise.

Node	Task type	Worker	Metric	Dir	Purpose in evaluation
resnet_trigger	Waveform classification	Claw	val_auc	max	Primary case study; memory ablation substrate
lr_synthetic	Logistic regression	Local	val_score	max	Cross-node transfer check; dependency-free
openml_credit_g	Tabular classification	Local	val_auc	max	Governance transfer to public tabular node
openml_bank_mkt	Tabular classification	Local	val_auc	max	Governance transfer with visible task improvement
mlagentbench_vec.	NumPy convolution	Local	speed_score	max	Benchmark compatibility probe
autoresearch_linux	LM pretraining	Claw	val_bpb	min	Real-task stress case; governance under high failure rate

Table 4: Trial lifecycle state machine. The pending guard is a file written before EXECUTING and deleted on any terminal transition; its presence at campaign start raises PendingTrialError, preventing concurrent writes to the append-only ledger.

From state	To state	Condition
PENDING	EXECUTING	Pending guard written; worker starts
EXECUTING	KEPT	Metric parses; beats current best; patch diff retained
EXECUTING	DISCARDED	Metric parses; does not beat current best
EXECUTING	FAILED_INVALID	Scope violation, runtime error, missing metric, no-op patch, or precondition failure
EXECUTING	FAILED_INVALID	Mid-trial crash detected via stale pending guard at next campaign start

Terminal states: KEPT, DISCARDED, FAILED_INVALID. All three produce a complete ledger record. Recovery tool: scripts/recover_pending.py.

We also instantiate two public OpenML [25] tabular nodes: openml_credit_g (dataset 31, the Statlog German Credit data [10]) and openml_bank_marketing (dataset 1461, the Bank Marketing data [19]). They expose only bounded config.yaml hyperparameters and model choices; dataset identity, target column, split logic, preprocessing, parser, and NodeSpec metadata are frozen. These are governance-transfer substrates, not competitive AutoML benchmarks.

The mlagentbench_vectorization node wraps a NumPy-convolution speed task as a bounded external-benchmark adapter. Benchmark code, correctness checking, and parser logic are frozen; only config.yaml edits over implementation choice and workload size are exposed. This is a compatibility probe, not a full MLAgentBench benchmark claim.

The sixth node, autoresearch_linux, is a GPT pretraining task on NVIDIA L40S hardware, adapted from Karpathy's autoresearch setup [16] through its macOS/CPU fork [1]: a small GPT, with training code that simplifies nanochat [15], trained on the climbmix-400b-shuffle data shards that the setup downloads. Its metric is val_bpb (bits per byte, minimise), the paper's only minimise-direction task. Workers may edit only named hyperparameters in train.py, including the model's depth, aspect ratio, head dimension and attention window pattern as well as batch sizes and learning rates; data preparation, data, tokenizer and parser are frozen. Each trial takes approximately 6 minutes, so only none and summary arms are run: 4 independent seeds × 30 trials per arm = 240 trials.

4.2 Campaign Protocol

Controlled ablations are designed to start each campaign from a reset node state. The canonical scientific-node campaign uses prompt_manager with structured state-aware hyperparameter proposals and claw_style_worker with the deterministic patch path. The memory ablation runs equal budgets across three memory modes (none, summary, and rationale-augmented summary) under the same node, budget, parser, and reset discipline. The OpenML campaigns use LangGraphManager, summary memory, and 20- and 30-trial budgets on each public node. The MLAgentBench adapter uses LocalWorker and a 30-trial budget. Stress campaigns exercise forced failure paths at 100% rate. For every trial, the control plane records proposal, patch reference, raw log reference, parsed metrics, validity status, failure category, decision rationale, provenance IDs, and reproducibility hashes.

Implementation and hardware environment. All cross-node campaigns use deepseek/deepseek-v4-flash [6] as the proposal manager via LangGraph [17], holding the model fixed to reduce the capability confound across nodes. Model-agnostic findings include provenance completeness, lifecycle correctness, and scope enforcement; proposal-quality findings may be manager-specific. ClawWorker campaigns ran on a server-class NVIDIA L40S host (CUDA 12.8, Ubuntu 22.04, PyTorch 2.x); LocalWorker campaigns ran on a CPU-only host under Python 3.11. The corrected ResNet L40S replication resets all editable node files to a verified frozen baseline before trial 1; older autoresearch_linux reset semantics are disclosed separately in Sections 5.6 and 6. Patch and run-log references are recorded with provenance IDs in the ledger; the ledgers needed to recompute the governance metrics are released with the code (Appendix A.1), while raw execution-host artifacts are not required for the paper's primary claims. Full environment details and campaign IDs are given in Appendices A and B.

4.3 Memory Ablation Design

As a diagnostic probe, we test the behavioral hypothesis

$$\text{RBR(none)} > \text{RBR(summary)} > \text{RBR(rationale)}$$

where RBR is *repeated-bad rate*: the fraction of proposals repeating the same edit target and mechanism as a prior rejected trial without a new corrective rationale. Rationale-linked memory should reduce repeated-bad proposals more than raw summaries because it provides targeted failure signals rather than history noise [4, 18, 23, 27].

Two manager configurations are tested on ResNet-trigger: a deterministic prompt manager (design-constrained negative — round-robin makes memory structurally irrelevant) and a LangGraph manager calling qwen2.5-coder:7b [12] at temperature 0.7 with an explicit AVOIDANCE RULE. All LangGraph proposals are parsed into structured parameter-change fields and applied via the deterministic patch bridge, eliminating coding-agent dependencies. Three ResNet-trigger replicates are run (two at budget 10, one at budget

Table 5: Governance metric suite. Given N trials, these metrics characterise process reliability independent of task-metric outcome.

Metric	Formal definition	What it detects
Acceptance rate (AR)	$\lvert\text{kept}\rvert/N$	Valid improving edits
Invalid rate (IR)	$\lvert\text{failed_invalid}\rvert/N$	Unevaluable trials
Repeated-bad rate (RBR)	Repeated-bad proposals $/N$	Redundant failure modes
Provenance completeness	Trials with all five IDs $/N$	Records lacking the ID chain
Artifact-reference coverage	Patch/log refs and hashes for materialized runs; classified records for pre-run failures	Ledger-level evidence traceability
Failure taxonomy	Counts per `FailureCategory`	Classified, not hidden, failures

20); a four-arm verbosity ablation with 50-token rationales tests whether rationale length drives the result. The same three-arm LangGraph ablation is repeated on `lr_synthetic` as a cross-node transfer check.

Large-scale replication (NVIDIA L40S). To obtain variance estimates and test the hypothesised ordering, a third ablation configuration runs on an institutional GPU cluster: 3 memory modes $\times$ 5 independent seeds $\times$ 15 trials = 225 total trials. The manager is `deepseek-v4-flash`; the worker is `qwen2.5-coder:7b` at temperature 0.2. Each seed resets `train.py` before trial 1. Per-trial SHA-256 context hashes verify that memory accumulates across the 225-trial ledger.

4.4 Governance Metrics

Failure labels are control-plane-assigned. The categories observed in the reported evidence include: `runtime_error`, `invalid_edit_scope`, `no_op_patch`, `proposal_precondition_failed`, and node-level `invalid_config`. The schema also supports `metric_missing` and `edit_failed`, although those labels are not part of the primary 1,445-trial evidence set. Valid-but-non-improving trials receive `degraded_metric` and are classified `discarded`, not `failed_invalid`. A campaign with high invalid rate but 100% provenance completeness is more interpretable than one with unexplained score changes and no audit trail.

4.5 Metric Validation Probes

We run three lightweight validation probes for the governance metrics. First, two reruns of the MLAgentBench vectorization adapter test whether the same NodeSpec contract continues to produce complete lifecycle records under repeated runs. Second, a re-execution script rebuilds kept-trial state from cumulative patch references and reruns selected kept trials where the underlying execution artifacts are available. Third, we estimate structural relabeling agreement by computing Cohen's κ between control-plane-assigned labels and an independent rule-based relabeler applied to structurally observable fields only (scope-validation flags, patch presence, training-initiation signals, execution status). On a stratified sample of 50 `failed_invalid` trials spanning the observed validation categories,

Table 6: Real-worker scientific-node case study (`kdd_resnet_scientific_20`, ResNet-trigger, `prompt_manager`, `claw_style_worker`, budget=20).

Field	Value
Kept / disc. / failed-invalid	4 / 5 / 11
Non-kept categories	5 × `degraded_metric`; 11 × `proposal_precondition_failed`
Acceptance / invalid rate	0.20 / 0.55
Provenance completeness	1.00 (20/20)
Decision record completeness	1.00 (20/20)
Failure evidence completeness	1.00 (11/11 precondition failures classified)
Patch refs/hashes for materialized valid runs	1.00 (9/9 valid runs)
Initial / best `val_auc`	0.7747 / 0.7827 (+0.008022)
Baseline seed noise (5 seeds)	mean 0.785, bootstrap 95% CI of mean [0.775, 0.798]

the structural relabeler achieves $\kappa = 0.830$ (95% CI [0.703, 0.958]). This is a structural consistency check, not a human inter-rater study. The remaining discordant records are edge cases where the ledger-level structural evidence does not distinguish the precise worker-side failure without consulting the referenced artifact file. Full human inter-rater validation remains future work.

5 Results

Results are presented in governance-first order. All numbers for the 1,445 primary trials derive from the primary campaign ledgers in `experiments/ledgers/`, which are released with the code (Appendix A.1); development and smoke ledgers are excluded from the paper totals. All governance metrics (acceptance rate, invalid rate, repeated-bad rate, provenance completeness) are recomputable from these ledgers without re-running any training; no GPU is required to verify the primary claims. The primary paper-relevant campaign set comprises 1,445 governed trials, indexed in Appendix B.1. Across the primary set, the harness runs one real-worker scientific case study (`resnet_trigger`), one dependency-free synthetic transfer node (`lr_synthetic`), two public OpenML tabular nodes, one MLAgentBench adapter, and one real LM training node (`autoresearch_linux`); supplemental validation adds two one-trial stress campaigns. Together these campaigns exercise kept, discarded, failed-invalid, invalid-scope, no-op, invalid-configuration, and repeated-bad outcomes under a single ledger schema.

The section is staged as follows: the single-node case study (§5.1); the six-node governance view (§5.2); stress campaigns and the governed-vs-ungoverned counterfactual (§5.3); the memory-ablation diagnostic and its large-scale replication (§§5.4–5.5); and the real LM-training stress case (§5.6). Throughout, AR denotes acceptance rate, IR invalid rate, and RBR repeated-bad rate (Table 5).

5.1 Real-Worker Scientific Node Case Study

Table 6 summarises the 20-trial real-worker campaign on ResNet-trigger. Four valid improving edits are kept, five valid non-improvements are discarded with patch/log references and metric evidence retained in the ledger, and eleven impossible/no-op structured proposals are classified as `proposal_precondition_failed` before they can corrupt node state. Provenance is complete for all 20 records (100%).

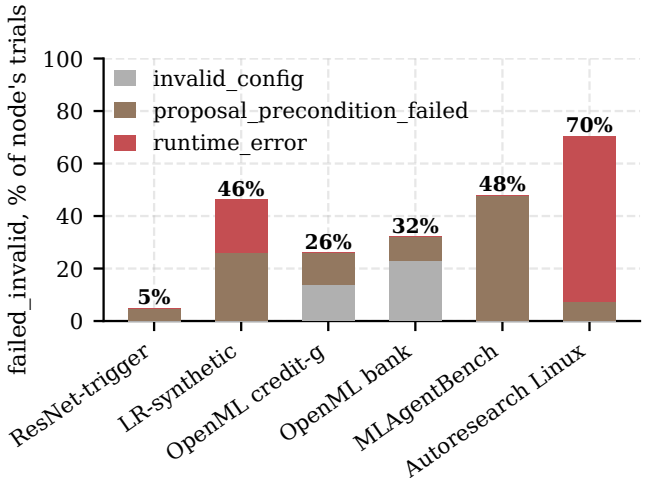


Figure 3: Failure taxonomy by node: `failed_invalid` trials as a share of each node's trials, split by `FailureCategory`. Composition is node-specific: `autoresearch_linux` concentrates in `runtime_error`; tightly-constrained nodes accumulate `proposal_precondition_failed`.

The net AUC gain (+0.008022) remains within the baseline seed interval and should be treated as secondary evidence that the loop runs correctly, not as an optimisation claim.

5.2 Cross-Node Governance Transfer

Table 7 and Figure 1 consolidate all six nodes into a single governance view; Figure 3 breaks the failed-invalid trials down by failure category. The central finding is that provenance completeness is 100% on every node without exception — including the 120-trial none-arm total failure and the s2 anomaly on `autoresearch_linux`, and across all three lifecycle outcomes (kept, discarded, failed-invalid).

Public tabular nodes. `openml_credit_g` and `openml_bank_marketing` test whether governed keep/discard/fail decisions transfer to public tabular substrates (LangGraph manager, summary memory, b20 and b30 budgets). Both nodes expose only bounded `config.yaml` edits; dataset, split, preprocessing, and parser code are frozen. Provenance completeness is 1.00 across all 460 trials combined. `openml_bank_marketing` shows visible task-metric improvement (mean best AUC 0.9295 at b20, 0.9341 at b30; 95% CI $[0.9337, 0.9345]$ at b30, σ=0.0005), while `openml_credit_g` is a neutral transfer case (mean best AUC 0.7679 at b30). Invalid rates increase with budget as bounded candidates are exhausted or violate node-level configuration constraints; all such trials are classified as `proposal_precondition_failed` or `invalid_config` rather than silently dropped.

MLAgentBench adapter. The 30-trial main campaign records 2 kept, 4 discarded, and 24 failed-invalid trials with 1.00 provenance completeness. The high failed-invalid count reflects exhausted bounded candidates (`proposal_precondition_failed`), not benchmark failure; the control plane classifies them explicitly rather than inventing out-of-scope edits. Two additional 10-trial reruns add kept and discarded outcomes under the same NodeSpec contract, bringing the adapter evidence set to 50 trials with complete provenance.

5.3 Stress Campaigns and Ungoverned Counterfactual

Table 8 documents forced failure paths and the synthetic-node governance validation.

The scope-violation and no-op cases are backed by real ledger records. The `lr_synthetic` campaigns confirm that the governance protocol operates on a different ML task and worker type without code changes.

Ungoverned counterfactual. Every "the harness makes X visible" claim is only as strong as evidence that X would be invisible without governance. Across the 1,445-trial primary campaign set, governance classification acts on the majority of outcomes: 242 trials (16.7%) produce `runtime_error` records that the pending-trial guard and append-only ledger make auditable; 220 trials (15.2%) are caught by precondition checks before reaching a valid run; 85 trials (5.9%) expose node-level `invalid_config` failures; 662 valid-but-degrading trials (45.8%) are discarded by the control-plane decision rule. Without the pending-trial guard and append-only ledger, each `runtime_error` would be a silent filesystem event with no trial identifier, failure taxonomy, or provenance chain.

Table 9 reports matched governed/ungoverned pairs on two nodes chosen to span failure rates; the visual companion is Fig. S8. The ungoverned arm is an independent matched campaign, not a reconstruction from the governed ledger. It disables the pending-trial guard, failure-record appends, and blocking scope enforcement (Table S3). A separate observation log (`*_ungoverned_obs.jsonl`) records every executed trial and is the denominator for ledger completeness. Because failure records are not appended, the trials missing from the ungoverned ledger are, by construction, exactly that arm's failed trials: the comparison shows what such a loop loses, not an independent estimate of loss.

The two nodes bracket the failure spectrum. The `autoresearch_linux` row is an extremal, definitional case: when every trial fails and failure appends are disabled, completeness is necessarily zero. The informative mixed-outcome comparison is `openml_bank_mkt`, where the ungoverned ledger retains only the 7 valid trials and silently omits all 13 failures. Combined, 43 of 50 ungoverned trials leave no ledger record versus 0 of 50 governed; because the two nodes were chosen to span failure rates, this combined fraction describes these campaigns, not a typical loss rate. The 95% bootstrap CI for ungoverned completeness lies strictly below 1.00 on both nodes.

5.4 Memory Ablation as Diagnostic Probe

Memory context is evaluated as a diagnostic probe for governance behaviour, not as a proposed optimisation method. The behavioral hypothesis is:

$$\mathrm{RBR}(\text{none}) > \mathrm{RBR}(\text{summary}) > \mathrm{RBR}(\text{rationale})$$

where RBR is *repeated-bad rate*: the fraction of proposals repeating the same edit target and mechanism as a prior rejected trial without a new corrective rationale.

Table 7: Cross-node governance summary. All nodes achieve 100% provenance completeness regardless of task type, worker class, or failure rate. AR = acceptance rate, IR = invalid rate. ResNet-trigger appears as two rows because the campaign configurations differ fundamentally (prompt_manager vs. DeepSeek LangGraph; different memory modes). The autoresearch arms are listed separately because their AR/IR differ categorically. The `lr_synthetic` row pools its three memory modes.

Node	Trials	AR%	IR%	Lifecycle outcomes
`resnet_trigger` (case study)	20	20.0	55.0	kept, disc., failed_inv.
`resnet_trigger` (L40S, all arms)	225	46.7	0.4	kept, disc., failed_inv. (5 clean seeds verified)
`lr_synthetic`	450	8.4	46.2	kept, disc., failed_inv.
`openml_credit_g`	230	13.5	26.1	kept, disc., failed_inv.
`openml_bank_marketing`	230	13.0	32.2	kept, disc., failed_inv.
`mlagentbench_vec.`	50	28.0	48.0	kept, disc., failed_inv.
`autoresearch` (none arm)	120	0.0	100.0	failed_inv. only (total failure)
`autoresearch` (summary)	120	11.7	40.8	kept, disc., failed_inv.
Provenance completeness: 100% on every row above				

Table 8: Stress campaigns and synthetic-node validation. Each row is a silent failure in a naive unguarded loop that the harness makes visible.

Scenario	Campaign	Governed outcome	Evidence
Out-of-scope edit (frozen file touched)	`kdd_stress_scope`	Patch rejected before training; classified `invalid_edit_scope`; git state unchanged	1/1 trials; patch ref retained
No-op patch (byte-identical edit)	`kdd_stress_noop`	Training skipped; classified `no_op_patch`; empty patch hash recorded	1/1 trials; ledger record present
Synthetic baseline: `lr_synthetic` (logistic regression, LocalWorker)	`lr_synth_baseline`	All three lifecycle outcomes in 5 trials; provenance complete; NodeSpec YAML portable	2 kept, 1 discarded, 2 failed-invalid; best val_score = 0.923
Synthetic LangGraph transfer	`lr_synth_lg_*`	Same manager/memory ablation runs on a different node and worker type; all data-generation constants remain frozen	30/30 valid trials; RBR none/summary/rationale = 0.70/0.70/0.80

Table 9: Governed vs. ungoverned counterfactual comparison across two nodes. Ledger completeness = (trials with a ledger record) / (trials that ran). Governed completeness is 1.00 by structural invariant (Section 3.4); ungoverned completeness falls below 1.00 because `failed_invalid` trials are silently absent from the ledger; by construction it equals one minus the ungoverned arm's failure rate. IR = invalid rate of the governed arm. 95% bootstrap CIs: 10,000 resamples, seed 42.

Node (profile)	N	Governed	Ungoverned	Drops	Interpretation
`autoresearch_linux` IR = 100%	30	1.00 [1.00, 1.00]	0.00 [0.00, 0.00]	30 (100%)	Total failure: ungoverned ledger is empty; campaign indistinguishable from one that never ran
`openml_bank_mkt` IR = 55%	20	1.00 [1.00, 1.00]	0.35 [0.15, 0.55]	13 (65%)	Mixed outcomes: ungoverned ledger retains only 7 valid trials; 13 classified failures leave no trace

Combined: 43 / 50 ungoverned trials absent, i.e. all failed trials of these two campaigns, which were chosen to span failure rates; 0 / 50 governed trials absent (100% completeness by structural invariant).

Full arm-level results appear in Table S4. These ResNet-trigger campaigns are not among the 1,445 primary trials; their ledgers are in the archive of commit `62e93a8` rather than in the tagged release (Appendix A.1). We use this ablation only as a diagnostic example for repeated-bad rate, not as a proposed memory method.

Primary diagnostic finding. On ResNet-trigger, the summary arm has the lowest repeated-bad rate of the three arms in 2 of 3 replicates (rep1 and rep3, where the no-memory and rationale arms tie); in rep2 the rationale arm is lowest, so the full hypothesised ordering holds only in rep2 (Table S4). The node-state hashes show, however, that the rep2 arms were not reset between campaigns: each started from the final `train.py` of the arm run before it, so the rep2 rationale arm began from the summary arm's best state (`val_auc` 0.7903 after its first trial). Rep2 is therefore not an independent replicate, and the one replicate in which the full ordering holds does not support it; rep1 and rep3 began every arm from the same node state. The verbosity arm C, run just before rep2, likewise began from the baseline plus one `BATCH_SIZE` edit, which has no effect in this mode. These are descriptive counts over three short replicates, not a significance test. These ResNet-trigger runs used the node's CPU fast-search mode, in which the training-set size, trace length, batch size and number of epochs are taken from the environment; edits to `N_SIGNAL`, `N_NOISE`, `TRACE_LEN`, `BATCH_SIZE` or `EPOCHS` in `train.py` therefore cannot change the result. Arms that kept proposing such edits ("stuck on `BATCH_SIZE`" and "stuck on `TRACE_LEN`" in Table S4) repeated edits without effect, which raises their repeated-bad rate, so the replicate comparison partly reflects this interface rather than the memory mode. A small cross-node `lr_synthetic` diagnostic (10 trials per memory mode) completes 30/30 valid LangGraph trials but does not reproduce the hypothesised RBR ordering: rationale (0.80) exceeds both none and summary (0.70), the reverse of the hypothesis. One possible reading of this reversal is an anchoring effect — richer rationale context gives the manager more prior failed patterns to fixate on rather than redirect from; either way, the result supports treating repeated-bad rate as a per-node diagnostic rather than as evidence of a general memory-mode effect.

Verbosity ablation. In this single run, truncating rationale to 50 tokens does not restore summary-level performance (RBR(C) = 0.80 = RBR(D) ≠ RBR(B) = 0.50), so rationale *length* alone does not

explain the difference. The full-rationale arm of this run was stuck on `BATCH_SIZE`, an edit without effect in the fast-search mode, and the truncated arm ended at the same starting AUC, so this comparison is weak evidence.

5.5 Large-Scale Memory Replication (NVIDIA L40S, 225 Trials)

The large-scale L40S replication (design and models in Section 4.3: 3 memory modes × 5 independent seeds × 15 trials = 225 trials) tests the full hypothesised ordering with variance estimates. Each seed reset `train.py` to a verified clean baseline before its first trial; reset integrity is confirmed by a shared `node_state_hash` at trial 1 across all 15 campaigns.

Starting point and configuration. All arms start from the same level: trial-1 `val_auc` lies between 0.886 and 0.890 in every arm and seed (Table S5, Fig. S4). This level is well above the case study's initial 0.7747, and the configurations differ: the L40S campaigns train `resnet_trigger` in its full mode on the GPU (4,000 signal and 4,000 noise traces of 8,000 samples; batch size 32 and 5 epochs before any edit), starting from the `train.py` committed on 30 April 2026 (harness commit `d589d88`), which reproduces the shared trial-1 `node_state_hash`, whereas the case study, the LangGraph ablation and the baseline-seed interval use the node's CPU fast-search mode (1,000 signal and 1,000 noise traces cropped to 4,096 samples; batch size 64 and 3 epochs, fixed by the mode). L40S AUCs are therefore comparable across arms but not with the other ResNet results. From the common start, the mean best AUC rises by about 0.011 without memory and by about 0.05 with either memory mode.

L40S replication: memory arms outperform none. Across all five clean seeds, mean best AUC follows none < summary ≈ rationale (0.899 vs. 0.939 and 0.941); the 0.002 difference between the two memory arms is not treated as meaningful. Cross-seed variability is lower for both memory arms, with $\sigma_{\text{none}} = 0.009$ vs. $\sigma_{\text{summary}} = \sigma_{\text{rationale}} = 0.002$, a 4–5× smaller standard deviation; with five seeds per arm this is descriptive. Memory arms diverge from the no-memory arm by trial 5 and hold the gap through trial 15 (Figs. S4 and S9). In an earlier run of these 15 campaigns the reset of `train.py` failed: within seeds s1–s3 each arm started from a different file, and s4–s5 started from a modified baseline. Those ledgers were discarded (deleted by the rerun scripts, outside the control plane), and all 15 campaigns were rerun under the same campaign IDs from a verified clean baseline; the numbers reported here come from the rerun. Provenance completeness: 100% across all 225 trials; the memory effect remains node- and manager-specific.

5.6 Real LM Training Node: Governance Under High Failure Rate

The `autoresearch_linux` node (GPT pretraining on the autoresearch data shards, NVIDIA L40S, metric `val_bpb`) is reported as a governance stress case under high failure rate. The `none` arm produces IR = 1.00 across all 120 trials (four seeds); three of four summary seeds produce valid improving trials (mean `val_bpb` reduction −0.062), while summary seed s2 also suffers total worker failure; one of its 30 records was written by the pending-trial recovery for a trial interrupted by a crash (Appendix B.7). Unlike the corrected ResNet replication, these campaigns were not run with the later frozen-baseline restore protocol, and the archived trial-1 hashes do not verify a common start state across all arms and seeds. We therefore report both the `none` arm's total failure and summary seed s2 as observed failures of undetermined cause (an infrastructure or protocol anomaly cannot be excluded), not as evidence for a memory effect. Across 240 trials including 151 runtime-error crashes, provenance completeness is 240/240 (100%): every `runtime_error` event is machine-readable with zero silent failures. The governance finding is classification fidelity under failure, not a general claim about memory-free managers.

6 Discussion and Limitations

Evidence strength. Table 10 separates demonstrated properties from case-study findings and non-claims.

Table 10: Evidence strength by claim.

Claim	Evidence level	Where
Control plane classifies kept / discarded / failed_invalid	Demonstrated	Secs 5.1–5.6
Append-only ledger: 100% provenance	Demonstrated	All nodes; Sec 5.2
Guard failures visible in audit record	Demonstrated	Stress, Sec 5.3
Protocol transfers across node types	Demonstrated, 6 nodes	Secs 5.2–5.6
Repeated-bad rate diagnoses memory sensitivity	Mixed case study	ResNet: summary lowest in 2/3; `lr_synthetic` fails
Memory ordering none < summary ≈ rationale (AUC)	Observed, 5 clean seeds	Sec 5.5; not on `lr_synthetic`
Smaller cross-seed spread with memory	Observed, 5 seeds (descriptive)	σ_{none}=0.009 vs. 0.002 (4–5×)
Rationale reduces repeated-bad rate	Not claimed	Non-monotonic; verbosity ablation not supported
No-memory arm: 100% invalid, cause undetermined (start state unverified)	Observed, 4 seeds	Sec 5.6
Governance classifies total worker failure	Demonstrated	Sec 5.6

Limitations. Four limitations bound the claims. *Scope of evidence:* six nodes broaden evidence beyond a single case study, but budgets are short, the summary arm has the lowest RBR in only 2 of 3 ResNet replicates (the full hypothesised ordering holds only in rep2, whose arms were not reset between campaigns), and the ordering reverses on `lr_synthetic`. *Label validation:* structural relabeling agreement (κ=0.830, Section 4.5) is a consistency check, not a human study; two-rater validation and holdout-node replication remain future work. *Task metric and reset scope:* the ResNet best AUC falls within the 5-seed baseline min–max range [0.771, 0.811] (bootstrap 95% CI of mean: [0.775, 0.798]), so AUC gain is secondary evidence; the `autoresearch_linux` campaigns predate the frozen-baseline restore protocol, so neither their `none`-arm total failure nor their s2 anomaly is used as memory-effect evidence. *Human utility:* the paper shows that governance properties can be enforced and measured, and what evidence is lost without them (Section 5.3) — not that governed ledgers improve human audit outcomes; a human-auditor utility study (do reviewers with governed ledgers detect seeded errors faster?) is the natural next step.

7 Conclusion

Across 1,445 governed trials on six ML nodes, the control plane classifies every outcome without silent corruption; provenance is complete by construction and is observed in all 1,445 trials, as the invariant requires. Autonomous ML experimentation should

be evaluated on governance metrics—provenance, lifecycle classification, scope enforcement, and failure taxonomy—not only task metrics.

Acknowledgments

The author is deeply grateful to Sudhir Shetty for the many inspiring discussions, and for his generous support and patience along the way. The author also warmly thanks Tony Chen and Girish Hoogar for their encouragement and for the thoughtful conversations that have enriched his perspective and helped shape his research interests.

AI Usage Statement

Apart from the LLMs studied here, the author used AI assistants, including Anthropic's Claude, to help write the harness code and to edit and check this manuscript, and reviewed all of their output.

References

[1] Artem Andreenko. 2026. *autoresearch-macos.* Retrieved November 8, 2026 from https://github.com/miolini/autoresearch-macos

[2] Lukas Biewald. 2020. *Experiment Tracking with Weights and Biases.* https://wandb.ai

[3] Joschka Birk, Gregor Kasieczka, Siddharth Mishra-Sharma, Benjamin Nachman, Dennis Noll, and Tanvi Wamorkar. 2026. *A Scientific Human-Agent Reproduction Pipeline.* arXiv:2604.18752 [hep-ph] doi:10.48550/arXiv.2604.18752

[4] Birgitta Böckeler. 2026. *Harness Engineering for Coding Agent Users.* Retrieved May 2026 from https://martinfowler.com/articles/harness-engineering.html

[5] Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, Lilian Weng, and Aleksander Mądry. 2024. *MLE-bench: Evaluating Machine Learning Agents on Machine Learning Engineering.* arXiv:2410.07095 [cs.CL] doi:10.48550/arXiv.2410.07095

[6] DeepSeek. 2026. *DeepSeek V4 Preview Release.* DeepSeek API Docs. Retrieved November 8, 2026 from https://api-docs.deepseek.com/news/news260424

[7] Pengfei Du. 2026. *Memory for Autonomous LLM Agents: Mechanisms, Evaluation, and Emerging Frontiers.* arXiv:2603.07670 [cs.AI] doi:10.48550/arXiv.2603.07670

[8] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. 2019. Neural Architecture Search: A Survey. *Journal of Machine Learning Research* 20, 55 (2019), 1–21. https://www.jmlr.org/papers/v20/18-598.html

[9] Mitchell Hashimoto. 2026. *My AI Adoption Journey.* Retrieved May 2026 from https://mitchellh.com/writing/my-ai-adoption-journey

[10] Hans Hofmann. 1994. *Statlog (German Credit Data).* UCI Machine Learning Repository. doi:10.24432/C5NC77

[11] Qian Huang, Jian Vora, Percy Liang, and Jure Leskovec. 2024. MLAgentBench: Evaluating Language Agents on Machine Learning Experimentation. In *Proceedings of the 41st International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 235).* PMLR, Vienna, Austria, 20271–20309. https://proceedings.mlr.press/v235/huang24y.html

[12] Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, Kai Dang, Yang Fan, Yichang Zhang, An Yang, Rui Men, Fei Huang, Bo Zheng, Yibo Miao, Shanghaoran Quan, Yunlong Feng, Xingzhang Ren, Xuancheng Ren, Jingren Zhou, and Junyang Lin. 2024. *Qwen2.5-Coder Technical Report.* arXiv:2409.12186 [cs.CL] doi:10.48550/arXiv.2409.12186

[13] Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren (Eds.). 2019. *Automated Machine Learning: Methods, Systems, Challenges.* Springer, Cham. doi:10.1007/978-3-030-05318-5

[14] Zhengyao Jiang, Dominik Schmidt, Dhruv Srikanth, Dixing Xu, Ian Kaplan, Deniss Jacenko, and Yuxiang Wu. 2025. *AIDE: AI-Driven Exploration in the Space of Code.* arXiv:2502.13138 [cs.AI] doi:10.48550/arXiv.2502.13138

[15] Andrej Karpathy. 2025. *nanochat.* Retrieved November 8, 2026 from https://github.com/karpathy/nanochat

[16] Andrej Karpathy. 2026. *autoresearch.* Retrieved November 8, 2026 from https://github.com/karpathy/autoresearch

[17] LangChain. 2026. *LangGraph.* Retrieved November 8, 2026 from https://github.com/langchain-ai/langgraph

[18] Ryan Lopopolo. 2026. *Harness Engineering: Leveraging Codex in an Agent-First World.* Retrieved May 2026 from https://openai.com/index/harness-engineering/

[19] Sérgio Moro, Paulo Cortez, and Paulo Rita. 2014. A Data-Driven Approach to Predict the Success of Bank Telemarketing. *Decision Support Systems* 62 (2014), 22–31. doi:10.1016/j.dss.2014.03.001

[20] Deepak Nathani, Lovish Madaan, Nicholas Roberts, Nikolay Bashlykov, Ajay Menon, Vincent Moens, Amar Budhiraja, Despoina Magka, Vladislav Vorotilov, Gaurav Chaurasia, Dieuwke Hupkes, Ricardo Silveira Cabral, Tatiana Shavrina, Jakob Foerster, Yoram Bachrach, William Yang Wang, and Roberta Raileanu. 2025. *MLGym: A New Framework and Benchmark for Advancing AI Research Agents.* arXiv:2502.14499 [cs.CL] doi:10.48550/arXiv.2502.14499

[21] Ollama. 2026. *Ollama.* Retrieved November 8, 2026 from https://github.com/ollama/ollama

[22] Prithvi Rajasekaran. 2026. *Harness Design for Long-Running Application Development.* Retrieved May 2026 from https://www.anthropic.com/engineering/harness-design-long-running-apps

[23] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language Agents with Verbal Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS '23, Vol. 36).* Curran Associates, Inc., New Orleans, LA, USA, 8634–8652. https://proceedings.neurips.cc/paper_files/paper/2023/hash/1b44b878bb782e6954cd888628510e90-Abstract.html

[24] Vivek Trivedy. 2026. *The Anatomy of an Agent Harness.* Retrieved May 2026 from https://www.langchain.com/blog/the-anatomy-of-an-agent-harness

[25] Joaquin Vanschoren, Jan N. van Rijn, Bernd Bischl, and Luis Torgo. 2013. OpenML: Networked Science in Machine Learning. *SIGKDD Explorations* 15, 2 (2013), 49–60. doi:10.1145/2641190.2641198

[26] Weco AI. 2026. *AIDE ML: The Machine Learning Engineering Agent.* Retrieved November 8, 2026 from https://github.com/WecoAI/aideml

[27] Chunlong Wu, Ye Luo, Zhibo Qu, and Min Wang. 2025. *Meta-Policy Reflexion: Reusable Reflective Memory and Rule Admissibility for Resource-Efficient LLM Agent.* arXiv:2509.03990 [cs.AI] doi:10.48550/arXiv.2509.03990

[28] Justin Young. 2025. *Effective Harnesses for Long-Running Agents.* Retrieved May 2026 from https://www.anthropic.com/engineering/effective-harnesses-for-long-running-agents

[29] Matei Zaharia, Andrew Chen, Aaron Davidson, Ali Ghodsi, Sue Ann Hong, Andy Konwinski, Siddharth Murching, Tomas Nykodym, Paul Ogilvie, Mani Parkhe, Fen Xie, and Corey Zumar. 2018. Accelerating the Machine Learning Lifecycle with MLflow. *IEEE Data Engineering Bulletin* 41, 4 (Dec. 2018), 39–45. https://people.eecs.berkeley.edu/~matei/papers/2018/ieee_mlflow.pdf

[30] Matei Zaharia, Omar Khattab, Lingjiao Chen, Jared Quincy Davis, Heather Miller, Chris Potts, James Zou, Michael Carbin, Jonathan Frankle, Naveen Rao, and Ali Ghodsi. 2024. The Shift from Models to Compound AI Systems. https://bair.berkeley.edu/blog/2024/02/18/compound-ai-systems/

[31] Qiran Zou, Hou Hei Lam, Wenhao Zhao, Yiming Tang, Tingting Chen, Samson Yu, Tianyi Zhang, Chang Liu, Xiangyang Ji, and Dianbo Liu. 2025. *FML-bench: Benchmarking Machine Learning Agents for Scientific Research.* arXiv:2510.10472 [cs.CL] doi:10.48550/arXiv.2510.10472

These appendices are the reproducibility material of the paper (the supplement of the workshop version). They provide the ledger schema, NodeSpec contract examples, control-plane pseudocode, extended experiment tables, the failure taxonomy with concrete ledger examples, campaign identifiers, the hardware and model environment, and notebook details. No central claim in the main text depends only on these appendices.

A Reproducibility and Artifact Description

A.1 Artifact and Code Availability

The harness code is public at https://github.com/dowlingwong/autoresearch_harness under the MIT licence. The trial ledgers behind the numbers in this paper are released as a tagged GitHub release of that repository, archived on Zenodo: release `kdd-etaai-arxiv-v1` (commit `0455171`), DOI 10.5281/zenodo.24174991. The release holds the 60 primary ledgers of Table S2, the validation and counterfactual ledgers of Table S3 (the ungoverned `autoresearch_linux` arm recorded no valid trial, so it has no ledger), the two `*_ungoverned_obs.jsonl` observation logs, and an `artifact_manifest.json` that lists each file with its SHA-256 digest and trial count. Each released ledger and observation log is byte-identical to its copy in `autoresearch-harness-artifact.zip`, which was committed to the same repository on 25 May 2026 (commit `62e93a8`) and remains in its history. The tagged commit is the code version from which the reported numbers are recomputed; it supersedes commit `ff3717b`, which the workshop version cited and which does not contain the ledgers. The ledgers of the ResNet-trigger LangGraph memory ablation (all rows of Table S4 except the `lr_synthetic` rows) are not in the tagged release; they are in the same archive of commit `62e93a8`, from which those rows are recomputed.

The release archive contains: `src/` (control-plane source), `configs/nodes/` (node YAML specs), `scripts/` (campaign runners), `tests/` (test suite), `experiments/ledgers/` (the ledgers listed above), `nodes/` (per-node training code), `paper_figures.ipynb` (figure notebook), `artifact_manifest.json` (ledger index), and `ARTIFACT_README.md` (step-by-step instructions). It excludes large or private execution artifacts: the two ResNet HDF5 data files and the execution-host patch and run-log files are not needed to recompute the reported governance metrics from the ledgers.

Execution-host patch and run-log files are referenced in each ledger record by path fields and provenance IDs; the files themselves are not part of the release.

Each trial record contains a trial identifier, campaign identifier, proposal summary/rationale, patch reference, raw-log reference, parsed metrics, terminal decision, failure category where applicable, node-state hash, and provenance IDs for proposal, patch, run, metric, and decision objects. Readers can map each paper table to its ledgers under `experiments/ledgers/` and from there to the artifact references recorded in each ledger row.

Table S1: Core ledger fields. Terminal lifecycle states are `kept`, `discarded`, and `failed_invalid`; failure labels are attached only where a control-plane guard or worker failure produces a failure condition.

Field	Meaning
`trial_id`	Stable identifier for one budget slot.
`campaign_id`	Links rows to the campaign ID inventory in Section B.1.
`budget_index`	Trial index within the fixed campaign budget.
`decision`	Terminal state: `kept`, `discarded`, or `failed_invalid`.
`failure_category`	Control-plane label such as `runtime_error`, `invalid_edit_scope`, `no_op_patch`, `proposal_precondition_failed`, `invalid_config`, `metric_missing`, or `edit_failed`; valid-but-worse trials use `degraded_metric` and are `discarded`.
`proposal_summary`, `proposal_rationale`	Manager proposal text recorded before worker execution.
`patch_ref`, `patch_hash`	Patch artifact path and content hash when a patch is materialised.
`raw_log_ref`	Worker run-log artifact path.
`parsed_metrics`	Machine-parsed task metrics used by the decision rule.
`provenance`	Stable IDs for the trial's proposal, patch, run, metric, and decision objects, each derived by hashing the trial ID.
`node_state_hash`	Hash of node-local state used for reset and divergence auditing.

A.2 Ledger Schema

A.3 Dry-Run Commands

The harness can be exercised without Ollama or a GPU using dry-run mode:

```
# Smoke test --- 3 trials, synthetic metrics, no LLM or node
      required
python3 scripts/run_campaign.py \
    --node resnet_trigger --campaign-id smoke \
    --budget 3 --dry-run

# LangGraph manager smoke test (stub LLM, no Ollama)
python3 scripts/run_campaign.py \
    --node resnet_trigger --campaign-id lg_smoke \
    --budget 2 --manager langgraph_manager \
    --dry-run --llm-stub

# Memory ablation dry run (all three modes)
python3 scripts/run_memory_ablation.py \
    --node resnet_trigger --budget 3 --dry-run
```

For a real campaign with a local Ollama model:

```
python3 scripts/run_campaign.py \
    --node resnet_trigger --campaign-id main \
    --budget 5 --manager prompt_manager \
    --memory-mode append_only_summary_with_rationale \
    --node-root nodes/ResNet_trigger \
    --model qwen2.5-coder:7b \
    --host http://localhost:11434
```

A.4 Hardware and Model Environment

Large-scale ResNet and `autoresearch_linux` campaigns ran on an institutional GPU cluster (NVIDIA L40S, CUDA 12.8, Ubuntu 22.04, PyTorch 2.x). Cross-node LangGraph campaigns use `deepseek/deepseek-v4-flash`, the Flash model of the DeepSeek V4 preview released on 24 April 2026 [6], through the OpenAI-compatible DeepSeek API endpoint (`DEEPSEEK_BASE_URL`, default https://api.deepseek.com; authenticated via `DEEPSEEK_API_KEY`). By their ledger timestamps (UTC), all campaigns in Tables S2 and S3 ran between 9 and 22 May 2026; the corrected L40S rerun ran on 21 May and the counterfactual campaigns on 22 May. The ResNet L40S replication uses `qwen2.5-coder:7b` at temperature 0.2. The stochastic memory ablation uses `qwen2.5-coder:7b` at temperature 0.7 with the explicit avoidance prompt and deterministic patch bridge. LocalWorker campaigns run under Python 3.11 on a CPU-only host.

Hosted API proposals may vary if the provider changes the served model behind a fixed identifier. Exact proposal text is most reproducible from the released ledgers. Governance metrics remain recomputable from archived trial records without API access. No API keys, account identifiers, or private endpoint secrets are included in the artifact.

A.5 Bootstrap Confidence Intervals

Seed-level bootstrap confidence intervals use 10,000 resamples with random seed 42. Resampling is performed over seed-level best metrics, not individual trial records, so each interval reflects between-seed variability.

A.6 Node-State-Hash Reset Verification

Each trial record contains a `node_state_hash` computed from the SHA-256 digest of the node's editable file set immediately before the worker executes. This field allows independent verification that every seed of a multi-seed ablation campaign started from an identical node state — a necessary condition for fair memory-mode comparisons. For the L40S replication (Table S5), all five seeds produce hash prefix `9c0118f33fcb` at `budget_index=1`:

```
python3 -c "
import json
for l in open('experiments/ledgers/<campaign_id>_trials.
     jsonl'):
    r = json.loads(l)
    if r['budget_index'] == 1:
        print(r['trial_id'], r['node_state_hash'][:12])
"
```

A divergence in the prefix across seeds indicates a reset failure and invalidates the ablation comparison for that seed. The hash is stored in every ledger record, so the check is replayable from the released ledgers without re-running any training code.

A.7 Provenance and Artifact Completeness

Figure S1 shows per-node provenance completeness and artifact-reference capture rates by decision type. Panel (a) confirms 100% provenance completeness on all six nodes. Panel (b) shows that lower artifact-reference capture rates for `failed_invalid` trials are expected by design: trials that fail at the precondition stage (before worker execution) produce no patch diff, and trials that fail at runtime may produce no parsed metric. The provenance chain itself (proposal → patch → run → metric → decision IDs) is always complete regardless of the terminal outcome.

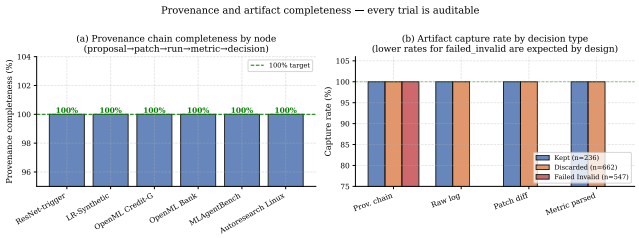


Figure S1: Provenance and artifact-reference completeness across all 1,445 governed trials. (a) Provenance chain completeness is 100% on every node. (b) Artifact-reference capture rates by decision type: lower patch and metric rates for `failed_invalid` are expected (pre-run failures produce no patch; runtime failures produce no parsed metric).

A.8 Analysis Notebook

The release includes `paper_figures.ipynb`. The notebook loads the campaign ledgers from `experiments/ledgers/`, treats the append-only ledgers as the source of truth, and writes derived outputs to `paper/tables/` and `paper/figures/`. Install dependencies and run:

```
uv pip install -e ".[dev,notebook]"
# then: Kernel -> Restart & Run All
```

The notebook is an analysis convenience layer; the ledgers, not rendered plots, are the evidence source.

Selected notebook outputs (after final L40S data synchronisation): `fig1_governance_overview`, `fig3_ir_taxonomy`, `fig5_cross_node`, `fig6_provenance`, `fig7_openml`, and `fig9_resnet_linux_ablation`.

Retired or exploratory outputs (including the MLP synthetic diagnostic plot) should not be cited as reported evidence unless the main text explicitly restores that node to the evidence set.

B Extended Experiment Tables

B.1 Campaign ID Inventory

Table S2 lists the 60 primary campaign ledgers behind the 1,445-trial evidence set; Table S3 lists the validation and counterfactual campaigns reported outside that total.

B.2 LangGraph Memory Ablation Detail

Table S4 gives the arm-level results of the LangGraph memory ablation summarised in Section 5.4.

Table S2: Primary campaign identifiers for the 1,445-trial governed evidence set. One ledger file exists per listed seed. Campaign IDs with ranges (e.g. `s1–s5`) denote five separate JSONL ledgers.

Campaign ID(s)	Node	Arm / mode	Seed(s)	Budget
kdd_resnet_scientific_20	resnet_trigger	case study	–	20
deepseek_resnet_none_s1–s5	resnet_trigger	none	s1–s5	15
deepseek_resnet_append_only_summary_s1–s5	resnet_trigger	summary	s1–s5	15
deepseek_resnet_append_only_summary_with_rationale_s1–s5	resnet_trigger	rationale	s1–s5	15
deepseek_lr_none_b30_s1–s5	lr_synthetic	none	s1–s5	30
deepseek_lr_append_only_summary_b30_s1–s5	lr_synthetic	summary	s1–s5	30
deepseek_lr_append_only_summary_with_rationale_b30_s1–s5	lr_synthetic	rationale	s1–s5	30
deepseek_openml_cg_s1–s4	openml_credit_g	summary	s1–s4	20
deepseek_openml_cg_b30_s1–s5	openml_credit_g	summary	s1–s5	30
deepseek_openml_bm_s1–s4	openml_bank_marketing	summary	s1–s4	20
deepseek_openml_bm_b30_s1–s5	openml_bank_marketing	summary	s1–s5	30
mlagentbench_vectorization_main_30	mlagentbench_vec.	adapter main	–	30
deepseek_mlagentbench_s1, s2	mlagentbench_vec.	adapter rerun	s1–s2	10
deepseek_autoresearch_linux_none_s1–s4	autoresearch_linux	none	s1–s4	30
deepseek_autoresearch_linux_append_only_summary_s1–s4	autoresearch_linux	summary	s1–s4	30

Table S3: Validation and counterfactual campaign identifiers reported outside the 1,445-trial headline total. Ungoverned counterfactual arms append only valid-trial records to the ledger by design; the full trial denominators are recorded in the paired `*_ungoverned_obs.jsonl` observation logs (Table 9).

Campaign ID(s)	Node	Arm / mode	Seed(s)	Budget
kdd_stress_scope	resnet_trigger	scope stress	–	1
kdd_stress_noop	resnet_trigger	no-op stress	–	1
lr_synth_baseline	lr_synthetic	baseline manager	–	5
lr_synth_lg_none	lr_synthetic	none	–	10
lr_synth_lg_summary	lr_synthetic	summary	–	10
lr_synth_lg_rationale	lr_synthetic	rationale	–	10
kdd_cf_openml_gov	openml_bank_marketing	governed (matched)	–	20
kdd_cf_openml_ung	openml_bank_marketing	ungoverned (matched)	–	20
kdd_cf_arlinux_gov	autoresearch_linux	governed (matched)	–	30
kdd_cf_arlinux_ung	autoresearch_linux	ungoverned (matched)	–	30

Table S4: LangGraph memory ablation (qwen2.5-coder:7b, temp=0.7, avoidance prompt, deterministic patch bridge). RBR = repeated-bad rate, the repeated-bad count divided by the budget N (Appendix C.7), recomputed for every row from the ledgers. The workshop version reported, for the ResNet-trigger rows, the count over the first $N-1$ trials divided by $N-1$ (e.g., 0.78 = 7/9); the order of the arms within each replicate is the same under both. The rep2 arms were not reset between campaigns (Section 5.4). Arm C tests whether rationale *length* drives Arm D's result. The `lr_synthetic` rows use validation score, not AUC.

Arm	Memory mode	RBR	Kept	Disc.	Best metric	Observation
Primary ablation — rep1, budget=10						
A	none	0.80	1	9	0.7747	Baseline
B	summary	0.50	4	6	0.8212	Discovers TRAIN_FRACTION
D	full rationale	0.80	1	9	0.7747	Stuck on BATCH_SIZE
Verbosity ablation — same node/budget/model						
C	short rationale (50 tok)	0.80	1	9	0.7747	Matches D, not B: length alone does not explain D
Seed replication — rep2, budget=10; arms not reset between campaigns						
A	none	0.80	1	9	0.7747	
B	summary	0.70	2	8	0.7903	Summary not lowest in this replicate
D	full rationale	0.60	3	7	0.8220	Lowest RBR in rep2; started from B's final state
Budget replication — rep3, budget=20						
A	none	0.85	2	18	0.8031	
B	summary	0.80	3	17	0.8150	Summary lowest; finds TRAIN_FRACTION at T5
D	full rationale	0.85	1	17	0.7747	Stuck on TRACE_LEN; also 2 failed-invalid (no_op_patch)
Cross-node transfer — lr_synthetic, budget=10						
A	none	0.70	2	8	0.9230	Clean LocalWorker run
B	summary	0.70	2	8	0.9228	No RBR reduction on this node
D	full rationale	0.80	1	9	0.9227	Slightly higher RBR

B.3 Cross-Node Acceptance and Invalid Rates

Figure S2 summarises acceptance rate (AR) and invalid rate (IR) across all six nodes with 95% seed-level bootstrap confidence intervals. AR and IR vary substantially by node and manager interface. The governance claim is not that these rates are uniform, but that the lifecycle classification and provenance record are correct on every node regardless of the rates observed.

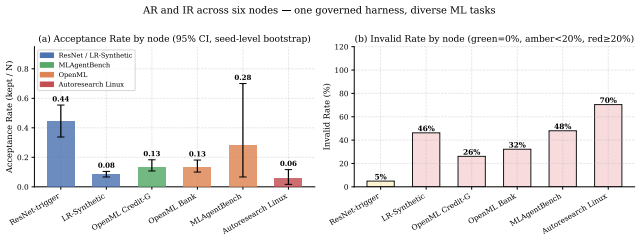


Figure S2: Cross-node acceptance rate (AR) and invalid rate (IR). Bootstrap CIs use 10,000 seed-level resamples. Colour denotes task family. AR and IR reflect proposal quality and task difficulty, not governance failures.

B.4 LR-Synthetic Memory Ablation

Figure S3 shows the memory ablation on `lr_synthetic` (450 trials: three arms, five seeds per arm, 30 trials per seed; Table S2). Acceptance is low in every arm. The rationale arm has the highest mean acceptance rate, and its seed-level bootstrap interval lies above that of the no-memory arm; with five seeds per arm this difference is descriptive and untested. The control plane classifies every trial in all three arms, so governance transfers to this node whatever its memory response, which differs from the ResNet response (Section 5.5).

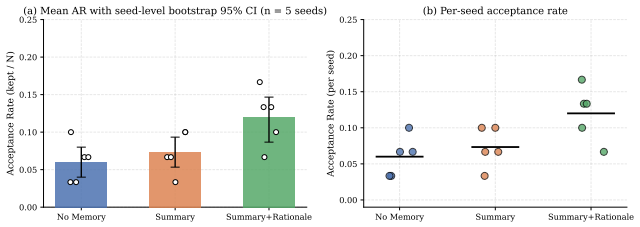


Figure S3: LR-synthetic memory ablation (150 trials per arm: 5 seeds of 30 trials). Left: mean AR with seed-level bootstrap 95% CIs; right: per-seed AR. The rationale arm's interval lies above the no-memory arm's (descriptive; five seeds per arm).

B.5 ResNet L40S Best-AUC Trajectory

Figure S4 shows the best-AUC trajectory of each arm and seed of the L40S replication (Section 5.5).

B.6 Large-Scale L40S Replication

Table S5 summarises the L40S replication by memory mode.

B.7 Autoresearch Linux Per-Seed Results

Table S6 gives the per-seed outcomes of the `autoresearch_linux` campaigns (Section 5.6). The ledger of summary seed s2 is the only

Table S5: Large-scale memory ablation (ResNet-trigger, NVIDIA L40S, 75 trials per arm = 3 modes × 5 seeds × 15 trials). Reset integrity confirmed by shared `node_state_hash` (`9c0118f33fcb`) at trial 1 across all 15 campaigns. Trial-1 `val_auc` is given per seed, s1 to s5, from the first record of each campaign's ledger; it is not comparable with the case-study values, which use the CPU fast-search mode (Section 5.5). σ is per-seed best-AUC standard deviation. 95% CIs use 10,000 seed-level bootstrap resamples (seed 42). Manager: `deepseek-v4-flash`; worker: `qwen2.5-coder:7b` at temperature 0.2.

Memory mode	Seeds	N	Acc%	Inv%	Trial-1 AUC (s1, s2, s3, s4, s5)	Mean best AUC	95% CI	σ
none	5	75	17.3	0.0	0.8878, 0.8891, 0.8890, 0.8872, 0.8873	0.899	[0.892, 0.907]	0.009
summary	5	75	64.0	1.3	0.8885, 0.8878, 0.8877, 0.8865, 0.8874	0.939	[0.937, 0.940]	0.002
rationale	5	75	58.7	0.0	0.8885, 0.8878, 0.8883, 0.8892, 0.8893	0.941	[0.940, 0.943]	0.002

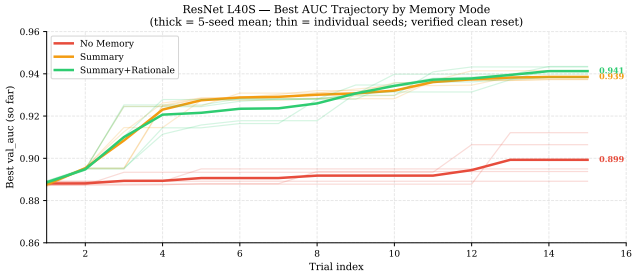


Figure S4: ResNet L40S best-AUC trajectory by memory mode (3 modes × 5 seeds, verified reset). Thick: 5-seed mean; thin: individual seeds. Memory arms diverge from the no-memory arm by trial 5 and hold the gap through trial 15, consistent with the 4–5× smaller cross-seed standard deviation reported in Section 5.5.

primary ledger with irregular budget indices: index 1 has no record, and index 2 has two records with the same trial ID, one for each time the slot was opened. The first was appended by the pending-trial recovery after an earlier run crashed while that trial was pending; the second records the trial executed after the restart. All 30 records are counted, so one of the 1,445 primary records was written by the recovery tool rather than at the end of a completed trial.

B.8 OpenML Per-Budget Summary

Table S7 summarises the OpenML campaigns by node and budget.

B.9 OpenML Dataset Statistics

The class imbalance in the bank-marketing dataset (88%) makes acceptance rate a more informative primary metric than raw accuracy; the control plane's keep rule (candidate > current best) is applied to `val_AUC`, which is robust to imbalance. Both datasets are public on OpenML (IDs 31 and 1461), where their licence is listed as "Public", and originate from the UCI Machine Learning Repository.

B.10 OpenML Running-Best Trajectories

Figure S5 shows running-best `val_AUC` trajectories for both OpenML nodes. The step function is monotonically non-decreasing because the control plane accepts only valid improving trials — the keep rule is enforced structurally, not by the manager. Individual seeds (thin blue lines) and the 5-seed mean (red) show the variance compressed by the governed campaign discipline.

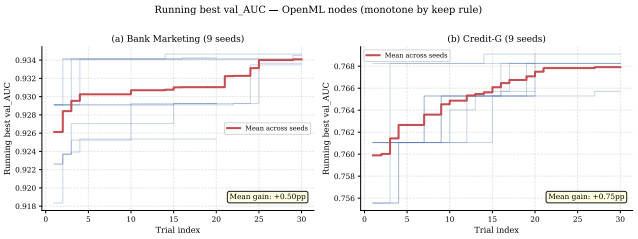


Figure S5: OpenML running-best val_AUC trajectories across all seeds. Monotone non-decrease is guaranteed by the control-plane keep rule. (a) Bank Marketing; (b) Credit-G.

C Control-Plane Algorithm

C.1 Trial Lifecycle

Figure S6 gives the pseudocode for one governed trial. The control plane is the only writer to the append-only ledger; the manager never transitions state directly.

1. Write a pending-guard file for this budget slot.
2. Call **Manager.propose**(memory context, node spec) → proposal.
3. **Scope-check:** if proposal touches any frozen path → reject.
4. **No-op check:** if patch is empty or byte-identical → reject.
5. **Precondition check:** validate proposal fields → reject if malformed.
6. Call **Worker.execute**(proposal, node spec) → patch + raw log.
7. If worker fails → record `failed_invalid/runtime_error` or `edit_failed` and go to step 10.
8. Parse task metric from raw log.
9. If metric missing → `failed_invalid/metric_missing`. If metric ≤ current best → `discarded`. Else → `kept`; update current best.
10. Append one immutable ledger record (fields: proposal, patch ref, log ref, metric, decision, failure category, provenance IDs, node-state hash).
11. Delete pending-guard file.
12. Update append-only memory context from completed record.

Figure S6: Trial lifecycle pseudocode (one budget slot).

Steps 3–5 are pre-run guards: they prevent invalid edits from corrupting node state and produce classified failure records before any training occurs. Steps 7–9 are post-run classification: a valid but non-improving trial becomes `discarded`, not `failed_invalid`. The manager is not consulted at any classification step.

C.2 Control-Plane Trial Decision Timeline

Figure S7 illustrates the one-trial-one-record discipline on a representative campaign (OpenML Bank Marketing, seed 1). Every trial

Table S6: Autoresearch Linux node (NVIDIA L40S, `deepseek-v4-flash` manager, `claw_style_worker`, 4 seeds × 30 trials per arm). AR = acceptance rate; IR = invalid rate; Gain = best − initial `val_bpb` (negative = improvement). Seed s2 of the summary arm suffered total worker failure; task-metric gain summaries use s1/s3/s4. AR uses the definition of Table 5, |kept|/N.

Arm	Seed	Kept	Disc	Fail	AR	IR	Gain
`none`	s1	0	0	30	0.00	1.00	—
	s2	0	0	30	0.00	1.00	—
	s3	0	0	30	0.00	1.00	—
	s4	0	0	30	0.00	1.00	—
`summary`	s1	4	17	9	0.13	0.30	−0.0616
	s2[†]	0	0	30	0.00	1.00	—
	s3	5	18	7	0.17	0.23	−0.0607
	s4	5	22	3	0.17	0.10	−0.0642

[†] s2 suffered total worker failure (27 `runtime_error` + 3 `proposal_precondition_failed`); no initial `val_bpb` was recorded. These campaigns predate the frozen-baseline restore protocol, and the archived trial-1 hashes do not verify a common start state across all arms and seeds; s2, like the total failure of the none arm, is therefore an observed failure of undetermined cause, not memory-effect evidence. Governance correctly classified all 30 trials as `failed_invalid`.

Table S7: OpenML campaign summary by public node and budget. Seed-level bootstrap confidence intervals use 10,000 resamples (seed 42).

Node	Budget	Seeds	Trials	AR%	IR%	Mean best AUC	95% CI
`credit_g`	b20	4	80	20.0	8.8	0.7671	[0.7655, 0.7687]
	b30	5	150	10.0	35.3	0.7679	[0.7667, 0.7688]
`bank_mkt`	b20	4	80	20.0	17.5	0.9295	[0.9263, 0.9330]
	b30	5	150	9.3	40.0	0.9341	[0.9337, 0.9345]

Table S8: OpenML dataset characteristics for the two public classification nodes used in Section 5.

ID	Name	N	Feats.	Task	Maj. cls.
31	credit-g	1,000	20	Binary	70%
1461	bank-marketing	45,211	16	Binary	88%

produces exactly one terminal record. The running best (step function) advances only when the control plane accepts a valid improving trial; it never decreases. This monotone property is structural: the keep rule `candidate > current_best` is evaluated and applied exclusively by the control plane.

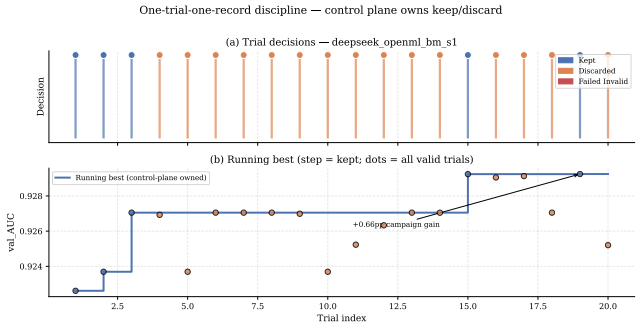


Figure S7: Trial decision timeline for one governed campaign (OpenML Bank Marketing, seed 1, 20 trials). Top: per-trial decision colour (kept / discarded / failed_invalid). Bottom: running best val_AUC step function and per-trial metric dots. The keep/discard decision is owned by the control plane, not the manager.

C.3 Governed vs. Ungoverned Counterfactual Visual

Figure S8 is the visual companion to the counterfactual comparison of Table 9.

C.4 ResNet L40S Per-Seed Panels

Figure S9 shows the per-seed acceptance and invalid rates and the per-seed best AUC of the L40S replication.

C.5 Decision Rule

The decision function `decide_trial()` is purely functional with no side effects. It receives the candidate metric, the current-best metric, and a validity flag, and returns exactly one terminal state:

C.6 Memory Context Algorithm

The memory context is built by `build_memory_context()` and injected into the manager at each trial. Three modes produce qualitatively different inputs:

C.7 Repeated-Bad Rate Detection

The repeated-bad rate (RBR) is a governance metric that measures how often the manager repeats the same edit target and mechanism as a prior rejected trial without providing a new corrective rationale.

Definition. Let $P_i = (t_i, d_i)$ be the edit target and direction of proposal i (e.g., target = `LEARNING_RATE`, direction = *increase*). Proposal i is *repeated-bad* if:

(1) A prior proposal P_j, $j < i$, had the same target $t_j = t_i$ and direction $d_j = d_i$; and
(2) P_j received decision `discarded` or `failed_invalid`; and
(3) Proposal i does not include a rationale that references the prior rejection.

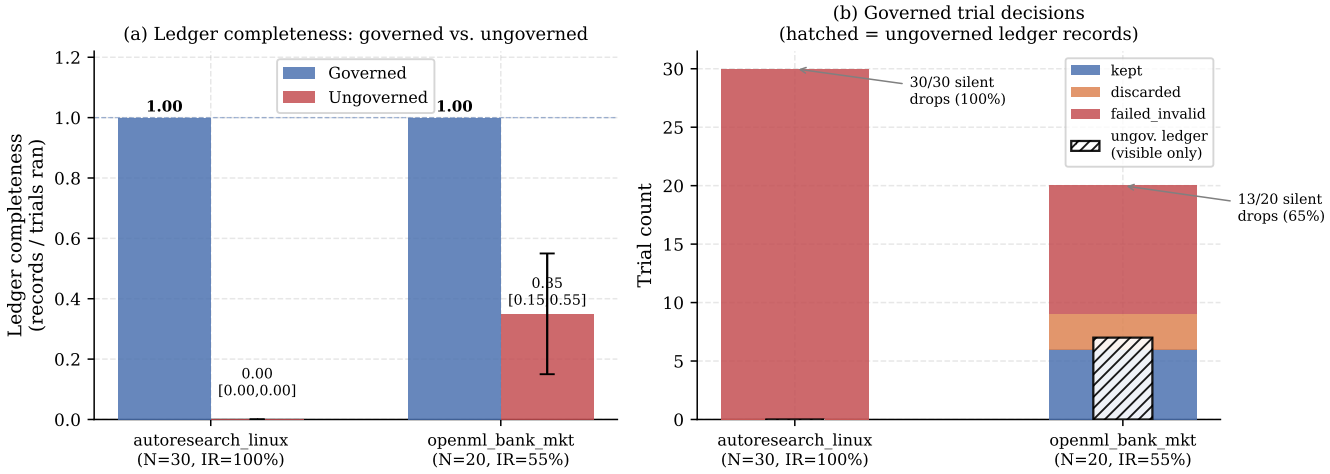


Figure S8: Governed vs. ungoverned counterfactual (visual companion to Table 9). Left: ledger completeness with 95% bootstrap CIs; the governed CI collapses to $[1.00, 1.00]$ by structural invariant, while the ungoverned CI lies strictly below 1.00 on both nodes. Right: governed trial decision timeline with silent-drop positions annotated. Combined: 43 / 50 ungoverned trials, all of them failed trials, leave no ledger record; the two campaigns were chosen to span failure rates.

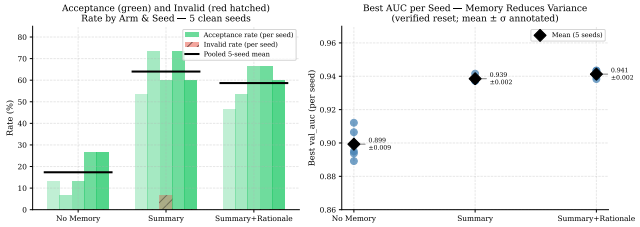


Figure S9: ResNet L40S memory ablation across five clean seeds with verified reset. Left: per-seed acceptance (green) and invalid (red) rates. Right: per-seed best AUC, showing $\sigma_{\text{none}} = 0.009$ vs. $\sigma_{\text{summary}} = \sigma_{\text{rationale}} = 0.002$ (a 4–5× smaller standard deviation; descriptive, five seeds per arm).

if validity flag is FALSE (any pre-run or post-run guard fired) → return `failed_invalid` with the guard's `FailureCategory`.

elif `candidate_metric` > `current_best_metric` → return `kept`; update current best to `candidate_metric`.

else → return `discarded` with `failure_category` = `degraded_metric`.

The function never raises an exception and never reads or writes the ledger. All side effects (ledger append, memory update, best-metric update) happen in the control-plane campaign loop after `decide_trial()` returns.

Figure S10: Decision rule pseudocode (`control_plane/decision.py`).

RBR = $|\{i : P_i \text{ is repeated-bad}\}|/N$ where N is the campaign budget.

Implementation. The detector in `memory/similarity.py` uses Jaccard similarity on unigram token sets of proposal summaries to identify near-duplicate proposals, then checks whether the candidate's summary token set S_i overlaps with any prior rejected proposal S_j above a threshold $\theta = 0.6$:

$$\text{Jaccard}(S_i, S_j) = \frac{|S_i \cap S_j|}{|S_i \cup S_j|} \geq \theta$$

Parameter direction is extracted by regex from the proposal rationale field (keywords: *increase*, *decrease*, *reduce*, *raise*, *lower*). Proposals that lack a detectable direction are not flagged.

Mode: none

Return a context containing only the NodeSpec parameters (editable symbols and their current values). No prior trial history is included. The manager has no information about past proposals or outcomes.

Mode: summary

Append one line per completed trial in chronological order:

```
Trial {i}: {decision} - {proposal_summary}
  metric: {parsed_metrics}
```

The manager receives the full history of proposals and outcomes, but with no explanation of *why* each trial succeeded or failed.

Mode: rationale (summary + rationale)

Same as summary, plus per-trial failure rationale lines:

```
Trial {i}: {decision} - {proposal_summary}
  metric: {parsed_metrics}
  rationale: {proposal_rationale}
  [AVOID: repeated {target} change - prior rejection]
```

The `[AVOID: ...]` warning is appended whenever the repeated-bad detector flags the current trial as semantically repeating a prior failed edit (see Section C.7). Rationale tokens can be truncated to a configurable `rationale_max_tokens` limit (50 tokens in the verbosity ablation).

Figure S11: Memory context algorithm (`memory/summarizer.py`).

C.8 Manager Prompt Template

The LangGraph manager (`manager/langgraph_manager.py`) builds a single-turn prompt from five components at each trial. Below is the template structure; bracketed tokens are filled at runtime.

```
You are a research manager proposing one bounded
hyperparameter change.

Node: [node_spec.name]
Description: [node_spec.description]
Metric to optimize: [metric_name] (direction: [direction])
Files you may instruct the worker to edit: [editable_paths]
Budget index (trial number): [budget_index]
Current best [metric_name]: [current_best_metric]

Editable values currently in [editable_file]:
  PARAM_A = current_value_A
  PARAM_B = current_value_B
  ...

# --- if memory_context is non-empty: ---
Prior trial memory:
[memory_context.context_text]

AVOIDANCE RULE (mandatory): Study the trial history
above carefully. Identify every hyperparameter that
was already tried in the same direction (e.g.
increasing BATCH_SIZE, reducing LEARNING_RATE) and
produced no improvement or a failure. You MUST NOT
propose the same parameter change in the same
direction again. Choose a DIFFERENT hyperparameter
or a DIFFERENT direction. If every direction for a
parameter has been explored, switch to a different
parameter entirely.

# --- if memory_context is empty: ---
No prior trial memory available.

Propose exactly one bounded experiment. Respond with
ONLY a JSON object (no markdown, no extra text):
  "summary"  -- short slug e.g. "reduce-lr-2e-4"
  "rationale" -- one sentence on why this might help
  "objective" -- full worker instruction:
                "In FILE, change PARAM from X to Y.
                 Edit only FILE."
  "param"    -- exact editable value name
  "old_value" -- current value as listed above
  "new_value" -- proposed replacement value

Constraints:
  - edit only: [editable_paths]
  - change exactly one listed value
  - param must be one of the editable values above
  - do not run the experiment; the harness runs it
```

The AVOIDANCE RULE is injected only when prior trial memory is non-empty (i.e., after the first trial). It operates on the memory context produced by the mode selected for that campaign arm. In the `none` mode the rule is absent; in the `summary` and `rationale` modes it is present and the manager is expected to honour it. The RBR metric (Section C.7) measures empirical compliance.

C.9 Scope Validation

The control plane's scope validator checks the file set touched by the patch against the NodeSpec's `editable_paths` whitelist and `frozen_paths` blocklist. A patch is rejected if it touches any file outside the whitelist or any file in the blocklist, regardless of the nature of the change. This fires before the worker executes training, so frozen files (dataset pipelines, split logic, evaluation harness) are never modified.

C.10 Crash-Recovery Path

If the process crashes mid-trial (e.g. OOM kill, host reboot), the pending-guard file persists on disk. At the start of the next campaign run, `PendingTrialError` is raised. The operator runs:

```
python3 scripts/recover_pending.py --list
python3 scripts/recover_pending.py --inspect \
    experiments/ledgers/<id>_pending.json
python3 scripts/recover_pending.py --mark-failed <id> \
    --reason "OOM during training"
```

This appends a `failed_invalid` record for the stale trial before any new trial begins, preserving the invariant that every opened budget slot produces exactly one terminal ledger record. The guard file is deleted only after the failure record is written.

D NodeSpec Contract Specification

D.1 Schema

A NodeSpec is a JSON-serialisable YAML file under `configs/nodes/`. The control plane reads it once at campaign start and treats it as immutable for the duration of the campaign. Table S9 lists the required and optional fields.

Table S9: NodeSpec YAML field schema. All fields are required unless marked optional.

Field	Type	Meaning
`name`	string	Unique node identifier used as campaign ledger key.
`description`	string	Human-readable node purpose.
`editable_paths`	list	Files the worker may edit; scope validator enforces this whitelist.
`frozen_paths`	list	Files and directories the worker must never touch; scope validator enforces this blocklist.
`run_command`	string	Shell command executed by the worker to train and produce output.
`metric_name`	string	Key expected in parsed metric output (e.g. `val_auc`).
`metric_direction`	string	`maximize` or `minimize`; determines the keep/discard rule.
`metric_parser`	string	Python import path to the parser function (`module:function`).
`acceptance_rule`	string	Expression evaluated by decision module (e.g. `candidate_metric > current_best_metric`).
`default_budget`	object	Default `trials` and `max_wall_clock_hours` (both optional at invocation time).
`setup_command`	string	Optional pre-campaign setup (e.g. `uv sync`).
`editable_symbols`	list	Optional: specific named symbols inside editable files (used by LocalWorker).
`failure_categories`	list	Exhaustive list of valid `FailureCategory` values for this node.

D.2 Annotated Example: `resnet_trigger`

This is the primary scientific node. The worker may only edit `train.py`; all other files including datasets, split logic, and architecture are frozen.

```
name: resnet_trigger
description: >
  Near-threshold detector waveform binary-classification
      benchmark
```

```
  with frozen H5 signal/noise traces and a single editable 1
       D ResNet
  training entrypoint.
editable_paths:
  - train.py
frozen_paths:
  - prepare.py
  - resnet_1d.py
  - signal_vacuum_sum_crop_4000x8000.h5
  - noise_traces_4000x8000.h5
  - artifacts/
  - data/
run_command: >
  ~/anaconda3/bin/python3 train.py > run.log 2>&1
metric_name: val_auc
metric_direction: maximize
metric_parser: >
  autoresearch.nodes.resnet_trigger.metric_parser:
       parse_val_auc
acceptance_rule: candidate_metric > current_best_metric
default_budget:
  trials: 50
  max_wall_clock_hours: 12
```

The metric parser converts `val_bpb` to `val_auc` via $1 - \texttt{val_bpb}$ when both values are present, preferring the `val_bpb` reading. This indirection allows the same governance infrastructure to handle both maximise- and minimise-direction nodes.

D.3 Annotated Example: `lr_synthetic`

This is the dependency-free transfer node. The worker edits `train.py` but is further restricted to four named symbols; the data-generation constants are frozen by scope enforcement and by the `editable_symbols` hint used by LocalWorker.

```
name: lr_synthetic
description: >
  Synthetic binary-classification node (NumPy logistic
       regression).
  Dependency-free governance-protocol validation.
editable_paths:
  - train.py
editable_symbols:
  - LEARNING_RATE
  - REGULARIZATION
  - N_EPOCHS
  - BATCH_SIZE
frozen_paths:
  - artifacts/
run_command: python3 train.py
metric_name: val_score
metric_direction: maximize
metric_parser: >
  autoresearch.nodes.lr_synthetic
  .metric_parser:parse_val_score
acceptance_rule: candidate_metric > current_best_metric
default_budget:
  trials: 10
  max_wall_clock_hours: 1
```

Adding a new node requires only a YAML file in `configs/nodes/` and a registration entry in `src/autoresearch/nodes/registry.py`. The control-plane contract (lifecycle, guards, ledger format) is identical.

E Failure Taxonomy with Concrete Examples

The control plane assigns a `FailureCategory` label at the terminal transition that writes the ledger record. The taxonomy is exhaustive for the current NodeSpec contract: every `failed_invalid` outcome maps to exactly one of the categories below. Table S10 lists each category, its detection point, and a representative ledger excerpt.

E.1 Concrete Ledger Records

The following records are drawn verbatim from the campaign ledgers (provenance hash suffixes truncated for readability; full records are in `experiments/ledgers/`). Each record shows the seven fields that identify a trial's terminal outcome. The `metric_missing` category has not been observed in any real campaign to date; all zero-metric outcomes in the released ledgers are classified as `runtime_error` (the worker exits non-zero before the metric parser runs).

Kept (campaign `kdd_resnet_scientific_20`, trial-001):


```
{"trial_id": "kdd_resnet_scientific_20-trial-001",
 "campaign_id": "kdd_resnet_scientific_20",
 "budget_index": 1,
 "decision": "kept",
 "failure_category": null,
 "proposal_summary": "lower-weight-decay",
 "parsed_metrics": {"val_auc": 0.774711}}
```


Discarded (campaign `kdd_resnet_scientific_20`, trial-002):


```
{"trial_id": "kdd_resnet_scientific_20-trial-002",
 "campaign_id": "kdd_resnet_scientific_20",
 "budget_index": 2,
 "decision": "discarded",
 "failure_category": "degraded_metric",
 "proposal_summary": "small-dropout",
 "parsed_metrics": {"val_auc": 0.773778}}
```


invalid_edit_scope (campaign `kdd_stress_scope`, trial-001):


```
{"trial_id": "kdd_stress_scope-trial-001",
 "campaign_id": "kdd_stress_scope",
 "budget_index": 1,
 "decision": "failed_invalid",
 "failure_category": "invalid_edit_scope",
 "proposal_summary": "lower-learning-rate",
 "parsed_metrics": null}
```


The scope guard fires before the worker executes; node state is never modified and training never runs.

no_op_patch (campaign `kdd_stress_noop`, trial-001):


```
{"trial_id": "kdd_stress_noop-trial-001",
 "campaign_id": "kdd_stress_noop",
 "budget_index": 1,
 "decision": "failed_invalid",
 "failure_category": "no_op_patch",
 "proposal_summary": "lower-weight-decay",
 "parsed_metrics": null}
```


proposal_precondition_failed (campaign `deepseek_autoresearch_linux_append_only_summary_s1`, trial-012):

Table S10: Failure taxonomy: detection point, guard type, and representative ledger evidence. All examples are drawn from real campaign ledgers in the released evidence; fields are formatted for readability.

Category	Fires at	Guard	Representative ledger evidence
`invalid_edit_scope`	Pre-run	Scope validator	Campaign `kdd_stress_scope`, trial-001: proposal attempts to edit `resnet_1d.py` (frozen); scope guard rejects; decision is `failed_invalid`; patch is never applied; node state is unchanged.
`no_op_patch`	Pre-run	No-op guard	Campaign `kdd_stress_noop`, trial-001: worker produces a patch that is byte-identical to the current `train.py`; no-op guard rejects; decision is `failed_invalid`; training is never executed.
`proposal_precondition_failed`	Pre-run	Precondition check	Campaign `deepseek_autoresearch_linux_append_only_summary_s2`, trials 28–30: manager produces malformed proposals after repeated `runtime_error` failures; precondition guard fires before worker call.
`runtime_error`	Post-run	Metric/exit check	Campaigns `deepseek_autoresearch_linux_none_s1–s4`, all 120 trials: the worker exits non-zero (cause undetermined; see Section 5.6); provenance IDs complete in all 120 records.
`invalid_config`	Post-run	Node validator	OpenML nodes reject hyperparameter values outside their bounded configuration domains, e.g. excessive tree depth or estimator counts. The run emits a machine-readable `invalid_config` label, which the control plane records as `failed_invalid`.
`metric_missing`	Post-run	Metric parser	Occurs when a run completes (exit zero) but the log file contains no parseable metric line, e.g. a training script that writes to `stderr` only or silently skips the validation epoch. Decision is `failed_invalid`; the raw log ref is retained for inspection.
`edit_failed`	Post-run	Worker path	ClawWorker only: the coding agent produces a patch file but fails to apply it cleanly (merge conflict, syntax error introduced by the patch). This is a worker-path artifact; it does not appear in LocalWorker or LangGraph deterministic-patch campaigns.
`degraded_metric`	Post-run	Decision rule	*Not* `failed_invalid`: valid run, metric parses, but does not beat current best. Decision is `discarded`. E.g. campaign `kdd_resnet_scientific_20`, trial-002: patch runs successfully, `val_auc` = 0.7738 vs. current best 0.7747; trial is discarded with full artifact retention.


```
{"trial_id":
   "deepseek_autoresearch_linux_append_only_summary_s1-trial-012",
 "campaign_id":
   "deepseek_autoresearch_linux_append_only_summary_s1",
 "budget_index": 12,
 "decision": "failed_invalid",
 "failure_category": "proposal_precondition_failed",
 "proposal_summary": "reduce-total-batch-size-32768",
 "parsed_metrics": null}
```


The precondition check fires when the manager produces a proposal whose field values fail validation (e.g. `new_value` is identical to `old_value`, or a required field is absent). No worker call is made.

runtime_error

(campaign `deepseek_autoresearch_linux_none_s1`, trial-001):


```
{"trial_id":
   "deepseek_autoresearch_linux_none_s1-trial-001",
 "campaign_id": "deepseek_autoresearch_linux_none_s1",
 "budget_index": 1,
 "decision": "failed_invalid",
 "failure_category": "runtime_error",
 "proposal_summary": "reduce-embedding-lr-0.01",
 "parsed_metrics": null}
```


All 120 trials in the `none`-memory arm of the autoresearch-linux node share this pattern: the worker exits non-zero and the control plane records `runtime_error`. The cause is undetermined because the start state of these campaigns is unverified (Section 5.6).

invalid_config

(campaign `deepseek_openml_bm_s2`, trial-005):


```
{"trial_id":
   "deepseek_openml_bm_s2-trial-005",
 "campaign_id": "deepseek_openml_bm_s2",
 "budget_index": 5,
 "decision": "failed_invalid",
 "failure_category": "invalid_config",
 "proposal_summary": "increase-n-estimators-700",
 "parsed_metrics": null}
```


The `invalid_config` category covers node-level constraint violations where the edited file is in scope but the proposed value is outside the NodeSpec's bounded domain. It appears in the primary OpenML evidence set and is therefore included in the taxonomy of Section 3.5.

E.2 Taxonomy Reliability Probe

To check that failure categories are structurally assignable from observable ledger evidence, we applied an independent rule-based relabeler to a stratified sample of 50 `failed_invalid` records spanning the observed validation categories. The relabeler uses only scope-validation flags, patch presence, training-initiation signals, and execution status — never the `failure_category` field itself. Cohen's κ between the control-plane assignments and the structural relabeler is $\kappa = 0.830$ (95 % CI $[0.703, 0.958]$), Landis–Koch "almost perfect" agreement. This is a structural agreement check, not a human inter-rater study. Perfect agreement ($\kappa = 1.0$) is achieved for `invalid_config`, `invalid_edit_scope`, and `proposal_precondition_failed`; the six discordant cases are `runtime_error` and `no_op_patch` edge cases where the distinguishing evidence (training seed, no-op message) is absent from the JSONL record but present in the referenced artifact files. Full human inter-rater validation remains future work.

E.3 Distinguishing Failed from Discarded

Key design choice: valid non-improving trials are `discarded`, not `failed_invalid`. This matters for audit:

- `failed_invalid`: a process problem (scope violation, coding failure, missing output) the operator should inspect.
- `discarded` trials indicate a valid experiment that did not improve the metric — a normal outcome in exploration campaigns with high discard rates.

Treating discard as failure would inflate the invalid rate and make it impossible to distinguish bad proposals from bad outcomes.

E.4 New Category Protocol

To add a new failure category:

(1) Add the enum label to `FailureCategory` in `memory/schemas.py`.
(2) Add detection logic in `control_plane/campaign.py` at the appropriate guard or post-run classification point.
(3) Add the name to `failure_categories` in the NodeSpec YAML(s).
(4) Add a test in `tests/test_trial_schema.py`.

The new category must map to `kept`, `discarded`, or `failed_invalid` before the ledger append.