arXiv:2611.19037v1 [cs.LG] 14 Nov 2026

THE CRITIC CAN HURT: ON THE SAFETY OF CRITICS IN LLM-GUIDED AUTOML SEARCH

Gaoyuan Du[1] **Xiaoyang Liu**[1] **Rex Zhou**[1] **Dengdu Jiang**[1] **Rong Fu**[2] **Meng Ding**[3] **Defu Cao**[4]

[1]Amazon [2]Independent Researcher
[3]University of Massachusetts Boston [4]University of Southern California

ABSTRACT

LLM critics increasingly steer iterative search in agentic AutoML, prompt optimization, and self-refinement—evaluated for whether they help, but rarely for whether they can hurt. We isolate the critic's marginal effect through a four-arm audit and characterize the loop's tolerance through controlled fault injection. We show a naive 8B critic drives search below the no-critic random-search floor on 7 of 8 tabular datasets, losing up to 0.245 AUC in the worst case; the magnitude varies strongly with dataset, frontier, and budget. *When*: stronger critics attenuate the loss on the evaluated suite, without establishing a universal capability threshold. *Why*: controlled fault injection shows that non-directional errors are nearly costless under best-so-far search, whereas directional errors consume finite search opportunities; an idealized effective-sample calculation yields a necessary condition, $q(m-1) > \beta$, not a deployment certificate. *How*: a simple observable validation gate and a learned detector can reduce realized loss, but neither is guaranteed to preserve the full-budget no-critic baseline. We will release CRITFAULT, to our knowledge the first benchmark for detecting harmful critic feedback in AutoML. A critic is not a free addition: audit before trusting, and gate when capability or verification is insufficient.

Keywords LLM critics · critic safety · iterative LLM-guided search · automated machine learning · LLM-as-a-judge

1 Introduction

Automated machine learning (AutoML) searches over preprocessing, feature transformations, model families, and hyperparameters to build predictive pipelines with minimal human effort [10, 12, 42, 52]. Recent agentic systems [4, 17, 26, 54] organize this search as a *propose–evaluate–correct* loop: a proposer selects a candidate pipeline, the candidate is scored on held-out data, and an LLM *critic* converts the score and search history into recommendations that condition the next proposal, while the system retains a best-so-far incumbent. In principle, critique buys sample efficiency [19]. **We show that it can do the opposite: a naive LLM critic can drive search *below* the floor of using no critic at all.**

We study the critic rather than the proposer because the two fail asymmetrically. A weak proposal is evaluated, scores poorly, and is discarded by the best-so-far rule; the loop verifies proposals by construction. A critique receives no such check: it is typically neither scored nor filtered, yet it conditions every subsequent proposal. Its cost is therefore not a degraded incumbent—elitist selection protects what has already been found—but sustained misdirection of a finite evaluation budget away from configurations that unguided search would otherwise reach.

Evidence from adjacent settings suggests this risk is real but leaves it unmeasured. LLMs often fail to self-correct without an external signal [22], reported gains concentrate where a verifier is available [27], and LLM-as-judge scores are biased and unstable [51, 63]; verbal self-refinement reports similarly uneven results [37, 48]. That evidence, however, concerns *one-shot* judgment, whose cost is bounded by a single decision. Closed-loop evaluations, in turn, ablate feedback to measure how much it *helps* (e.g., SELA's search-policy ablation [4]); existing evaluations rarely test the implicit harmlessness assumption directly against an otherwise identical loop with the critic removed.

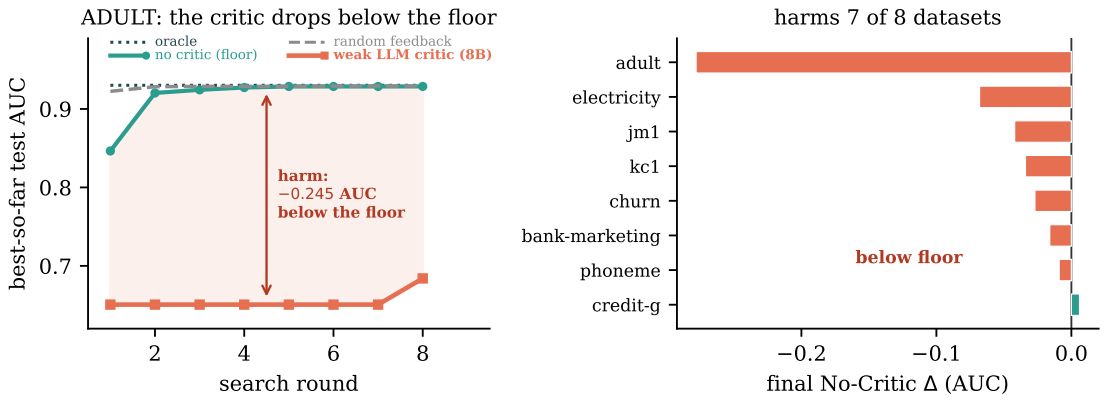


Figure 1: **A weak LLM critic actively harms search. Left:** worst case on ADULT—an 8B critic (coral, AUC 0.68) pins search far below no-critic random search (teal, 0.93) and an independently seeded second random search (gray). **Right:** the below-floor sign holds on $7/8$ datasets. Bars show the *No-Critic* Δ, the critic-guided minus no-critic final test AUC, so negative means below floor (used throughout).

To measure that margin, we introduce a *four-arm audit—no critic* (floor), a second independently seeded random search, *real critic*, and *oracle* (ceiling)—running an identical budget over the same cached configuration frontier (§3). We formalize the underlying question as *critic safety*: whether, and under what conditions, adding an LLM critic improves, preserves, or degrades the outcome of an otherwise identical search. The audit reveals a substantial failure mode. A naive 8B critic falls below the no-critic floor on 7 of 8 binary tabular datasets (Fig. 1). Typical losses are $\sim$0.03 AUC, reaching 0.245 in a deterministic worst case.

Two empirical factors modulate this harm (§4). First, *critic capability*: on a targeted six-dataset subset, a 14B critic improves on the 8B critic on five of six datasets, but stronger critics do not eliminate negative point estimates (a 70B critic is negative on five of six), so we find attenuation rather than a capability threshold. Second, the *search setting*: when the frontier grows from 96 to 250 configurations and the horizon from 8 to 20 rounds, the same critic's below-floor effect largely disappears; because both factors change together, we do not attribute this to headroom alone. The resulting trade-off is asymmetric: on these frontiers even an ideal critic gains at most $\sim$0.005 AUC over the floor, while the worst-case loss is roughly fifty times larger, and a practitioner cannot know *ex ante* which regime a given critic and task fall into.

Why does the loop tolerate some errors and not others? Controlled fault injection separates error rate from error structure (§5): because the loop keeps a best-so-far incumbent, non-directional errors act like unguided exploration and remain nearly costless even at high rates, whereas systematically adversarial guidance wastes effective search rounds and accumulates cost. An idealized effective-horizon analysis condenses this into the condition $q(m-1) > \beta$ (competence q, lookahead m, adversarial rate β), which is necessary for net benefit in the model and is the leading-order boundary when improvements over the incumbent are rare; it is not sufficient and does not characterize real critics exactly. Audits in text, forecasting, retrieval, and a second tabular task are consistent with the hypothesis that harm tracks mismatch between the critic's priors and the task.

The condition's one actionable variable is β. A training-free validation gate and a learned bad-feedback detector target it without retraining the critic (§6). In deployable end-to-end reruns that use only observed quantities, the simple gate reduces the raw critic's average loss relative to uniform search, while neither gate guarantees the full-budget no-critic baseline on every dataset.

This paper makes three contributions:

- **Critic safety and a marginal-effect audit (*when*).** We formalize *critic safety* and introduce the four-arm audit that isolates a critic's marginal effect on a shared frontier (§3–§4). The audit uncovers a below-floor failure mode that attenuates with critic capability and varies with frontier and budget, and observes below-floor results in the published CAAFE system (§6).
- **An error-structure explanation (*why*).** Controlled fault injection and an effective-horizon analysis show the loop is resilient to non-directional errors but not to systematically adversarial ones, yielding the necessary condition $q(m-1) > \beta$ and its tail approximation (§5). Cross-domain audits are consistent with its qualitative predictions.

- **Prevention and a benchmark (*how*).** A validation gate and a learned bad-feedback detector target β using only observed quantities and are evaluated end to end (§6). We also contribute CRITFAULT, to our knowledge the first benchmark for detecting harmful critic feedback in tabular AutoML, to be released with all artifacts upon publication.

An LLM critic is thus not a free addition but a component whose marginal effect must be measured. Before trusting one, audit it against a no-critic control at the intended budget; when critic capability or verifier reliability is insufficient, gate or omit the critique.

2 Related Work

Classical and LLM-based AutoML. Classical pipeline optimizers [10, 12, 13, 33, 42, 52] search configurations via Bayesian, evolutionary, or ensembling methods, with sample-efficient hyperparameter tuning [11, 24] and standardized suites [14, 55]; strong tabular learners [3, 15, 18, 20, 28] set the ceiling they chase. These optimize a validation metric directly and never reason about *why* a configuration failed. More recent systems add an LLM to drive pipeline construction: LLM-orchestrated search with targeted refinement [4, 40, 62], retrieval/planning with iterative execution [17, 21], search over plans or code [26, 54], and RL-trained MLE agents [36, 58]; LLMs also serve as feature engineers [19] or direct predictors [38], evaluated end-to-end by agentic benchmarks [2, 23, 46]. All use the LLM primarily as a *proposer*, assessed at the outcome level; our object is the complementary *critic* role—the *conditions* under which an added critic turns from helpful to harmful, which these works do not isolate.

Verbal feedback and self-critique. Reflexion [48], Self-Refine [37], ReAct [59], self-correction training [57], and trained critics [39] show verbal critique can help on coding/reasoning, while a consistent line of negative results shows intrinsic self-correction can be net-*negative*: it degrades accuracy once oracle labels are removed [22, 27], often turns correct answers incorrect [53], and self-verification is itself unreliable [50]. Closest to our thesis, an imperfect in-context evaluator can induce *reward hacking* [43]—a proxy-over-optimization inherent to inference-time selection such as best-of-N [29], with *self-preference* a concrete driver [44]. This evidence is mostly on math/code with intrinsic self-evaluation; we study an *externally anchored* critic on tabular data, sharpening "fails without external grounding" into a measured harm *below the do-nothing floor* with an idealized necessary condition for when a critic can help.

LLM-as-judge reliability and verifier-gated selection. A growing literature documents biases and instability in LLM-as-judge setups [16, 51, 63], while trained *verifiers* [7, 35, 60], self-taught evaluators [56], and prover–verifier games [30] gate generation. Gating on verifier reliability is now an active theme [31, 32, 47], but these operate as candidate reranking or task-level verification. Our Verification-Gated Critique differs in setting and mechanism: it gates *online, inside* the search loop from observed validation utilities, motivated by our net-benefit condition. Itkin [25] similarly show that *delayed* verification destabilizes multi-agent belief—independent evidence that verification, wrongly placed, can do harm. Relatedly, the generation–verification gap widens with compute [9, 49]; the capability attenuation we observe (§4) is a consistent but distinct phenomenon, measuring net harm to an external search rather than intrinsic self-improvement.

Noise tolerance in feedback-driven optimization. Robust preference optimization gives explicit noise thresholds—rDPO [5] tolerates flip rates below $1/2$, symmetric losses add robustness [41]. Our resilience threshold (§5) is structurally different: it arises from *elitist best-so-far* dynamics rather than gradient preference training, and decomposes bad feedback into random (nearly free) versus systematically adversarial (costly)—a distinction a symmetric flip-rate model cannot capture.

3 Problem Setup and Evaluation Protocol

We ask whether feedback from an LLM critic improves—or degrades—finite-budget search over a fixed space of tabular ML pipelines. This section defines the search loop and the two instruments the rest of the paper builds on: a *four-arm audit* that measures a real critic's end-to-end effect (§4) and *controlled fault injection* that varies feedback-error rate and structure independently, motivating the theory of §5.

3.1 Critic-guided search

A critic-guided AutoML agent tunes a tabular pipeline over T rounds (Fig. 2; Alg. 1 in App. C). Each round, a *proposer* selects an untried configuration a_t from a finite space $\mathcal{A}$—each a a complete pipeline (model family, feature-transformation set, hyperparameter preset)—which is trained and scored on a frozen validation split; an LLM *critic* then turns the score and history into a short natural-language correction (e.g., *"boosting underperforms here; try a neural family next"*) conditioning the next proposal. The agent returns the best-so-far validation incumbent $\hat{a}_T$, whose quality

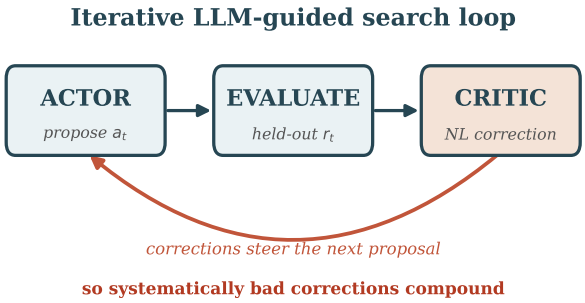


Figure 2: The iterative LLM-guided search loop we study.

is the hidden test utility $s_{\text{test}}(\hat{a}_T)$. For dataset D and condition r, we write $J_D(r) = \mathbb{E}[s_{\text{test}}(\hat{a}_T^{(r)})]$, the expected test utility over search randomness. Systems such as SELA [4] and AIDE [26] instantiate this propose–evaluate–correct loop on the premise that critique makes search more sample-efficient; we ask whether confidently-wrong critique can instead steer search *below* unguided exploration.

Two design choices isolate the quantity of interest: the critic acts *only* through the feedback channel (it never trains or executes code, alters the candidate space, or sees test utilities), so any effect is attributable to its guidance rather than tool-use failures; and the object of evaluation is therefore the critique's effect on the full trajectory, not the correctness of any individual critique. Formal history and feedback maps are deferred to App. C.

3.2 Instantiation and the shared cached frontier

Data flow. Each dataset is partitioned once into train/validation/test splits, and every configuration in a 96-configuration pipeline space $\mathcal{A}$ (design in App. C) is scored *before* search, with validation and test AUC cached. During search, selecting a configuration reveals only its cached validation AUC; test AUC stays hidden and evaluates only the final incumbent. Because all conditions on a dataset share this same frozen frontier—candidate space, split, budget, outcomes—they differ only in the trajectory-selecting mechanism, isolating search-policy differences from training randomness, execution failures, and pool differences (access rules in App. C.1).

Critic and proposer. In the experiments reported here, critique and proposal are realized by a single LLM call per round. The LLM receives the evaluated history (configuration names and their validation AUCs) and the menu of untried configurations, states its assessment, and names the next configuration. If the response names no untried configuration, the next configuration is drawn uniformly from the menu (a fallback). Without feedback the loop reduces to uniform sampling. The audited object is therefore this feedback-conditioned LLM selection policy as a whole, including its parser and fallback; we do not separate the effect of the critique from that of the mapping to a configuration. The primary naive critic is Llama-3.1-8B; the capability analysis of §4 evaluates seven tiers (3B–70B open-weight models and Claude Haiku 4.5) across four providers, all with deterministic decoding through a provider-agnostic interface (identifiers and settings in App. C).

3.3 Four-arm audit and controlled fault injection

Four-arm critic-safety audit. We compare four budget-matched conditions on the shared frontier that differ only in their feedback or selection mechanism: (i) *no critic*—no critique ($c_t = \varnothing$), uniform sampling over untried configurations; the unguided-search floor; (ii) *random feedback*—in the reported runs, an independently seeded uniform random search that does not pass feedback through the LLM; it is a second draw of the floor, not a semantic placebo; (iii) *real critic*—feedback from the LLM under evaluation, accepted without verification unless stated; and (iv) *oracle*—bypasses the language channel and directly selects the highest-validation-utility untried configuration; an unattainable ceiling. We report two paired contrasts,

$$\Delta_{\text{NC}}(D) = J_D(\text{real}) - J_D(\text{no-critic}),$$
$$\Delta_{\text{RF}}(D) = J_D(\text{real}) - J_D(\text{random-feedback}),$$

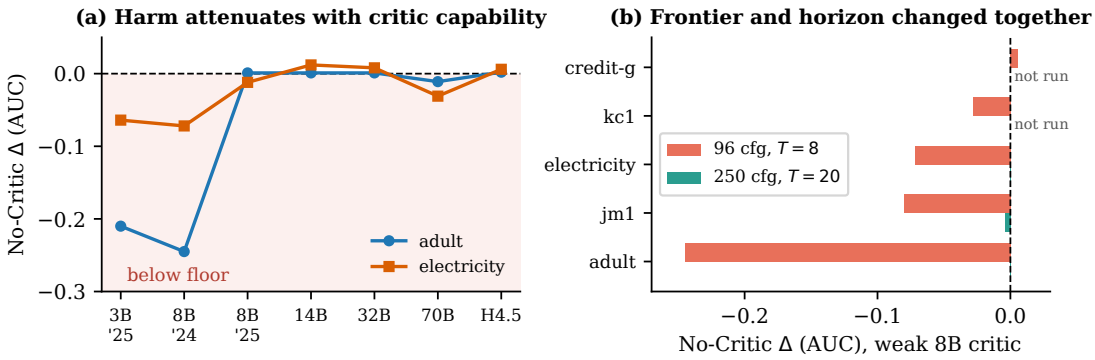


Figure 3: ***When*. (a)** Capability ladder (No-Critic Δ, 8 seeds): ADULT falls from -0.245 (2024-gen 8B) to $|\Delta| \leq 0.011$ at every $\geq$14B tier (Table 1). **(b)** The same 8B critic on the 96-configuration frontier with T=8 (coral) and, for three datasets, on a 250-configuration frontier with T=20 (teal); frontier and horizon change together, so the contrast does not isolate either factor.

where Δ_{NC} measures the system-level value of critic guidance over unguided search. Because arm (ii) is a random search, Δ_{RF} mirrors Δ_{NC} under a different seed and does not isolate the value of critique content; all conclusions below rest on Δ_{NC}. We say a critic is *below floor* when $\Delta_{\mathrm{NC}}(D) < 0$; prevalence statements ("7 of 8 datasets") refer to signs of dataset-level point estimates, while the headline inferential claim rests on a suite-level sign-flip permutation test, with per-dataset bootstrap intervals reported alongside (details in App. C). Because all outcomes are cached, auditing an additional critic requires no repeated pipeline training.

Controlled fault injection. The audit cannot separate how *often* a critic errs from how its errors are *structured*, entangled in a real trajectory. We therefore corrupt a clean reference policy $\pi^\star$: each round follows $\pi^\star$ with probability $1 - p$ and a corrupted choice with probability p, sweeping p. Three structures span the error space—RANDOM (uniform untried; non-directional), WORST (lowest-utility untried; globally adversarial), and MISLEAD (lowest-utility incumbent neighbor; locally plausible but directionally harmful). These are diagnostic stress tests with privileged frontier information, not deployable policies; they realize the limiting error structures of §5. Degradation curves and the resilience threshold are defined in App. C.

3.4 Datasets and experimental scope

We evaluate on 12 public OpenML classification datasets spanning sample size, dimensionality, class imbalance, and domain (Table 12). We report in-space headroom $h_D = J_D(\text{oracle}) - J_D(\text{no-critic})$ (AUC, T=8, 96-configuration frontier) as a continuous quantity: it is below 0.01 on all 12 datasets, ranging from 0.0013 (ADULT) to 0.0077 (ELECTRICITY) on the eight binary datasets and LETTER, and ≤ 0.0003 on MFEAT-FACTORS, SEGMENT, and SPLICE. We refer to the latter three as the *near-zero-headroom* datasets; this is a descriptive grouping, not a pre-specified threshold. Each analysis uses an explicitly defined subset (Table 3): the audit the eight binary datasets, fault injection the other nine, the capability ladder a targeted six, and CRITFAULT all twelve. Unless stated, the primary audit uses $T = 8$, eight matched seeds, and deterministic decoding; exact settings in App. C.

4 *When*: Capability and Search Setting

We now characterize when critic guidance becomes harmful and when below-floor harm disappears. We examine two empirical factors: the search setting (frontier size and horizon) and the capability of the critic. All experiments use the shared configuration frontiers and evaluation protocol of §3.

4.1 Reference policies establish limited upside

We first quantify the improvement available from informative guidance. The clean hill-climbing reference achieves a positive but small No-Critic delta on each of the nine datasets outside the near-zero-headroom group ($+0.001$ to $+0.005$ AUC over 30 seeds, Table 15; suite-level Wilcoxon $p = 0.001$). This limited upside is not a weak search space:

Table 1: Capability ladder: No-Critic Δ (AUC, 8 seeds) across seven tiers spanning four providers and two generations (2024–2025). The 3B model was evaluated on four of the six datasets.

	3B[25]	8B[24]	8B[25]	14B	32B	70B	H4.5
adult	$-.210$	$-.245$	$+.001$	$+.001$	$+.001$	$-.011$	$+.002$
electr.	$-.064$	$-.072$	$-.012$	$+.012$	$+.008$	$-.031$	$+.006$
jm1	$-.008$	$-.080$	$-.003$	$-.002$	$-.000$	$-.011$	$+.002$
kc1	$+.000$	$-.028$	$-.000$	$-.001$	$-.001$	$-.023$	$-.000$
bank-mkt		$-.007$	$+.000$	$+.000$	$+.001$	$-.007$	$+.001$
credit-g		$+.006$	$-.005$	$-.012$	$-.007$	$+.006$	$-.001$

3B[25]=Ministral-3B (Mistral '25); 8B[24]=Llama-3.1 (Meta '24); 8B[25]=Ministral-3-8B (Mistral '25); 14B=Ministral-14B; 32B=Qwen3; 70B=Llama-3.3; H4.5=Haiku 4.5 (Anthropic).

the in-space oracle matches or beats AutoGluon-Tabular on 5 of 8 datasets (Table 6), so the floors we measure harm against are SOTA-competitive.

In-space headroom is small on every dataset (0 to 0.0077 AUC, §3.4); on the three near-zero-headroom datasets the oracle and the unguided-search reference are nearly indistinguishable. The reference results therefore establish an asymmetric opportunity: the attainable upside of guidance is small on these fixed frontiers, whereas the downside of systematically misleading guidance can be considerably larger.

4.2 Weak critics can fall below the search floor

Replacing the clean reference with a real LLM critic changes the outcome substantially. A naive Llama-3.1-8B critic (temperature 0) has a negative No-Critic delta on 7 of the 8 binary datasets, ranging from -0.009 AUC on PHONEME to -0.245 on ADULT; only CREDIT-G has a positive point estimate ($+0.006$; Fig. 1). In the representative ADULT trajectory the critic emits fluent rationales while its picks stay below the incumbent (App. G; §5.4).

An exact one-sided sign-flip test over all 2^8 sign assignments (§3.3) gives $p = 0.0078$ for a negative suite-level mean, so the negative direction is not driven by a single task. It persists under stochastic decoding and additional seeds, though the largest magnitudes occur under deterministic decoding; per-run logs are released with the code.

4.3 Harm attenuates with the search setting and critic capability

Harm attenuates on a larger frontier with a longer horizon. We evaluate the same 8B critic on an expanded 250-configuration frontier with a longer horizon of $T{=}20$. In this joint setting the below-floor effect decreases by one to two orders of magnitude: ADULT $-0.245 \rightarrow +0.0003$, ELECTRICITY $-0.072 \rightarrow +0.0003$, JM1 $-0.080 \rightarrow -0.004$. Because frontier and horizon change together, this contrast does not identify whether headroom, budget, or their interaction drives the attenuation. A separate stress test shows that enlarging the space does not mechanically reduce harm for a systematically adversarial critic and can amplify it (Table 7).

Harm attenuates with critic capability. We repeat the real-critic arm across a seven-tier capability ladder over four providers and two generations on six datasets (Table 1; Fig. 3a). The weak 2024 8B critic is negative on 5 of 6 datasets (up to 0.245 AUC); at the 14B tier every point estimate is within 0.012 of zero, although three of six remain negative. Relative to the 8B critic, the 14B critic improves on five of six datasets (exact one-sided paired sign-flip test, $p = 0.047$). The pattern is not monotone in size or release date: a 2025 3B critic stays harmful on 3 of 4 (ADULT -0.210), and Llama-3.3-70B is negative on five of six datasets (up to -0.031). Because these six datasets were chosen partly for the 8B critic's effect (App. C), the ladder supports task-dependent attenuation, not a capability threshold.

Robustness. The harm is neither a parsing-failure artifact (harm and parse-failure are anti-correlated across the suite; App. D) nor a naive-prompting one—chain-of-thought and self-consistency mitigate but do not eliminate it—a point we develop when motivating prevention (§6.1).

4.4 Magnitude and operational consequences

The harm is asymmetric relative to the available upside. On the fixed frontier, the largest gain from clean reference guidance is approximately 0.005 AUC (Table 15), whereas the naive 8B critic loses approximately 0.03 AUC on a typical harmed dataset and 0.245–0.28 in the worst case (Fig. 3a)—roughly six times the largest observed reference-policy gain typically, and up to fifty times in the worst case.

Below-floor harm is thus largest for a weak critic on these small, fixed frontiers at short horizons, and because upside is bounded while downside is not, the risk is asymmetric. The next section asks why the loop tolerates some critic errors but not others.

5 *Why*: Error Structure, Not Error Rate

The capability results identify when below-floor harm is likely, but not why some critic errors accumulate while others do not. We isolate error structure through controlled fault injection and derive its consequences from the loop's best-so-far dynamics. The intuition: the loop reports the best configuration seen so far, so a random error merely spends a round like the no-critic loop and cannot lower the incumbent, whereas a systematic error that steers away from the high-value region denies the incumbent its chance to improve round after round. A critic can thus err often yet stay harmless, or err rarely but directionally and do real damage. We make this precise across four levels of claim: exact structural results, a probabilistic bound, a first-order criterion, and empirical tests.

5.1 Fault injection reveals a wide but asymmetric safety envelope

The controlled interventions of §3.3 reveal two empirical regularities (Table 14, 30 seeds). First, RANDOM corruption is nearly costless: even at p=1, final degradation is below 0.005 AUC on every tested dataset. Second, systematically adversarial corruption (WORST/MISLEAD) exhibits a broad but finite resilience region: the measured threshold is $p^\star$=0.6–0.9 on the primary binary datasets, with LETTER at the boundary value 1.0. Error frequency alone therefore does not determine harm: errors at similar rates differ sharply depending on whether they preserve exploration or systematically redirect it—unlike the symmetric flip-rate thresholds of robust preference optimization (rDPO's $\varepsilon < 1/2$ [5]). We explain both regularities from the loop's structure.

Model. Let $\mathcal{A} = \{a_1, \ldots, a_N\}$ be a finite configuration space with fixed cached scores $s(a) \in [0, 1]$; a T-round run selects without replacement and reports the incumbent $b_T = \max_{t \le T} s(a_{\pi(t)})$. The elitist update (Alg. 1) gives the structural fact we exploit:

$$b_t \;=\; \max(b_{t-1},\, s(a_{\pi(t)})) \;\ge\; b_{t-1} \qquad \text{(monotone, "elitist").} \tag{1}$$

A corruption process replaces the selected configuration with probability p, independently across rounds: RANDOM substitutes a uniform untried configuration, ADVERSARIAL (WORST/MISLEAD) an argmin-type one. The no-critic floor selects uniformly at random every round.

5.2 Best-so-far search and random corruption

Proposition 1 (Uniform corruption preserves the random-search floor)**.** *Suppose every selected configuration is uniformly distributed over the remaining untried configurations, regardless of whether the round is labeled clean or corrupted. Then for every $p \in [0, 1]$ and horizon T,*

$$\mathbb{E}\big[b_T^{\mathrm{uniform}(p)}\big] \;=\; \mathbb{E}\big[b_T^{\textit{no-critic}}\big].$$

The proof (App. B) follows from without-replacement exchangeability: both loops induce the same distribution over the visited set, on which b_T depends alone.

This is an exact floor-preservation result for the uninformative regime a naive critic degenerates to; it does not claim that randomizing an *informative* critic is costless. It explains why non-directional corruption barely degrades search (random-mode cost within noise of zero on all 12 datasets, Table 14): random selections still explore the same pool and cannot lower an observed incumbent. Falling below the random-search floor therefore requires selection that departs from uniform in a harmful direction; the controlled WORST/MISLEAD interventions show that such directional error is sufficient to do so.

5.3 Adversarial corruption and the resilience threshold

For adversarial corruption we need only that a corrupted round contributes no improvement.

Assumption 1 (Wasted adversarial rounds)**.** An adversarial round selects a configuration a with $s(a) \le b_{t-1}$, i.e. it never raises the incumbent. (WORST picks the global argmin among untried; MISLEAD the worst neighbor of the incumbent; both satisfy this once b_{t-1} exceeds the low-score region.)

Empirically 90.5% of adversarial picks fail to raise the incumbent (per-dataset 87–93%), so the assumption holds for most rounds and the bound degrades gracefully otherwise. Let T_{sat} be the *saturation budget*: the number of effective rounds after which the clean loop is within δ of its ceiling ($\mathbb{E}[b^{\mathrm{clean}}_{T_{\mathrm{sat}}}] \geq b^\star - \delta$, $b^\star = \max_a s(a)$); empirically $T_{\mathrm{sat}} \approx$ 1–3 (§4.1).

Proposition 2 (Resilience under adversarial corruption)**.** *Let $K \sim \mathrm{Binomial}(T, 1-p)$ be the number of effective rounds. Under Assumption 1, if the clean loop is within δ of its ceiling whenever $K \geq T_{\mathrm{sat}}$, then, since scores lie in $[0,1]$,*

$$\mathbb{E}\,[\,b^\star - b_T\,] \;\leq\; \delta \;+\; \Pr\big[K < T_{\mathrm{sat}}\big].$$

Consequently, for a tolerated additional failure probability ϵ, the loop is $(\delta+\epsilon)$-robust whenever $\Pr[\mathrm{Binomial}(T, 1-p) < T_{\mathrm{sat}}] \leq \epsilon$, and we write $p^\star(\epsilon)$ for the largest such p. The simpler expression

$$p^\star \;\approx\; 1 - \frac{T_{\mathrm{sat}}}{T}$$

is the expectation-level approximation to this probabilistic threshold.

The proof (App. B) splits on $\{K \geq T_{\mathrm{sat}}\}$, using that corrupted rounds leave the monotone incumbent unchanged. With T=8 and $T_{\mathrm{sat}} \in [1,3]$ the approximation predicts $p^\star \in [0.63, 0.88]$, matching the measured 0.6–0.9 (Table 14): the loop is robust not because the critic is accurate, but because elitist monotonicity converts its job from never erring into supplying T_{sat} good rounds out of T.

5.4 The adversarial signature in real critics

Before building a criterion on the adversarial rate, we examine whether *real* critics show the systematically adversarial structure of Assumption 1. These diagnostics are consistent with directional mismatch but do not show that every real-critic step is adversarial. First, *family-selection divergence*: harmed trajectories under-select the families that dominate the cached frontier, with mean KL to the oracle family distribution of 8.4 for the harmful Llama-3.1-8B versus 3.2 for Claude Haiku 4.5, tracking the harm rate (Table 8); we read this as a diagnostic of directional mismatch, not a direct estimate of β. Second, a representative trajectory (App. G): after finding a strong gradient-boosted configuration in round one, the critic verbally endorses it every subsequent round yet proposes steadily worse alternatives—Assumption 1 realized by a real critic. Because the loop retains the incumbent, their cost is the forgone exploration unguided search would have reached.

5.5 A net-benefit criterion from effective search horizon

The following model is deliberately idealized (i.i.d. scores, independent round types); its purpose is to explain the observed qualitative asymmetry and derive intervention targets, not to predict the sign of every dataset-level effect. Model a critic by two per-round rates: with probability q it is *competent*, proposing the best of an effective m-configuration lookahead ($m \geq 1$, a proxy for how sharply it concentrates on good regions); with probability β it is *adversarial* (Assumption 1); otherwise it selects uniformly. The no-critic loop is the all-uniform special case (q=β=0).

Lemma 5.1 (Conditional effective sample size)**.** *Conditioned on K_c competent and K_u uniform rounds, the final incumbent has the same distribution as the maximum of $N_{\mathrm{eff}} = mK_c + K_u$ i.i.d. score draws.*

The proof (App. B) deletes adversarial rounds (they never enter the maximum) and treats each competent round as m exchangeable uniform draws.

Lemma 5.1 is exact. Taking expectations, $\mathbb{E}[N_{\mathrm{eff}}] = T\big(1 + q(m-1) - \beta\big) = T_{\mathrm{eff}}$, so $T_{\mathrm{eff}} > T$ exactly when $q(m-1) > \beta$. A per-round comparison with random search makes the role of this condition precise:

Proposition 3 (Effective-horizon criterion: necessary condition and tail boundary)**.** *Let F be the score distribution of a uniform draw and $\varepsilon(b) = 1 - F(b)$ the probability that a uniform round improves on incumbent b. Under the model above, a critic round improves on b with higher probability than a uniform round if and only if $\beta < q\sum_{j=1}^{m-1} F(b)^j$. Consequently the critic can be net-beneficial ($\mathbb{E}[b_T^{\mathit{critic}}] > \mathbb{E}[b_T^{\mathit{no\text{-}critic}}]$) only if*

$$\boxed{\,q\,(m-1) \;>\; \beta\,,}$$

equivalently $T_{\mathrm{eff}} > T$, and in the rare-improvement regime the per-round advantage is $\varepsilon\,[\,q(m{-}1) - \beta\,] + O(\varepsilon^2)$ as $\varepsilon \to 0$.

Table 2: Cross-domain test of the mismatch hypothesis (same 8B critic, four-arm protocol, 8 seeds; No-Critic Δ, 95% bootstrap CI). Credit scoring is an *in-domain control* (same domain, same 96-config space) yet neutral: only the original tabular pipeline-search setting is harmed.

Setting (space size)	Δ	95% CI	Verdict
Tabular AutoML (96)	-0.245	$[-0.27, -0.22]$	harm
Credit scoring (96)	$+0.000$	$[-0.001, +0.002]$	neutral
20NG text-pipe (36)	-0.003	$[-0.008, +0.001]$	neutral
20NG weak→strong (50)	$+0.003$	$[-0.005, +0.010]$	neutral
M4 forecasting (13)	$+0.001$	$[+0.001, +0.002]$	helps
RAG retrieval (54)	$+0.007$	$[+0.006, +0.008]$	helps

The condition is necessary, not sufficient: away from the tail, the higher-order terms matter. With Bernoulli$(1/2)$ scores and q=0.3, m=3, β=0.3, $q(m-1) = 0.6 > \beta$, yet a critic round reaches the top score with probability $0.4625 < 0.5$, so the critic is worse at every budget T (App. B). In the idealized Monte-Carlo simulation of App. B at T=8, the sign of $q(m-1) - \beta$ matches the sign of net benefit on $43/48$ grid cells.

The criterion organizes both axes: capability enters through q, m, error structure through β. Consistent with this, the family-selection divergence of §5.4 falls with critic tier, tracking the harm rate—though we do not directly identify (q, m, β) for real critics, treating the family KL only as a proxy for directional mismatch. Its single actionable variable, β, motivates the gates of §6, which we evaluate end to end; because gating also changes exploration, those experiments do not manipulate β in isolation. It also unifies the folklore heuristics as single-variable moves on one condition—stronger critic raises q, m; verification lowers β; breadth raises m (App. B).

The criterion does not include search headroom or the frontier–budget setting of §4, both of which modulate how strongly directional guidance appears relative to the no-critic floor (CREDIT-G illustrates the gap: a positive delta with a β-proxy indistinguishable from the harmed datasets; further scope discussion in §7).

The criterion also predicts that search *breadth* is a partial structural defense—parallel k-wide search dilutes any single adversarial pick and falls below floor on only $1/12$ datasets versus $9/12$ for sequential, tree-search, and multi-agent backbones (corollary and derivation in App. B; Table 9 in App. D).

5.6 A cross-domain test of the prior–reality mismatch hypothesis

If prior–task mismatch drives directionally adversarial guidance, then, holding the critic fixed, harm should track where its inductive preferences conflict with the task's performance landscape. We test this with the same 8B critic and four-arm protocol (8 seeds) across five further settings—credit scoring, text pipelines, forecasting, retrieval (Table 2). Only the original ADULT pipeline-search is clearly negative (-0.245); the in-domain credit-scoring control and both text settings are neutral, and forecasting and retrieval are slightly positive. Holding the critic fixed, model size cannot explain this variation; the pattern is consistent with the mismatch hypothesis, which we do not test directly (e.g., by anonymizing configuration names).

Broader implication. The analysis suggests a hypothesis for other systems that retain a best-so-far candidate or use feedback to accept, reject, or redirect future proposals: non-directional evaluator noise may be substantially less harmful than systematic bias. Extending the quantitative boundary $q(m-1) > \beta$ to self-refinement, verifier-guided decoding, or preference-based optimization requires setting-specific assumptions and remains future work.

6 *How*: Verification, Filtering, and Selective Abstention

6.1 Motivation: the harm appears in a published system and survives prompting

Before asking how to prevent harm, we test two natural objections: that it is specific to our cached frontier, and that it can be prompted away.

The failure reproduces in a published system. Running the published CAAFE system [19] with the same Llama-3.1-8B on the eight frozen splits, the LLM-guided system is negative on three datasets (ADULT -0.030, BANK-MARKETING -0.054, CREDIT-G -0.090), mean -0.021, no gain above $+0.004$. Because CAAFE differs in feature-engineering space, dynamics, and headroom, we read this as evidence that below-floor harm is not specific to our cached frontier,

not as a per-dataset replication. A stress test on two public regression tasks likewise goes negative, so the mechanism is not specific to classification or ROC-AUC (App. D).

Prompting and added compute do not reliably remove it. Chain-of-thought and self-consistency shrink the harm but do not reliably flip its sign: on ADULT, chain-of-thought lifts the delta from -0.278 to -0.041 yet three of four datasets stay below floor, and five-sample self-consistency spends $5\times$ the generations while doing *worse* than plain chain-of-thought (Table 10). Compute does not monotonically buy safety—the zero-call no-critic condition beats naive, chain-of-thought, and 40-generation variants alike (Fig. 4). This sets the prevention bar: remove the loss without extra critic generations, by targeting the directional-error mechanism of §5 directly.

6.2 Prevention objective and evaluation protocol

The prevention objective is deliberately asymmetric. Because the attainable upside on these frontiers is small (§4), directional error is costly (§5), and added reasoning does not remove the loss (§6.1), we prioritize non-inferiority to the no-critic floor over extracting small gains from every critic proposal. Filtering cannot create search headroom: on a frontier with little headroom, successful prevention may return the system to the floor without measurable positive gain.

Both methods in this section instantiate one gated policy. At each round a binary gate g_t decides whether the newly generated critique may condition the next proposal:

$$a_{t+1} \;=\; \begin{cases} \pi(H_t,\, c_t), & g_t = 1, \\ \pi_0(H_t), & g_t = 0, \end{cases} \tag{2}$$

where π is the proposer of §3.1 and π_0 is a critique-free fallback policy. The gate g_t is produced by a training-free validation rule (**VGC**, §6.3) or by a learned detector (§6.5). When the validation anchor is unstable, the policy may abstain from further critic-conditioned decisions and use the uniform fallback. This prevents additional directional interventions after the switch, but it does not recover evaluations already consumed by the critic-guided prefix. The two gates trade data for immediacy: the **learned detector** (§6.5) applies when historical labeled runs are available to supervise it offline; **VGC** (§6.3) needs no training data and gates online from the current run's validation signal—the cold-start default. Low-stability abstention is an empirical fallback, not a construction-level guarantee relative to full-budget random search.

These methods are motivated by $q(m{-}1) > \beta$, but we do not claim they manipulate β in isolation: gating also changes the history the proposer conditions on and the effective exploration policy, so the experiments evaluate the entire resulting policy end to end. We report three dimensions: *safety* (No-Critic delta of the final incumbent), *retained utility* (relative to the raw critic), and *cost* (configuration evaluations and critic generations); a policy has *no observed below-floor effect* when $\Delta_{\mathrm{NC}}(D) \geq 0$, an empirical description rather than a safety guarantee. The end-to-end evaluations (§6.6) use $T{=}8$, eight paired seeds, and a uniform reject branch.

6.3 Training-free verification-gated critique

VGC sets g_t from the most recent realized transition. The proposal $a_t = \pi(H_{t-1}, c_{t-1})$ is evaluated once (counted against the budget); VGC sets $g_t = 1$ if $s_{\mathrm{val}}(a_t)$ is within a tolerance τ_t of the incumbent's validation utility, and $g_t = 0$ otherwise, discarding c_t so round $t{+}1$ follows π_0. This is a transition-level trust update: it stops a poorly performing feedback-conditioned trajectory from propagating, rather than certifying an individual critique's correctness before use. Rejected rounds consume budget like any other round; the gate never triggers an additional same-round evaluation. configurations and the search history; it never accesses test utilities or the validation utilities of untried configurations.

6.4 Reliability-aware abstention

The gate overrules the critic using the validation anchor, so it should defer when that anchor is itself unstable. We measure anchor stability by $\widehat{\rho}_t$, the bootstrap stability of the validation-induced ranking over the configurations tried so far: validation examples are resampled, each tried configuration's utility is recomputed on the resample, and $\widehat{\rho}_t$ summarizes the rank agreement across replicates. Computing $\widehat{\rho}_t$ requires the per-example validation predictions of tried configurations; it uses neither test labels nor the scores of untried configurations observable proxy for whether the anchor can support gating, not as a calibrated estimate of test-set reliability. The tolerance interpolates

$$\tau_t \;=\; \tau_{\min} \;+\; (1 - \widehat{\rho}_t)\,(\tau_{\max} - \tau_{\min}), \tag{3}$$

so that when the ranking is stable the gate is strict ($\tau_t \to \tau_{\min}$) and when stability is low the system becomes less willing to overrule the critic with an unreliable signal. Because a widened tolerance still gates with an untrustworthy signal, a low-stability policy can abstain and use uniform proposals for the remaining rounds. This limits further exposure but does not make the complete trajectory distributionally equivalent to no-critic search.

6.5 A transferable harmful-feedback detector and CRITFAULT

With historical trajectories available, controlled intervention outcomes can supervise a learned gate. We package the fault-injection trajectories as CRITFAULT (∼3,840 labeled loop transitions across 12 datasets; labels, features, and a separate real-harm relabeling detailed in App. F), to our knowledge the first benchmark for detecting harmful critic feedback in tabular AutoML. The deployable detector reads only evaluated validation utility, incumbent-relative change, and round fraction. The two structural indicators (untried-utility rank and neighborhood membership) are retained only for offline diagnosis. Like VGC it produces g_t after evaluation and before the critique conditions the next proposal (Eq. (2)); unlike post-hoc failure monitors [45, 61], it prevents propagation online at no added LLM cost (a gradient-boosted model over the three deployable features).

With all five features, the leave-one-dataset-out AUC is 0.95 over the 12-dataset fault-injection benchmark and transfers across critic models (leave-one-model-out mean 0.988). On six datasets and 384 realized critic transitions with the same labels and folds, the fully deployable f0–f2 detector attains 0.982; adding the two offline structural features raises this to 0.992 (App. D). Classification AUC alone, however, does not establish a safe filter; the end-to-end test follows.

6.6 End-to-end prevention results

Paired VGC audit (T=8). We reran fixed and reliability-adaptive tolerances on eight datasets and eight paired seeds, using a uniform rather than validation-greedy reject branch. Adaptive minus fixed is -0.00006 AUC with a paired 95% CI of $[-0.00109, 0.00090]$; five of eight point estimates are positive. Relative to the uniform reference, fixed minus uniform is -0.00032 $[-0.00140, +0.00072]$ and adaptive minus uniform is -0.00038 $[-0.00166, +0.00078]$; both lower bounds exceed -0.002, which supports exploratory average non-inferiority at a 0.002 margin, not dataset-wise safety. Thus the data do not support adaptive superiority or a suite-wide safety guarantee. We retain VGC as an exploratory gate and treat the fixed observable rule as the primary training-free baseline.

Detector and simple gate online. In a deployable end-to-end rerun, both gates use only evaluated utility, incumbent-relative change, and round fraction; a trigger sends the next proposal to uniform search. Across six datasets, the raw policy averages -0.00541 AUC relative to uniform, the simple non-improvement gate $+0.00046$, and the learned detector -0.00015. The simple gate significantly improves over raw on bank-marketing and kc1; the detector significantly improves over raw on bank-marketing and electricity. Neither learned detection nor classification AUC alone establishes floor preservation.

Summary. Observable gating reduces the average loss in this suite, with the simple validation rule the most stable end-to-end baseline. Adaptive tolerance and learned detection can help individual datasets, but neither supports a universal non-inferiority claim. The operating thresholds reported here should be recalibrated when the critic, search space, or data regime changes.

7 Limitations and Ethical Considerations

Scope. The primary estimates concern finite-budget tabular search on small, fixed frontiers (96 configurations, T=8), where in-space headroom is below 0.01 AUC on every dataset. The below-floor effect shrinks when frontier and horizon grow together (§4.3), and we do not claim harm at long horizons; Proposition 2 predicts that directional corruption is tolerated once the budget is large relative to T_{sat}. The capability ladder uses six datasets chosen partly by the 8B critic's outcome, and per-dataset intervals quantify search randomness conditional on one frozen split.

What the audit isolates. In the reported experiments, critique and proposal share one LLM call, and the random-feedback arm is an independently seeded random search (§3). The audit therefore measures the complete feedback-conditioned selection policy, including its parser and fallback, rather than the content of the critique in isolation. The condition $q(m-1) > \beta$ is derived in an idealized model; (q, m, β) are not identified for real critics.

Broader impact. Tabular ML is the production default in high-stakes domains—credit scoring [1, 34], fraud detection [8], clinical risk [6]. In such threshold-driven applications, even a ~ 0.03 AUC drop can shift a non-trivial slice of the accept/reject boundary. The work is defensive—high-stakes critics still need human oversight; all data are public OpenML benchmarks.

8 Conclusion

Adding an LLM critic to an iterative search loop is not guaranteed to be harmless. We formalized *critic safety* and characterized when adding a critic helps or hurts on finite-budget tabular search: **when**—harm attenuates with critic capability and varies with frontier and budget, without a universal capability threshold; **why**—best-so-far search tolerates non-directional errors but not directional ones, and an idealized model gives the necessary condition $q(m-1) > \beta$; and **how**—observable gates that target β reduce the average loss in deployable reruns, without guaranteeing the no-critic baseline on every dataset. Cross-domain audits are consistent with the hypothesis that harm tracks mismatch between the critic's priors and the task. The practical takeaway: audit a critic against a no-critic control at the intended budget before trusting it, and gate or omit it when capability or verification is insufficient. We will release CRITFAULT and all artifacts upon publication.

References

[1] Bart Baesens, Tony Van Gestel, Stijn Viaene, Maria Stepanova, Johan Suykens, and Jan Vanthienen. Benchmarking state-of-the-art classification algorithms for credit scoring. *Journal of the Operational Research Society*, 54(6): 627–635, 2003. doi: 10.1057/palgrave.jors.2601545.

[2] Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, et al. MLE-bench: Evaluating Machine Learning Agents on Machine Learning Engineering. *arXiv preprint arXiv:2410.07095*, 2024.

[3] Tianqi Chen and Carlos Guestrin. XGBoost: A Scalable Tree Boosting System. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2016.

[4] Yizhou Chi et al. SELA: Tree-Search Enhanced LLM Agents for Automated Machine Learning. *arXiv preprint arXiv:2410.17238*, 2024.

[5] Sayak Ray Chowdhury, Anush Kini, and Nagarajan Natarajan. Provably Robust DPO: Aligning Language Models with Noisy Feedback. *arXiv preprint arXiv:2403.00409*, 2024.

[6] Evangelia Christodoulou, Jie Ma, Gary S. Collins, Ewout W. Steyerberg, Jan Y. Verbakel, and Ben Van Calster. A systematic review shows no performance benefit of machine learning over logistic regression for clinical prediction models. *Journal of Clinical Epidemiology*, 110:12–22, 2019. doi: 10.1016/j.jclinepi.2019.02.004.

[7] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, et al. Training Verifiers to Solve Math Word Problems. In *arXiv preprint arXiv:2110.14168*, 2021.

[8] Andrea Dal Pozzolo, Olivier Caelen, Yann-Aël Le Borgne, Serge Waterschoot, and Gianluca Bontempi. Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Systems with Applications*, 41(10): 4915–4928, 2014. doi: 10.1016/j.eswa.2014.02.026.

[9] Tim R. Davidson, Anja Surina, and Caglar Gulcehre. The Future of Facts: Tracing the Factual Generation–Verification Gap. *arXiv preprint arXiv:2605.27564*, 2026.

[10] Nick Erickson, Jonas Mueller, Alexander Shirkov, Hang Zhang, Pedro Larroy, Mu Li, and Alexander Smola. AutoGluon-Tabular: Robust and Accurate AutoML for Structured Data. *arXiv preprint arXiv:2003.06505*, 2020.

[11] Stefan Falkner, Aaron Klein, and Frank Hutter. BOHB: Robust and Efficient Hyperparameter Optimization at Scale. In *International Conference on Machine Learning (ICML)*, 2018.

[12] Matthias Feurer, Aaron Klein, Katharina Eggensperger, Jost Springenberg, Manuel Blum, and Frank Hutter. Efficient and Robust Automated Machine Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.

[13] Matthias Feurer, Katharina Eggensperger, Stefan Falkner, Marius Lindauer, and Frank Hutter. Auto-Sklearn 2.0: Hands-free AutoML via Meta-Learning. In *Journal of Machine Learning Research (JMLR)*, 2022.

[14] Pieter Gijsbers, Erin LeDell, Janek Thomas, Sébastien Poirier, Bernd Bischl, and Joaquin Vanschoren. An Open Source AutoML Benchmark. In *6th ICML Workshop on Automated Machine Learning (AutoML)*, 2019.

[15] Léo Grinsztajn, Edouard Oyallon, and Gaël Varoquaux. Why do Tree-based Models Still Outperform Deep Learning on Typical Tabular Data? In *Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*, 2022.

[16] Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, et al. A Survey on LLM-as-a-Judge. *arXiv preprint arXiv:2411.15594*, 2024.

[17] Siyuan Guo, Cheng Deng, Ying Wen, Hechang Chen, Yi Chang, and Jun Wang. DS-Agent: Automated Data Science by Empowering Large Language Models with Case-Based Reasoning. *International Conference on Machine Learning (ICML)*, 2024.

[18] Noah Hollmann, Samuel Müller, Katharina Eggensperger, and Frank Hutter. TabPFN: A Transformer That Solves Small Tabular Classification Problems in a Second. *International Conference on Learning Representations (ICLR)*, 2023.

[19] Noah Hollmann, Samuel Müller, and Frank Hutter. Large Language Models for Automated Data Science: Introducing CAAFE for Context-Aware Automated Feature Engineering. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

[20] Noah Hollmann, Samuel Müller, Lennart Purucker, Arjun Krishnakumar, Max Körfer, Shi Bin Hoo, Robin Tibor Schirrmeister, and Frank Hutter. Accurate predictions on small data with a tabular foundation model. *Nature*, 637: 319–326, 2025.

[21] Sirui Hong, Yizhang Lin, Bang Liu, Bangbang Liu, Binhao Wu, Ceyao Zhang, Chenxing Wei, Danyang Li, Jiaqi Chen, Jiayi Zhang, et al. Data Interpreter: An LLM Agent For Data Science. *arXiv preprint arXiv:2402.18679*, 2024.

[22] Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large Language Models Cannot Self-Correct Reasoning Yet. In *International Conference on Learning Representations (ICLR)*, 2024.

[23] Qian Huang, Jian Vora, Percy Liang, and Jure Leskovec. MLAgentBench: Evaluating Language Agents on Machine Learning Experimentation. *International Conference on Machine Learning (ICML)*, 2024.

[24] Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. Sequential Model-Based Optimization for General Algorithm Configuration. In *Learning and Intelligent Optimization (LION)*, 2011.

[25] Igor Itkin. Delayed Verification Destabilizes Multi-Agent LLM Belief: Instability Thresholds and Optimal Corrector Placement. *arXiv preprint arXiv:2606.27409*, 2026.

[26] Zhengyao Jiang, Dominik Schmidt, Dhruv Srikanth, Dixing Xu, Ian Kaplan, Deniz Jacenko, and Yuxiang Wu. AIDE: AI-Driven Exploration in the Space of Code. *arXiv preprint arXiv:2502.13138*, 2025.

[27] Ryo Kamoi, Yusen Zhang, Nan Zhang, Jiawei Han, and Rui Zhang. When Can LLMs Actually Correct Their Own Mistakes? A Critical Survey of Self-Correction of LLMs. *Transactions of the Association for Computational Linguistics (TACL)*, 2024.

[28] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, et al. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[29] Hadi Khalaf, Claudio Mayrink Verdun, Alex Oesterling, Himabindu Lakkaraju, and Flavio du Pin Calmon. Inference-Time Reward Hacking in Large Language Models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.

[30] Jan Hendrik Kirchner, Yining Chen, Harri Edwards, Jan Leike, Nat McAleese, and Yuri Burda. Prover-Verifier Games Improve Legibility of LLM Outputs. In *arXiv preprint arXiv:2407.13692*, 2024.

[31] Shayan Kiyani et al. When to Trust the Cheap Check: Weak and Strong Verification for Reasoning. *arXiv preprint arXiv:2602.17633*, 2026.

[32] Jacky Kwok, Shulu Li, and Pranav Atreya. LLM-as-a-Verifier: A General-Purpose Verification Framework. *arXiv preprint arXiv:2607.05391*, 2026.

[33] Erin LeDell and Sebastien Poirier. H2O AutoML: Scalable Automatic Machine Learning. *7th ICML Workshop on Automated Machine Learning (AutoML)*, 2020.

[34] Stefan Lessmann, Bart Baesens, Hsin-Vonn Seow, and Lyn C. Thomas. Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *European Journal of Operational Research*, 247(1):124–136, 2015. doi: 10.1016/j.ejor.2015.05.030.

[35] Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, et al. Let's Verify Step by Step. In *International Conference on Learning Representations (ICLR)*, 2024.

[36] Zexi Liu, Jingyi Chai, Xinyu Zhu, Shuo Tang, Rui Ye, Bo Zhang, Lei Bai, and Siheng Chen. ML-Agent: Reinforcing LLM Agents for Autonomous Machine Learning Engineering. *arXiv preprint arXiv:2505.23723*, 2025.

[37] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, et al. Self-Refine: Iterative Refinement with Self-Feedback. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

[38] Hariharan Manikandan, Yiding Jiang, and J. Zico Kolter. Language Models are Weak Learners. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

[39] Nat McAleese et al. LLM Critics Help Catch LLM Bugs. *arXiv preprint arXiv:2407.00215*, 2024.

[40] Jaehyun Nam, Jinsung Yoon, Jiefeng Chen, Jinwoo Shin, Sercan Ö. Arık, and Tomas Pfister. MLE-STAR: Machine Learning Engineering Agent via Search and Targeted Refinement. *arXiv preprint arXiv:2506.15692*, 2025.

[41] Soichiro Nishimori, Yu-Jie Zhang, Thanawat Lodkaew, and Masashi Sugiyama. On Symmetric Losses for Robust Policy Optimization with Noisy Preferences. *arXiv preprint arXiv:2505.24709*, 2025.

[42] Randal S. Olson, Nathan Bartz, Ryan J. Urbanowicz, and Jason H. Moore. Evaluation of a Tree-based Pipeline Optimization Tool for Automating Data Science. In *Genetic and Evolutionary Computation Conference (GECCO)*, 2016.

[43] Jane Pan, He He, Samuel R. Bowman, and Shi Feng. Spontaneous Reward Hacking in Iterative Self-Refinement. *arXiv preprint arXiv:2407.04549*, 2024.

[44] Arjun Panickssery, Samuel R. Bowman, and Shi Feng. LLM Evaluators Recognize and Favor Their Own Generations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

[45] Divya Pathak, Harshit Kumar, Anuska Roy, Felix George, Mudit Verma, and Pratibha Moogi. Detecting Silent Failures in Multi-Agentic AI Trajectories. *arXiv preprint arXiv:2511.04032*, 2025.

[46] Rushi Qiang, Yuchen Zhuang, Yinghao Li, Dingu Sagar V K, Rongzhi Zhang, Changhao Li, Ian Shu-Hei Wong, Sherry Yang, Percy Liang, Chao Zhang, and Bo Dai. MLE-Dojo: Interactive Environments for Empowering LLM Agents in Machine Learning Engineering. *arXiv preprint arXiv:2505.07782*, 2025.

[47] Jon Saad-Falcon et al. Shrinking the Generation-Verification Gap with Weak Verifiers. *arXiv preprint arXiv:2506.18203*, 2025.

[48] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language Agents with Verbal Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

[49] Yuda Song et al. Mind the Gap: Examining the Self-Improvement Capabilities of Large Language Models. *arXiv preprint arXiv:2412.02674*, 2024.

[50] Kaya Stechly, Karthik Valmeekam, and Subbarao Kambhampati. On the Self-Verification Limitations of Large Language Models on Reasoning and Planning Tasks. *arXiv preprint arXiv:2402.08115*, 2024.

[51] Rickard Stureborg, Dimitris Alikaniotis, and Yoshi Suhara. Large Language Models are Inconsistent and Biased Evaluators. *arXiv preprint arXiv:2405.01724*, 2024.

[52] Chris Thornton, Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. Auto-WEKA: Combined Selection and Hyperparameter Optimization of Classification Algorithms. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2013.

[53] Guiyao Tie, Zenghui Yuan, Zeli Zhao, Chaoran Hu, Tianhe Gu, Ruihang Zhang, Sizhe Zhang, Junran Wu, Xiaoyue Tu, Ming Jin, Qingsong Wen, Lixing Chen, Pan Zhou, and Lichao Sun. Can LLMs Correct Themselves? A Benchmark of Self-Correction in LLMs. *arXiv preprint arXiv:2510.16062*, 2025.

[54] Patara Trirat, Wonyong Jeong, and Sung Ju Hwang. AutoML-Agent: A Multi-Agent LLM Framework for Full-Pipeline AutoML. *arXiv preprint arXiv:2410.02958*, 2024.

[55] Joaquin Vanschoren, Jan N. van Rijn, Bernd Bischl, and Luis Torgo. OpenML: Networked Science in Machine Learning. *ACM SIGKDD Explorations Newsletter*, 15(2), 2014.

[56] Tianlu Wang, Ilia Kulikov, Olga Golovneva, Ping Yu, Weizhe Yuan, Jane Dwivedi-Yu, Richard Yuanzhe Pang, Maryam Fazel-Zarandi, Jason Weston, and Xian Li. Self-Taught Evaluators. *arXiv preprint arXiv:2408.02666*, 2024.

[57] Sean Welleck, Ximing Lu, Peter West, Faeze Brahman, Tianxiao Shen, et al. Generating Sequences by Learning to Self-Correct. In *International Conference on Learning Representations (ICLR)*, 2023.

[58] Sherry Yang, Joy He-Yueya, and Percy Liang. Reinforcement Learning for Machine Learning Engineering Agents. *arXiv preprint arXiv:2509.01684*, 2025.

[59] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct: Synergizing Reasoning and Acting in Language Models. In *International Conference on Learning Representations (ICLR)*, 2023.

[60] Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh Agarwal. Generative Verifiers: Reward Modeling as Next-Token Prediction. In *International Conference on Learning Representations (ICLR)*, 2025.

[61] Shaokun Zhang, Ming Yin, Jieyu Zhang, Jiale Liu, Zhiguang Han, Jingyang Zhang, Beibin Li, Chi Wang, Huazheng Wang, Yiran Chen, and Qingyun Wu. Which Agent Causes Task Failures and When? On Automated Failure Attribution of LLM Multi-Agent Systems. In *International Conference on Machine Learning (ICML)*, 2025.

[62] Shujian Zhang, Chengyue Gong, Lemeng Wu, Xingchao Liu, and Mingyuan Zhou. AutoML-GPT: Automatic Machine Learning with GPT. *arXiv preprint arXiv:2305.02499*, 2023.

[63] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, et al. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

A Metric Definitions

For dataset D, critic model M, and search policy Π, let $\mathrm{AUC}_{\Pi}(D, M)$ be the final best-so-far out-of-sample ROC-AUC (mean over seeds). We report:

$$\text{No-Critic delta:} \quad \Delta(D, M) = \mathrm{AUC}_{\text{real}}(D, M) - \mathrm{AUC}_{\text{no-critic}}(D), \tag{4}$$

$$\text{Second-replicate delta:} \quad \Delta_{\text{RF}}(D, M) = \mathrm{AUC}_{\text{real}}(D, M) - \mathrm{AUC}_{\text{rand-fb}}(D), \tag{5}$$

$$\text{Harm rate:} \quad \mathrm{HR}(M) = \tfrac{1}{|\mathcal{D}|} \textstyle\sum_D \mathbb{1}[\Delta(D, M) < 0 \wedge \text{CI excl. } 0], \tag{6}$$

$$\text{Guard gain:} \quad G(D, M) = \mathrm{AUC}_{\text{vgc}}(D, M) - \mathrm{AUC}_{\text{real}}(D, M), \tag{7}$$

$$\text{Resilience thr.:} \quad p^{\star}(D) = \min\{p : \mathrm{AUC}(p{=}0) - \mathrm{AUC}(p) > 0.01\}. \tag{8}$$

In the reported runs the random-feedback arm is an independently seeded random search, so Δ_{RF} is a second-seed version of Δ rather than a semantic placebo. All CIs are 10,000-sample nonparametric bootstrap over seeds; paired significance is Wilcoxon signed-rank across seeds (per dataset) or across datasets (suite level).

B Proof Details

Statements of Proposition 1, Proposition 2, Lemma 5.1, and Proposition 3 are in §5; the full proofs follow.

Proof of Proposition 1 (uniform corruption preserves the floor). Selection without replacement of T configurations uniformly at random induces the same distribution over the *set* $S_T = \{\pi(1), \ldots, \pi(T)\}$ regardless of the order or of which rounds were labeled "corrupted," because each round's draw is uniform over the remaining configurations in both the corrupted and uncorrupted case. Since $b_T = \max_{a \in S_T} s(a)$ depends only on the set S_T, its distribution, and hence its expectation, is identical to the no-critic loop, which draws S_T the same way. The claim holds for all p because the mixing weight between two identical sampling processes is irrelevant. The proposition requires only sampling-without-replacement exchangeability and that b_T be a function of the visited set; it makes no assumption on $s(\cdot)$. As stated in §5.2, the hypothesis is that selection is uniform in *every* round; the result is a floor-preservation statement for an uninformative policy, not a claim that randomly corrupting an informative critic is costless.

Proof of Proposition 2 (resilience under adversarial corruption). By Eq. (1) the incumbent is non-decreasing, and by Assumption 1 corrupted rounds leave it unchanged; hence b_T equals the incumbent obtained from the subsequence of effective rounds alone. Let $K \sim \mathrm{Binomial}(T, 1-p)$ denote the (random) number of effective rounds, and let $E = \{K \geq T_{\text{sat}}\}$. On E, the definition of the saturation budget gives $\mathbb{E}[\, b^{\star} - b_T \mid E \,] \leq \delta$. On the complement E^c, scores lie in $[0, 1]$, so $b^{\star} - b_T \leq 1$ pointwise. Splitting the expectation,

$$\begin{aligned} \mathbb{E}[\, b^{\star} - b_T \,] &= \mathbb{E}[\, b^{\star} - b_T \mid E \,] \Pr[E] \;+\; \mathbb{E}[\, b^{\star} - b_T \mid E^c \,] \Pr[E^c] \\ &\leq \delta \;+\; \Pr\big[K < T_{\text{sat}}\big], \end{aligned}$$

which is the stated bound. For a tolerated additional failure probability ϵ, the loop is $(\delta + \epsilon)$-robust whenever $\Pr[\mathrm{Binomial}(T, 1-p) < T_{\text{sat}}] \leq \epsilon$; since this tail probability is nondecreasing in p, the threshold $p^{\star}(\epsilon) = \sup\{p : \Pr[\mathrm{Binomial}(T, 1-p) < T_{\text{sat}}] \leq \epsilon\}$ is well-defined. The expectation-level approximation $p^{\star} \approx 1 - T_{\text{sat}}/T$ arises from replacing K by its mean $(1-p)T$ and requiring $(1-p)T \geq T_{\text{sat}}$; at the paper's scale ($T{=}8$, $T_{\text{sat}} \in [1, 3]$) it predicts $p^{\star} \in [0.63, 0.88]$, and the binomial tail at these parameters is small enough that the probabilistic and expectation-level thresholds agree to the granularity of the swept grid $\mathcal{P}$. Two remarks on the hypotheses. Assumption 1 (adversarial rounds do not raise the incumbent) holds once b_{t-1} exceeds the low-score region, empirically after round 1 on all nine fault-injection datasets, and the 90.5% empirical satisfaction rate reported in §5.3 means the deleted-round argument applies to the large majority of adversarial rounds, with the bound degrading gracefully in the remainder. The event E conditions on the count K only, not on which rounds are effective; this is licensed by the exchangeability of uniform-without-replacement selection over the effective subsequence.

Proof of Lemma 5.1 (conditional effective sample size). By elitist monotonicity (Eq. (1)) the incumbent depends only on the multiset of realized values, and adversarial rounds (Assumption 1) leave it unchanged, so they may be deleted from the multiset without changing b_T. Condition on $K_c = k_c$ competent and $K_u = k_u$ uniform rounds. A uniform round contributes one i.i.d. score draw. A competent round proposes the best of an effective m-configuration lookahead, so its realized value is distributed as the maximum of m i.i.d. draws; for the purpose of the *outer* maximum, such a round is exchangeable with m independent uniform draws (the maximum of maxima over disjoint blocks equals the maximum of the pooled draws). Hence, conditioned on (k_c, k_u), the final incumbent is distributed as the maximum of $N_{\text{eff}} = mk_c + k_u$ i.i.d. draws, as claimed. The lemma is exact under the model's assumptions (i.i.d. scores, independent round types).

Proof of Proposition 3 (necessary condition and tail boundary). Taking expectations over round types in Lemma 5.1 gives $\mathbb{E}[N_{\text{eff}}] = T(1 + q(m-1) - \beta) = T_{\text{eff}}$, so $T_{\text{eff}} > T$ if and only if $q(m-1) > \beta$. Under the model, rounds are independent and each uniform round contributes one draw from F, so it suffices to compare the score distribution of a single round. Treating an adversarial round as contributing no improvement (Assumption 1), the per-round CDF of a critic round is

$$G(x) \;=\; \beta + (1 - q - \beta)\, F(x) + q\, F(x)^m.$$

For an incumbent b, a uniform round improves with probability $1 - F(b)$ and a critic round with probability $1 - G(b)$. Their difference is

$$F(b) - G(b) \;=\; q\big(F(b) - F(b)^m\big) - \beta\big(1 - F(b)\big) \;=\; \big(1 - F(b)\big)\Big[\, q \sum_{j=1}^{m-1} F(b)^j - \beta \,\Big],$$

which is positive exactly when $\beta < q \sum_{j=1}^{m-1} F(b)^j$. Because $\sum_{j=1}^{m-1} F(b)^j < m-1$ whenever $F(b) < 1$, the condition $q(m-1) \le \beta$ implies $G \ge F$ everywhere: each critic round is stochastically dominated by a uniform round, and so is the final incumbent, giving $\mathbb{E}[b_T^{\text{critic}}] \le \mathbb{E}[b_T^{\text{no-critic}}]$ for every T. Hence $q(m-1) > \beta$ is necessary for net benefit. Writing $\varepsilon = 1 - F(b)$ and expanding $F(b)^j = (1-\varepsilon)^j$ gives $F(b) - G(b) = \varepsilon\,[\,q(m-1) - \beta\,] + O(\varepsilon^2)$, so $q(m-1) > \beta$ is the leading-order boundary in the rare-improvement regime $\varepsilon \to 0$.

The condition is not sufficient. Take Bernoulli$(1/2)$ scores, initial incumbent 0, and q=0.3, m=3, β=0.3, so $q(m-1) = 0.6 > \beta$. A critic round reaches score 1 with probability $0.3\,(1 - 2^{-3}) + 0.4 \cdot \frac{1}{2} = 0.4625$, versus 0.5 for a uniform round, so $\mathbb{E}[b_T^{\text{critic}}] = 1 - 0.5375^T < 1 - 0.5^T = \mathbb{E}[b_T^{\text{no-critic}}]$ for every $T \ge 1$. Here $\varepsilon = 0.5$ is far from the tail and the higher-order terms dominate; for this distribution the actual boundary is $\beta = 0.225$ at every T. Expected final utility therefore depends on the full score distribution, and we use $q(m-1) > \beta$ only as a necessary condition and tail approximation.

The criterion subsumes existing critic heuristics. The folklore practices for "when to trust an LLM critic" are each a single-variable move on the same condition $q(m-1) > \beta$: *use a stronger critic* raises q, m; *verify before trusting* (self-consistency, LLM-as-judge gating) lowers β; *search broadly* (parallel/tree policies) raises the effective lookahead m (corollary below). This is why none of them is universally safe: raising capability fails if β stays high (a fluent but confidently wrong critic), and verifying fails if the verifier is itself unreliable (§6, low-$\widehat{\rho}$). The condition indicates which move is needed when; our prevention (§6) targets β.

Corollary: breadth is a structural defense. Proposition 3 suggests that *parallel* search resists harm. Expanding k candidates per round and keeping the best raises the effective lookahead ($m \to m' \ge m$), which widens the range of adversarial rates for which net benefit is possible to $\beta < q(m'-1)$; a k-wide round also dilutes a single adversarial pick to weight $1/k$. The boundary thus shifts toward net-benefit as breadth grows, predicting that breadth-oriented policies are hard to harm, matching the measured backbones (Table 9: parallel k-wide search below floor on $1/12$ datasets versus $9/12$ for sequential, tree-search, and multi-agent; §5.5).

Monte-Carlo check of the criterion. We cannot check $q(m-1) > \beta$ per real dataset (sequential critics violate the independent-round assumption), so we simulate the idealized loop. For q=0.3, m=3 ($q(m-1) = 0.60$), the measured sign-flip boundary in our simulation is 0.42, 0.53, 0.56, and 0.60 at T=8, 32, 128, 512, and at T=8 the sign of $q(m-1) - \beta$ matches the sign of net benefit on $43/48$ grid cells. This agreement is specific to the simulated score distribution: the Bernoulli example above has boundary 0.225 at every T, so the simulation illustrates the tail approximation rather than a general convergence result.

C Implementation Details

In the reported experiments, ACTOR and CRITIC are realized by a single LLM call per round that reads the evaluated history and the menu of untried configurations and names the next configuration; a response that names no untried configuration triggers a uniform fallback (§3.2).

Search space: 8 model families × 4 transform sets × 3 hyperparameter presets (= 96). The four feature-transformation sets are generated by applying the same fixed construction rule separately to each dataset, using its two highest-variance numeric features as transformation anchors; the complete configuration list is released with the code. Loop: T=8 rounds (legacy VGC runs: T=10), greedy incumbent, no explicit patience. LLM critics served via Amazon Bedrock Converse, temperature 0, ≤ 1536 max tokens; models are the published open-weights Llama-3.1-8B / Llama-3.3-70B inference profiles and Claude Haiku 4.5. All tree models pinned single-threaded (`OMP_NUM_THREADS=1`) to avoid oversubscription. Detector: gradient-boosted trees, 100 estimators, depth 3, evaluated leave-one-dataset-out.

Algorithm 1 Actor–Critic AutoML loop

Input: $D_{\text{tr}}, D_{\text{val}}$, rounds T, patience p
$f^* \leftarrow \textsc{Train}(D_{\text{tr}}, D_{\text{val}})$; $r^* \leftarrow \textsc{Metric}(f^*)$; $c \leftarrow 0$
for $t = 1$ to T **do**
  $a_t \leftarrow \textsc{Actor}(s_t, h_{t-1}, \text{corr}_{t-1})$ ▷ propose
  $f_t \leftarrow \textsc{Pipeline}(D_{\text{tr}}, D_{\text{val}}; a_t)$ ▷ execute
  $r_t \leftarrow \textsc{Metric}(f_t)$
  $\text{corr}_t \leftarrow \textsc{Critic}(a_t, r_t, h_{t-1})$ ▷ critique
  $h_t \leftarrow h_{t-1} \cup \{(a_t, r_t, \text{corr}_t)\}$
  if $r_t > r^*$ **then** $f^* \leftarrow f_t$; $r^* \leftarrow r_t$; $c \leftarrow 0$
  else $c \leftarrow c{+}1$; **if** $c \geq p$ **break**
  end if
end for
return f^*

C.1 Formal setup details (moved from §3)

History and feedback maps. Let $H_{t-1} = \{(a_i, s_{\text{val}}(a_i))\}_{i=1}^{t-1}$ denote the evaluated configurations and their observed validation utilities, and c_{t-1} the most recent feedback. At round t the proposer selects $a_t = \pi(H_{t-1}, c_{t-1})$; after evaluation and the incumbent update $\hat{a}_t = \arg\max_{a \in \{a_1,\ldots,a_t\}} s_{\text{val}}(a)$, the critic produces $c_t = C(a_t, s_{\text{val}}(a_t), H_{t-1})$, which may influence the next proposal. In the no-critic condition $c_t = \varnothing$ reduces selection to uniform exploration over untried configurations; the oracle bypasses the language channel entirely and upper-bounds any feedback-mediated policy on the same frontier.

Terminology. Because the critic only shapes which configurations are explored—it holds no value function and receives no reward—we refer to the system as a *proposer–critic loop* rather than an actor–critic method with a reinforcement-learning interpretation. The loop is externally grounded (the critic sees held-out validation utilities), unlike intrinsic self-critique in which a model evaluates its own output without an external task signal [48]; the validation signal both makes critic guidance auditable and supplies the evidence used to gate harmful feedback in §6.

Cached-frontier access rules. Before any search condition runs, every configuration in $\mathcal{A}$ is evaluated once on the frozen split and its validation and test utilities cached. During a run, the search may read $s_{\text{val}}(a)$ only after selecting a; test utilities remain hidden until the final incumbent is fixed. Observed differences between arms therefore cannot be attributed to pipeline-training randomness, execution failures, resampling, or access to different candidate pools.

Resilience threshold for fault injection. For dataset D, corruption mode c, and rate p, let $J_{D,c}(p)$ denote the expected test utility of the final validation incumbent. The mode-specific resilience threshold is the first evaluated corruption rate at which degradation exceeds a fixed tolerance,

$$p^{\star}_{D,c} = \min\left\{p \in \mathcal{P} : \; J_{D,c}(0) - J_{D,c}(p) > \delta_{\text{res}}\right\},$$

with $\delta_{\text{res}} = 0.01$ AUC, where $\mathcal{P} \subset [0, 1]$ is the swept grid. Two configurations are neighbors (for MISLEAD) when they differ in exactly one pipeline component. The tolerance δ_{res} defines meaningful degradation under corruption.

Statistical assessment of below-floor prevalence. Evidence for the below-floor direction is assessed at two levels: per-dataset bootstrap confidence intervals quantify task-level uncertainty, while an exact suite-level sign-flip test over all 2^8 sign assignments evaluates whether negative Δ_{NC} occur systematically across datasets. Prevalence statements such as "below floor on 7 of 8 datasets" refer to the signs of dataset-level point estimates; the headline inferential claim relies on the suite-level test rather than requiring every individual interval to exclude zero.

Analysis-specific dataset subsets. The primary real-critic audit uses the eight binary-classification datasets. Controlled fault injection uses the nine datasets other than the three near-zero-headroom datasets. The capability ladder is a targeted follow-up on six binary datasets selected partly on the reference 8B critic's outcome; it is therefore used only to study how an established effect changes with capability, not to estimate prevalence of harm. The CRITFAULT detector is evaluated on the full 12-dataset suite, including the three near-zero-headroom datasets, where harmful transitions are less sharply defined.

Table 3: Experimental scope by analysis. Exact dataset names and model settings are reported in the appendix.

Analysis	Datasets	Purpose
Real-critic audit	8 binary	Below-floor harm
Fault injection	9 (see text)	Error structure
Capability ladder	6 targeted	Capability attenuation
CRITFAULT	12 total	Feedback detection

C.2 Analysis-specific experimental settings

Table 3 summarizes the datasets, search budgets, seed counts, and decoding settings used in each analysis. The primary four-arm audit and capability ladder use $T = 8$ and eight matched seeds. Controlled fault injection uses 30 seeds. The end-to-end prevention reruns of §6.6 use $T = 8$ and eight paired seeds; the legacy prevention results in App. D use the budgets stated in their captions. The expanded-frontier experiment uses $T = 20$, and the wide-space stress test uses $T = 16$. Unless otherwise stated, critic decoding uses temperature 0; stochastic-decoding robustness experiments use temperature 0.7.

Critics are accessed through a provider-agnostic Amazon Bedrock Converse interface. Exact model identifiers, providers, and release generations are listed with the capability ladder (Table 1); all critics use their default inference profiles with a 1536-token generation cap. All prompts, cached frontiers, frozen splits, configuration definitions, and evaluation scripts will be released with the paper.

D Additional Experiments

Per-dataset detector generalization (leave-one-dataset-out). Table 4 gives the per-dataset breakdown behind the mean AUC 0.982 reported in §6: on injection labels with all five features, the bad-feedback signature transfers to each of the nine datasets outside the near-zero-headroom group (each ≥ 0.96, mean 0.982) and is lower (mean 0.851) on the three near-zero-headroom datasets, where harmful transitions are less sharply defined. This 0.982 is distinct from the 0.982 of the deployable f0–f2 detector on real-harm labels (§6.5).

Table 4: Bad-feedback detector, **leave-one-dataset-out** AUC (train on 11 datasets, test on the held-out one), injection labels, all five features (including the two offline structural features). Mean 0.982 on the nine datasets outside the near-zero-headroom group (every one ≥ 0.96) and 0.851 on the three near-zero-headroom datasets (segment/mfeat/splice), where nearly every configuration scores near the oracle.

Nine datasets (mean 0.982, all ≥ 0.96)			
bank-marketing	0.990	letter	0.983
credit-g	0.989	churn	0.980
phoneme	0.988	jm1	0.974
adult	0.987	kc1	0.966
electricity	0.985		
Near-zero headroom (mean 0.851)			
mfeat-factors	0.962	segment	0.732
splice	0.860		
Overall mean 0.950 (12 datasets, 3,840 transitions)			

Capability ladder, full per-dataset numbers. Table 5 gives the complete seven-tier, four-provider ladder behind Fig. 3a. The 2024 8B critic is negative on $5/6$ datasets with losses up to 0.245; every $\geq$14B critic stays within 0.031 of zero, but negative point estimates remain (three of six at 14B, five of six at Llama-3.3-70B). The ladder supports attenuation of the large losses, not a capability threshold.

Real SOTA baseline. Table 6 compares our loop against AutoGluon-Tabular on identical frozen public splits, confirming our search space is SOTA-competitive (in-space oracle matches or beats AutoGluon on 5/8 datasets).

Space size alone is not protective. Two observations should be kept apart. *(i) For the real 8B critic*, harm is largest on the 96-config frontier at T=8 and largely disappears on the 250-config frontier at T=20 (§4: ADULT

Table 5: Capability ladder across a *seven-tier, four-provider* span: real-critic No-Critic Δ (AUC, 8 seeds, identical frozen splits) on six discriminating datasets (full per-dataset version of Table 1). The 2024-generation 8B critic (Meta) is negative on $5/6$; the larger critics are within 0.031 of zero but not uniformly non-negative (Llama-3.3-70B reaches -0.031 on ELECTRICITY and -0.023 on KC1). The pattern is not monotone in size, provider, or release date.

Dataset	$3B^{25}$	$8B^{24}$	$8B^{25}$	14B	32B	70B	H4.5
adult	$-.210$	$-.245$	$+.001$	$+.001$	$+.001$	$-.011$	$+.002$
electricity	$-.064$	$-.072$	$-.012$	$+.012$	$+.008$	$-.031$	$+.006$
jm1	$-.008$	$-.080$	$-.003$	$-.002$	$-.000$	$-.011$	$+.002$
kc1	$+.000$	$-.028$	$-.000$	$-.001$	$-.001$	$-.023$	$-.000$
bank-marketing		$-.007$	$+.000$	$+.000$	$+.001$	$-.007$	$+.001$
credit-g		$+.006$	$-.005$	$-.012$	$-.007$	$+.006$	$-.001$

$3B^{25}$=Ministral-3B (Mistral '25); $8B^{24}$=Llama-3.1 (Meta '24); $8B^{25}$=Ministral-3-8B (Mistral '25); 14B=Ministral-14B; 32B=Qwen3; 70B=Llama-3.3; H4.5=Haiku 4.5 (Anthropic).

Table 6: Competitive AutoML setup. Test ROC-AUC on identical frozen public splits: our no-critic floor and in-space oracle vs. **AutoGluon-Tabular** [10] (real SOTA, 120s matched budget, `medium_quality`). The in-space oracle matches or beats AutoGluon on 5 of 8 datasets; the harm we measure is thus relative to a SOTA-competitive baseline.

Dataset	no-critic	oracle	AutoGluon
electricity	.9411	**.9542**	.9316
credit-g	.8017	**.8085**	.7849
churn	.9412	**.9463**	.9264
adult	.9287	**.9302**	.9236
bank-marketing	.9297	**.9328**	.9309
kc1	.8279	.8314	**.8378**
phoneme	.9433	.9483	**.9549**
jm1	.7290	.7320	**.7484**

$-0.245 \rightarrow +0.0003$); because frontier and horizon change together, this does not identify the cause. *(ii) For a* modeled *systematically-adversarial critic* (not the real LLM), harm grows as the space enlarges $2.6\times$ ($96 \rightarrow 250$ configs) on a wide weak-to-strong space with real headroom (Table 7). Enlarging the space is therefore not by itself a defense against a directional critic.

Table 7: Enlarging the space does not protect against a modeled systematically-adversarial critic. Wide weak-to-strong space with real headroom (1.7–5.6% oracle-over-mean AUC); enlarging the space $2.6\times$ ($96 \rightarrow 250$ configs) with a modeled systematically-adversarial critic (No-Critic Δ, AUC$\times 10^3$, 20 seeds, T=16, 95% bootstrap CI). Harm grows with space size on all five datasets.

Dataset	headroom	Δ @96	Δ @250
electricity	.056	-30.8 $[-43.6, -18.9]$	-36.0 $[-48.1, -24.4]$
credit-g	.042	-22.2 $[-29.0, -15.1]$	-28.3 $[-38.8, -19.7]$
kc1	.030	-14.3 $[-19.7, -8.9]$	-25.4 $[-36.3, -15.9]$
jm1	.028	-14.9 $[-19.6, -10.2]$	-15.9 $[-24.0, -8.6]$
adult	.017	-5.8 $[-9.0, -3.2]$	-13.2 $[-18.1, -8.9]$

Legacy prevention protocol. Several prevention results in this appendix (Table 8, the last row of Table 10, the VGC column of Table 11, Fig. 4, and the second-backbone paragraph) come from an earlier protocol whose reject branch selects the untried configuration with the best cached validation utility; its learned filter also uses the two offline structural features. Both read validation utilities of untried configurations, which a deployed loop cannot observe, so we report these results as exploratory. The deployable reruns in §6.6 use a uniform reject branch and only observed features.

Detector-as-online-filter and the harm mechanism (legacy protocol). Table 8 reports the family-selection KL between the critic's picked model-family distribution and the oracle-optimal one: a harmful raw critic over-commits to weak families (high KL_{raw}). Under the legacy protocol above, the filtered policy ends above the no-critic floor on ADULT, JM1, and KC1 and below it on CREDIT-G, with lower KL.

Table 8: Detector-as-online-filter under the legacy protocol (validation-best reject branch, five features) and the family-selection KL (Llama-3.1-8B critic, 5 seeds, T=10, LODO-trained, reject threshold τ=0.6). KL is between the critic's picked model-family distribution and the oracle-optimal one.

Data	raw−nc	filter−nc	KL_{raw}	KL_{filter}
adult	−.0002	+.0015	8.96	0.15
jm1	−.0028	+.0036	8.70	0.37
kc1	−.0045	+.0018	7.32	5.17
credit-g	−.0036	−.0017	8.44	1.65

Harm generalizes across search backbones (full data). Reimplementing four search paradigms over the identical frontier and contrasting no-critic vs. a systematically-adversarial critic (30 seeds), the harmful-critic delta is negative for every paradigm (Table 9), largest for sequential (which commits hardest) and smallest for breadth-oriented policies—so the harm is a property of critic-guided search itself, not of any one control flow, and breadth is a partial structural defense.

Table 9: Harm generalizes across search backbones. Mean harmful-critic Δ (AUC vs. no-critic) over 12 datasets, 30 seeds; $n_{<0}$ = datasets with $\Delta < -0.002$. A systematically adversarial critic degrades sequential, tree-search, and multi-agent search alike; parallel k-wide expansion is nearly immune—breadth hedges against an unreliable critic.

Search policy	mean harmful Δ	$n_{<0}$/12
sequential	−0.0183	9
multi-agent (pooled)	−0.0095	9
tree-search (UCB1)	−0.0080	9
parallel k-wide	−0.0004	1

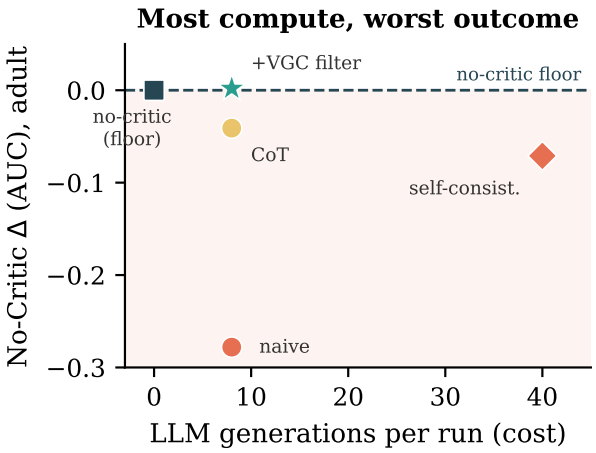


Figure 4: Accuracy–cost trade-off for the weak 8B critic (ADULT). Naive, chain-of-thought, and 5×-cost self-consistency all spend real LLM generations to land *below* the zero-cost no-critic floor. The "+VGC filter" point uses the legacy learned filter (validation-best reject branch; see above).

Prompting family. Chain-of-thought and self-consistency (§6.1) reduce but do not eliminate harm (Table 10): both stay negative on $3/4$ discriminating datasets, and added compute is not monotone help—self-consistency at $5\times$ the calls scores -0.071 on ADULT, *worse* than plain CoT's -0.041, so the failure is not merely insufficient compute. The last row uses the legacy learned filter (validation-best reject branch) and is not a deployable result.

Harm is bad advice, not bad formatting (full data). A natural worry is that the 8B harm is merely a format-compliance artifact—a failed parse triggering a random fallback. The data rule this out: harm and parse-failure are *anti*-correlated across the suite. ADULT, the most harmed dataset (-0.245), has *zero* fallbacks in 64 critic rounds, so its harm is entirely substantive; CREDIT-G, with the *most* fallbacks ($18/64$), is not harmed at all ($+0.006$), because a

Table 10: Harm is not a naive-prompting artifact. No-Critic Δ (AUC, 8 seeds, single controlled run) by critic *method family*, Llama-3.1-8B. Chain-of-thought and self-consistency (#calls = LLM generations/run) *reduce* but do not eliminate harm (negative on 3/4; self-consistency worse than CoT on ADULT at $5\times$ cost). The last row uses the legacy filter protocol.

Family (#calls)	ADULT	JM1	CREDIT-G	KC1
no-critic (0)	.000	.000	.000	.000
naive (8)	$-.278$	$-.018$	$+.006$	$-.019$
CoT (8)	$-.041$	$-.014$	$+.001$	$-.006$
self-consistency (40)	$-.071$	$+.001$	$+.003$	$-.017$
+ legacy filter (8)	$+.002$	$+.004$	$-.002$	$+.002$

Table 11: Supervised fine-tuning vs. VGC under *leave-one-dataset-out* (No-Critic Δ, AUC$\times 10^3$, 8 seeds, 95% bootstrap CI). For each row the critic is QLoRA-fine-tuned on oracle picks from the other seven datasets and evaluated on the held-out one. Fine-tuning averts the collapse on all eight but only returns to the floor (no CI excludes zero positively) and needs oracle labels. The VGC column uses the legacy protocol.

Held-out dataset	raw (§4)	Fine-tuned (LOO)	VGC
ADULT	-245	$+0.5$ $[-0.5, +1.4]$	$+1.0$ $[0.0, +1.9]$
ELECTRICITY	-72	$+1.6$ $[-2.1, +7.0]$	$+3.5$ $[-1.5, +8.2]$
JM1	-80	-0.8 $[-3.4, +1.7]$	$+1.7$ $[-1.7, +4.5]$
KC1	-28	-3.1 $[-9.4, +1.2]$	-1.2 $[-3.9, +1.9]$
CHURN	-2	$+1.7$ $[-1.1, +4.5]$	-1.8 $[-6.1, +1.9]$
BANK-MARKETING	-7	-0.7 $[-1.4, +0.0]$	-0.7 $[-6.6, +3.6]$
PHONEME	-9	-0.5 $[-2.0, +0.7]$	$+1.2$ $[-1.0, +3.4]$
CREDIT-G	$+6$	$+0.9$ $[-5.3, +6.8]$	$+4.9$ $[+0.7, +10.2]$

failed parse degenerates to random search (the floor), never below it. Format reliability does improve with scale, but the capability effect is driven by proposal *quality*: a small critic that parses perfectly still steers confidently into low-value regions.

Second backbone (legacy protocol). On a Llama-3.3-70B critic (8 seeds), legacy VGC moves JM1 from a raw -0.0010 to $+0.0036$ and leaves ADULT essentially unchanged.

Can supervised fine-tuning replace the guard? A natural objection is to fix the weak critic by *fine-tuning* it rather than gating it. We test this directly under *leave-one-dataset-out*: for each dataset we QLoRA-fine-tune Llama-3.1-8B on oracle picks (the truly-best untried config each round) from the *other seven* datasets, then evaluate the fine-tuned critic on the held-out eighth—which it never saw in training—in the exact $T{=}8$ setup of §4 (Table 11, 8 seeds, 95% bootstrap CI). Fine-tuning does avert the collapse: on every held-out dataset the No-Critic Δ returns to within ± 0.003 AUC of the floor (e.g. ADULT $-0.245 \to +0.0005$, JM1 $-0.080 \to -0.0008$), with no dataset below floor beyond noise. But it only recovers the *floor*—no held-out dataset shows a positive delta whose CI excludes zero—and it requires oracle supervision unavailable at deployment (the oracle label is precisely what the critic cannot see) plus per-suite GPU training. The VGC column is from the legacy protocol (validation-best reject branch) and is shown for reference only; it is not a deployable comparison.

Harm generalizes to regression. The main study is binary classification; harm is not specific to it. On public regression datasets (California housing, diabetes; weak-to-strong regressor space scored by test R^2, 20 seeds), a systematically-adversarial critic remains harmful: No-Critic Δ (in R^2) is -0.034 $[-0.045, -0.023]$ on California and -0.024 $[-0.033, -0.016]$ on diabetes, both CIs excluding zero. The failure mode is a property of elitist critic-guided search, not of the classification metric.

The harm is not free: full cost accounting. The four-arm comparison (§6.1) is not accuracy-neutral. The no-critic arm issues zero LLM calls; a naive critic issues one per round ($T{=}8$) and a self-consistency critic issues $k{\cdot}T$ ($k{=}5$, i.e. 40) generations per run, each a multi-hundred-token reasoning trace, yet on the discriminating datasets both land below the free no-critic floor (Table 10). The critic thus spends the most compute to achieve the worst outcome—for the weak 8B critic, adding a critic dominates in cost while being dominated in quality, inverting the usual accuracy–cost trade-off and providing an independent argument against deploying an unverified critic.

Cross-domain discriminative test (full per-domain results). We run the same 8B critic on four further search domains, all with the identical four-arm protocol (8 seeds). Per-domain No-Critic Δ (bootstrap CI):

- 20-Newsgroups text-pipeline (36 variants): -0.003 $[-0.008, 0.001]$ (contains zero).
- 20-Newsgroups text-pipeline (50 weak-to-strong pool): $+0.003$ $[-0.005, 0.010]$.
- Forecasting-method selection (13 methods, 40 real M4-monthly series): $+0.001$ $[+0.001, 0.002]$ — *significant positive*.
- RAG retrieval-config selection (SQuAD, 54 chunk/top-k/retriever variants): $+0.007$ $[+0.006, 0.008]$ — *significant positive*.

Tabular AutoML (-0.245) is the sole clearly negative setting. One hypothesis consistent with the case study (App. G) is that the critic's priors about model families conflict with a frontier dominated by gradient boosting; we do not test this directly.

Detector validity controls: real-harm labels and leakage ablation (full protocols). *(i) Real-harm labels:* we relabel transitions from a *real* Llama-3.1-8B critic—$y{=}1$ iff the step both fails to raise the incumbent and leaves ≥ 0.01 AUC of achievable pool value on the table (a genuinely wasteful real step, not an injected one)—and retrain leave-one-dataset-out; the detector attains mean AUC 0.992 (6 datasets, balanced 52% positive). *(ii) Leakage-feature ablation:* dropping the two structural features that partly encode the injection (rank-among-untried, neighbor-of-incumbent) and keeping only verifier-view scores (val, claimed-Δ, round) still gives mean AUC 0.982 on the real-harm labels (and only 0.76 on the injection labels), confirming the signature is not an artifact of label generation.

Detector signature transfers across critic models (leave-one-model-out). The detector is not tied to the critic it was trained on. Under a stricter leave-one-model-out protocol—train on the trajectories of two critics and test on a *held-out third*, spanning two providers (Llama-3.1-8B, Claude Haiku 4.5, Claude-3-Haiku; 960 transitions)—it attains mean AUC 0.988, and ≥ 0.967 for every held-out model (Llama-3.1-8B 0.967, Claude Haiku 4.5 0.998, Claude-3-Haiku 0.999). A detector trained on one set of critics thus flags harmful feedback from an unseen critic of a different family.

CRITFAULT scale and baselines. The benchmark packages ∼3,840 loop transitions across 12 datasets, each labeled clean/adversarial with three decision-time features (f0–f2) and two offline structural features (f3, f4), under a fixed leave-one-dataset-out protocol. Reference baselines show the task is non-trivial yet learnable: random 0.50, best single feature 0.92, our detector 0.95 (Table 13); the full datasheet is in App. F.

Why a cheap pre-deployment red-flag fails (detail). We tested a one-shot, oracle-free screen: run a few pilot rounds and flag a critic if its proposals fall below the incumbent too often. It fails on these low-headroom frontiers. A *strong* critic reaches a near-ceiling incumbent within a round or two, after which almost any further proposal (including healthy exploration) scores below that high incumbent; a *weak* critic keeps a lower incumbent, so its below-incumbent rate is not higher. The screen therefore measures "exploration versus an already high incumbent," not "harmful versus safe," and does not separate the two. This is why our prevention (§6) gates online rather than relying on a static pre-deployment test.

E Dataset Suite

Table 12 lists the 12 public OpenML classification datasets used throughout. They span roughly an order of magnitude in size (700–8,400 training rows), a mix of native-binary and binarized multiclass targets (minority vs. rest), and several application domains (credit, software defect prediction, marketing, telecom churn, electricity-market drift). All train/test splits are frozen once and reused across every arm and seed, so any measured difference reflects the critic, not resampling.

F CRITFAULT Datasheet

Motivation. CRITFAULT is, to our knowledge, the first benchmark for detecting harmful critic feedback in a tabular AutoML loop (to be released publicly upon publication). It exists to make the diagnostic finding of this paper—that bad feedback has a dataset- and model-agnostic signature—reproducible and comparable.

Composition. ∼3,840 instances across 12 public OpenML datasets. Each instance is one proposed loop transition (a critic-selected next configuration), labeled $y{=}1$ (adversarial: injected WORST-neighbor pick) or $y{=}0$ (clean: ideal-critic pick). The classes are balanced by construction (1,920 each).

Table 12: Public OpenML suite. All splits frozen; multiclass targets binarized (minority vs. rest).

Dataset	n_{train}	n_{test}	Type
credit-g	700	300	binary
kc1	1476	633	binary (defect)
mfeat-factors	1400	600	multiclass→bin
segment	1617	693	multiclass→bin
splice	2233	957	3-class→bin
churn	3500	1500	binary
phoneme	3782	1622	binary
jm1	7619	3266	binary (defect)
adult	8400	3600	binary
bank-marketing	8400	3600	binary
electricity	8400	3600	binary (drift)
letter	8400	3600	26-class→bin

Features. Five features with no access to the held-out test outcome: (f0) proposal validation score; (f1) claimed improvement over the incumbent; (f2) round fraction t/T; (f3) normalized rank among the cached validation utilities of untried configurations; (f4) neighbor-of-incumbent indicator. Only f0–f2 are observable at decision time; f3 uses unobserved utilities of untried configurations and f4 the grid structure, so both are offline diagnostic features. The deployable detector of §6.5 uses f0–f2 only.

Task and protocol. Binary detection of adversarial transitions. The canonical split is *leave-one-dataset-out* (train on 11 datasets, test on the held-out one; 12 folds), scored by ROC-AUC and averaged. This measures generalization to an *unseen* dataset, not within-dataset fit.

Reference baselines (Table 13). New detectors should report LODO mean AUC against these anchors.

Table 13: CRITFAULT reference baselines (leave-one-dataset-out mean AUC, 12 folds). The task is non-trivial (random/majority $= 0.50$); the bad-feedback signature is strong enough that a single decision-time feature reaches 0.92, and the full detector tops the panel.

Detector	LODO mean AUC
random	0.502
majority (base rate)	0.500
best single feature	0.916
logistic (5 feat.)	0.921
gradient-boosted (ours)	**0.950**

Limitations. Labels come from controlled fault injection on one 8B critic; real-critic and cross-model trajectories are a natural extension (the harness supports both). "Adversarial" is ill-defined on the three near-zero-headroom datasets (segment/mfeat/splice), where all baselines degrade—these are retained for completeness but flagged in the card.

Access. To be released publicly with the harness upon publication, under a permissive license: the instance table (in CSV form), the LODO evaluator, and the baseline panel above.

Table 14: Fault-injection (30 seeds). $p^\star$ = smallest corruption rate with > 0.01 AUC drop; cost_1 = AUC lost at full corruption. Adversarial modes tolerate 60–100% bad feedback; random corruption is nearly free ($\text{cost}_1 \leq 0.005$, no threshold)—see §5.

Dataset	$p^\star_{\text{worst}}$	$\text{cost}_1^{\text{worst}}$	$\text{cost}_1^{\text{random}}$
credit-g	0.6	0.093	0.005
adult	0.6	0.084	0.001
churn	0.6	0.078	0.004
jm1	0.9	0.061	0.004
electricity	0.9	0.051	0.002
kc1	0.9	0.048	0.004
bank-marketing	0.9	0.024	0.002
phoneme	0.9	0.023	0.002
letter	1.0	0.019	0.001

Table 15: Four-arm ablation with a deterministic *ideal* critic (30 seeds). No-Critic Δ = real$-$no-critic, 95% bootstrap CI (paired). Small but robust: all 9 listed datasets have a CI excluding 0 (paired Wilcoxon $p < 0.005$); suite-level Wilcoxon $p = 0.001$. The three near-zero-headroom datasets are omitted. The consequential quantity is not this upside but the downside of a bad critic (§4).

Dataset	no-crit	oracle	No-Critic Δ [95% CI]
electricity	.9465	.9542	$+.0053$ [.0027, .0085]
churn	.9412	.9463	$+.0050$ [.0035, .0066]
credit-g	.8031	.8085	$+.0049$ [.0027, .0073]
jm1	.7281	.7321	$+.0038$ [.0025, .0053]
bank-marketing	.9298	.9328	$+.0028$ [.0022, .0034]
phoneme	.9444	.9483	$+.0025$ [.0016, .0034]
letter	.9933	.9959	$+.0023$ [.0015, .0034]
kc1	.8284	.8314	$+.0023$ [.0014, .0032]
adult	.9289	.9302	$+.0013$ [.0008, .0019]

G Qualitative Case Study: a Real Harmful Trajectory

To make the failure concrete, Table 16 shows a real Llama-3.1-8B critic trajectory on ADULT (one seed, reasoning-mode prompt so the critic states its rationale). The pattern is striking: the critic *verbally* recognizes the strong family—every round it repeats that "LightGBM has already shown strong performance (0.930)"—yet its *actions* walk steadily away from it, proposing configurations of broadly decreasing validation AUC ($0.930 \to 0.927 \to 0.926 \to 0.922 \to \ldots \to 0.651$; every post-initial pick below the incumbent), ending on a logistic-regression config near chance. Each proposal after round 1 scores *below the incumbent*—the "systematically adversarial" signature of §5 (Assumption 1) in a real critic. Elitist selection saves the final *incumbent* (it keeps round 1's good value), but the critic's guidance stream is actively misleading: a loop that trusted the latest proposal, or a non-elitist variant, would track the critic down to 0.65. The trajectory illustrates one route to aggregate harm: confident, consistent misdirection rather than random noise. The artifact to be released contains full transition logs.

Table 16: A real Llama-3.1-8B critic trajectory on ADULT (oracle: lightgbm, 0.930). The critic endorses the strong family in words while proposing steadily worse configurations; every post-initial pick is below the incumbent (Assumption 1).

Round	Proposed family	Val AUC	Below incumbent?
1	lightgbm	0.930	— (initial)
2	lightgbm	0.927	yes
3	lightgbm	0.926	yes
4	gradient_boosting	0.922	yes
5	hist_gbdt	0.928	yes
6	extra_trees	0.902	yes
7	extra_trees	0.880	yes
8	logistic_regression	0.651	yes