Peak-Then-Collapse and the Four Interface Channels of Knowledge-Graph Tool Use

Anonymous ACL submission

Abstract

We test the standard RLVR tool-use recipe (GRPO on Qwen2.5-7B-Instruct) on a deliberately minimal knowledge-graph tool API: four Freebase navigation verbs over Complex WebQuestions. Under a self-verifiable retrieval reward, the policy's correct-via-tool rate climbs from 3.8% to 9.6% over 250 steps, then collapses to 0% within a single 50-step window, a *peak-then-collapse* pattern replicated across four seeds. Across seven reward designs, we find four recurring failure modes: adding denser or more targeted proxy rewards shifts the failure mode rather than eliminating it. We argue that a key difference from Python interpreters, web search, and JSON APIs is interface feedback: their failures often leak natural-language signal the model saw in pretraining. A Python traceback names the failing line; an empty Freebase result `[]` does not. Stripping away that surface exposes a degradation regime that same-family reward redesigns do not fix. A direct oracle ablation rules out relation selection: injecting gold relations at every retrieval call lifts exact-match accuracy by only $+0.20$ pp, and 95.4% of retrieval-dependent errors are retrieval-composition failures rather than answer-extraction failures. As a mitigation, one-iteration self-distillation reaches 40.0% EM at 7B; the best 14B model we train reaches 40.2% (a gap within sampling noise) and initialization barely matters, so within the 7B–14B range tested the $\sim$40% ceiling appears interface-bound rather than capacity-bound.

1 Introduction

Large language models increasingly act as *agents*: they call external tools (code interpreters, web search, function APIs) and are trained to use them by reinforcement learning with verifiable rewards (RLVR), where a checkable outcome (did the code run? did the answer match?) supplies the reward (Li et al., 2025; Feng et al., 2025; Jin et al., 2025; Qian et al., 2025). The dominant recipe (dense process rewards, GRPO (Shao et al., 2024), a 7B-scale policy) is treated as portable across tool families. We test that portability on a knowledge graph (KG), whose interface is by construction the opposite of the pretraining-aligned surfaces RLVR usually enjoys. Our benchmark is Complex WebQuestions (CWQ) (Talmor and Berant, 2018), a standard multi-hop question-answering task over the Freebase (Bollacker et al., 2008) knowledge graph, which we re-cast from its usual gold-subgraph/SPARQL form into a *deliberately minimal tool-use interface*: the policy must reach answers by composing four navigation verbs (the outgoing or incoming relations of an entity, and the entities reachable from it along a relation), where every entity is an opaque machine ID (`m.0d3k14`) and every miss returns an uninformative empty list. Python tool use leaks natural-language signal the model saw in pretraining (a traceback names the failing line), whereas this KG interface leaks almost none. We ask whether the standard RLVR-tool recipe still converges once that surface is removed.

Across seven reward designs trained on Qwen2.5-7B-Instruct (Team, 2025) + CWQ, we find four mechanistically distinct failure modes, none solved by the next reward design in the same family. The signature finding is a *peak-then-collapse* (§4.3): a self-verifiable retrieval reward (one that checks a returned entity appears verbatim in the answer, gold-free) lifts *grounded retrieval* (the correct-via-tool rate, CvT) from 3.8% to 9.6% over 250 training steps, then collapses to 0% in a single 50-step window (Figure 1, replicated across four seeds).

Agentic-RL collapse phenomena have been documented (Wang et al., 2025a; Jin et al., 2025; Wang et al., 2026), but at the level of action-distribution entropy and reward variance. The specific mechanism we observe, a one-shot quote-and-stop attractor with mean tool calls per ques-

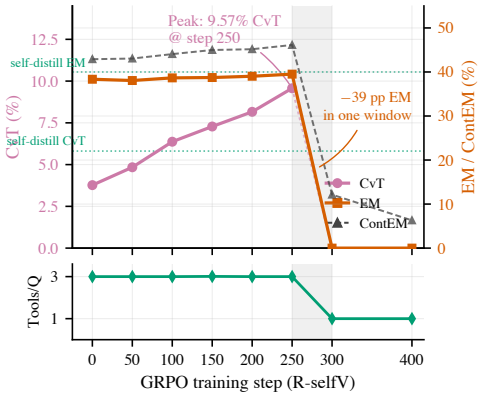

Figure 1: **Peak-then-collapse under self-verifiable retrieval (R-selfV).** Top: CvT (correct-via-tool rate; left, blue) climbs 3.77→9.57% (steps 0–250) then catastrophically collapses to 0 by step 300; EM (exact match; right, red) follows. Bottom: Tools/Q drops 3.0→1.0 across the shaded 50-step window. self-distill (dashed) reaches 40.0% EM / 5.81% CvT in the same reward family. Mechanism: §4.3.

tion (Tools/Q) dropping 3.0→1.0 in a 50-step window, is a call-volume signature distinct from these precedents and from the standard RLHF reward-hacking framework (Gao et al., 2023; Pan et al., 2022; Skalse et al., 2022), which reports proxy-overoptimisation at the level of learned reward models. Stripping that pretraining-aligned surface exposes a degradation regime no reward redesign in the same family can fix.

We run a controlled comparison: one model, one task, one algorithm (GRPO), one tool API (the 4-verb Freebase interface), seven reward designs read as a pragmatic ladder (§3.4; two rungs also vary the KL tether or initialisation). To explain the convergent failures we formalise the tool-API-to-policy-gradient pathway as four channels along which a tool API supplies (or fails to supply) RL signal: silent failure, symbolic schema, opaque compositional state, and absent pretraining prior ($L_{\text{sig/lang/comp/prior}}$). The KG case is degraded along all four relative to natural-language tool surfaces (§5). We test the framework's dominant-bottleneck attribution via a direct oracle ablation: gold-relation injection at the strongest in-family checkpoint lifts EM (exact match) by only +0.20 pp, ruling out relation selection as the binding constraint.

Contributions. **(C1)** A controlled comparison of seven reward designs on a single model × task × algorithm × tool API, organised into a four-mode failure taxonomy (§4.3). **(C2)** A signal-theoretic framework (four channels $L_{\text{sig/lang/comp/prior}}$) supported by a direct oracle ablation ruling out relation selection as the bottleneck, plus a step-level error decomposition where 95.4% of 1,958 *retrieval-dependent* errors localise to retrieval composition, not answer extraction (§5.3). **(C3)** A one-iteration ReST-EM-style self-distillation recipe (self-distill) reaching EM=40.0% at 7B; the 5K rule-based pre-SFT prefix is dispensable (raw-base init reaches the same ≈ 40%) and 14B matches it (40.2%, within sampling noise). Concurrent richer pipelines reach higher EM via interface enrichments along channels identified by the framework (§4.5, §2). **(C4)** Step-resolved (50-step cadence, full CWQ test set) peak-then-collapse on a KG-tool agent: a call-volume signature (Tools/Q 3.0→1.0) novel relative to the action-distribution entropy and reward-variance diagnostics of prior agentic-RL collapse (Wang et al., 2025a; Jin et al., 2025; Wang et al., 2026) (§4.3).

We release all code, reward implementations, per-checkpoint full-test evaluations, and the seven-category trajectory classifier at `https://anonymous.4open.science/r/KG_GRPO-D47D`.[1]

2 Related Work

Process-reward RL for tools and knowledge graphs. Process-reward recipes for tool-using LLMs (Jin et al., 2025; Chen et al., 2025; DeepSeek-AI, 2025; Wang et al., 2025b; Qian et al., 2025; Zhang et al., 2025c) succeed on tool APIs with pretraining-aligned surfaces; we test whether the same recipe transfers to KG tools deliberately stripped of natural-language surface. KG-specific variants using KGs to *define* rewards (fs1 (Zhang et al., 2025b); HyperGraphPro (Park et al., 2026); KG-Implicit-RM (Kansal and Jha, 2026)) differ from our setup, which uses the KG to evaluate the policy (under RoG's (Luo et al., 2024) subgraph filter; strict EM motivated by Zhang et al. (2025a)). On the same Qwen2.5-7B-Instruct + CWQ benchmark, richer-pipeline concurrent systems reach higher numbers by enlarging the interface (GraphWalker (Xu et al., 2026) adds ∼ 21K-trajectory SFT + BM25 + LLM reranker, 79.6% EM; Graph-RFT (Song et al., 2025) adds web fallback over a 40%-incomplete KG, 67.2% Hits@1), confirming that interface enrichment, not capacity, is the differentiating axis. Our minimal-interface scope is a deliberate complement, isolating what

[1] URL anonymised for review.

gradient-bearing signal a stripped KG tool affords.

Cold-start SFT-then-RL. Our self-distillation recipe sits in the cold-start-SFT-then-RL family (Feng et al., 2025; Li et al., 2025; Chen et al., 2023); the methodological extension is that our SFT corpus is *self-distilled* from our own RL checkpoint with a strict EM-correct ∩ tool-productive ∩ format-valid filter, connecting to ReST-EM (Singh et al., 2024). The no-rule-prefix variant (§4.5) shows init-from-base alone (no rule-SFT, no re-distillation) suffices to reach the 40% ceiling, corroborating ReST-EM's init-from-base prescription on a different task family.

Reward hacking and agentic-RL collapse. The peak-then-collapse mode (§4.3) extends a known reward-hacking pattern. RLHF (Skalse et al., 2022; Pan et al., 2022; Gao et al., 2023) catalogues proxy-overoptimisation at the learned-reward-model level; agentic-RL diagnostics include action-distribution entropy (Yu et al., 2025), reward variance (Echo Trap Wang et al., 2025a; Wang et al., 2026; Deng et al., 2025), and metric-without-capability gain (Shao et al., 2025) (closely related to our R-stepwise tool-free shortcut, §4.3); Yue et al. (2025) (RL on math does not push pass@k beyond base) is our closest non-KG analogue. Our peak-then-collapse contribution is the Tools/Q call-volume signature (3.0→1.0) and step-resolved cadence on a KG-tool agent.

3 Methodology

3.1 Task and KG tool interface

We use Complex WebQuestions (CWQ) (Talmor and Berant, 2018) over Freebase (Bollacker et al., 2008), following the standard RoG (Luo et al., 2024) splits (27,639 train / 3,531 test).[2] Answers are normalised (lowercase, strip articles, strip punctuation, whitespace-collapse) before exact-match scoring. The RoG subgraph filter yields 2.59M entities, 7,058 relations, 8.3M triples accessible at inference time.

The model accesses the KG through a *four-verb* tool interface:

- `get_tail_relations(e)`, `get_head_relations(e)`: list relations outgoing from / incoming into entity e.
- `get_tail_entities(e, r)`, `get_head_entities(e, r)`: list entities reachable from e along relation r in either direction.

Tool calls are emitted as `<search>`*`verb(args)`*`</search>` tokens; the KG response is inlined verbatim into the context and generation continues. Up to five tool-calling turns are permitted per question. The interface is deliberately minimal: four verbs, one symbol per query, responses that list entities or relations line-by-line. Unlike a web-search RAG surface, every query response can be objectively classified as successful, empty, or off-schema, the property that makes step-resolved reward analysis tractable. Figure 7 (App. E) illustrates the re-cast.

3.2 Shared SFT initialisation

All reward-family ablations share a single SFT-initialised Qwen2.5-7B-Instruct (Team, 2025) checkpoint with LoRA (r=64, α=128, DoRA + rslora) on 5,000 rule-based trajectories whose `<search>` calls are constructed *from the gold triple chain* (no LLM teacher: no GPT-4, no Claude, no Qwen-72B distillation). The rule-based generator and the 14,082-trajectory G-variant corpus are in Appendix C. Capacity (Qwen2.5-14B; Table 2) and cross-family (Llama-3.1-8B-Base; App. F) replications use model-specific SFT pre-warms with the same recipe.

3.3 GRPO training

GRPO (Shao et al., 2024) on the SFT base: group size 8, `low_var_kl` loss, KL coefficient 0.05 unless noted (Table 2 for per-variant overrides), lr 3e−7, batch 128, 500 steps; `verl` (Sheng et al., 2024) + vLLM (Kwon et al., 2023) rollouts on NVIDIA GH200. Reward implementations are unit-tested (18/18 passing; Appendix B).

3.4 Reward ladder

Intuitively, the ladder first tests whether final-answer reward alone preserves the tool format, then adds verifiable tool-call structure, then rewards productive tool use, then stabilises the policy with a stronger KL tether. R-selfV tests whether a gold-free retrieved-entity signal can improve grounding, while self-distill tests whether successful trajectories can stabilise the best in-family policy. Table 1 summarises the six-rung design.

[2] 27,639 is RoG's row-expanded count over (question, gold-path) pairs; the canonical `rmanluo/RoG-cwq` release reports 16,894 unique training questions. Talmor & Berant's original v1.1 release uses 27,734 / 3,480 / 3,475; the residual differences are Freebase-reachability filtering.

Rung	Name	Patch over predecessor	EM	CvT	Mode emerged
1	R-binary	r_{out} only	0.000	—	Mode 1 (format)
2	R-stepwise	$+\,r_{\text{valid}} + r_{\text{path}} + r_{\text{coh}}$	0.325	0.03%	Mode 2 (ritual)
3	R-toolverbs	$+\,r_{\text{tool-type}}$ bonus	0.322	3.03%	Mode 3 (drift)
4	R-toolverbs·KL	KL $\times 5$ (SFT-prior tether)	0.384	3.77%	stable to 400+
5	R-selfV*	$+\,r_{\text{retrv}}$ (entity in answer)	0.395	**9.57%**	Mode 4 (cliff)
6	**self-distill***	R-toolverbs + self-distilled SFT pre-warm	**0.400**	5.81%	stable (best EM)

Table 1: **Reward-ladder design.** Each rung patches a named failure of its predecessor (§4.3). *Rungs 5–6 fork from R-toolverbs·KL@400 (R-selfV initialises from it; self-distill self-distills from its rollouts then re-inits from base). EM/CvT in Table 2; formulas in App. B.

We organise reward design as a six-rung ladder (Table 1); each rung is introduced as a patch for a specific named failure of its predecessor (failure modes in §4.3; full reward formulas in App. B). Rungs 4 and 6 also vary non-reward ingredients (KL coefficient; initialisation/data regime), so we read the ladder *pragmatically* (does each patch fix its predecessor's failure?) rather than as strictly single-variable reward ablations.[3] **(1) R-binary**: outcome-only $r{=}0.5\,\text{EM}{+}0.5\,\text{F1}$ (Search-R1 (Jin et al., 2025), DeepSeek-R1-zero (DeepSeek-AI, 2025) family). **(2) R-stepwise**: equal-weight $r_{\text{out}}{+}r_{\text{valid}}{+}r_{\text{path}}{+}r_{\text{coh}}$ (r_{valid} checks call syntax, r_{path} rewards gold-relation queries, r_{coh} scores `<think>`-vs-call coherence; KG analogue of StepSearch (Wang et al., 2025b)). **(3) R-toolverbs**: $r{=}0.25\,r_{\text{out}}{+}0.50\,r_{\text{tool-type}}{+}0.25\,r_{\text{tool-usage}}$, rewarding all 4 verbs and non-empty responses (no gold supervision). **(4) R-toolverbs·KL**: R-toolverbs with KL coef raised $5\times$ ($0.05{\to}0.25$); extends R-toolverbs's stable window step $150{\to}400$+. **(5) R-selfV**: $r{=}0.25\,r_{\text{ans}}{+}0.50\,r_{\text{tool-type}}{+}0.25\,r_{\text{retrv}}$, $r_{\text{retrv}}{=}|\mathcal{T}_{\text{prod}}|/|\mathcal{T}|$ where $t{\in}\mathcal{T}_{\text{prod}}$ iff response non-empty and $\geq$1 returned KG entity appears verbatim in `<answer>`. Gold-free at GRPO training. **(6) self-distill**: roll out R-toolverbs·KL@400 greedily, filter EM-correct/tool-productive/format-valid trajectories (14,082, 50.9% yield), SFT from base, then R-toolverbs GRPO. One-iteration ReST-EM (Singh et al., 2024) variant; init-from-base ablation (§4.5) drops the 5K rule-based pre-SFT prefix and reaches the same $\sim 40\%$, corroborating ReST-EM's init-from-base prescription.

3.5 Evaluation

Unless otherwise noted, numbers are reported on the full 3,531-question CWQ test set with greedy decoding, `max_new_tokens`=512, and up to five tool-calling turns; pass@k and the closed-model control instead use the smaller subsets and sampling temperatures noted inline.

Held-out checkpoint selection. To confirm that reported checkpoints (e.g., the R-selfV step-250 peak) are not test-set-informed, we additionally hold out a hop-stratified dev split ($n{=}800$, seed 42) and re-select every reported checkpoint on dev alone: the dev peak is step 250 (matching the body), and held-out test EM ($n{=}2{,}731$) lies within the Wilson 95% half-width (± 1.85 pp) of full-set EM for all ten checkpoints (App. B.10).

Scalar metrics.

- **EM**: exact match after normalisation.
- **CvT**, the *correct-via-tool* rate: EM-correct *and* ≥ 1 returned KG entity appears verbatim in the final `<answer>`.
- **ContEM**, *containment* EM: gold answer is contained as a substring of the normalised prediction (looser than EM; reported for closed-model controls where strict gold-form normalisation is lossy, e.g., GPT-4o in App. D, and for collapse-trajectory diagnostics, e.g., Figure 1).
- **hard-partition CvT**: CvT restricted to the Cat B subset ($n{=}2{,}556$ structurally-hard questions; pass@10=0 on base Qwen; App. B.8).
- **Tools/Q**: mean number of `<search>` calls per question.

Seven-category trajectory classification. Every trajectory is assigned exactly one of: `correct-via-tool` / `correct-via-memory` / `correct-no-tool` / `wrong-no-tool` / `kg-incomplete` / `tool-misuse` / `wrong-answer`; full definitions + decision tree in Appendix B.

Statistics. Confidence intervals are Wilson 95% (Wilson, 1927); paired comparisons use McNemar's test (McNemar, 1947) (exact binomial for

[3] Three earlier variants (E2 heuristic, E4 random-reward control, E5a retrieval-grounded cold-start) were run but cut as redundant; see App. B.6.

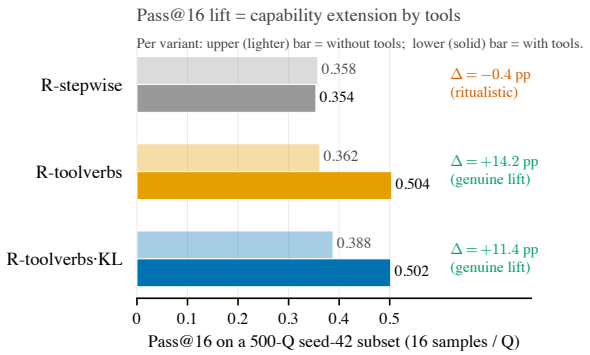

Figure 2: Pass@16 with vs. without tool access on a 500-Q seed-42 subset (16 samples per question). R-stepwise's gap is statistically zero (-0.4 pp): the policy has stopped using tools. R-toolverbs and R-toolverbs·KL each show $+11.4$ to $+14.2$ pp of genuine tool-lift. self-distill (EM-best) is measured separately (App. B.11): a $+11.5$ pp tool-lift confirms it genuinely consumes retrieval.

$n < 25$, continuity-corrected χ^2 otherwise). For pass@k (§4) we additionally evaluate 16 samples per question at temperature 1 on a 500-question seed-42 subset.

4 Results: A Reward Ladder Producing Four Failure Modes

4.1 The reward-ladder matrix

The strongest variant (self-distill) reaches EM=40.0%, and R-toolverbs·KL-14B reaches 40.2% (a difference within sampling noise), consistent with an interface-level bottleneck (§5).

4.2 Pass@16: tools as decoration vs. capability extension

On 500 seed-42 questions $\times$ 16 samples, R-stepwise's with-vs-without-tools gap is statistically zero (-0.4 pp; Figure 2), while R-toolverbs@100 and R-toolverbs·KL@400 show $+14.2$ and $+11.4$ pp; the EM-best self-distill shows a genuine $+11.5$ pp lift on a separate subset (App. B.11). The gap distinguishes ritualistic from capability-extending tool use, complementing Yue et al. (2025) (RL does not push pass@k beyond the base model): process rewards shift the answer source from parametric memory to retrieval consumption.

4.3 Four failure modes

The seven reward designs produce four mechanistically distinct failure modes, differing in *what fails*: format, tool-use ritual, format-after-tool-use, or the signal the reward optimises. Each successive design *shifts* the failure rather than eliminating it. Mode 4 (peak-then-collapse) is the strongest expression of this pattern and we give its mechanism in full here; Modes 1–3 follow the same shift-not-eliminate dynamic and are summarised below, with full mechanism prose in App. B.5. The seven-category trajectory classifier (App. B; exemplars in App. G) labels every test-set trajectory.

Mode 1: Sparse-Reward Format Collapse (R-binary, R-binary-SR). Both reach EM=0.000 at format-valid $0/3531$. Outcome-only reward fires only on EM-match; with L_{sig} producing identical `[]` early responses and L_{prior} no schema bias, the gradient never reaches the format-keeping subspace before low-entropy degeneracy locks in. R-binary-SR (Search-R1's exact `kl_coef=0.001` config; App. A) collapses identically, ruling out our hyperparameters as cause and locating collapse as intrinsic to outcome-only reward on this interface.

Mode 2: Verify-Then-Answer Goodhart (R-stepwise). R-stepwise reaches EM=32.5% but **CvT=**0.03% ($1/3531$): Tools/Q=1.0 with 100% of calls hitting `get_tail_relations` (the only verb that never returns empty). The 4-component verifiable-step reward incentivises structural existence checks over tool-output consumption, a tool-free shortcut (component-gaming Goodhart, cf. Gao et al., 2023). The Pass@16 with-vs-without-tools gap of -0.4 pp (Figure 2) confirms tools are decorative.

Mode 3: Format Drift Collapse (R-toolverbs, late stage). R-toolverbs briefly succeeds at step 100 (EM=32.2%, **CvT=**3.03%, Tools/Q=2.30, the first reward with non-trivial grounded retrieval) but by step 150 `<search>` tags drift inside `<think>` tags; by step 200 calls fail to parse; by step 250 EM $\rightarrow$ 0%. R-toolverbs rewards tool-type diversity but anchors no output format; format conventions drift from the SFT prior under loose KL. Resolution: R-toolverbs·KL raises KL $5\times$ ($0.05\rightarrow0.25$), extending the stable window to step 400+. A fixed-reward KL sweep suggests that stronger tethering alone is not the full explanation: the broader instability arises because malformed or failed KG interactions provide little corrective signal (App. B.5).

Mode 4: Specification Gaming Cliff (R-selfV). R-selfV, the R-toolverbs reward augmented with a self-verifiable retrieval-contribution term (a returned KG entity must appear verbatim in the answer; full reward in App. B), exhibits the most

Model	Step	EM	CvT count	CvT %	Wilson 95% CI	Tools/Q
SFT base (pre-GRPO)	—	0.001	—	—	—	4.96
R-binary	1250	0.000	—	—	—	0–1
R-stepwise	500	0.325	1 / 3531	0.03%	[0.00, 0.16]	1.00
R-toolverbs	100	0.322	107 / 3531	3.03%	[2.51, 3.65]	2.30
R-toolverbs·KL	400	0.384	133 / 3531	3.77%	[3.19, 4.45]	3.00
R-selfV @ 250 (peak)	**250**	**0.395**	**338 / 3531**	**9.57%**	**[8.65, 10.59]**	**3.00**
R-selfV @ 300 (collapsed)	300	0.000	—	collapsed	—	1.00
init-from-iterate	500	0.394	162 / 3531	4.59%	[3.95, 5.33]	3.00
self-distill	**500**	**0.400**	**205 / 3531**	**5.81%**	**[5.08, 6.63]**	**3.00**
R-binary-SR (Search-R1 reimpl.)[‡]	best	0.000	0 / 3531	0.00%	[0.00, 0.11]	0.00
R-toolverbs·KL-14B	400	0.402	226 / 3531	6.40%	[5.64, 7.26]	3.99

Table 2: **Reward-ladder results on full CWQ test (N=3,531), Qwen2.5-7B-Instruct unless noted. self-distill @ 500** is the EM-best 7B variant (40.0%); **R-selfV @ 250** reaches the 9.57% CvT peak before the peak-then-collapse (§4.3). Wilson 95% CIs reported for CvT; for EM at n=3531 the half-width is $\pm$1.6 pp, so EM differences below $\sim$ 3 pp are within sampling noise. McNemar paired tests: R-toolverbs·KL@400 vs. R-stepwise@500, $p<10^{-4}$; self-distill vs. init-from-iterate on CvT, $p<10^{-2}$. self-distill retrained from scratch under three additional seeds (4 total, mean EM=40.21%, std 0.46 pp; CIs all overlap); other rows are single-seed point estimates with Wilson 95% CIs (the R-selfV cliff signature replicates across 4 seeds, §4.3). [‡]R-binary-SR replicates Search-R1's `kl_coef=0.001` config; format-valid 0/3531, so EM=0 is format collapse not retrieval failure (App. A).

distinctive trajectory and is the paper's signature finding.

Climb. Initialised from R-toolverbs·KL@400 (CvT=3.77%), CvT climbs monotonically across five 50-step windows: 3.77 $\rightarrow$ 4.84 $\rightarrow$ 6.37 $\rightarrow$ 7.28 $\rightarrow$ 8.16 $\rightarrow$ 9.57% (peak at step 250; Figure 1, top). EM moves in lockstep at low amplitude (0.384 $\rightarrow$ 0.395). Mean Tools/Q is steady at 3.00, so the gain is entirely a query-precision improvement at constant call volume. The signature replicates across four seeds (peak step 200–250, full collapse to EM=0 by step 300 in every seed).

Cliff. Between checkpoints 250 and 300 (a single 50-step window), the policy collapses: EM 0.3951 $\rightarrow$ 0.0000, ContEM 46.13% $\rightarrow$ 12.14%, Tools/Q 3.00 $\rightarrow$ 1.00 (Figure 1, bottom). By step 400 ContEM is 6.30%; recovery does not occur. *No training hyperparameter changes between step 250 and step 300*; the collapse is entirely policy-internal.

Mechanism (hypothesis). The retrieval-contribution term r_{retrv} fires whenever a returned entity appears verbatim in the answer. Once the policy has learned to quote returned entities (the productive climb), a *locally optimal degenerate* strategy becomes available: call any single tool, copy any entity from its response, emit that entity as a terse answer, stop. This guarantees a productive call at minimum sampling budget, maximising r_{retrv} while sacrificing EM. Two analyses confirm the mechanism (Figure 3). First, per-checkpoint reward decomposition shows r_{outcome} dies first (0.41 $\rightarrow$ 0.0007 across step 250–300) while $r_{\text{tool-type}}$ and r_{retrv} *briefly saturate* at $\approx$ 0.51 at step 300 (the per-call productivity denominator trivially hits 1.0 once Tools/Q= 1) before collapsing to $\approx$ 0.09 by step 400. Second, token-level entropy inside `<search>` tags drops 0.35 $\rightarrow$ 0.21 nats across step 200–400 while outside-tag entropy stays at floor ($\sim$ 0.01–0.04 nats), a *selective* collapse on the action subspace, complementing prior agentic-RL entropy diagnostics (Yu et al., 2025; Wang et al., 2025a; Deng et al., 2025).

Why this matters. This proxy-overoptimisation shape is catalogued in RLHF and agentic-RL (Gao et al., 2023; Wang et al., 2025a; Deng et al., 2025); our contribution is the Tools/Q call-volume signature (3.0$\rightarrow$1.0) characterising it at a behavioural level on a KG-tool agent (§2).

4.4 Trajectory composition

The seven-category breakdown (full table in App. B.3) reveals two patterns. *kg-incomplete* is the largest non-success bucket for every rung-3-or-better checkpoint, falling monotonically across rungs 3–6 (1,435 $\rightarrow$ 1,201 $\rightarrow$ 824 $\rightarrow$ 818) as the policy reaches the KG more productively. The *wrong-answer* bucket rises in lockstep: retrieval-grounded but the answer is wrong.

4.5 Self-distillation: a one-iteration recipe reaches the 40% ceiling

self-distill and **self-distill (init-from-base)** share an identical 14,082-trajectory self-distillation SFT (strict-filtered R-toolverbs·KL@400 rollouts, 50.9% yield) then R-toolverbs GRPO; they differ

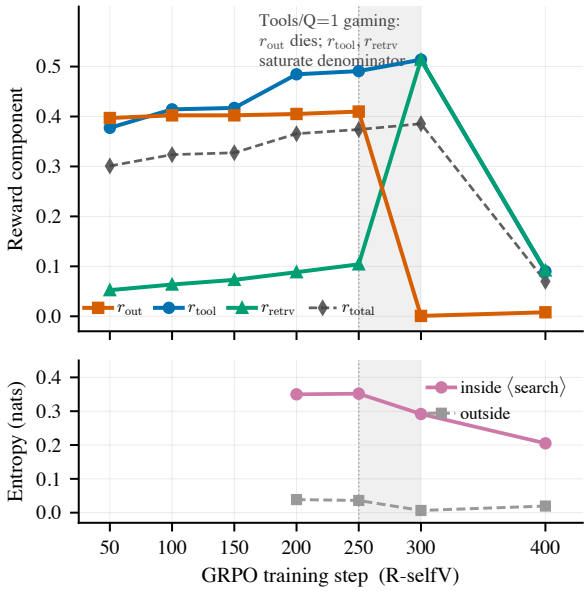

Figure 3: **Mode 4 mechanism evidence.** *Top:* per-checkpoint reward components (R-selfV; HPC seed-42). r_{out} dies in the collapse window (shaded); r_{tool} and r_{retrv} briefly saturate at ≈ 0.51 at step 300 (denominator gaming with Tools/Q$=1$) before themselves collapsing by step 400. The weighted total r_{total} peaks *post-collapse* (step 300): specification-gaming Goodhart made visible. *Bottom:* token entropy in nats, partitioned by inside vs. outside `<search>` spans. Inside-tag entropy drops $0.35 \rightarrow 0.21$ across step 200–400; outside-tag entropy stays at floor, a selective action-subspace collapse.

only in whether a 5K rule-based gold-path pre-SFT prefix is prepended (self-distill: yes; init-from-base: no, from raw base). self-distill reaches EM=40.0%, CvT=5.81% at step 500 (Table 2; Tools/Q=3.00), replicated across four seeds (mean EM=40.21%, std 0.46 pp); init-from-base matches at $\approx$ 40% across step 100–500; the 5K rule-based prefix is dispensable. **self-distill (init-from-iterate)** (no SFT; init from R-toolverbs·KL@400 iterate) reaches EM=39.4%: the self-distillation SFT, not the rule-based prefix, is the operative ingredient, consistent with ReST-EM's (Singh et al., 2024) init-from-base prescription.

4.6 External validity

Three controls rule out simple alternative explanations. *(i) KG coverage*: bounded BFS marks 99.5% of CWQ and 69.14% of KGQAGen-10k (Zhang et al., 2025a) as 2-hop-solvable, both well above the 3–10% CvT our models reach (coverage is not the binding constraint). *(ii) Cross-benchmark*: 8 base Qwen variants give Spearman ρ=0.976 between CWQ and KGQAGen-10k EM orderings (P4). *(iii) GPT-4o*: under the same 4-verb interface, near-100% format-validity but only 6.2% ContEM and 2.1% F1 on n=500 with $\sim$ 78% of trajectories self-reporting "unable to retrieve X" (App. D), a schema-unfamiliarity control, not a head-to-head baseline.

Reward	Failure mode	Dominant channel
R-binary / R-binary-SR	Mode 1: format collapse	$L_{\text{sig}}+L_{\text{prior}}$
R-stepwise	Mode 2: ritual tool use	L_{lang} (gameable)
R-toolverbs	Mode 3: format drift	L_{sig} (silent fail)
R-toolverbs·KL	stable, low CvT	$L_{\text{sig}}+L_{\text{lang}}$
R-selfV	**Mode** 4: peak-then-collapse	L_{sig} (gameable)
self-distill (stable)	$\sim$40% interface ceiling	all 4 channels bind

Figure 4: **Synthesis: reward $\rightarrow$ failure mode $\rightarrow$ dominant interface channel.** Each ladder rung (§3.4) produces a distinctive failure mode (§4.3) with channel attribution from §5.3. R-toolverbs·KL is stable but plateaus at low CvT (a non-Mode-numbered intermediate); self-distill attains the $\sim$ 40% EM ceiling but is bound by all four channels jointly.

5 A Signal-Theoretic Account of Tool-Interface Difficulty

The four failure modes characterised in §4.3 share a common substrate: each is a degradation of the gradient signal GRPO needs to converge on a tool-grounded policy. We propose this substrate has structure: four interface-level channels along which a tool API supplies (or fails to supply) usable RL signal (Figure 4), yielding four channel-level claims. These channels are an organising account of the substrate: §5.3 validates one of them directly (the relation oracle), and cleanly isolating each of the others is future work.

5.1 Four interface-level channels

Why do small ($\leq$7B) Qwen-class models learn Python (Li et al., 2025; Feng et al., 2025), web-search (Jin et al., 2025), and JSON-API (Qian et al., 2025) tool use under similar GRPO recipes, yet fail at compositional KG tool use under all seven reward designs? Four structural properties of the interface jointly determine RL trainability; the KG case is degraded along all four channels relative to natural-language-surface tools (diagram + tool-family table in App. E):

- L_{sig} **(silent failure).** A KG miss returns `[]` with no information about *why* (wrong entity, wrong relation, or both), in contrast to a Python traceback or HTTP `400` with named parameter. On our 7-reward matrix, $\geq 99\%$ of erroneous tool calls return identically `[]` and contribute no credit-assignment signal.
- L_{lang} **(symbolic schema).** KG queries use opaque MIDs (`m.0d3k14`) and dotted relation names (`people.person.place_of_birth`); pretraining corpora contain essentially zero exposure to this schema, in contrast to Python (10–100 B tokens) or natural-language search queries.
- L_{comp} **(opaque compositional state).** Multi-hop KG reasoning carries entity IDs across turns; these IDs are semantically opaque and any paraphrasing breaks the chain.
- L_{prior} **(no pretraining prior).** Zero pretraining examples of Freebase schema navigation, versus extensive Python / search / JSON-API exposure.

5.2 Four channel-level claims

We treat P1–P4 (Table 3) as joint *consistency checks* of the descriptive structure (per-claim caveats in App. E.4).

Claim	Evidence	Reading
P1 Outcome-only collapse	R-binary+R-binary-SR EM=0%, format-valid 0/3531 (Table 2)	confounds L_{sig} with L_{prior} *(observational)*
P2 Capacity invariance	14B vs 7B self-distill hard-partition CvT 4.61% vs. 4.19% (CIs overlap; §4.5)	capacity does not resolve the bottleneck *(ablation)*
P3 Self-distillation sufficient	raw-base and rule-SFT-prefixed both $\approx$ 40% EM, $\geq$ init-from-iterate's 39.4%	matches ReST-EM's init-from-base *(ablation)*
P4 Cross-benchmark transfer	KGQAGen-10k Spearman ρ=0.976 on 8 base Qwen variants (§4)	not CWQ-specific *(interpretive)*

Table 3: Four channel-level claims (P1–P4) as joint consistency checks; per-claim evidence status in *italics* (observational / ablation-supported / interpretive). The stronger evidence is the direct decomposition in §5.3.

5.3 Mechanism evidence chain

We characterise self-distill's 2,119 strict-EM=0 error trajectories in two stages.

Schema-surface failures are common. A 4-channel attribution rubric on these 2,119 has L_{sig} dominant at 55.2%; a decision-boundary analysis (App. B.4) finds 70.4% of kg-incomplete first-search calls within Lev$\leq$ 1 of gold (4.7$\times$ null; held-out-unseen 77.7% rules out memorisation), 89.1% entity-correct/relation-wrong (90.4% same-entity-different-relation vs. R-toolverbs·KL).

But relation selection is not the EM bottleneck. We inject the gold-chain relation at every `<search>` call (replacing the model's relation argument with the ground-truth relation while keeping its entity argument) and re-evaluate. This oracle lifts EM by only $+\mathbf{0.20}$ pp (40.19% vs. 39.99%), within Wilson 95% CI half-width ($\pm$1.6 pp at n=3,531): relation selection alone cannot move EM. The residual gap must lie either upstream of relation selection (entity threading across hops, L_{comp}) or downstream (answer extraction). Step-level classification of the 1,958 *retrieval-dependent* errors (Table 9) localises $\mathbf{95.4}\%$ to *retrieval-composition failures* vs. 4.6% *answer-extraction failures*: the residual gap is upstream of answer extraction, not in relation selection.

A Llama-3.1-8B-Base cross-family check (App. F) also fails to train a tool-grounded policy under this interface, but via a distinct parametric-memory mode-collapse (Tools/Q=0), indicating the specific collapse mode is model-family-dependent.

6 Conclusion

Within a deliberately controlled scope (one model family, one task, one RL algorithm), seven GRPO-trained reward designs on Qwen2.5-7B-Instruct + CWQ reveal four interface-level failure modes ($L_{\text{sig/lang/comp/prior}}$); an oracle ablation rules out relation selection as the EM bottleneck, locating the residual gap upstream of answer extraction (the specific channel is unresolved). A one-iteration self-distillation recipe reaches 40.0% EM at 7B (self-distill), matched by 14B; richer concurrent pipelines exceed this via interface enrichments. Advances need interface-aware training, not larger models or denser proxy rewards. Methodologically, our 50-step cadence caught the peak-then-collapse mode; step-resolved instrumentation is essential for process-reward RL. Open: whether a non-gameable self-verifiable signal exists; how far interface enrichments push the ceiling is the richer-pipeline line's domain, where our channel attribution interprets their gains.

Limitations

Our empirical scope is bounded in several ways. All main-body results use Qwen2.5-7B-Instruct, consistent with the concurrent KGQA-RL evalua-

tion norm (GraphWalker, Graph-RFT, StepSearch, ToolRL, and Search-R1 all evaluate on Qwen2.5-7B-Instruct); cross-family replication on Llama-3.1-8B-Base (App. F) reaches comparable EM but exhibits a parametric-memory mode-collapse distinct from Qwen's peak-then-collapse, indicating the specific collapse mode is model-family-dependent. On scale, the R-toolverbs·KL-14B replication (Table 2: EM=40.24%, CvT=6.40%; Table 8: hard-partition CvT 4.61%, within noise of 7B self-distill's 4.19%) is positive evidence that the $\sim$40% ceiling on this reward family is interface-bound rather than capacity-bound, but the capacity-invariance argument (§5.2, P2) anchors only on the 7B–14B comparison, with $\geq$32B replication left to follow-up.

Our GRPO training is on CWQ only. KGQAGen-10k provides cross-benchmark validation: comparable bounded-BFS solvability coverage and a strong EM-ordering correlation across base Qwen variants (Spearman ρ=0.976; §4.6) indicate the findings are not CWQ-specific. We do not, however, retrain or evaluate our trained policy on a second benchmark, which we leave to future work, and we make no claims about web-search-style RAG or question answering over non-KG textual corpora.

Finally, per-step monitoring (50-step cadence) and final reporting both use the full 3,531-question CWQ test set, so specific checkpoint labels such as the R-selfV peak at step 250 are test-set-informed selections. We frame this as a controlled diagnostic study rather than a leaderboard claim: the peak-then-collapse signature replicates across four independent seeds, peaking in steps 200–250 and fully collapsing by step 300 in every seed (§4.3), and the comparative claims that anchor the framework (capacity invariance, the oracle $+0.20$ pp, the 95.4% retrieval-composition split) are robust to a $\pm$50-step window. A held-out dev/test re-selection (§3.5; App. B.10) confirms the reported checkpoints are not test-set artefacts: the dev peak is step 250 and all ten checkpoints fall within ±1.85 pp, so we retain full-set reporting for comparability.

Acknowledgments

In preparing this work, the authors used AI-based assistants for language editing of the manuscript (drafting and revising prose and LaTeX formatting) and for coding assistance (implementing and debugging the training and evaluation pipelines), under the authors' direction. All research questions, experimental design, analyses, and claims are the authors' own; the authors verified all AI-assisted text and code for correctness and take full responsibility for the content of this paper.

References

Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: A collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, pages 1247–1250. ACM.

Baian Chen, Chang Shu, Ehsan Shareghi, Nigel Collier, Karthik Narasimhan, and Shunyu Yao. 2023. FireAct: Toward language agent fine-tuning. *Preprint*, arXiv:2310.05915.

Mingyang Chen, Linzhuang Sun, Tianpeng Li, Haoze Sun, Yijie Zhou, Chenzheng Zhu, Haofen Wang, Jeff Z. Pan, Wen Zhang, Huajun Chen, Fan Yang, Zenan Zhou, and Weipeng Chen. 2025. ReSearch: Learning to reason with search for LLMs via reinforcement learning. *Preprint*, arXiv:2503.19470.

DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *Preprint*, arXiv:2501.12948.

Wenlong Deng, Yushu Li, Boying Gong, Yi Ren, Christos Thrampoulidis, and Xiaoxiao Li. 2025. On group relative policy optimization collapse in agent search: The lazy likelihood-displacement. *arXiv preprint arXiv:2512.04220*.

Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang, Jinxin Chi, and Wanjun Zhong. 2025. ReTool: Reinforcement learning for strategic tool use in LLMs. *Preprint*, arXiv:2504.11536.

Leo Gao, John Schulman, and Jacob Hilton. 2023. Scaling laws for reward model overoptimization. *Preprint*, arXiv:2210.10760.

Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. Search-R1: Training LLMs to reason and leverage search engines with reinforcement learning. *Preprint*, arXiv:2503.09516.

Yuval Kansal and Niraj K. Jha. 2026. Knowledge graphs are implicit reward models: Path-derived signals enable compositional reasoning. *Preprint*, arXiv:2601.15160.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with PagedAttention. In *Proceedings of the 29th Symposium on Operating Systems Principles (SOSP)*.

Xuefeng Li, Haoyang Zou, and Pengfei Liu. 2025. ToRL: Scaling tool-integrated RL. *Preprint*, arXiv:2503.23383.

Linhao Luo, Yuan-Fang Li, Gholamreza Haffari, and Shirui Pan. 2024. Reasoning on graphs: Faithful and interpretable large language model reasoning. In *International Conference on Learning Representations (ICLR)*.

Quinn McNemar. 1947. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157.

Alexander Pan, Kush Bhatia, and Jacob Steinhardt. 2022. The effects of reward misspecification: Mapping and mitigating misaligned models. In *International Conference on Learning Representations (ICLR)*.

Jinyoung Park, Sanghyeok Lee, Omar Zia Khan, Hyunwoo J. Kim, and Joo-Kyung Kim. 2026. HyperGraphPro: Progress-aware reinforcement learning for structure-guided hypergraph RAG. *Preprint*, arXiv:2601.17755.

Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiusi Chen, Dilek Hakkani-Tür, Gokhan Tur, and Heng Ji. 2025. ToolRL: Reward is all tool learning needs. *Preprint*, arXiv:2504.13958.

Rulin Shao, Shuyue Stella Li, Rui Xin, Scott Geng, Yiping Wang, Sewoong Oh, Simon Shaolei Du, Nathan Lambert, Sewon Min, Ranjay Krishna, Yulia Tsvetkov, Hannaneh Hajishirzi, Pang Wei Koh, and Luke Zettlemoyer. 2025. Spurious rewards: Rethinking training signals in RLVR. *arXiv preprint arXiv:2506.10947*.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. DeepSeekMath: Pushing the limits of mathematical reasoning in open language models. *Preprint*, arXiv:2402.03300. Introduces GRPO.

Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2024. HybridFlow: A flexible and efficient RLHF framework. *Preprint*, arXiv:2409.19256. `verl` framework.

Avi Singh, John D Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Xavier Garcia, Peter J Liu, James Harrison, Jaehoon Lee, Kelvin Xu, Aaron Parisi, Abhishek Kumar, Alex Alemi, Alex Rizkowsky, Azade Nova, Ben Adlam, Bernd Bohnet, Gamaleldin Elsayed, Hanie Sedghi, and 22 others. 2024. Beyond human data: Scaling self-training for problem-solving with language models. *Preprint*, arXiv:2312.06585.

Joar Skalse, Nikolaus H. R. Howe, Dmitrii Krasheninnikov, and David Krueger. 2022. Defining and characterizing reward hacking. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Yanlin Song, Ben Liu, Víctor Gutiérrez-Basulto, Zhiwei Hu, Qianqian Xie, Min Peng, Sophia Ananiadou, and Jeff Z. Pan. 2025. Plan then retrieve: Reinforcement learning-guided complex reasoning over knowledge graphs. *arXiv preprint arXiv:2510.20691*.

Alon Talmor and Jonathan Berant. 2018. The web as a knowledge-base for answering complex questions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 641–651.

Qwen Team. 2025. Qwen2.5 technical report. *Preprint*, arXiv:2412.15115.

Zihan Wang, Chi Gui, Xing Jin, Qineng Wang, Licheng Liu, Kangrui Wang, Shiqi Chen, Linjie Li, Zhengyuan Yang, Pingyue Zhang, Yiping Lu, Jiajun Wu, Li Fei-Fei, Lijuan Wang, Yejin Choi, and Manling Li. 2026. RAGEN-2: Reasoning collapse in agentic RL. *arXiv preprint arXiv:2604.06268*.

Zihan Wang, Kangrui Wang, Qineng Wang, Pingyue Zhang, Linjie Li, Zhengyuan Yang, Xing Jin, Kefan Yu, Minh Nhat Nguyen, Licheng Liu, Eli Gottlieb, Yiping Lu, Kyunghyun Cho, Jiajun Wu, Li Fei-Fei, Lijuan Wang, Yejin Choi, and Manling Li. 2025a. RAGEN: Understanding self-evolution in LLM agents via multi-turn reinforcement learning. *arXiv preprint arXiv:2504.20073*.

Ziliang Wang, Xuhui Zheng, Kang An, Cijun Ouyang, Jialu Cai, Yuhang Wang, and Yichao Wu. 2025b. StepSearch: Igniting LLMs search ability via step-wise proximal policy optimization. *Preprint*, arXiv:2505.15107.

Edwin B. Wilson. 1927. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212.

Shuwen Xu, Yao Xu, Jiaxiang Liu, Chenhao Yuan, Wenshuo Peng, Jun Zhao, and Kang Liu. 2026. GraphWalker: Agentic knowledge graph question answering via synthetic trajectory curriculum. *arXiv preprint arXiv:2603.28533*.

Zhaochen Yu, Ling Yang, Jiaru Zou, Shuicheng Yan, and Mengdi Wang. 2025. Demystifying reinforcement learning in agentic reasoning. *Preprint*, arXiv:2510.11701.

Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao Huang. 2025. Does reinforcement learning really incentivize reasoning capacity in LLMs beyond the base model? *arXiv preprint arXiv:2504.13837*.

Liangliang Zhang, Zhuorui Jiang, Hongliang Chi, Haoyang Chen, Mohammed Elkoumy, Fali Wang, Qiong Wu, Zhengyi Zhou, Shirui Pan, Suhang Wang, and Yao Ma. 2025a. Diagnosing and addressing pitfalls in KG-RAG datasets: Toward more reliable benchmarking. *Preprint*, arXiv:2505.23495.

Mike Zhang, Johannes Bjerva, and Russa Biswas. 2025b. Follow the path: Reasoning over knowledge graph paths to improve LLM factuality. *Preprint*, arXiv:2505.11140.

Yaocheng Zhang, Haohuan Huang, Zijun Song, Yuanheng Zhu, Qichao Zhang, Zijie Zhao, and Dongbin Zhao. 2025c. CriticSearch: Fine-grained credit assignment for search agents via a retrospective critic. *Preprint*, arXiv:2511.12159.

A Search-R1 reimplementation rationale

We re-implement the Search-R1 recipe using our verl pipeline: outcome-only binary-EM reward (our `outcome_em_only` reward type, unit-tested), Search-R1's published hyperparameters (`kl_coef` $= 0.001$, lr $= 10^{-6}$, batch 256, 500 steps). One configuration change to our pipeline, zero dependency changes. We label this R-binary-SR throughout the paper and treat it as a *reimplementation*, not a reproduction. Table 2 reports format-valid $0/3531$ and EM$= 0$ for R-binary-SR at convergence — a format-collapse outcome under outcome-only reward, consistent with the sparse-reward format-collapse mechanism (§4.3).

B Reward formulas and trajectory classifier

B.1 Reward type definitions

All reward types are implemented in a shared `verl_reward` module with $18/18$ unit tests passing. Notation: a is the final `<answer>` text after normalisation, y is the gold answer, $\mathcal{T}$ is the ordered sequence of `<search>` calls with their KG responses.

- $r_{\text{EM}}(a, y) = \mathbf{1}[a = y]$, $r_{\text{F1}}(a, y)$ token-overlap F1.
- $r_{\text{out}} = 0.5\, r_{\text{EM}} + 0.5\, r_{\text{F1}}$.
- r_{valid}: fraction of calls in $\mathcal{T}$ that parse as valid `verb(args)` matching the 4-verb schema.
- r_{path}: $\mathbf{1}[\exists t \in \mathcal{T} : \text{rel}(t)$ on gold triple chain$]$.
- r_{coh}: sequence-match between `<think>` spans and their following `<search>` calls.
- $r_{\text{tool-type}}$: number of distinct tool verbs covered in $\mathcal{T}$, scaled to $[0, 1]$.
- $r_{\text{tool-usage}}$: fraction of calls in $\mathcal{T}$ with non-empty response.
- r_{retrv} (I-Self): $|\mathcal{T}_{\text{prod}}|/|\mathcal{T}|$, where $t \in \mathcal{T}_{\text{prod}}$ iff (i) the response is non-empty and (ii) at least one returned KG entity appears verbatim in a.
- r_{ans} (I-Self): identical to r_{out} but with r_{EM} replaced by a length-normalised form that treats multi-token gold answers leniently.

B.2 Seven-category trajectory classifier

For each trajectory τ on test question q with gold answer y, the classifier proceeds:

1. Parse τ into `<think>`, `<search>`, `<tool_response>`, `<answer>` spans.
2. If $r_{\text{EM}}(a, y) = 1$:
   - if no `<search>` call: **correct-no-tool**;
   - else if $\exists t \in \mathcal{T}$ with entity-in-answer: **correct-via-tool**;
   - else: **correct-via-memory**.
3. Else (EM=0):
   - if no `<search>` call: **wrong-no-tool**;
   - if any `<search>` call parses invalid: **tool-misuse**;
   - if all responses empty or disjoint from gold sub-graph: **kg-incomplete**;
   - otherwise: **wrong-answer**.

Full implementation: `scripts/task16_classify.py` in the release bundle. Edge cases (degenerate loops, overlong outputs, malformed `<answer>` envelopes) are logged separately for manual audit.

B.3 Seven-category trajectory breakdown across the reward ladder

Category	R-stepwise @500	R-toolverbs @100	R-tlvb ·KL @400	R-selfV @250	init-from-itr @500	self-distill @500
correct-via-tool	1	107	133	**338**	162	205
correct-via-memory	1136	1027	1218	1054	1226	1203
wrong-no-tool	14	15	6	2	0	0
kg-incomplete	209	1435	1201	824	954	**818**
tool-misuse	833	216	173	205	196	165
wrong-answer	1338	731	800	1108	993	1140

Table 4: Seven-category trajectory breakdown across reward variants on full CWQ test (N=3,531). The R-selfV column is at its peak step 250 (§4.3). *kg-incomplete* falls monotonically across rungs 3-6 while *wrong-answer* rises (R-stepwise has only 209 because it almost never calls a tool; *correct-no-tool* is uniformly near zero and omitted). Category counts use normalised answer matching, so for self-distill the error rows sum to 2,123 vs the strict EM$= 0$ count of 2,119 (4 trajectories are strict-EM-correct but normalised-wrong).

B.4 Decision-boundary bucket distribution

The bucket distribution underlying the decision-boundary analysis (§5.3):

Failure mode	self-distill@500 (n=818)
1. Relation typo ($\leq$1 edit)	**70.4%** (n=576)
2. Correct entity, wrong relation	18.7% (n=153)
3. Wrong entity, near-miss relation	8.6% (n=70)
4. Format / genuine miss	2.3% (n=19)
Entity-correct (1+2)	**89.1%**

Table 5: Bucket distribution underlying the decision-boundary analysis (§5.3). Bucket 1+2 is the entity-correct, relation-wrong subset (89.1% of n=818).

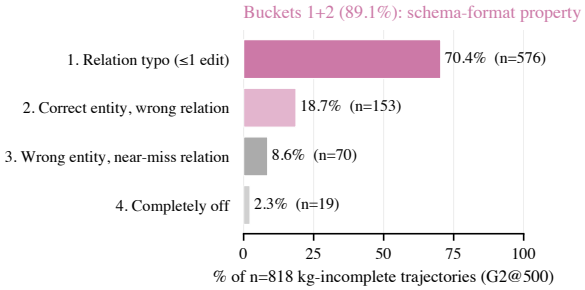

Figure 5: **Decision-boundary view of the schema-format property.** Visual of Table 5: 70.4% of 818 kg-incomplete trajectories from self-distill@500 are *exact-relation-typo* ($\leq$ 1 edit from a gold relation); buckets 1+2 (**89.1%**) are entity-correct, relation-wrong. The schema-format property is a robust surface signal but not the EM-gating mechanism (§5.3).

B.5 Failure-mode mechanism details (Modes 1–3)

Main body §4.3 presents Modes 1–3 in condensed form; full mechanism prose is preserved here. Mode 4 (the signature peak-then-collapse finding) is presented in full in the main body.

Mode 1: Sparse-Reward Format Collapse (R-binary, R-binary-SR). R-binary (outcome-only) reaches EM=0.000, F1=0.119; R-binary-SR at Search-R1's exact hyperparameters (`kl_coef=0.001`, lr=1e-6, batch=256) on our pipeline collapses identically (EM=0.000, format-valid 0/3531, ContEM=36.45%; App. A). R-binary's 0% strict EM differs qualitatively from GPT-4o's 0% on the same interface: R-binary's 0/3531 is a genuine format collapse, whereas GPT-4o is $\sim$ 100% format-valid with surface-form drift. Mechanism: outcome-only reward fires only on EM-match; with L_{sig} producing identical `[]` responses on most early-exploration calls and L_{prior} giving no schema bias, the gradient never reaches the format-keeping subspace before tail thinness drives the policy into low-entropy degeneracy. The R-binary-SR control rules out our hyperparameters as the cause; collapse is intrinsic to outcome-only reward on this interface.

Mode 2: Verify-Then-Answer Goodhart (R-stepwise). R-stepwise reaches EM=32.5% with CvT=0.03% (1/3531) — one trajectory of 3,531 actually grounds its answer in a tool response. Tools/Q is 1.0 and 100% of those calls are `get_tail_relations`, the only verb that never returns empty: tool calls are ritualistic, the answer is generated from parametric memory. Mechanism: the 4-component verifiable-step reward incentivises the easier component (call a tool that satisfies a structural existence check) over the harder one (consume the tool's output) — a tool-free shortcut (component-gaming Goodhart, cf. Gao et al., 2023); unlike RLHF proxy decorrelation, here the proxy-gold correlation r=0.95–0.97 is *maintained* throughout training. Falsification via Pass@16 with-vs-without-tools: R-stepwise's gap is -0.4 pp (Figure 2) — even with 16 attempts and tool access the policy does not solve more questions than parametric sampling; tools are decorative.

Mode 3: Format Drift Collapse (R-toolverbs, late stage). R-toolverbs briefly succeeds: at step 100 the policy reaches EM=32.2% with CvT=3.03% (Tools/Q=2.30) — the first reward design with non-trivial grounded retrieval. By step 150 `<search>` tags drift inside `<think>` tags; by step 200 tool calls stop being parsed; by step 250 EM falls to $\sim$ 0% in a repetition loop. Mechanism: R-toolverbs rewards tool-type diversity but anchors no output format; under GRPO's KL-loose configuration, format conventions drift from the SFT prior, the executor cannot parse drifted tags, the reward signal disappears, and the policy converges to a low-entropy attractor. This is not an L_{sig} failure (early-step trajectories *do* retrieve) but a format-anchor absence interacting with L_{lang}. Resolution: raising KL 5$\times$ to 0.25 (R-toolverbs·KL) extends the stable window to step 400+. A fixed-reward KL ablation (kl $\in \{0.001, 0.005, 0.01, 0.05, 0.25\}$) separates the two: collapse is *non-monotonic* in KL — only kl=0.005 collapses by step 400, while 0.001, 0.01, and 0.25 stay stable — so L_{sig} silent-failure, not KL-tether, is the unified mechanism.

B.6 Cut variants and rationale

For completeness, three reward variants were run but excluded from Table 2 because they are redundant with retained variants:

Cut variant	EM (%)	CvT (%)	Reason cut
E2 heuristic	19.91	0	Tool abandonment; redundant with R-stepwise
E4 random	5.04	0	Spurious-rewards control; single-seed sanity check
E5a retr-ground	31.49	0	Cold-start failure; redundant with R-binary

Table 6: Reward variants run but cut from Table 2, with rationale.

B.7 Reward-name to HPC tracking-code mapping

For reproducibility, Table 7 maps the paper's reward names to the HPC training-code identifiers and the corresponding verl YAML configuration files used in our release bundle. The verl configs are the canonical specification; the paper names are pedagogical labels for the body.

B.8 Cat A / Cat B stratified results

B.9 Mechanism evidence chain (full table)

B.10 Held-out dev/test checkpoint-selection robustness

To rule out test-set-informed checkpoint selection (the evaluation-protocol limitation; §3.5), we draw a hop-count-stratified split (seed 42) of the 3,531-question CWQ test set into a dev split (n=800) and a held-out test split (n=2,731), re-select each reported checkpoint on dev EM, and compare held-out test EM against the full-set EM used in the body. At n=2,731, $p\approx0.4$, the Wilson 95% half-width is ±1.85 pp. The R-selfV dev-EM peak is step 250 (scanning $\{200, 250, 300\}$) — identical to the step reported throughout. Table 10 gives the triple-column comparison: all ten reported checkpoints lie within ±1.85 pp (test$-$full), so the body's full-set numbers are not artefacts of selection on test.

B.11 Self-distill pass@16 tool-dependence

The main-body Pass@16 comparison (Figure 2) used a 500-question seed-42 subset at temperature 1 and did not originally include self-distill (the EM-best 7B system). We separately measure self-distill (loaded from the merged self-distill@500 checkpoint) on a 200-question seed-42 subset with 16 samples per question at temperature 0.7; Table 11 reports pass@$\{1, 4, 8, 16\}$ with and without tool access. The gap is $+11.5$ pp at pass@16 (48.0% vs. 36.5%) and positive at every k — self-distill genuinely consumes retrieval, unlike R-stepwise's ornamental -0.4 pp (Figure 2). Because this subset (n=200) and sampling temperature (0.7) differ from the Figure 2 protocol (n=500, temperature 1), the absolute magnitude is not directly comparable to the Figure 2 bars; the sign and scale establish genuine tool-dependence for the winning system.

C SFT corpora

Rule-based shared SFT corpus (5,000 trajectories). The shared SFT initialisation used by every reward variant is constructed from the CWQ training split as follows. (i) For each question, the gold triple chain is extracted from the WebQuestionsSP annotation. (ii) `<think>` content is filled in from one of three fixed templates with variable substitution (entities and relations drawn from the gold chain). (iii) `<search>` calls are constructed directly from the gold chain, with the tool response obtained by executing the query against the Freebase sub-graph — so every call returns by construction a response that contains the next entity on the gold path. (iv) The final `<answer>` is the gold answer. *No LLM teacher is used anywhere in this pipeline.*

Enhanced variant (6,000 trajectories). An enhanced variant `sft_trajectories_enhanced.jsonl` uses the same pipeline with additional diversity in the `<think>` template and richer tool-call ordering. Both corpora are used interchangeably for the shared SFT base; results are not sensitive to which is used.

Self-distillation corpus (14,082 trajectories). For init-from-iterate and self-distill (§4.5) we roll out the R-toolverbs·KL@400 checkpoint (HPC code: E5b-stab) greedily on the full 27,639-question CWQ training split, then filter to trajectories satisfying *simultaneously*:

- strict EM$=1$ against the normalised gold answer;
- ≥ 1 productive tool call (in the I-Self sense — non-empty response and ≥ 1 returned entity in the final `<answer>`);
- valid `<answer>` envelope format.

Yield: 14,082 trajectories (50.9% of the 27,639). This corpus is LLM-generated but LLM-*self*-generated — no external teacher.

Methodological positioning. Most 2025–26 agentic-RL papers use LLM teachers (DeepSeek-R1, Qwen-72B, GPT-4) for SFT data generation, introducing teacher-model bias confounds (Qian

Paper name	HPC code	verl yaml	One-line description
R-binary	E1	`e1_outcome.yaml`	Outcome-only (0.5 EM+0.5 F1); Search-R1 / DeepSeek-R1-zero family
R-binary-SR	E1′	`e1prime_searchr1.yaml`	R-binary with Search-R1's exact hyperparameters (kl_coef=0.001)
R-stepwise	E3	`e3_verifiable.yaml`	Equal-weight $r_{\text{out}}+r_{\text{valid}}+r_{\text{path}}+r_{\text{coh}}$
R-toolverbs	E5b	`e5b_tooltype.yaml`	$r_{\text{tool-type}}$ bonus (4 verbs + non-empty responses)
R-toolverbs·KL	E5b+KL	`e5b_kl05.yaml`	R-toolverbs with KL coef raised 5× (0.05→0.25); a.k.a. E5b-stab
R-selfV	E5b+SelfV	`e5b_selfv.yaml`	Self-verifiable retrieval reward (r_{retrv}); the peak-then-collapse variant
self-distill	G2	`g2_selfdistill_rlfromsft.yaml`	Rule-SFT prefix → self-distilled SFT (14,082 traj.) → R-toolverbs GRPO
self-distill (init-from-base)	G3	`g3_selfdistill_nopred.yaml`	No rule-SFT prefix; otherwise identical to self-distill
self-distill (init-from-iterate)	G1	`g1_init_from_e5bkl400.yaml`	Init from R-toolverbs·KL@400; no self-distillation SFT

Table 7: Paper reward names ↔ HPC training-code identifiers ↔ verl YAML configurations. The body uses paper names exclusively; HPC operators / replicators should consult HPC-code and yaml columns. Cut variants (E2 / E4 / E5a) documented separately in App. B.6.

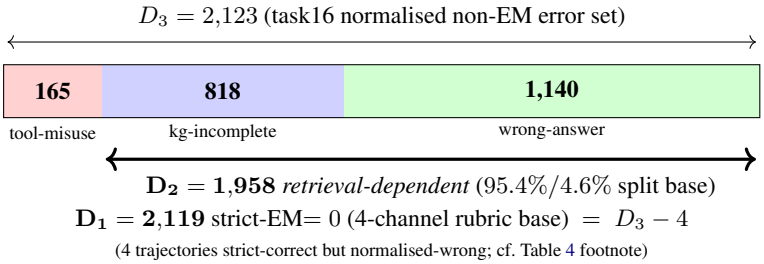

Figure 6: **Three error denominators, set-decomposition view.** The task16-normalised non-EM error set D_3=2,123 partitions into *tool-misuse* 165 + *kg-incomplete* 818 + *wrong-answer* 1,140. The retrieval-vs-extraction split (95.4%/4.6%, §5.3) is computed over D_2=1,958=D_3−165 (*kg-incomplete* ∪ *wrong-answer*). The 4-channel rubric (Table 9 row 1) is computed over D_1=2,119 strict-EM=0 errors, which differs from D_3 by the 4 strict-correct/normalised-wrong trajectories. The three bases are not additive.

Variant	EM (%)		CvT (%)	
	Cat A (n=975)	Cat B (n=2556)	Cat A	Cat B
R-stepwise @ 500	69.64	18.27	0.00	0.04
R-toolverbs·KL @ 400	73.13	25.08	5.85	2.97
self-distill @ 500	74.36	26.88	10.05	4.19
R-toolverbs·KL-14B @ 400	76.62	26.37	8.91	4.61

Table 8: Cat A (n=975, pass@10>0): parametric-recoverable. Cat B (n=2,556, pass@10=0): structurally hard. The R-stepwise→self-distill lift on Cat B EM (+8.61 pp) and Cat B CvT (+4.15 pp) shows the gain is on the hard partition. 14B and 7B self-distill agree within 0.5 pp on Cat B CvT (P2 capacity invariance, §5.2). The Cat B fraction is 2,556/3,531 = 72.4% of CWQ — the share of questions structurally beyond parametric memory.

et al., 2025). Our shared SFT is teacher-free; the G-variant is self-distilled from our own RL checkpoint. Both choices are deliberate.

D GPT-4o closed-model schema-unfamiliarity control

We include a closed-model GPT-4o (`gpt-4o-2024-08-06`) baseline as a *schema-unfamiliarity control*, not as a head-to-head comparator. The substantive findings are GPT-4o's **6.2**% ContEM at 1.46 tools/Q on n=500, with a 200-Q sub-eval showing ∼ 78% of trajectories self-describing as "unable to retrieve X" (verbose-correct without strict gold-form match): a frontier closed model under the same four-verb Freebase interface produces parseable answers at near-100% format-validity, but cannot productively retrieve from this schema.

Why we include this control. Schema-unfamiliarity is itself a meaningful obstacle even at frontier closed-model scale on this interface, so the ∼ 10% CvT ceiling reached by our trained models is not trivially dismissed as "the model is too small". The ∼ 78% self-reported "unable to retrieve X" rate from GPT-4o trajectories is direct evidence that *schema unfamiliarity, not reasoning capability*, is the binding constraint at frontier scale on this interface.

Why this is a control, not a baseline. GPT-4o is not RL-trained on our reward; it does not see the four-verb interface during pretraining; and we do not present any "GPT-4o vs ours" claim in the main body. The control's role is to bracket the schema-

Analysis	Method	Key finding
Failure-mode classification	Deterministic 4-channel rubric on self-distill's 2,119 strict-EM=0 error trajectories (Appendix B)	L_{sig} **55.2%** (1,170), L_{lang} **23.6%** (500), L_{comp} **7.8%** (165), L_{prior} **13.4%** (284) — silent failure dominates the surface
Decision-boundary analysis	Edit-distance bucketing of the first `<search>` call in self-distill's 818 *kg-incomplete* trajectories (full bucket table App. B.4)	**70.4%** within Levenshtein distance 1 of a gold relation; **89.1%** entity-correct, relation-wrong — a measurable schema-format property
Behavioral query diff	Compare 500 *kg-incomplete* trajectories shared between self-distill and R-toolverbs·KL	**90.4%** same-entity-different-relation — the two checkpoints disagree on relation but agree on entity

Table 9: Three error decompositions characterising self-distill's failure distribution. **Three denominators (kept distinct, not additive).** (i) The task16 *normalised* error set is $2{,}123$ = *tool-misuse* $165+$ *kg-incomplete* $818+$ *wrong-answer* $1{,}140$ (Table 4). (ii) The retrieval-vs-answer-extraction split (§5.3; $95.4\%/4.6\%$) is over the $\mathbf{1{,}958}$ *retrieval-dependent* errors $= 2{,}123 - 165$ *tool-misuse* (the *kg-incomplete* $\cup$ *wrong-answer* categories, $818+1{,}140$). (iii) The 4-channel rubric (row 1) is over the $\mathbf{2{,}119}$ *strict-EM=0* errors ($= 3{,}531 - 1{,}412$ correct at EM$= 40.0\%$; differs from the $2{,}123$ normalised set by the 4 strict-correct/normalised-wrong trajectories noted in Table 4). The 95.4% split (base $1{,}958$) and the $L_\bullet$ shares (base $2{,}119$) are over different denominators and are not additive. The schema-format surface property is real (rows 2–3) but oracle-injected gold relations only lift EM by $+0.20$ pp (§5.3); measurable but not the EM-gating mechanism.

Variant (step)	dev %	test %	full %	Δ pp
R-binary (1250)	0.25	0.18	0.20	−0.02
R-stepwise (500)	32.00	32.59	32.46	+0.13
R-toolverbs (100)	31.50	32.41	32.20	+0.21
R-toolverbs·KL (400)	36.88	38.78	38.35	+0.43
R-selfV (250, peak)	39.62	39.47	39.51	−0.03
R-selfV (300, collapsed)	0.00	0.00	0.00	+0.00
init-from-iterate (500)	39.00	39.55	39.42	+0.13
self-distill (500)	39.87	40.02	39.99	+0.03
R-binary-SR (300)	0.00	0.00	0.00	+0.00
R-toolverbs·KL-14B (400)	39.87	40.35	40.24	+0.11

Table 10: **Held-out dev/test re-aggregation.** Hop-stratified split (seed 42): dev n=800, held-out test n=2,731, full n=3,531. Δ = test−full EM (pp). All ten checkpoints fall within the Wilson 95% half-width at n=2,731 ($\pm$1.85 pp); the R-selfV dev peak is step 250, matching the body. Re-selecting checkpoints on dev reproduces every reported checkpoint.

pass@k	with tools %	without tools %	tool-lift pp
@1	38.41	26.41	+12.0
@4	44.17	31.97	+12.2
@8	46.28	34.32	+12.0
@16	**48.00**	**36.50**	**+11.5**

Table 11: **Self-distill pass@k tool-dependence** on a 200-question seed-42 subset (16 samples/Q, temperature 0.7), loaded from merged self-distill@500. The with-vs-without-tools gap is positive at every k ($+11.5$ pp at pass@16). Subset size and temperature differ from Figure 2, so magnitudes are not directly comparable to its bars.

unfamiliarity factor in our framework's L_{lang} channel: a model with no exposure to Freebase MIDs and dotted relations during training cannot productively use these tools, regardless of reasoning ability.

E Framework illustration and per-claim caveats

E.1 Setup: standard multi-hop KGQA re-cast as a minimal tool interface

E.2 Tool-family comparison: why the KG case is uniquely degraded

E.3 Four-channel pathway diagram

E.4 Per-claim caveats for P1–P4

P1–P4 are joint consistency checks of the descriptive framework; the per-channel decomposition in §5.3 (oracle ablation, behavioural diff, decision-boundary analysis on $1{,}958$ self-distill errors) is the primary evidence. The per-claim scope statements:

- **P1 scope.** R-binary/R-binary-SR format collapse evidences a joint L_{sig}+L_{prior} degradation typical of outcome-only KG RL.
- **P2 scope.** 7B vs. 14B is a two-point falsification

Model	n	EM	ContEM	F1	Tools/Q
GPT-4o (5-shot ReAct)	500	*0.6*†	**6.2**	2.1	1.46
R-stepwise @500 (ours)	3,531	**32.5**	45.3	38.1	1.00
self-distill @500 (ours)	3,531	**40.0**	49.5	42.8	3.00

Table 12: Closed-model schema-unfamiliarity control. All numbers are percentages except n and Tools/Q. The substantive GPT-4o findings are the bolded ContEM 6.2% and the low Tools/Q 1.46 (vs. self-distill's 3.00). †Strict EM reflects entity-casing / gold-form normalisation drift on closed-model output; ContEM is the comparable metric for schema-unfamiliarity assessment.

Tool type	Failure mode	Recovery signal
Python	`ZeroDivisionError`	traceback (file:line)
Web search	noisy but readable docs	extract anyway
Function API	`400` + named param	fix parameter
Browser/UI	page state visible	DOM/visual
KG tool	wrong entity *or* relation → `[]`	**none**

Table 13: Failure mode and recovery signal across tool families. KG tools are uniquely degraded: a single empty-list response cannot distinguish wrong entity from wrong relation from both — the L_{sig} silent-failure that the four-channel framework formalises.

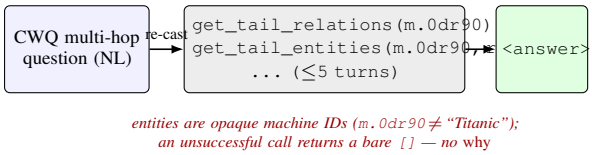

Figure 7: **Setup: a standard multi-hop KGQA benchmark re-cast as a minimal tool-use interface.** CWQ questions over Freebase are reached not by SPARQL or gold-subgraph injection but by composing four navigation verbs; entities are opaque machine IDs and an unsuccessful call returns a bare `[]` — the $L_{\text{sig}}/L_{\text{lang}}$ degradations our framework formalises (§5).

test rather than a multi-point scaling regression; a clean L_{lang}-isolation control would require a non-KG synthetic-schema task with matched reward density and is left to follow-up work.

- **P3 scope.** The init-from-iterate variant's SFT is computed against a model already exposed to R-toolverbs·KL trajectories, so it absorbs less new signal per SFT step; the init-from-base ablation (§4.5) isolates init source by removing the rule-based SFT pre-pass, with init-from-base ≈ self-distill ≈ 40% EM establishing *sufficiency* of init-from-base.
- **P4 scope.** The KGQAGen-10k cross-benchmark Spearman ρ=0.976 is evaluated on *base* Qwen variants, testing parametric-capability transfer of the benchmark rather than retrieval-bottleneck transfer of our trained policies.

F Cross-family replication: Llama-3.1-8B-Base

We replicate the self-distill protocol on Llama-3.1-8B-Base, following Search-R1's choice for the same multi-turn verl-based setup (Jin et al., 2025); the 8B-Instruct variant's tool-use limitations are documented in Meta's model card. Pipeline, hyperparameters, and SFT corpus are identical to self-distill (App. C); only the base model changes.

Stage	EM	Tools/Q	Tool-involving
SFT-only	0.284	0.65	44.4%
GRPO step 50	0.327	**0.00**	0.1%
GRPO step 100	**0.345**	**0.00**	0.0%

Table 14: Llama-3.1-8B-Base on full CWQ test (n=3531). "Tool-involving" = share of trajectories with ≥ 1 tool call under the 7-category classifier (App. B). Wilson 95% CIs ±1.6 pp at n=3531.

Cross-architecture interpretation. Llama-Base GRPO reaches comparable EM (Table 14: 0.345 at step 100; vs Qwen self-distill 0.400 at step 500) but exhibits a qualitatively distinct collapse: Tools/Q drops to 0 by step 50 and remains there; SFT-only's 44.4% tool-involving share collapses to $\sim$ 0% under GRPO. The policy bifurcates into pure parametric-memory operation (32.5% correct-no-tool + 67.5% wrong-no-tool), distinct from Qwen's tool-retaining peak-then-collapse (§4.3). The four-channel framework (§5) predicts model-family-specific channel weights: Llama-Base appears to escape L_{sig} silent failures by abandoning the tool entirely, whereas Qwen continues to query through them up to and including the peak-then-collapse.

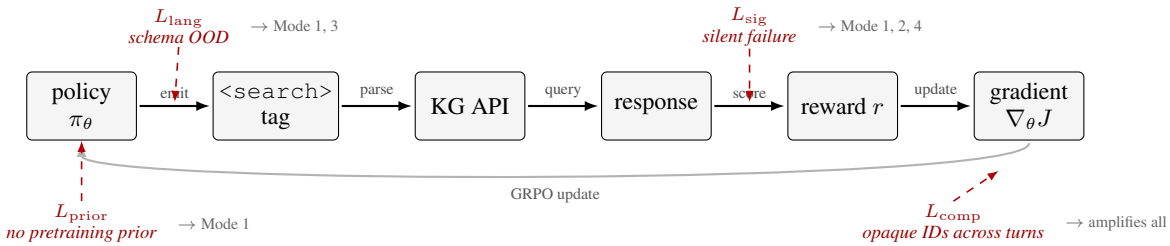

Figure 8: **Four interface-level channels along the policy → tool → reward pathway.** Each red dashed arrow marks where a structural property of the KG interface attenuates RL gradient signal: L_{sig} (silent failure on `[]` responses), L_{lang} (out-of-distribution Freebase MIDs and dotted relations), L_{comp} (opaque entity IDs threading across multi-turn calls), L_{prior} (no pretraining inductive bias for schema navigation). Annotations point to the failure modes (§4.3) each channel drives.

G Trajectory exemplars

Mode 2 (R-stepwise): *tool-free shortcut*

Q: "What religion does Ovadia Yosef lead?"
Gold: Judaism

<think> Ovadia Yosef is a religious leader; let me check.</think>

```
<search>get_tail_relations(
  Ovadia_Yosef)</search>
<resp>people.person.religion,
  ...</resp>
```

(response not consumed)

<answer>Judaism</answer> ✓ *from memory*

---

1 tool call · entity in answer: **no**

self-distill: *grounded multi-call*

Q: "What religion does Ovadia Yosef lead?"
Gold: Judaism

```
<search>get_tail_relations(
  Ovadia_Yosef)</search>
<resp>people.person.religion,
  ...</resp>
<search>get_tail_entities(
  Ovadia_Yosef,

people.person.religion)</search>
<resp>Judaism</resp>
```

<answer>Judaism</answer> ✓ *retrieved*

2–3 tool calls · entity in answer: **yes**

Figure 9: **Two contrasting behavioural patterns from real trajectory dumps.** (Left) Mode 2 *tool-free shortcut* from R-stepwise: the model issues a single `get_tail_relations` ritual call and answers from parametric memory (sample WebQTest-1384 from `e3_step500` trajectory dump). (Right) self-distill healthy multi-call grounded answer on the same WebQTest-1384 question: the model issues two grounded calls and the answer comes from the retrieved KG content (`39g2_step500`). The Mode 4 *quote-and-stop* attractor (copy the first returned entity verbatim and stop) is characterised quantitatively in §4.3 and Figure 1.