

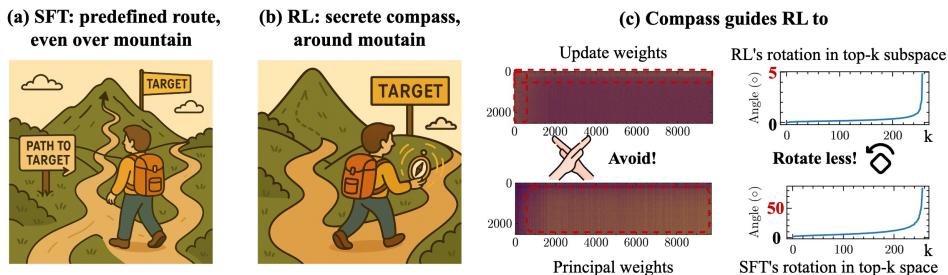
000 001 002 003 004 005 006 007 008 009 THE PATH NOT TAKEN: RLVR PROVABLY LEARNS OFF THE PRINCIPALS

010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
Anonymous authors

Paper under double-blind review

ABSTRACT

010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
Reinforcement learning (RL) reliably improves LLM reasoning while appearing to change only a small fraction of parameters. We revisit this paradox and argue that the visible sparsity is not the phenomenon itself but the trace of a *optimization bias*, where RLVR stubbornly commits updates to preferred regions that remain invariant across datasets and RL variants, *as if guided by an implicit compass*. We propose a Three-Gate Theory to formalize this mechanism, where the Anchor Gate I shows RL induces a one-step policy-KL leash that keeps updates proximal to the base policy; This constrained update is then steered by Gate II (Model Geometry) towards lower-curvature, spectra-preserving directions, a data-invariant feature; and finally, it is filtered by Gate III (Precision), where the bfloat16 format acts as a lens that amplifies the bias by hiding micro-updates, making the underlying pattern visible as apparent sparsity. Empirically, we validate this theory with a comprehensive suite of experiments. We show that RL preserves the model’s spectral structure and avoids its principal weights, in sharp contrast to SFT, which alters spectra and mainly targets those weights. Causal interventions confirm that this bias is destroyed when the model’s geometry is disrupted, proving that the geometry is the steering core of the “compass”. By providing the first parameter-level account of RLVR’s training dynamics: **RLVR learns off-principal directions in weight space**, our work not only demystifies its optimization bias but also provides a new perspective of understanding RLVR. **Crucially, we show that RL operates in a distinct optimization regime from SFT, directly adapting SFT-era parameter-efficient fine-tuning (PEFT) methods can be flawed, as evidenced by our case studies on advanced sparse fine-tuning and LoRA variants**, motivating the design of efficient **geometry-aware, RLVR-native learning** algorithms.



042
043
044
045
046
Figure 1: SFT and RL update in different manners. (a) SFT follows a predefined (guided externally) route, even over the mountain, to reach the target; (b) RLVR, without an explicit guide, behaves as if steered by an *implicit compass (optimization bias)*, taking a detour around the mountain, with two (c) Evidences: RLVR avoids principal weights during updates (left) and rotates less in the top- k subspace with preserved spectra (right).

1 INTRODUCTION

047
048
049
050
051
052
053
Large reasoning models (LRMs), such as OpenAI-o3 (Jaech et al., 2024) and DeepSeek-R1 (Guo et al., 2025), have advanced the ability of large language models to solve complex mathematical and programming tasks. A key driver is large-scale *Reinforcement Learning with Verifiable Rewards* (RLVR), which uses simple, easy-to-verify rewards to incentivize complex, multi-step reasoning. Yet, despite these advances, the mechanisms by which RL shapes model representations and behavior remain poorly understood. Given the substantial computational resources devoted to RL (relative

054 to SFT) and the emergence of striking new behaviors, it is natural to assume that realizing these be-
 055 haviors requires substantial parameter changes. However, recent evidence points in the opposite
 056 direction: RL induces *sparse* parameter updates, whereas SFT yields *dense* ones (Mukherjee et al.,
 057 2025). This counterintuitive observation reveals a paradox: *a high-cost, high-gain process that relies*
 058 *on surprisingly minimal modifications.*

060 **Our thesis.** We resolve this paradox by uncovering a deeper mechanism behind the apparent spar-
 061 **RLVR has a stubborn persistent optimization bias.** It consistently routes visible weight up-
 062 *dates into a narrow, reproducible subset of parameters, a pattern that remains strikingly invariant*
 063 *across diverse algorithms and datasets.* This inherent selectivity is then amplified by the precision
 064 limit of bfloat16, which produces the apparent sparsity as *the symptom of a persistent optimization*
 065 *bias.* We refer to the organizing principle behind this optimization bias as an *implicit RL's compass*.
 066 As illustrated in Fig. 1, while SFT is pulled toward an explicit external target, RLVR, despite having
 067 no such teacher, is secretly guided.

068 These observations prompt two central questions:

069 **Where does this optimization bias originate, and which parameters does it preferentially update?**

070 In this paper, we present a **Three-Gate Theory** that formalizes this mechanism. We show that
 071 an RL update is first constrained by Gate I (Anchoring), where an on-policy KL leash keeps the
 072 policy proximal to its base. This update is then steered by Gate II (Geometry), with the intuition
 073 that the unlike a random initialized model, the pretrained model's structured landscape routes the
 074 change away from high-curvature principal subspaces under the implicit KL leash. *This geometry*
 075 *gate serves as a key bridge to understand why the optimization bias is data- and algorithm-invariant.*
 076 Finally, we show that the Gate III: precision of bfloat16 storage acts as a realization filter, amplifying
 077 the optimization bias by hiding minimal micro updates and leading to the apparent sparsity.

078 We validate this theory with a comprehensive suite of experiments, especially confirming its pre-
 079 *diction on which parameters it preferentially updates.* We show that RL (1) preserves the model's
 080 spectral structure, in sharp contrast to SFT; (2) avoids updating the model's principal weights (the
 081 core pathways identified by the rank-k reconstruction of the weights, defined as a key driver for
 082 parameter-efficient SFT Liu et al. (2025c); and (3) we establish causality via a geometry interven-
 083 tion, showing that disrupting the model's geometry with orthogonal rotations destroys the optimiza-
 084 tion bias, confirming that the geometry is the steering compass. Finally, to validate the hypothesis,
 085 we construct an "safe mask", without any training, that can closely recover the training dynamics of
 086 a dense RLVR model in terms of KL divergence wrt. base model Shenfeld et al. (2025), showing
 087 less optimization intervention, while training on the "principal weights" yields the worst trend.

088 Our main contributions are:

- 089 • **Observation.(Sec. 2).** For the first time, we identify a *persistent, data/algorithm-invariant op-*
 090 *timization bias* in RLVR fine-tuning, an *implicit optimization compass* that shapes the training
 091 *behaviors.*
- 092 • **Theory (Sec. 3).** We propose a Three-Gate Theory (Anchor, Geometry, Precision) that provides
 093 a mechanistic account of this bias, showing how an RL update is jointly constrained, steered, and
 094 filtered.
- 095 • **Evidence (Sec. 4).** We provide strong empirical validation, consistently contrasting RL with SFT.
 096 Our evidence includes near-invariant layer spectra, sub-random overlap with principal weights,
 097 and causal interventions that confirm the role of model geometry in guiding the optimization.
- 098 • **Insight (Sec. 5).** We show that SFT-era sparse and low-rank priors (e.g., principal-targeted vari-
 099 *ants) are misaligned with RLVR's off-principal dynamics, motivating geometry-aware, RLVR-*
 100 *native learning algorithms.*

102 Our results, to our knowledge, provide the first systematic link between RL training dynamics and
 103 weight-space changes, complementing concurrent analyses that remain at the abstract level (policy
 104 distribution, output KL loss) (Wu et al., 2025; Shenfeld et al., 2025). This parameter-level account
 105 also explains why RL preserves pretrained capabilities more faithfully than SFT (Wang et al., 2024;
 106 Chu et al., 2025) in large reasoning models from a fresh perspective. Furthermore, it provides
 107 a critical insight for developing efficient RL algorithms: parameter-efficient fine-tuning (PEFT)
 108 recipes must be rethought for RL, as SFT and RL show disjoint training dynamics (See Sec. 4).

108

109
110
111
Table 1: **Update sparsity in SFT vs. RLVR.** Higher sparsity_{bf16} indicates more weights unchanged. RLVR
is consistently much sparser than SFT. \dagger *Mixed* denotes a diverse data source combining math, coding, STEM,
logic puzzles, and instruction-following Liu et al. (2025a).

Base Model	Finetuned (FT) Model	Algorithm	Data	sparsity _{bf16}
Qwen-1.5B	DS-R1-Distill-Qwen-1.5B	SFT	Mixed	2.8%
DS-R1-Distill-Qwen-1.5B	DeepScaleR-1.5B-Preview	GRPO	Math	53.8%
DS-R1-Distill-Qwen-1.5B	DeepCoder-1.5B-Preview	GRPO	Code	45.5%
DS-R1-Distill-Qwen-1.5B	Archer-Code-1.5B	GRPO	Code	52.5%
DS-R1-Distill-Qwen-1.5B	NV-ProRL	GRPO	Mixed \dagger	38.4%
DS-R1-Distill-Qwen-1.5B	NV-ProRL-v2	Reinforcement++	Mixed \dagger	36.3%
Qwen3-8B-Base	Klear-Reasoner-8B-SFT	SFT	Math+Code	0.6%
Klear-Reasoner-8B-SFT	Klear-Reasoner-8B	GRPO	Math+Code	69.5%
Qwen3-8B-Base	GT-Qwen3-8B-Base	GRPO	Math	79.9%
Qwen3-8B-Base	OURS	DAPO	Math	79.7%
Qwen3-14B-Base	UniReason-Qwen3-14B-think-SFT	SFT	Math	18.8%
Qwen3-14B-Base	UniReason-Qwen3-14B-RL	GRPO	Math	68.3%
Qwen3-4B	Polaris-4B-Preview	DAPO	Math	79.3%
DS-R1-Distill-Qwen-7B	Polaris-7B-Preview	DAPO	Math	61.7%
Qwen3-30B-A3B	UloRL-A3B	GRPO	Math	91.7%

124
125
2 A STUBBORN OPTIMIZATION BIAS IN RLVR126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
We revisit the observation: *RL induces sparse parameter updates*, but move beyond quantification by analyzing where RL localizes these changes. Our analysis uncovers a deep “**optimization bias**” phenomenon, which we demonstrate *RL exhibits a stubborn, structured optimization bias: it consistently routes visible changes to specific regions of the network*. The observed sparsity is a readout of this bias, amplified by `bfloat16`, rather than intrinsically sparse gradients.130
131
132
133
134
135
Model suite. We analyze publicly released checkpoints, as shown in Tab. 1. The suite spans multiple RLVR variants (e.g., GRPO, DAPO, Reinforcement++), diverse data domains (math, coding, instruction), and several model families and types (dense and Mixture-of-Experts). We place particular emphasis on DeepSeek-R1-Distill-Qwen-1.5B (DS-Qwen-1.5B), for which a long-horizon RL checkpoint is available (Liu et al., 2025a). This model serves as a robust case study given its extensive training for over 3,000 steps on a diverse data mixture.136
137
2.1 A ROBUST, `BFLOAT16`-AWARE ANALYSIS OF UPDATE SPARSITY138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
A `bfloat16`-aware probe for unchanged weights. `bfloat16` (bf16) is standard in modern RL frameworks like verl (Sheng et al., 2024), to improve throughput without compromising performance. However, analyzing parameter changes under bf16 requires a careful probe. Its unique numerical format, with only 7 mantissa bits for precision, means that the smallest representable difference between two numbers scales with their magnitude. Consequently, a fixed absolute-tolerance check as used in (Mukherjee et al., 2025), is *unreliable*, which can over- or under-report (see Appendix D.1).156
157
158
159
160
161
To ensure a rigorous report, we adopt a numerically robust, `bfloat16`-aware probe to define the update sparsity sparsity_{bf16} as the fraction of parameters that remain unchanged.156
157
158
159
160
161
Definition 2.1 (Unchanged Weight in bf16). *Let $w_i, \widehat{w}_i \in \mathbb{R}$ be scalars stored in bf16 (finite, nonzero). We say w_i is unchanged with respect to \widehat{w}_i iff*

156
157
158
159
160
161
$$|\widehat{w}_i - w_i| \leq \eta \max(|w_i|, |\widehat{w}_i|), \quad \eta = 10^{-3}. \quad (1)$$

156
157
158
159
160
161
Choosing $\eta = 10^{-3} < 2^{-9}$ makes equation 1 equivalent to bitwise equality (See Appendix D.2.).156
157
158
159
160
161
Definition 2.2 (`bfloat16`-aware Update Sparsity). *Write $x \approx_{\eta}^{\text{bf16}} y$ for Def. 2.1. Define the bf16 change count $\|\theta^1 - \theta^0\|_{0,\eta}^{\text{bf16}} := |\{i : \theta_i^1 \not\approx_{\eta}^{\text{bf16}} \theta_i^0\}|$ and the corresponding sparsity*

156
157
158
159
160
161
$$\text{sparsity}_{\text{bf16}}(\theta^0, \theta^1; \eta) := 1 - \|\theta^1 - \theta^0\|_{0,\eta}^{\text{bf16}} / n. \quad (2)$$

156
157
158
159
160
161
where n is the total number of parameters. Values near 1 indicate few stored changes, while values near 0 indicate dense apparent change.156
157
158
159
160
161
RLVR update sparsity results. As shown in Tab. 1, our analysis confirms that RL yields substantially higher update sparsity than SFT. Across models, SFT sparsity is consistently low (typically 0.6%–18.8%), whereas RL sparsity is an order of magnitude higher, ranging from 36% to 92%. However, absolute levels on recent checkpoints are lower than earlier reports (Mukherjee et al., 2025), underscoring the need for bf16-aware probes and re-evaluation on current models.

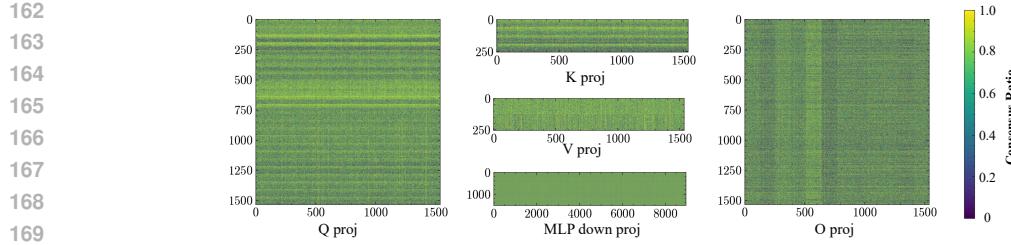


Figure 2: **Consensus ratio of weight updates** across five RLVR runs on the 13th layer’s projection (Q/K/V/O) and the MLP down projection (zoom in for structures). Lighter bands indicate coordinates updated in most runs, revealing a stable, stripe-like routing pattern rather than random scatter.

2.2 RLVR EXHIBITS A SURPRISING UPDATE BIAS

Magnitude alone does not reveal *where* changes occur, impeding the deep analysis on *how* sparse changes arise. We therefore examine the *updated subnetwork*. We use 5 independent RLVR checkpoints from the same DS-Qwen-1.5B in Tab. 1, trained on diverse data and different RLVR algorithms. For each layer ℓ and run r , we first form the bf16-aware *changed* mask $M_{\ell}^{(r)} := \mathbf{1}[W_{\ell}^{(r)} \neq_{bf16} W_{\ell}^0]$ (Def.2.2) against the base weights W_{ℓ}^0 .

Stability across runs. We first analyze their spatial agreement using *Jaccard Overlap*. For runs r, s , let $A = \{(i, j) : M_{\ell,ij}^{(r)} = 1\}$ and $B = \{(i, j) : M_{\ell,ij}^{(s)} = 1\}$. We report the mean off-diagonal of the pairwise Jaccard matrix $J(A, B) = \frac{|A \cap B|}{|A \cup B|}$ and compare it to the independent Bernoulli baseline $\mathbb{E}[J] = \frac{pq}{p+q-pq}$. As summarized in Tab. 2, Jaccard

is consistently high across runs, confirming a shared footprint when trained from the same base model, with Jaccard matrix shown in Fig. 8.

Consensus ratio (where updates land). Stability alone does not indicate *where* updates land. We therefore visualize and analyze the consensus ratio $C_{\ell,ij} = \frac{1}{R} \sum_{r=1}^R M_{\ell,ij}^{(r)}$, the fraction of runs realizing a *weight update* at coordinate (i, j) . Values near 1 indicate that *all* runs consistently change that weight; values near 0 indicate that none do. As shown in Fig. 2, consensus maps reveal contiguous row/column bands, stripe-like, localized routing rather than scattered noise. Especially, there are obvious *row-wise stripes* in Q/K/V projections and *column-wise stripes* in O projections. This exposes a clear **optimization bias**: RLVR consistently concentrates updates in specific regions of the parameter matrices, even though the five runs use disjoint data and RL variants.

Temporal stability (how the bias emerges). To examine *within-run* dynamics, we track the row-wise ratio $\rho_{\ell,i}(t) = \frac{1}{n_{\ell}} \sum_j M_{\ell,ij}(t)$ and column-wise ratio $\kappa_{\ell,j}(t) = \frac{1}{m_{\ell}} \sum_i M_{\ell,ij}(t)$ across checkpoints at t steps. On DS-Qwen-1.5B (training setting in Appendix C.1), the *relative* profiles $\rho_{\ell,\cdot}(t)$ and $\kappa_{\ell,\cdot}(t)$ remain aligned while overall density grows as shown in Fig. 3: peaks and troughs persist. The routing bias *emerges early* and is *reinforced* over training, indicating a temporally stable phenomenon rather than a transient artifact. Moreover, the peak is consistent with the bias structure shown in Fig. 2. We also show their remaining column-wise (Q) and row-wise (O) update ratio dynamics in Fig. 10, without a clear trend, *indicating the bias is indeed structured, not random*.

Other model families (whether only on Qwen). We observe similar stripe-structured footprints on Llama and Mistral (Fig. 9 in Appendix), suggesting the routing bias is generic to RLVR. We further examine the update consensus on the Llama model across RLVR variants in Appendix G.1.

Table 2: Cross-run stability for 13th block.

Layer	Jaccard Overlap	Random Baseline
Q	0.580	0.430
K	0.580	0.413
V	0.597	0.467
O	0.552	0.373
MLP-down	0.585	0.453
MLP-up	0.578	0.443
MLP-gate	0.575	0.437

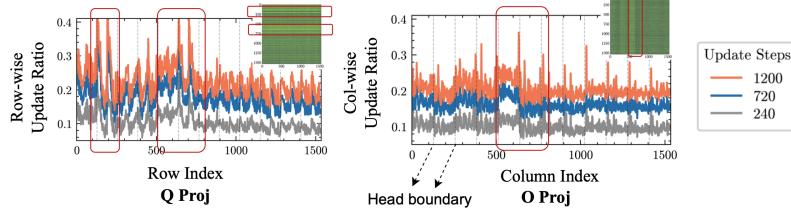


Figure 3: **Temporal emergence of the optimization bias** with row and column-wise update ratios for the 13th attention block across gradient update steps ($t \in \{240, 720, 1200\}$), smoothed with a 3-step window. The row-dominant (Q) and column-dominant (O) patterns are consistent with the bias structures in Fig. 2. We visualize the head boundaries with grey dashed lines. The bias appears not only across heads but also within heads.

216 2.3 SPARSITY IS A SYMPTOM, NOT THE PHENOMENON
217

218 The stable footprint of *where updates land*, persisting both throughout training and in the final
219 model, suggests the focus should move from *sparsity* itself to the *underlying optimization bias*.
220

221 We find that sparsity is actually the *readout* of this optimization bias, whose visibility is amplified by
222 the precision limits of `bf16` storage. Because `bf16` has a limited mantissa, changes smaller than
223 the unit-in-the-last-place (ULP) threshold (Lemma D.2) are not representable. *Therefore, if RLVR*
224 *consistently routes sub-ULP updates toward a particular subset of parameters, the stored values*
225 *will not change, and the result appears as sparsity.*

226 We test this hypothesis by increasing the learning rate to scale otherwise sub-ULP updates above the
227 representable threshold. As predicted, the apparent update sparsity largely disappears. This directly
228 challenges the interpretation of (Mukherjee et al., 2025) that sparsity stems from zero gradients.
229 Instead, our results point to sparsity as a byproduct of an optimization bias interacting with finite
230 precision. Consistent with this view, concurrent work observes that sparsity mostly vanishes under
231 `fp32` storage (Shenfeld et al., 2025), even though task performance does not improve.

232 **Remark on precision.** One natural confusion is treated the `bf16` as the final cause, while it is
233 important to note that in verl, optimizer states and gradient reductions/accumulation are maintained
234 in `float32`¹. So the sparsity cannot show up unless the RL process is consistently biased toward
235 where to assign visible changes throughout the training.

236 *Aha Finding! — RLVR exhibits a patterned, rather than random, optimization bias toward
237 where the visible changes land. The apparent sparsity is a direct readout of this underlying
238 bias, an effect amplified by `bf16`’s precision.*

239 3 A MECHANISTIC THEORY OF THE RL OPTIMIZATION BIAS
240

241 In the post-training era, reinforcement learning (RL) has been the most compute-intensive yet pow-
242 erful stage xAI (2025). Paradoxically, as shown in Sec. 2, these large gains arise not through broad
243 updates, but through selective, patterned edits that reveal a persistent optimization bias. Understanding
244 this bias is essential to demystify how RL achieves its improvements with two central questions:
245

246 *Where does this optimization bias originate, and what does RL preferentially optimize?*

247 We answer this question with a **Three-Gate Theory**. First, on-policy RL introduces a KL constraint
248 that **anchors** the fine-tuned policy nearby. Second, the pretrained model **geometry** steers updates
249 towards specific regions, which is finally visualized through the *lens of precision*.

250 **Notations.** We consider a large language model with parameters θ , defining a conditional distribu-
251 tion $\pi_\theta(y | x)$ over possible output token sequences $y = (y_1, \dots, y_T) \in \mathcal{Y}$ given a prompt $x \in \mathcal{X}$
252 from the space \mathcal{X} . Each sequence y is composed of tokens from a vocabulary \mathcal{V} of size N .

253 **RLVR objective.** Various RLVR algorithms including PPO, GRPO, DAPO, and REINFORCE++,
254 learn a policy π_θ by optimizing variants by optimizing variants of a KL-regularized objective:
255

$$256 \max_{\theta} \mathbb{E}_{y \sim \pi_\theta(\cdot | x), x \sim \mathcal{X}} [R(x, y) - \beta \text{KL}(\pi_\theta(\cdot | x) \| \pi_{\text{ref}}(\cdot | x))]. \quad (3)$$

257 where π_{ref} is a fixed reference policy and $\beta \geq 0$ controls the KL regularization ($\beta = 0$ recovers
258 the clip-only variants such as DAPO). Rewards $R(x, y)$ are *verifiable* and (after normalization)
259 *bounded* (e.g., pass/fail or execution scores). Moreover, the surrogate typically uses the token-wise
importance ratio $w_t = \frac{\pi_\theta(y_t | x, y_{\leq t})}{\pi_{\text{old}}(y_t | x, y_{\leq t})}$ with clipping relative to π_{old} .

260 3.1 GATE I: ANCHORING VIA AN ON-POLICY KL LEASH
261

262 We first show that online policy gradient updates yield a per-step *policy* KL bound (an **anchoring**
263 effect), which in turn limits parameter movement during the RLVR update.

264 **One-step surrogate.** With equation 3, a standard sequence-level online policy-gradient surrogate
265 is

$$266 \mathcal{L}_{\text{PG}}(\theta) = -\mathbb{E}_{x \sim \mathcal{X}, y \sim \pi_\theta(\cdot | x)} [A^\perp(x, y) \log \pi_\theta(y | x)], \quad (4)$$

267 where A^\perp is a (normalized) advantage estimate, optionally *shaped* by a reference-KL log-ratio term.
268 In practice, updates are performed over mini-batches, with a collected batch of data, not in a fully
269 on-policy manner. But the resulting error after a small step size $\Delta\theta$ is $O(\|\Delta\theta\|^2)$ (Lemma E.1).

¹verl mixed-precision settings with `{reduce_type, buffer_dtype}=float32`.

270 **Implicit KL leash.** The KL leash emerges as policy gradient methods can be understood as a
 271 conservative projection, keeping new policy close to its starting point while reweighting it toward
 272 higher-reward outcomes, not pulling it toward a potentially distant external distribution like SFT:
 273

274 **Proposition 3.1** (One-step policy-KL leash). *Let $q(\cdot | x)$ be a full-support reference and let $\tilde{q}_\beta(\cdot |$
 275 $x) \propto q(\cdot | x) \exp(R/\beta)$ denote the soft-regularized improvement oracle. Let θ^+ be the parametric
 276 fit obtained by the M -projection of \tilde{q}_β onto the policy class, $\theta^+ \in \arg \min_\theta D_{\text{KL}}(\tilde{q}_\beta \| \pi_\theta)$. Then, for
 277 a sufficiently small one-step update,*

$$278 D_{\text{KL}}(\pi_{\theta^+} \| \pi_\theta) \leq (1 + o(1)) D_{\text{KL}}(\tilde{q}_\beta \| \pi_\theta), \quad (5)$$

279 where the $o(1)$ term vanishes as $D_{\text{KL}}(\tilde{q}_\beta \| \pi_\theta) \rightarrow 0$.
 280

281 Notably, even when the explicit KL term is removed (e.g., in DAPO with $\beta = 0$), the ratio clip-
 282 ping trick still imposes a KL bound $O(\varepsilon^2)$ in the small-step regime (Appendix E.2.4), confirmed
 283 empirically with a bounded KL divergence change during a DAPO run (Fig. 11).

284 **Weight update constraint.** Now we show the KL leash puts a constraint on weight update ΔW

285 **Proposition 3.2** (Policy-KL leash \Rightarrow weight bound). *Assume $\log \pi_\theta$ is C^3 and let $F(\theta)$ denote the
 286 Fisher information. If a one-step update $\theta^+ = \theta + \Delta$ satisfies $D_{\text{KL}}(\pi_{\theta^+} \| \pi_\theta) \leq K$ and, on the update
 287 subspace, $F(\theta) \succeq \mu I$ for some $\mu > 0$, then for K sufficiently small*

$$288 \|\Delta\|_{F(\theta)} \triangleq \sqrt{\Delta^\top F(\theta) \Delta} \leq \sqrt{2K} (1 + o(1)), \quad \|\Delta\|_2 \leq \sqrt{\frac{2K}{\mu}} (1 + o(1)). \quad (6)$$

290 Consequently, for any weight matrix block $W \in \theta$, $\|\Delta W\|_F \leq \sqrt{2K/\mu} (1 + o(1))$.
 291

292 See a detailed proof for Proposition 3.1 in Appendix E.2.1 and Proposition 3.2 in Appendix E.2.2.

293 **Take-away 1: RL update imposes an implicit KL leash (anchor effect), ensuring that the per-
 294 step drift from the current policy is small.** This aligns with recent work arguing that even the
 295 final policy is KL-proximal Wu et al. (2025); Shenfeld et al. (2025). Our focus, however, is to
 296 understand how this leash affects the weight change dynamics.
 297

298 3.2 GATE II: PRETRAINED GEOMETRY DETERMINES Where A KL-BOUNDED STEP GOES

299 **From Gate I to location.** Gate I supplies a one-step KL leash, but it does not explain *where* the
 300 step lands. We propose Gate II: the Model Geometry Gate, where we argue that unlike a randomly
 301 initialized network, a well-pretrained model possesses a highly structured geometry, e.g., spectrum
 302 statics, high-curvature directions related to reasoning performance, acts as a “**key compass**” that
 303 determines where the update is favored to be applied.

304 **Layerwise norm bound from the KL leash.** Let W_0 be a pretrained linear block, $W_+ = W_0 + \Delta W$
 305 the post-step block, and let $F_W \succeq \mu_W I$ be a per-layer curvature proxy. If the per-layer KL budget
 306 satisfies $\frac{1}{2} \langle \text{vec } \Delta W, F_W \text{ vec } \Delta W \rangle \leq \delta_W$, then (Appendix E.10)

$$308 \|\Delta W\|_F \leq \sqrt{\frac{2\delta_W}{\mu_W}}, \quad \|\Delta W\|_2 \leq \sqrt{\frac{2\delta_W}{\mu_W}}. \quad (7)$$

310 We then show this conservative update yields three consequences making them preserve pretrained
 311 weight spectrum instead of destroying them based on weight perturbation theory Stewart (1998).

312 **Limited subspace rotation.** First, as shown in Theorem 3.3, the angle between the original and
 313 updated subspaces is quadratically bounded, meaning the fundamental directions are preserved.

314 **Theorem 3.3** (Constrained subspace rotation with Wedin’s sin– Θ theorem Wedin (1972).). *For any
 315 k with $\gamma_k > 0$,*

$$317 \max(\|\sin \Theta(U_k(W_0), U_k(W_+))\|_2, \|\sin \Theta(V_k(W_0), V_k(W_+))\|_2) \leq \frac{\|\Delta W\|_2}{\gamma_k} \leq \frac{\sqrt{2\delta_W/\mu_W}}{\gamma_k}. \quad (8)$$

319 **Singular value stability.** Second, the magnitudes of the principal components themselves are pre-
 320 served. The change in each singular value is bounded by the norm of the update.

321 **Corollary 3.4** (Singular-value stability). *For each k ,*

$$323 |\sigma_k(W_+) - \sigma_k(W_0)| \leq \|\Delta W\|_2 \leq \sqrt{\frac{2\delta_W}{\mu_W}}, \quad \sum_i (\sigma_i(W_+) - \sigma_i(W_0))^2 \leq \|\Delta W\|_F^2 \leq \frac{2\delta_W}{\mu_W}. \quad (9)$$

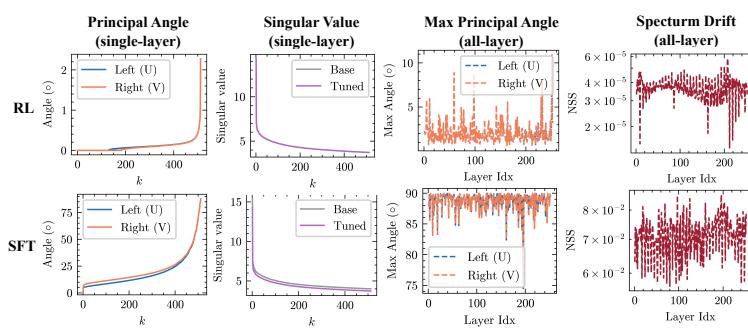


Figure 4: The spectrum probe results on the RL and SFT version on the Qwen3-8B Su et al. (2025) with the full top- k principal angles and singular value curve on one exemplar layer. RLVR exhibits a surprisingly stable top- k spectrum with minimal subspace rotation and changes in top- k eigenvalues. More visualization in Appendix F.2.

Top-k energy preservation. Finally, these effects combine to ensure the cumulative energy of the top- k components of the weights remains stable.

Corollary 3.5 (Top- k energy and Ky Fan norms). *Let $\|\cdot\|_{(k)} := \sum_{i=1}^k \sigma_i(\cdot)$ be the Ky Fan k -norm. Then*

$$|\|W_+\|_{(k)} - \|W_0\|_{(k)}| \leq \sum_{i=1}^k |\sigma_i(W_+) - \sigma_i(W_0)| \leq k \|\Delta W\|_2 \leq k \sqrt{\frac{2\delta_W}{\mu_W}}. \quad (10)$$

See a detailed proof in Appendix E.3.

Take-away 2: Under the KL leash, RL updates tend to preserve the model’s original weight structure rather than destroy it. This naturally favors updates in low-curvature directions of the optimization landscape, which avoids dramatic changes in model behavior. Since directly quantifying curvature in LRM with long CoTs is computationally prohibitive, we instead adopt a powerful and efficient proxy, principal weights Liu et al. (2025c), as detailed in Sec. 4.2.

3.3 GATE III: PRECISION AS A LENS THAT REVEALS THE COMPASS

Building on the optimization bias, the bfloat16 with limited precision acts as a *lens*: it hides those micro-updates that occur where the RL consistently holds a weak willingness to apply large changes.

Corollary 3.6 (Magnitude-dependent realization threshold). *A stored weight W_{ij} changes at a step iff $|\Delta W_{ij}| \gtrsim \frac{1}{2} \text{ULP}_{\text{bf16}}(W_{ij})$.*

The effect of this gate has been discussed aforementioned. We would emphasize again that precision is more an *amplifier* for visible sparsity, not the *cause* of optimization bias, as optimizer states, etc., are still in float32 (See 2.3).

4 EMPIRICAL VALIDATION OF THE OPTIMIZATION COMPASS

We now present theory-driven empirical evidence to validate our ‘‘Optimization Compass’’, the core claim of Gate II: that the model’s geometry is the ‘‘compass’’ that steers the KL-constrained updates. The following experiments confirm our theory’s predictions about where the optimization bias originates and what RL preferentially optimizes.

4.1 RLVR PRESERVES SPECTRAL STRUCTURE, WHILE SFT DOES NOT

First, we probe spectral changes to directly test the prediction of Gate II (Geometry): the model’s geometry steers the update, causing RLVR to preserve the underlying structure rather than destroy it. We examine checkpoints trained with a standard SFT-then-RLVR pipeline: one from Qwen3-8B-Base (Su et al., 2025) and another long-horizon RL run on DS-Qwen-1.5B (Liu et al., 2025a), both from industry with SOTA performance. Besides, we also analyze a setting where SFT and RL are applied separately to the Qwen3-14B-Base model, delivering comparable in-domain math performance (Huan et al., 2025). We compare the base weights W_0 with the finetuned weights W_+ :

- **Subspace rotation.** For the top- k left (U)/right(V) singular subspaces, we check the rotation using **principal angles** via $\cos \theta_i(U) := \sigma_i(U_{0,k}^\top U_{+,k})$ and $\cos \theta_i(V) := \sigma_i(V_{0,k}^\top V_{+,k})$.

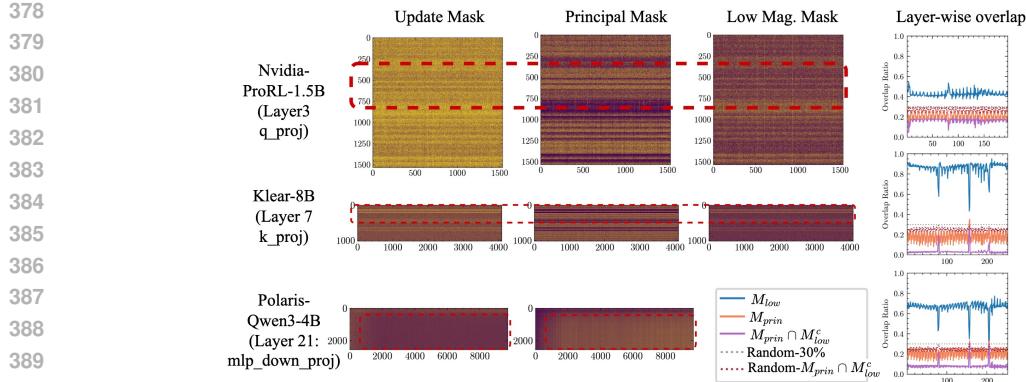


Figure 5: **RL avoids updating principal weights.** We compare the RL update mask with principal weight mask M_{princ} , low magnitude mask M_{low} , and the one $M_{princ} \cap M_{low}^c$. The layer-wise overlap between RL updates and principal weights is consistently *sub-random*, an effect more pronounced when removing its overlapped weights with M_{low} , i.e., $M_{princ} \cap M_{low}^c$.

• **Spectrum drift.** Beyond showing the singular value curve, we quantify singular-value change with a normalized ℓ_2 shift: $\text{NSS}(W) = \|\sigma(W_+) - \sigma(W_0)\|_2 / \|\sigma(W_0)\|_2$

Our findings: RLVR checkpoints show a surprisingly stable spectrum within their top principal components. As shown in Fig. 4, RLVR consistently exhibits low subspace rotation and low spectrum drift. In sharp contrast, SFT induces significantly higher rotation and drift across the same metrics. This provides the first direct evidence for our "compass" theory.

4.2 RLVR AVOIDS PRINCIPAL WEIGHTS WHILE SFT TARGETS THEM

Next, we examine which specific parameters the optimization compass targets or avoids beyond a macro-level spectral check.

Principal weights as a proxy for high-curvature directions. Directly identifying high-curvature directions is computationally prohibitive, especially given LRM with long CoTs. Instead, we adopt a powerful proxy from recent work Liu et al. (2025c), **principal weights**, which is defined as *the weights with the largest magnitude after low-rank approximation*, representing its most influential computational pathways. The validity of this proxy is confirmed by their perturbation studies, which show that modifying these specific weights causes sharp *reasoning performance degradation*. This degradation is directly linked to high-curvature regions via a Taylor expansion of the loss. The *principal mask*, $M_{princ}^{(k)} = \text{Top}_\alpha(s_{ij}^{(k)})$, is defined as the top- α fraction of weights with the highest score, $s_{ij}^{(k)} = |W_0^{(k)}(i, j)|$, where $W_0^{(k)}$ is the rank- k SVD reconstruction of W_0 .

Low-magnitude weights as low-resistance pathway. We further include the top- α lowest magnitude weights, as $M_{low} = \text{Bottom}_\alpha(|W_0|)$. The magnitude is also a bias from the model geometry (distribution prior), impacting how easily the weights can be updated based on our precision gate.

Metrics. Let M be the weight update *update mask* from an RLVR run. We report the overlap ratio between our identified mask M_\bullet with it, defined as $\text{Overlap}(M_\bullet, M) = \frac{|M_\bullet \cap M|}{|M|}$, with a random guess baseline overlap ratio as the density of M_\bullet itself, i.e., α .

Our findings. Fig. 5 visualizes the RL update mask M in relation to the principal mask M_{princ} and the low-magnitude mask M_{low} , reporting their layer-wise overlap against a random baseline as well.

The results show a clear dichotomy. RL updates exhibit a sub-random overlap with principal weights, indicating a strong tendency to avoid them. Conversely, the updates show a super-random overlap with low-magnitude weights due to their low-resistance to micro-updates. Besides, we found that the residual overlap between updates and principal weights is highly accounted for by weights that are both principal (defined by the rank- k approximation of W_0) and low-magnitude (original W_0). After excluding this intersection, i.e., $M_{princ} \cap M_{low}^c$, the overlap drops significantly.

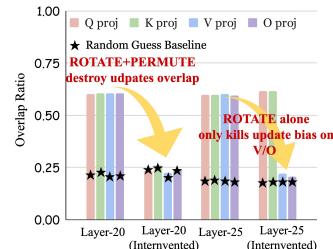


Figure 6: **Overlap ratio after intervention.**

432 **Remark.** This leads to a crucial and counter-intuitive finding: ***RLVR and SFT are driven by updates***
 433 ***to fundamentally different regions of a model***, though performance could be similar.
 434

435 **4.3 RLVR RELIES ON MODEL GEOMETRY, DISRUPTING GEOMETRY DESTROYS THE BIAS**
 436 Gate II posits that the pretrained model’s geometry steers RL updates. To test this causal link,
 437 we deliberately ”scramble” the geometry of specific layers in a Qwen3-4B-Base model using
 438 orthogonal rotations and head permutations (details in Appendix C.3) and compare the overlap
 439 $\text{Overlap}(M_\bullet, M) = \frac{|M_\bullet \cap M|}{|M|}$. between the base run with another independent run without inter-
 440 vention and one run with intervention.

441 **Findings.** We modify (i) layer 20 with ROTATE+PERMUTE, and (ii) layer 25 with ROTATE. As
 442 shown in Fig. 6, the update overlap collapsed to a random level in the intervened layers, while
 443 remaining high in all untouched layers. This provides strong causal evidence that the pretrained
 444 model’s geometry is the source of the optimization bias.

445 **4.4 RLVR SIGNATURES PERSIST IN AGENTIC TASKS, RLHF, ROBOTICS**

446 **Setup.** We analyze additional *agent*, *RLHF* (RL with human feedback), and *embodied* checkpoints
 447 and apply the same weight–space diagnostics as in Sec. 4.1 and Sec. 4.2: (i) principal-subspace
 448 rotation, (ii) spectral drift, and (iii) update–principal misalignment. The extended model suite
 449 is summarized in Tab. 5. **(i)Agents.** We evaluate policies from AGENTFLOW (Li et al., 2025b)
 450 and VERL-AGENT (Feng et al., 2025) on multi-turn and long-horizon tasks. We also assess tool-
 451 augmented agents from SKYRL (Cao et al., 2025) and VERL-TOOL (Jiang et al., 2025) on *Web-
 452 Search*, *DeepSearch*, and *SWE*. **(ii)RLHF.** We include preference-optimized models trained with
 453 DPO (Rafailov et al., 2023a) and SimPO (Meng et al., 2024c), primarily targeting instruction follow-
 454 ing. **(iii)Embodied AI.** We include embodied AI models with both language and vision backbones
 455 from SimpleVLA-RL (Li et al., 2025a) and Embodied R1 (Yuan et al., 2025).

456 **Our Findings.** **(i) Stable spectra, minimal rotation.** Across models, top- k subspaces rotate only
 457 slightly, and layer spectra remain near-identical to the base model (Fig. 15; Fig. 16; Fig. 17), matching
 458 the spectrum-preserving, off-principal regime observed earlier. **(ii) Off-principal updates.** Up-
 459 date masks in checkpoints consistently *avoid principal weights*: the most active bands are spatially
 460 misaligned with the principal mask (Fig. 18). **Takeaway.** RLVR’s optimization dynamics, *mini-
 461 mal rotation, spectrum preservation, off-principal routing*, persist beyond verifiable math/code to
 462 *agents, embodied AI RLHF*, indicating a common, model-conditioned optimization bias within a
 463 KL-anchored RL post-training game, consistent with our Three-Gate Theory.

464 **5 THEORY-GUIDED RETHINKING OF LEARNING ALGORITHMS FOR RL**

465 A good theory should not only explain a phenomenon, beautifully validated by observations, but
 466 also provide actionable insights. Our account shows that RLVR and SFT follow disjoint optimiza-
 467 tion dynamics in parameter space, which implies that many SFT-era PEFT methods, especially those
 468 aligned with principal directions through sparse or low-rank priors, transfer poorly to RLVR. This
 469 section *validates* our predictions and *demonstrates* how they guide the redesign of learning algo-
 470 rithms for RL.

472 **5.1 PROBING SPARSE FINE-TUNING IN RL**

473 Rather than judge success by final task accuracy, which is noisy and can reward “lucky” runs in RL,
 474 we instead track the token-wise **forward KL** drift $\text{KL}(\pi \parallel \pi_{\text{ref}})$ throughout training. This allows
 475 us to assess how closely a sparse run follows the dense baseline trajectory. Intuitively, if removing
 476 weights impedes the training, the policy cannot effectively shift away from the base policy.

477 **Masks.** We evaluate by performing RLVR on the DS-Qwen-1.5B with only the following weights,
 478 identified one-shot before training. (i) $U = M_{\text{princ}}$ (principal-only, sparsity 50%), (ii) $U = M_{\text{princ}}^c$
 479 (non-principal-only, sparsity 50%), (iii) $U = M_{\text{low}}$ (lowest weights with the threshold as the mean
 480 of magnitude), (iv) $U = M_{\text{low}} \cup M_{\text{princ}}^c$ (favor non-principal and low-magnitude), and (v) a random
 481 mask with the same layer-wise sparsity as (iv). We choose 50% for (i) as we want to isolate the
 482 effect of the number of parameters for a fair comparison to see the difference between (i) and (ii).

483 **Results (KL in Fig. 7 and accuracy in Tab. 3)** The **union mask** $M_{\text{low}} \cup M_{\text{princ}}^c$ tracks the dense
 484 run’s KL curve most closely and outperforms its random baseline, showing our theory indeed distin-
 485 guishes those *highly touchable weights* with a similar trend. **principal-only** is the worst with much
 slower increasing training KL loss and much lower accuracy, showing a clear training block.

486 **Takeaways and limits.** (1) RLVR’s effective updates concentrate away from *principal* directions
 487 (the ones SFT tends to favor), consistent with our
 488 theory. This reveals a fundamental contrast be-
 489 tween SFT and RL: SFT directly targets principal
 490 weights Liu et al. (2025c), whereas RL actively
 491 avoids them. This finding demonstrates that tra-
 492 ditional sparse fine-tuning algorithms designed for
 493 SFT may be a poor match for RL and motivates
 494 the need for a new class of RL-specific methods.
 495 (2) Freezing principal and large-magnitude weights
 496 while updating non-principal and low-magnitude
 497 ones can *approximate the dense KL trajectory* with
 498 *competitive* accuracy, demonstrating the guidance ef-
 499 fect of our theory, especially since the mask is pre-
 500 dicted without any additional training. (3) These
 501 are *one-shot* masks without schedule/retuning, with
 502 residual accuracy gaps expected. Future work com-
 503 bining our theory with dynamic mask refresh or
 504 PEFT schedules will be promising next steps (Zhao
 505 et al., 2024; Zhu et al., 2024; Liu et al., 2025c).

505 506 5.2 REVISITING LoRA THROUGH THE LENS OF OUR THEORY

507 A recent report (Schulman & Lab, 2025) finds that low-rank LoRA, even rank-1, can match full-
 508 parameter RL performance. Our theory offers an explanation: in full-parameter RL, effective up-
 509 dates lie *off* the principal directions and induce only small spectral changes. Low-rank adapters
 510 can approximate these off-principal updates, while freezing the base weights regularizes training
 511 and discourages moves toward principal directions. With an appropriately scaled learning rate, the
 512 limited adapter capacity is therefore sufficient to catch up to full-parameter performance at least in
 513 the short run.

514 **However**, the same report suggests principal-targeted variants such as **PiSSA** (Meng et al., 2024a)
 515 should yield further gains. Our geometry account disagrees: aligning updates to top- r principal
 516 directions enforces SFT-style behavior that is *misaligned* with RLVR’s off-principal bias.

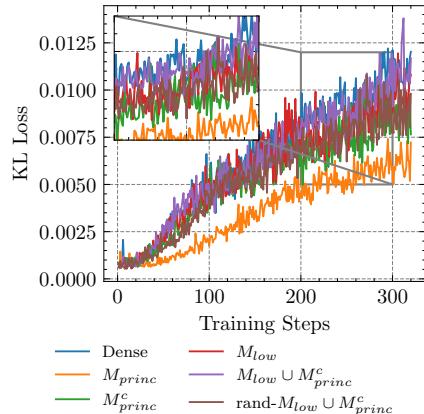
517 **Empirical test.** On **DS-Qwen-1.5B** with **DeepMath-103K** (He et al., 2025), we sweep ranks
 518 $\{8, 32, 64\}$ and learning rates $\{1 \times 10^{-4}, 5 \times 10^{-5}, 1 \times 10^{-5}\}$ for 200 steps, and report pass@1 (mean
 519 over 16 samples) on **AIME24** and **AMC23** (Fig. 19). To control for model effects, we repeat on
 520 **Llama-3.2-3B-Instruct** with a Math corpus and report pass@1 (mean over 4) on **MATH500**
 (Fig. 20).

521 **Our findings.** Across settings, the principal-targeted **PiSSA** provides no clear gain over LoRA. At
 522 the higher learning rates used for low-rank adapters to match full-parameter performance, **PiSSA**
 523 *often becomes unstable and collapses* earlier than LoRA. This occurs because scaling the learning
 524 rate in **PiSSA** enforces updates along principal directions, higher-curvature and spectrum-distorting,
 525 precisely the directions RLVR tends to avoid. The result is brittle optimization and early collapse,
 526 whereas LoRA’s off-principal updates remain better aligned with RLVR’s geometry.

527 **Insight.** These results support the geometry-based account: principal-aligned LoRA variants are
 528 *over-fit to SFT’s update geometry* and misaligned with RL’s training dynamics, so success in SFT
 529 does not transfer to RL.

530 531 6 CONCLUSION

532 In this work, we resolve the paradox of sparse but effective reinforcement learning (RL) updates
 533 by identifying a persistent, geometry-aligned optimization bias, an “implicit compass” that steers
 534 training. We propose a Three-Gate Theory that provides a mechanistic account for this phenomenon,
 535 showing how on-policy constraints, pretrained model geometry, and bfloat16 precision interact to
 536 guide updates. Our experiments, including causal interventions, confirm that this compass steers RL
 537 to preserve the model’s spectral structure by avoiding the principal weights targeted by Supervised
 538 Fine-Tuning (SFT). This parameter-level account not only demystifies its optimization bias but also
 539 *charts a path toward a white-box understanding of RLVR and the design of geometry-aware, RLVR-native* learning algorithms, rather than repurposed SFT-era heuristics.



505 506 507 **Figure 7: KL loss curve on**
DS-Qwen-1.5B under different masks. A
 508 Better visualization.

540 REFERENCES
541

542 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-
543 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical
544 report. *arXiv preprint arXiv:2303.08774*, 2023.

545 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
546 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are
547 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

548

549 Shiyi Cao, Sumanth Hegde, Dacheng Li, Tyler Griggs, Shu Liu, Eric Tang, Jiayi Pan, Xingyao
550 Wang, Akshay Malik, Kourosh Hakhamaneshi, Richard Liaw, Philipp Moritz, Matei Zaharia,
551 Joseph E. Gonzalez, and Ion Stoica. Skyrl-v0: Train real-world long-horizon agents via rein-
552 forcement learning. <https://novasky-ai.notion.site/skyrl-v0>, 2025. NovaSky
553 AI, Notion page. Accessed 2025-11-10.

554 Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V
555 Le, Sergey Levine, and Yi Ma. Sft memorizes, rl generalizes: A comparative study of foundation
556 model post-training. *arXiv preprint arXiv:2501.17161*, 2025.

557

558 Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li,
559 Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned lan-
560 guage models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.

561 DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning,
562 2025. URL <https://arxiv.org/abs/2501.12948>.

563

564 Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah Smith.
565 Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping.
566 *arXiv preprint arXiv:2002.06305*, 2020.

567 Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. Group-in-group policy optimization for llm
568 agent training. *arXiv preprint arXiv:2505.10978*, 2025.

569

570 Daya Guo, Dejian Yang, Haowei Zhang, and Junxiao Song. Deepseek-r1: Incentivizing reasoning
571 capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.07570*, 2025. URL
572 <https://arxiv.org/abs/2501.07570>.

573

574 Seungwook Han, Jyothish Pari, Samuel J Gershman, and Pulkit Agrawal. General reasoning requires
575 learning to reason from the get-go. *arXiv preprint arXiv:2502.19402*, 2025.

576

577 Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian
578 Yu, Zhenwen Liang, Wenxuan Wang, et al. Deepmath-103k: A large-scale, challenging, de-
579 contaminated, and verifiable mathematical dataset for advancing reasoning. *arXiv preprint
arXiv:2504.11456*, 2025.

580

581 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
582 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv
preprint arXiv:2103.03874*, 2021.

583

584 Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification.
585 *arXiv preprint arXiv:1801.06146*, 2018.

586

587 Yafei Hu, Quanting Xie, Vidhi Jain, Jonathan Francis, Jay Patrikar, Nikhil Keetha, Seungchan Kim,
588 Yaqi Xie, Tianyi Zhang, Hao-Shu Fang, et al. Toward general-purpose robots via foundation
589 models: A survey and meta-analysis. *arXiv preprint arXiv:2312.08782*, 2023.

590

591 Maggie Huan, Yuetai Li, Tuney Zheng, Xiaoyu Xu, Seungone Kim, Minxin Du, Radha Pooven-
592 dran, Graham Neubig, and Xiang Yue. Does math reasoning improve general llm capabilities?
593 understanding transferability of llm reasoning. *arXiv preprint arXiv:2507.00432*, 2025.

Aaron Jaech et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024.

594 Dongfu Jiang, Yi Lu, Zhuofeng Li, Zhiheng Lyu, Ping Nie, Haozhe Wang, Alex Su, Hui Chen, Kai
 595 Zou, Chao Du, et al. Verltool: Towards holistic agentic reinforcement learning with tool use.
 596 *arXiv preprint arXiv:2509.01055*, 2025.

597

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

602 Chunyuan Li, Zhe Gan, Zhengyuan Yang, Jianwei Yang, Linjie Li, Lijuan Wang, Jianfeng Gao, et al.
 603 Multimodal foundation models: From specialists to general-purpose assistants. *Foundations and
 604 Trends® in Computer Graphics and Vision*, 16(1-2):1–214, 2024.

605

606 Haozhan Li, Yuxin Zuo, Jiale Yu, Yuhao Zhang, Zhaoxiang Yang, Kaiyan Zhang, Xuekai Zhu, Yuchen
 607 Zhang, Tianxing Chen, Ganqu Cui, et al. Simplelevla-rl: Scaling vla training via reinforcement
 608 learning. *arXiv preprint arXiv:2509.09674*, 2025a.

609 Zhuofeng Li, Haoxiang Zhang, Seungju Han, Sheng Liu, Jianwen Xie, Yu Zhang, Yejin Choi, James
 610 Zou, and Pan Lu. In-the-flow agentic system optimization for effective planning and tool use.
 611 *arXiv preprint arXiv:2510.05592*, 2025b.

612

613 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan
 614 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. *arXiv preprint
 615 arXiv:2305.20050*, 2023.

616

617 Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero:
 618 Benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information
 619 Processing Systems*, 36:44776–44791, 2023.

620

621 Mingjie Liu, Shizhe Diao, Ximing Lu, Jian Hu, Xin Dong, Yejin Choi, Jan Kautz, and Yi Dong.
 622 Prorl: Prolonged reinforcement learning expands reasoning boundaries in large language models.
 623 *arXiv preprint arXiv:2505.24864*, 2025a.

624

625 Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee,
 626 and Min Lin. Understanding r1-zero-like training: A critical perspective. *arXiv preprint
 627 arXiv:2503.20783*, 2025b.

628

629 Zihang Liu, Tianyu Pang, Oleg Balabanov, Chaoqun Yang, Tianjin Huang, Lu Yin, Yaoqing Yang,
 630 and Shiwei Liu. Lift the veil for the truth: Principal weights emerge after rank reduction for
 631 reasoning-focused supervised fine-tuning. *arXiv preprint arXiv:2506.00772*, 2025c.

632

633 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint
 634 arXiv:1711.05101*, 2017.

635

636 Michael Luo, Sijun Tan, Roy Huang, Xiaoxiang Shi, Rachel Xin, Colin Cai, Ameen Patel, Alpay
 637 Ariyak, Qingyang Wu, Ce Zhang, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepcoder: A
 638 fully open-source 14b coder at o3-mini level. <https://pretty-radio-b75.notion.site/DeepCoder-A-Fully-Open-Source-14B-Coder-at-O3-mini-Level-1cf81902c14680b3bee5eb349a512a51>, 2025a.
 639 Notion Blog.

640

641 Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai,
 642 Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepscaler: Surpassing o1-preview
 643 with a 1.5b model by scaling rl, 2025b. Notion Blog.

644

645 MAA. American mathematics contest 12 (amc 12), November 2023. URL https://artofproblemsolving.com/wiki/index.php/AMC_12_Problems_and_Solutions.

646

647 MAA. American invitational mathematics examination (aime), February 2024. URL
 648 https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions.

648 MAA. American invitational mathematics examination (aime), February 2025. URL
 649 https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions.

650

651 Fanxu Meng, Zhaohui Wang, and Muhan Zhang. Pissa: Principal singular values and singular
 652 vectors adaptation of large language models. *Advances in Neural Information Processing Systems*,
 653 37:121038–121072, 2024a.

654

655 Yu Meng, Mengzhou Xia, and Danqi Chen. Simpo: Simple preference optimization with a
 656 reference-free reward. *Advances in Neural Information Processing Systems*, 37:124198–124235,
 657 2024b.

658

659 Yu Meng, Mengzhou Xia, and Danqi Chen. Simpo: Simple preference optimization with a
 660 reference-free reward. *Advances in Neural Information Processing Systems*, 37:124198–124235,
 661 2024c.

662

663 Sagnik Mukherjee, Lifan Yuan, Dilek Hakkani-Tur, and Hao Peng. Reinforcement learning fine-
 664 tunes small subnetworks in large language models. *Advances in Neural Information Processing
 Systems*, 2025.

665

666 Long Ouyang et al. Training language models to follow instructions with human feedback. In
 667 *Advances in Neural Information Processing Systems*, volume 35, pp. 27730–27744, 2022.

668

669 Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language under-
 670 standing by generative pre-training. *arXiv preprint arXiv:2303.08774*, 2018.

671

672 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 673 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 674 models from natural language supervision. In *International conference on machine learning*, pp.
 675 8748–8763. PMLR, 2021.

676

677 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 678 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances
 in neural information processing systems*, 36:53728–53741, 2023a.

679

680 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 681 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances
 in neural information processing systems*, 36:53728–53741, 2023b.

682

683 Negin Raoof, Etrash Kumar Guha, Ryan Marten, Jean Mercat, Eric Frankel, Sedrick Keh, Hritik
 684 Bansal, Georgios Smyrnis, Marianna Nezhurina, Trung Vu, Zayne Rea Sprague, Mike A Mer-
 685 rill, Liangyu Chen, Caroline Choi, Zaid Khan, Sachin Grover, Benjamin Feuer, Ashima Su-
 686 varna, Shiye Su, Wanja Zhao, Kartik Sharma, Charlie Cheng-Jie Ji, Kushal Arora, Jeffrey Li,
 687 Aaron Gokaslan, Sarah M Pratt, Niklas Muenninghoff, Jon Saad-Falcon, John Yang, Asad Aali,
 688 Shreyas Pimpalgaonkar, Alon Albala, Achal Dave, Hadi Pouransari, Greg Durrett, Sewoong Oh,
 689 Tatsunori Hashimoto, Vaishaal Shankar, Yejin Choi, Mohit Bansal, Chinmay Hegde, Reinhard
 690 Heckel, Jenia Jitsev, Maheswaran Sathiamoorthy, Alex Dimakis, and Ludwig Schmidt. Auto-
 691 matic evals for llms, 2025. URL <https://github.com/mlfoundations/evalchemy>.

692

693 Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and struc-
 694 tured prediction to no-regret online learning. In *Proceedings of the fourteenth international con-
 695 ference on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference
 696 Proceedings, 2011.

697

698 John Schulman and Thinking Machines Lab. Lora without regret. *Thinking Machines Lab: Con-
 699 nectionism*, 2025. doi: 10.64434/tml.20250929. <https://thinkingmachines.ai/blog/lora/>.

700

701 Zhihong Shao et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language
 702 models. *arXiv preprint arXiv:2402.03300*, 2024.

703

704 Idan Shenfeld, Jyothish Pari, and Pulkit Agrawal. RL's razor: Why online reinforcement learning
 705 forgets less. *arXiv preprint arXiv:2509.04259*, 2025.

702 Guangming Sheng et al. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint*
 703 *arXiv:2409.19256*, 2024.

704 Gilbert W Stewart. Perturbation theory for the singular value decomposition. 1998.

706 Zhenpeng Su, Leiyu Pan, Xue Bai, Dening Liu, Guanting Dong, Jiaming Huang, Wenping Hu, and
 707 Guorui Zhou. Klear-reasoner: Advancing reasoning capability via gradient-preserving clipping
 708 policy optimization. *arXiv preprint arXiv:2508.07629*, 2025.

709 Qwen Team. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.

710 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
 711 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
 712 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

713 Xinyi Wang, Antonis Antoniades, Yanai Elazar, Alfonso Amayuelas, Alon Albalak, Kexun Zhang,
 714 and William Yang Wang. Generalization vs memorization: Tracing language models' capabilities
 715 back to pretraining data. *arXiv preprint arXiv:2407.14985*, 2024.

716 Per-Åke Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972.

717 Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du,
 718 Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint*
 719 *arXiv:2109.01652*, 2021.

720 Fang Wu, Weihao Xuan, Ximing Lu, Zaid Harchaoui, and Yejin Choi. The invisible leash: Why rlvr
 721 may not escape its origin. *arXiv preprint arXiv:2507.14843*, 2025.

722 xAI. Grok: Ai assistant, 2025. URL <https://x.ai/grok>. Accessed: 2025-09-24, continuously
 723 updated.

724 Wei Xiong, Jiarui Yao, Yuhui Xu, Bo Pang, Lei Wang, Doyen Sahoo, Junnan Li, Nan Jiang, Tong
 725 Zhang, Caiming Xiong, et al. A minimalist approach to llm reasoning: from rejection sampling
 726 to reinforce. *arXiv preprint arXiv:2504.11343*, 2025.

727 An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu,
 728 Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu,
 729 Xingzhang Ren, and Zhenru Zhang. Qwen2.5-math technical report: Toward mathematical ex-
 730 pert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024.

731 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong
 732 Liu, Lingjun Liu, Xin Liu, et al. Dapo: An open-source llm reinforcement learning system at
 733 scale. *arXiv preprint arXiv:2503.14476*, 2025.

734 Yifu Yuan, Haiqin Cui, Yaoting Huang, Yibin Chen, Fei Ni, Zibin Dong, Pengyi Li, Yan Zheng,
 735 and Jianye Hao. Embodied-r1: Reinforced embodied reasoning for general robotic manipulation.
 736 *arXiv preprint arXiv:2508.13998*, 2025.

737 Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerl-
 738 zoo: Investigating and taming zero reinforcement learning for open base models in the wild. *arXiv*
 739 *preprint arXiv:2503.18892*, 2025.

740 Simon Zhai, Hao Bai, Zipeng Lin, Jiayi Pan, Peter Tong, Yifei Zhou, Alane Suhr, Saining Xie, Yann
 741 LeCun, Yi Ma, et al. Fine-tuning large vision-language models as decision-making agents via
 742 reinforcement learning. *Advances in neural information processing systems*, 37:110935–110971,
 743 2024.

744 Xiaojiang Zhang et al. Srpo: A cross-domain implementation of large-scale reinforcement learning
 745 on llm. *arXiv preprint arXiv:2504.14286*, 2025.

746 Jiawei Zhao, Zhenyu Zhang, Beidi Chen, Zhangyang Wang, Anima Anandkumar, and Yuandong
 747 Tian. Galore: Memory-efficient llm training by gradient low-rank projection. *arXiv preprint*
 748 *arXiv:2403.03507*, 2024.

756 Hanqing Zhu, Zhenyu Zhang, Wenyan Cong, Xi Liu, Sem Park, Vikas Chandra, Bo Long, David Z
 757 Pan, Zhangyang Wang, and Jinwon Lee. Apollo: Sgd-like memory, adamw-level performance.
 758 *arXiv preprint arXiv:2412.05270*, 2024.

760 Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul
 761 Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *arXiv*
 762 *preprint arXiv:1909.08593*, 2019.

764 A CLARIFICATION OF LLM USAGE

766 In this work, we employ LLMs to polish the writing throughout the paper and to assist in generating
 767 code for figure plotting. Besides, we use it for drawing the teaser figure.

769 B MORE RELATED WORKS

771 **Post-training** Large-scale models pre-trained on broad domains serve as general-purpose back-
 772 bones with extensive domain knowledge and notable zero-shot capabilities (Radford et al., 2021;
 773 Achiam et al., 2023; Touvron et al., 2023; Hu et al., 2023; Li et al., 2024; Radford et al., 2018; Brown
 774 et al., 2020). However, such pre-trained models often fail to meet the specific application require-
 775 ments or align with domain-specific constraints. *Post-training* methods address this gap by adapting
 776 foundation models to downstream tasks. Common approaches include supervised fine-tuning on
 777 curated datasets (Howard & Ruder, 2018; Dodge et al., 2020; Wei et al., 2021; Chung et al., 2024),
 778 reinforcement learning from human or automated feedback (Ziegler et al., 2019; Ouyang et al., 2022;
 779 Guo et al., 2025; Zhai et al., 2024), and other recent techniques (Rafailov et al., 2023b).
 780 Especially, the recent advances in LLM reasoning (DeepSeek-AI, 2025) highlight the effectiveness
 781 of *Reinforcement Learning with Verifiable Rewards* (RLVR), which replaces subjective human judg-
 782 ments with automatically verifiable signals. RLVR has been shown to significantly enhance reason-
 783 ing ability using policy optimization algorithms such as PPO (Ouyang et al., 2022) and GRPO (Shao
 784 et al., 2024). Building on these successes, a growing body of work (Yu et al., 2025; Liu et al., 2025b;
 785 Luo et al., 2025a; Zhang et al., 2025; Liu et al., 2025a; Xiong et al., 2025) continues to refine RL
 786 methods tailored for LLM reasoning.

787 **SFT versus RL.** Prior work comparing these paradigms has largely focused on downstream per-
 788 formance. A foundational result shows that on-policy RL can outperform offline SFT even with the
 789 same expert data (Ross et al., 2011). Recent empirical studies consistently reinforce this, finding
 790 that RL-tuned models often generalize better out-of-distribution (Han et al., 2025; Chu et al., 2025)
 791 and transfer more effectively to new tasks (Huan et al., 2025) than their SFT counterparts.

792 While these studies establish a performance hierarchy, our work investigates a different dimension:
 793 how these distinct methods affect the model’s internal structure. A recent study observed that RL
 794 fine-tunes only a fraction of the network’s parameters (Mukherjee et al., 2025), but this empirical
 795 finding left the underlying mechanism unexplored and did not characterize or predict the affected
 796 subnetwork. Our work aims to bridge this gap by providing a mechanistic explanation for this
 797 phenomenon.

798 C EXPERIMENTAL DETAILS

799 C.1 TRAINING SETTINGS

800 **Models & Datasets.** We run post-training experiments on three open models:
 801 **DeepSeek-R1-Distill-Qwen-1.5B** (Yang et al., 2024), **Qwen2.5-Math-7B** (Yang et al.,
 802 2024), and **Qwen3-Base** (Team, 2025). The maximum context length is set to 8192 for
 803 DeepSeek-R1-Distill-Qwen-1.5B and Qwen2.5-Math-7B, and to 20480 for Qwen3-Base.
 804 We evaluate primarily on mathematics using two training corpora to reduce dataset-specific
 805 confounds. (1) **DAPO+MATH (DM)**: a union of the DAPO-Math-17k set² and the MATH
 806 dataset (Hendrycks et al., 2021). (2) **DS+SR**: the 47k DeepScaler collection (Luo et al., 2025b) com-
 807 bined with high-difficulty (levels 3–5) problems extracted from SimpleRL (Zeng et al., 2025). We
 808 use the version from Huan et al. (2025).

809 ²DAPO-Math-17k

810 **Training details.** We implement RLVR on the VeRL pipeline (Sheng et al., 2024) and use
 811 vLLM (Kwon et al., 2023) for rollouts. We use FSDPv2 with the default mixed precision config-
 812 uration. All experiments run on NVIDIA H200 GPUs. Unless otherwise noted, we use DAPO (Yu
 813 et al., 2025) *without* an explicit reference-KL penalty (ratio clipping as in DAPO), a global batch
 814 size of 256 (mini-batch 64) with 4 gradient update per step.
 815 Per-model configurations without specific mention:

- 816 • **Qwen2.5-Math-7B on DM:** 16 rollouts per prompt; 8 x H200 GPUs; 300 training steps.
- 817 • **DeepSeek-R1-Distill-Qwen-1.5B on DS+SR:** 12 rollouts per prompt; 16 x H200 GPUs; 320
 818 steps.
- 819 • **Qwen3-4B-Base on DS+SR:** 16 rollouts per prompt; 32 x H200 GPUs; 150 steps.

820 We optimize the actor with AdamW (Loshchilov & Hutter, 2017) (constant learning rate 1×10^{-6} ,
 821 $\beta_1=0.9$, $\beta_2=0.999$). Rewards are *verifiable*: +1.0 if the extracted final answer is correct, otherwise
 822 -1.0 (no separate format score), following the detailed verifier implementation in Su et al. (2025).
 823 We enable an over-length penalty with an extra 1024-token budget and penalty factor 1.0.

824 C.2 EVALUATION SETTINGS

826 We evaluate models on four widely used benchmarks: AIME24 (MAA, 2024), AIME25 (MAA,
 827 2025), AMC23 (MAA, 2023), MATH-500 (Lightman et al., 2023), as we main train using math
 828 daastets. We used Eval-Chemy (Raouf et al., 2025) with their default temperature 0.7 and 0.8 as
 829 the top-p value. In our experiments, we used the **averaged accuracy**, i.e., $\text{pass}@1(\text{avg}@k)$ for
 830 all benchmarks. to evaluate the models' performance. Specifically, for AIME24 and AIME 25, we
 831 averaged accuracy on 64 samples, for AMC, we average accuracy on 32 samples, For MATH 500,
 832 our score is the average accuracy over 2 samples.

833 C.3 INTERVENTION DETAILS

834 **Intervention 1: loss-preserving V/O rotation.** Let D be the head dimension, H_q the number of
 835 query heads, H_{kv} the number of key/value heads, and $n_{\text{rep}} = H_q/H_{kv}$ (grouped GQA). Denote

$$837 \quad W_v \in \mathbb{R}^{d_{\text{model}} \times (H_{kv}D)}, \quad W_o \in \mathbb{R}^{d_{\text{model}} \times (H_qD)}.$$

838 Draw any orthogonal $R \in \mathbb{R}^{D \times D}$ (Haar/Hadamard) and form the block rotations

$$839 \quad R_{kv} = \text{diag}(\underbrace{R, \dots, R}_{H_{kv}}) \in \mathbb{R}^{(H_{kv}D) \times (H_{kv}D)}, \quad R_q = \text{diag}(\underbrace{R, \dots, R}_{n_{\text{rep}}}, \underbrace{R, \dots, R}_{n_{\text{rep}}}, \dots) \in \mathbb{R}^{(H_qD) \times (H_qD)}.$$

840 We edit the weights by right-multiplication along the head axis:

$$843 \quad W'_v = W_v R_{kv}, \quad W'_o = W_o R_q. \quad (11)$$

845 If b_v exists, reshape b_v per head and set $b'_v = b_v R_{kv}$.

846 **Proposition C.1** (Exact invariance). *Let $\text{Ctx} = \text{Attn}(Q, K, V) \in \mathbb{R}^{(\times (H_qD))}$. Under equation 11,*

$$847 \quad \text{out}' = \text{Attn}(Q, K, VR_{kv}) (W_o R_q)^\top = \text{Ctx} R_q R_q^\top W_o^\top = \text{Ctx} W_o^\top = \text{out}.$$

848 **Intervention 2: head shuffle (lossless).** Let P_{kv} be a permutation of the H_{kv} KV heads and P_q
 849 its grouped expansion to H_q heads. Apply

$$851 \quad \begin{aligned} & \text{rows of } (W_k, W_v) \leftarrow P_{kv}, & \text{rows of } W_q \leftarrow P_q, & \text{columns of } W_o \leftarrow P_q^{-1}. \\ & \text{cols of } (W_k, W_v) \leftarrow P_{kv}, & \text{cols of } W_q \leftarrow P_q, & \text{columns of } W_o \leftarrow P_q^{-1}. \end{aligned}$$

853 This relabels which head carries which subspace, while leaving the block function unchanged.

854 We show that after weight intervention, the model weights update position has a sub-random overlap
 855 while those untouched weights stay a high overlap.

856 C.4 EVALUTION SETTINGS

858 D EXAMPLES OF WHY PREVIOUS IDENTIFIED METHOD FAILS

859 D.1 FAILURES OF A FIXED ABSOLUTE TOLERANCE RULE

- 861 • **False positives at large scale.** Within $[2^{10}, 2^{11}] = [1024, 2048]$, the bf16 spacing is $\text{ULP}_{\text{bf16}} =$
 862 $2^{10-7} = 8$. Numbers like 1024.001 and 1024.002 differ by $10^{-3} > 10^{-5}$, hence would be flagged as
 863 “changed” by the 10^{-5} rule, yet both round to the same bf16 code (1024), i.e., *no storage-level
 864 change*.

864 • **False negatives at small scale.** Around $10^{-6} \approx 2^{-20}$, the bf16 spacing is $\text{ULP}_{\text{bf16}} = 2^{-27} \approx$
 865 7.45×10^{-9} . Weights $w=10^{-6}$ and $\widehat{w}=2 \times 10^{-6}$ differ by $10^{-6} \leq 10^{-5}$ and would be marked “equal”
 866 by the 10^{-5} rule, yet they are separated by ≈ 134 ULPs and quantize to *different* bf16 codes.
 867

868 **D.2 JUSTIFICATION OF OUR PROBE**

869 **Lemma D.1** (Gap between distinct bf16 representables). *If $x \neq y$ are normalized bf16 numbers in*
 870 *the same binade $[2^e, 2^{e+1}]$, then*

871
$$|x - y| \geq 2^{-7} \quad \text{and} \quad \frac{|x - y|}{\max(|x|, |y|)} > 2^{-8}.$$

 872
 873

874 *The strict inequality also holds across the binade boundary.*

875 **Lemma D.2** (ULP lens: magnitude-dependent threshold). *For normalized bf16 values x with $|x| \in$*
 876 *$[2^e, 2^{e+1}]$,*

877
$$\frac{\text{ULP}_{\text{bf16}}(x)}{|x|} \in (2^{-8}, 2^{-7}] = (0.390625\%, 0.78125\%).$$

 878

879 *Hence the minimal realized relative update at magnitude $|x|$ is $\gtrsim \frac{1}{2} \text{ULP}_{\text{bf16}}(x)/|x| \in$*
 880 *$(0.195\%, 0.391\%)$. In particular, larger $|x|$ requires a larger absolute step to register.* \square
 881

882 **Proposition D.3** (Soundness and completeness of the probe). *Let w_i, \widehat{w}_i be normalized bf16 values*
 883 *(finite, nonzero), and suppose $\eta < \frac{1}{2} \min_x \text{ULP}_{\text{bf16}}(x)/|x| = 2^{-9} \approx 1.953 \cdot 10^{-3}$. Then*

884
$$|\widehat{w}_i - w_i| \leq \eta \max(|w_i|, |\widehat{w}_i|) \iff \text{bf16}(w_i) = \text{bf16}(\widehat{w}_i).$$

 885

886 *Proof.* (\Rightarrow) If $w_i \neq \widehat{w}_i$, Lemma D.2 gives $|\widehat{w}_i - w_i|/\max(|w_i|, |\widehat{w}_i|) > 2^{-8} > 2\eta$, contradiction. Hence
 887 $w_i = \widehat{w}_i$ as bf16 numbers.

888 (\Leftarrow) If the stored bf16 values are equal, the difference is 0, which satisfies equation 1. \square \square

889 **Corollary D.4** (Choice $\eta = 10^{-3}$ is safe). *Since $10^{-3} < 2^{-9}$, Proposition D.3 applies: the test*
 890 *equation 1 passes iff the two bf16 entries are bit-wise identical (or both zero). Thus $\eta = 10^{-3}$ yields*
 891 *a scale-aware probe that flags equality only when storage is unchanged.*

892 **E MATH ANALYSIS**

893 **E.1 POLICY-GRADIENT FINE-TUNING (DAPO)**

894 Assume an *old* policy π_{old} that we use to sample G candidate completions $y^{1:G}$ for each prompt
 895 $x \in \mathcal{X}$. For a single token $y_{i,t}$ (token t in completion i) we define the *importance-weighted advantage*

896
$$w_{i,t} = \underbrace{\frac{\pi_{\theta}(y_{i,t}|x, y_{<t})}{\pi_{\text{old}}(y_{i,t}|x, y_{<t})}}_{\text{importance ratio}} \hat{A}_{i,t} \mathbb{I}_{\text{clip}} \in \mathbb{R}, \quad (1)$$

 897
 898

899 where $\hat{A}_{i,t}$ is the estimated advantage and $\mathbb{I}_{\text{clip}} \in \{0, 1\}$ implements the usual trust-region clipping.
 900

901 **Token-level objective.** The DAPO loss can be written as a sum of weighted log-probabilities

902
$$J_{\text{RL}}(\theta) = \mathbb{E}_{x \sim \mathcal{X}, y^{1:G} \sim \pi_{\text{old}}} \left[\frac{1}{\sum_i |y^i|} \sum_{i=1}^G \sum_{t=1}^{|y^i|} w_{i,t} \log \pi_{\theta}(y_{i,t} | x, y_{<t}^i) \right]. \quad (2)$$

 903
 904

905 **E.2 PROOF OF GATE I: ON-POLICY RL IMPLIES A ONE-STEP KL LEASH**

906 This appendix provides the standard tilting oracle and M -projection facts, local second-order ex-
 907 pansions, and the proof of the one-step policy-KL leash (Prop. 3.1 in the main text). *We keep the*
 908 *proof concise, otherwise too lengthy, especially for those has shown in some prior work Shenfeld*
 909 *et al. (2025); Wu et al. (2025).* Our one-step analysis is inspired by recent work Wu et al. (2025);
 910 Shenfeld et al. (2025), which uses a similar variational approach to show that even the final con-
 911 verged policy remains KL-proximal to the base policy. We also record a trust-region/clipping bound
 912 used when $\beta = 0$.

913 Throughout, x is fixed, $q(\cdot | x)$ has full support on \mathcal{Y} , and $\pi_{\theta}(\cdot | x)$ is a C^3 parametric family with
 914 log-density $\log \pi_{\theta}$ locally smooth. Expectations without explicit subscript are conditional on x .
 915 We first show useful lemmas here.

918 **Lemma E.1** (Frozen-policy surrogate is second-order tight). *Let $f(\theta) := \mathcal{L}_{\text{PG}}(\theta)$ in equation 4 and
919 $g(\theta) := \tilde{\mathcal{L}}_{\text{PG}}(\theta; \theta_t)$ be the frozen-policy surrogate with A_{θ_t} . Then $f(\theta_t) = g(\theta_t)$ and $\nabla f(\theta_t) =$
920 $\nabla g(\theta_t)$. If ∇f and ∇g are L -Lipschitz in a neighborhood of θ_t , then*

$$922 \quad |f(\theta_t + \Delta\theta) - g(\theta_t + \Delta\theta)| \leq \frac{L}{2} \|\Delta\theta\|^2.$$

923 Proof. At θ_t , both objectives evaluate to $-\mathbb{E}_{\pi_{\theta_t}}[A_{\theta_t} \log \pi_{\theta_t}]$. For the gradient, using the
924 log-derivative trick and the centering of A_{θ_t} , both yield $-\mathbb{E}_{\pi_{\theta_t}}[A_{\theta_t} \nabla \log \pi_{\theta_t}]$. Thus $f(\theta_t) = g(\theta_t)$
925 and $\nabla f(\theta_t) = \nabla g(\theta_t)$. The bound is the standard second-order Taylor remainder under Lipschitz
926 gradients. \square

928 1: Exponential tilting and M-projection

929 **Lemma E.2** (Gibbs variational principle / exponential tilting). *Fix $\beta > 0$ and a full-support reference
930 $q(\cdot | x)$. Then*

$$932 \quad \max_{\pi \ll q} \left\{ \mathbb{E}_{y \sim \pi}[R(x, y)] - \beta D_{\text{KL}}(\pi \| q) \right\}$$

933 is uniquely maximized by

$$935 \quad \tilde{q}_{\beta}(y | x) = \frac{q(y | x) \exp(R(x, y) / \beta)}{\mathbb{E}_{y \sim q}[\exp(R(x, y) / \beta)]}.$$

937 Proof. Consider $\mathcal{L}(\pi, \lambda) = \mathbb{E}_{\pi}[R] - \beta \mathbb{E}_{\pi}[\log \frac{\pi}{q}] + \lambda(\sum_y \pi(y) - 1)$. Stationarity in π gives $\log \frac{\pi}{q} =$
938 $R/\beta - \lambda - 1$, hence $\pi \propto q e^{R/\beta}$. Strict concavity in π yields uniqueness. \square

940 **Lemma E.3** (Policy Gradient Update as Parametric M-projection). *For fixed \tilde{q}_{β} ,*

$$942 \quad \arg \min_{\theta} D_{\text{KL}}(\tilde{q}_{\beta} \| \pi_{\theta}) = \arg \max_{\theta} \mathbb{E}_{y \sim \tilde{q}_{\beta}}[\log \pi_{\theta}(y | x)].$$

944 Proof. $D_{\text{KL}}(\tilde{q}_{\beta} \| \pi_{\theta}) = \mathbb{E}_{\tilde{q}_{\beta}}[\log \tilde{q}_{\beta}] - \mathbb{E}_{\tilde{q}_{\beta}}[\log \pi_{\theta}]$, where the first term is θ -independent. We omit
945 the full proof here, with one can be found in Shenfeld et al. (2025). \square

947 2: Local second-order identities

948 **Lemma E.4** (Local Pythagorean identity for the M-projection). *Let $f(\theta) := D_{\text{KL}}(\tilde{q}_{\beta} \| \pi_{\theta}) =$
949 $\mathbb{E}_{\tilde{q}_{\beta}}[-\log \pi_{\theta}] + \text{const}$. Assume $\log \pi_{\theta}$ is C^3 near θ , and let $\theta^+ \in \arg \min f$. Writing $\Delta := \theta^+ - \theta$, for
950 $\|\Delta\|$ small,*

$$952 \quad f(\theta) - f(\theta^+) = \frac{1}{2} \Delta^T H_{\tilde{q}}(\theta) \Delta + O(\|\Delta\|^3), \quad H_{\tilde{q}}(\theta) := -\mathbb{E}_{\tilde{q}_{\beta}}[\nabla^2 \log \pi_{\theta}].$$

954 Proof. Taylor-expand f at θ^+ : $f(\theta) = f(\theta^+) + \frac{1}{2} \Delta^T H_{\tilde{q}}(\theta^+) \Delta + O(\|\Delta\|^3)$ since $\nabla f(\theta^+) = 0$. Local
955 C^3 smoothness implies $H_{\tilde{q}}(\theta^+) = H_{\tilde{q}}(\theta) + O(\|\Delta\|)$, which is absorbed into the cubic remainder. \square

957 **Lemma E.5** (Quadratic expansion of policy KL). *Let $F(\theta) := -\mathbb{E}_{\pi_{\theta}}[\nabla^2 \log \pi_{\theta}]$ be the Fisher in-
958 formation. Then*

$$959 \quad D_{\text{KL}}(\pi_{\theta+\Delta} \| \pi_{\theta}) = \frac{1}{2} \Delta^T F(\theta) \Delta + O(\|\Delta\|^3).$$

961 Proof. Expand $\log \frac{\pi_{\theta+\Delta}}{\pi_{\theta}} = \Delta^T \nabla \log \pi_{\theta} + \frac{1}{2} \Delta^T \nabla^2 \log \pi_{\theta} \Delta + O(\|\Delta\|^3)$, take expectation under $\pi_{\theta+\Delta} =$
962 $\pi_{\theta} + O(\|\Delta\|)$, use $\mathbb{E}_{\pi_{\theta}}[\nabla \log \pi_{\theta}] = 0$ and $-\mathbb{E}_{\pi_{\theta}}[\nabla^2 \log \pi_{\theta}] = F(\theta)$. \square

963 3. Relating projection Hessian and Fisher under small tilt

965 **Lemma E.6** (Hessian–Fisher proximity). *Suppose $\|\nabla^2 \log \pi_{\theta}(y | x)\|_{\text{op}} \leq L$ uniformly near θ . Then*

$$967 \quad \|H_{\tilde{q}}(\theta) - F(\theta)\|_{\text{op}} \leq 2L \text{TV}(\tilde{q}_{\beta}, \pi_{\theta}) \leq L \sqrt{2 D_{\text{KL}}(\tilde{q}_{\beta} \| \pi_{\theta})}.$$

968 In particular, with $\kappa := D_{\text{KL}}(\tilde{q}_{\beta} \| \pi_{\theta}) \rightarrow 0$, we have $H_{\tilde{q}}(\theta) = (1 + O(\sqrt{\kappa})) F(\theta)$ as quadratic forms.

970 Proof. For bounded matrix-valued h , $\|\mathbb{E}_{\tilde{q}} h - \mathbb{E}_{\pi} h\|_{\text{op}} \leq 2\|h\|_{\infty} \text{TV}(\tilde{q}, \pi)$. Apply this with $h :=$
971 $-\nabla^2 \log \pi_{\theta}$ and Pinsker’s inequality $\text{TV}(p, q) \leq \sqrt{\frac{1}{2} D_{\text{KL}}(p \| q)}$. \square

972 **4. Remainder control**
 973

974 **Lemma E.7** (Cubic remainder is $o(f)$). *If $H_{\tilde{q}}(\theta) \succeq mI$ on the update subspace (local strong
 975 convexity), then for $\|\Delta\|$ small*

$$976 \quad \|\Delta\|^2 \leq \frac{2}{m} (f(\theta) - f(\theta^+)), \quad O(\|\Delta\|^3) = o(f(\theta)).$$

977 *Proof.* From Lemma E.4, $f(\theta) - f(\theta^+) \geq \frac{m}{2} \|\Delta\|^2 + O(\|\Delta\|^3)$. Rearranging yields $\|\Delta\|^2 = O(f(\theta) - f(\theta^+))$, so the cubic term is lower order. \square

980 **E.2.1 PROOF OF PROPOSITION 3.1**
 981

982 *Proof of Proposition 3.1.* Let $f(\theta) = D_{\text{KL}}(\tilde{q}_\beta \| \pi_\theta)$ and $\Delta = \theta^+ - \theta$. By Lemma E.4,

$$983 \quad f(\theta) - f(\theta^+) = \frac{1}{2} \Delta^\top H_{\tilde{q}}(\theta) \Delta + O(\|\Delta\|^3).$$

984 By Lemma E.5,

$$985 \quad D_{\text{KL}}(\pi_{\theta^+} \| \pi_\theta) = \frac{1}{2} \Delta^\top F(\theta) \Delta + O(\|\Delta\|^3).$$

986 By Lemma E.6 with $\kappa = f(\theta)$, $\Delta^\top F \Delta = (1 + O(\sqrt{\kappa})) \Delta^\top H_{\tilde{q}} \Delta$. Hence

$$987 \quad D_{\text{KL}}(\pi_{\theta^+} \| \pi_\theta) = (1 + O(\sqrt{\kappa})) (f(\theta) - f(\theta^+)) + O(\|\Delta\|^3).$$

988 Since $f(\theta^+) \geq 0$, $f(\theta) - f(\theta^+) \leq f(\theta) = \kappa$. By Lemma E.7, $O(\|\Delta\|^3) = o(f(\theta))$. Therefore

$$989 \quad D_{\text{KL}}(\pi_{\theta^+} \| \pi_\theta) \leq (1 + o(1)) f(\theta) = (1 + o(1)) D_{\text{KL}}(\tilde{q}_\beta \| \pi_\theta),$$

990 which is the desired inequality. \square

991 **E.2.2 PROOF OF PROPOSITION 3.2**
 992

993 *Proof of Proposition 3.2.* By the quadratic expansion of policy KL (Lemma E.5),

$$994 \quad D_{\text{KL}}(\pi_{\theta+\Delta} \| \pi_\theta) = \frac{1}{2} \Delta^\top F(\theta) \Delta + R(\Delta), \quad |R(\Delta)| \leq C \|\Delta\|^3 \quad (12)$$

995 for some local constant $C > 0$ (from C^3 smoothness). Let $a := \Delta^\top F(\theta) \Delta$. Using the spectral lower
 996 bound $F(\theta) \succeq \mu I$ on the update subspace,

$$997 \quad \|\Delta\|^2 \leq \frac{a}{\mu}. \quad (13)$$

998 Combining equation 12–equation 13 yields

$$999 \quad D_{\text{KL}}(\pi_{\theta+\Delta} \| \pi_\theta) \geq \frac{1}{2} a - C \left(\frac{a}{\mu} \right)^{3/2}.$$

1000 Since $D_{\text{KL}}(\pi_{\theta^+} \| \pi_\theta) \leq K$, we have

$$1001 \quad K \geq \frac{1}{2} a - C \mu^{-3/2} a^{3/2}. \quad (14)$$

1002 For a sufficiently small (equivalently, K small), the cubic term is dominated by the linear term:
 1003 choose $a_0 > 0$ so that $C \mu^{-3/2} \sqrt{a} \leq \frac{1}{4}$ whenever $0 < a \leq a_0$. Then from equation 14

$$1004 \quad K \geq \left(\frac{1}{2} - \frac{1}{4} \right) a = \frac{1}{4} a \Rightarrow a \leq 4K.$$

1005 Substituting $a \leq 4K$ back into equation 12 refines the remainder: $|R(\Delta)| \leq C \|\Delta\|^3 \leq C(a/\mu)^{3/2} =$
 1006 $O(K^{3/2}) = o(K)$, so $D_{\text{KL}}(\pi_{\theta+\Delta} \| \pi_\theta) = \frac{1}{2} a + o(K)$. Hence $a = 2 D_{\text{KL}}(\pi_{\theta+\Delta} \| \pi_\theta) + o(K) \leq$
 1007 $2K + o(K)$, i.e.

$$1008 \quad \Delta^\top F(\theta) \Delta \leq 2K(1 + o(1)).$$

1009 Taking square roots gives the Fisher-norm bound in equation 6: $\|\Delta\|_{F(\theta)} = \sqrt{\Delta^\top F(\theta) \Delta} \leq$
 1010 $\sqrt{2K}(1 + o(1))$. The Euclidean bound follows from equation 13:

$$1011 \quad \|\Delta\|_2 \leq \sqrt{\frac{\Delta^\top F(\theta) \Delta}{\mu}} \leq \sqrt{\frac{2K}{\mu}} (1 + o(1)).$$

1012 Finally, for any parameter block $W \subset \theta$, its Frobenius change is the ℓ_2 -norm of the corresponding
 1013 subvector of Δ ; therefore $\|\Delta W\|_F \leq \|\Delta\|_2$. \square

1026 E.2.3 ONE-STEP KL BUDGET (USED IN GATE II)
10271028 **Corollary E.8** (KL budget). *If $D_{\text{KL}}(\pi_{\theta^+} \| \pi_\theta) \leq K$, then*

1029
$$\frac{1}{2} \Delta^\top F(\theta) \Delta \leq K (1 + o(1)).$$

1030

1031 *Proof.* Apply Lemma E.5 and Lemma E.7. \square 1032 E.2.4 TRUST-REGION / CLIPPING BOUND (FOR $\beta = 0$)
10331034 **Lemma E.9** (Implicit KL leash from ratio clipping). *Let $r_t = \frac{\pi_{\theta^+}(y_t | x, y_{\leq t})}{\pi_\theta(y_t | x, y_{\leq t})}$ and suppose clipping
1035 enforces $r_t \in [1 - \varepsilon, 1 + \varepsilon]$ on the batch. Then*

1036
$$\widehat{D}_{\text{KL}}(\pi_{\theta^+} \| \pi_\theta) \leq \widehat{\mathbb{E}}[T(x)] \cdot \max\{-\log(1 - \varepsilon), \log(1 + \varepsilon)\} = O(\varepsilon) \cdot \widehat{\mathbb{E}}[T(x)],$$

1037

1038 and in the small-step regime (mean-zero advantage) this tightens to $O(\varepsilon^2)$.1039 *Proof.* Autoregressive factorization gives $D_{\text{KL}}(\pi_{\theta^+} \| \pi_\theta) = \mathbb{E}_{\pi_{\theta^+}}[\sum_t \log r_t]$. Because $\log r_t \in [0, \log(1 + \varepsilon)]$, we have $|\log r_t| \leq c(\varepsilon)$; summing over t and taking batch expectation
1040 yields the stated bound. Using $\log(1 \pm \varepsilon) = \pm \varepsilon + O(\varepsilon^2)$ and small-step arguments gives $O(\varepsilon^2)$. \square 1041 E.3 PROOFS FOR GATE II (SEC. 3.2)
10421043 **Setup (layer-conditioned budget).** Partition $\theta = (\text{vec}(W), \theta_{\neg W})$ and let the Fisher at $\theta = \theta_t$ be

1044
$$F(\theta) = \begin{bmatrix} F_{W,W} & F_{W,\neg W} \\ F_{\neg W,W} & F_{\neg W,\neg W} \end{bmatrix} \geq 0.$$

1045

1046 For a one-step update $\Delta\theta$, the global KL leash implies $\frac{1}{2} \Delta\theta^\top F(\theta) \Delta\theta \leq K$. Define the
1047 layer-conditioned curvature

1048
$$S_W := F_{W,W} - F_{W,\neg W} F_{\neg W,\neg W}^{-1} F_{\neg W,W} \geq 0,$$

1049

1050 and the per-layer budget $\delta_W := \frac{1}{2} \text{vec}(\Delta W)^\top S_W \text{vec}(\Delta W) \leq K$. Let $\mu_W := \lambda_{\min}(S_W) > 0$ on the
1051 update subspace.1052 **Lemma E.10** (Layer-conditioned Frobenius/operator bounds). $\|\Delta W\|_F \leq \sqrt{2\delta_W/\mu_W}$ and
1053 $\|\Delta W\|_2 \leq \|\Delta W\|_F$. \square 1054 *Proof.* Since $S_W \geq \mu_W I$, $\delta_W \geq \frac{1}{2} \mu_W \|\Delta W\|_F^2$. \square \square 1055 **Lemma E.11** (Wedin’s sin- Θ). *For $W_+ = W_0 + \Delta W$, the principal subspace angles satisfy
1056 $\|\sin \Theta(U_k(W_0), U_k(W_+))\|_2 \leq \|\Delta W\|_2/\gamma_k$ and similarly for V_k .* \square 1057 **Lemma E.12** (Weyl/Mirsky and Hoffman–Wielandt). $|\sigma_k(W_+) - \sigma_k(W_0)| \leq \|\Delta W\|_2$ and
1058 $\sum_i (\sigma_i(W_+) - \sigma_i(W_0))^2 \leq \|\Delta W\|_F^2$. \square 1059 **Corollary E.13** (Projection stability). *With the same assumptions,*

1060
$$\|U_k(W_0)U_k(W_0)^\top - U_k(W_+)U_k(W_+)^\top\|_2 = \|\sin \Theta(U_k(W_0), U_k(W_+))\|_2 \leq \frac{\sqrt{2\delta_W/\mu_W}}{\gamma_k}.$$

1061

1062 *The analogous bound holds for the right subspaces with V_k . Interpretation. The leading invariant
1063 subspaces rotate by at most $O(\sqrt{\delta_W/\mu_W}/\gamma_k)$; when the gap is moderate, the rotation is small.* \square 1064 F MORE VISUALIZATION
1065

1066 F.1 JACCARD MATRIX

1067 RL updates are highly consistent across independent training runs. Fig. 8 shows the pair-wise Jac-
1068 card similarity between the final update masks from five RLR runs on different data and algorithms.
1069 The high similarity scores demonstrate that the optimization process consistently targets the same
1070 subset of parameters, providing strong evidence for a deterministic, non-random optimization bias.1071 F.2 SPECTRUM SHIFT FOR DS-1.5B AND QWEN3-1
10721073 We also show the spectrum shift for DS-1.5B and Qwen3-1 here.
1074

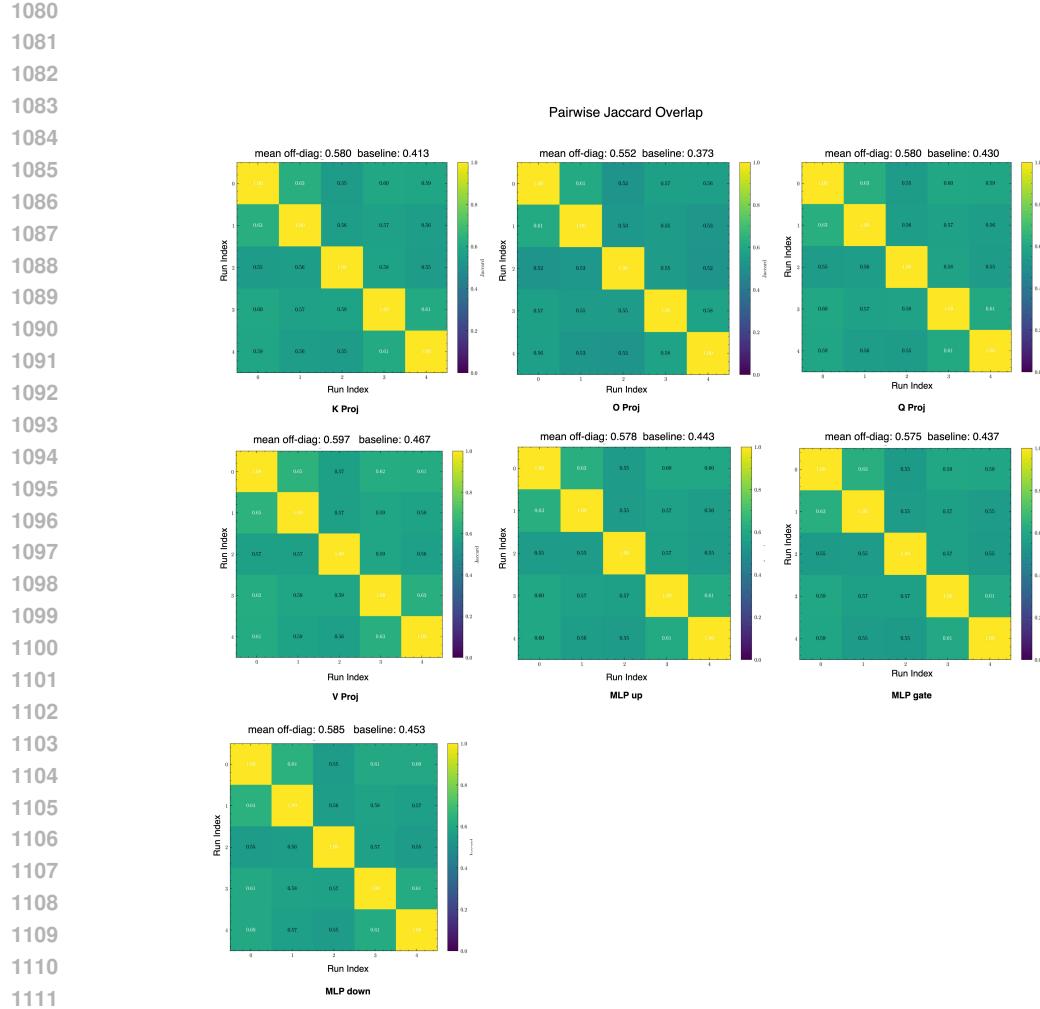


Figure 8: Pair-wise Jaccard similarity of update masks from five independent RLVR runs on Layer 13 of the DS-Distill-Qwen-1.5B model.

Table 3: Performance of DS-Qwen-1.5B with different masking strategies [with a extended training window to 500 steps](#). Parameter counts shown are for linear layers only, excluding the embedding and head layers. Detailed evaluation settings are available in Appendix C.2. *We observe that training only on principal weights M_{princ} results in a clear accuracy gap compared to both the dense baseline and its complement M_{princ}^c . The models using the M_{low} and $M_{princ}^c \cup M_{lowest}$ masks achieve performance closest to the dense baseline.*

Model	Mask	Math500	AMC23	AIME24	AIME25	Average	#params
DS-Qwen-1.5B	Dense	84.5	83.52	38.28	28.075	58.59	100%
	M_{princ}	83.60	78.83	34.06	25.63	55.44	50%
	M_{princ}^c	84.0	77.97	38.64	27.81	56.90	50%
	M_{low}	83.8	82.42	37.03	27.82	57.77	58.59%
	$M_{princ}^c \cup M_{low}$	84.10	81.41	40.30	27.70	58.37	74.02%
	Random- $M_{princ}^c \cup M_{low}$	84.10	81.72	34.69	27.34	56.89	74.02%

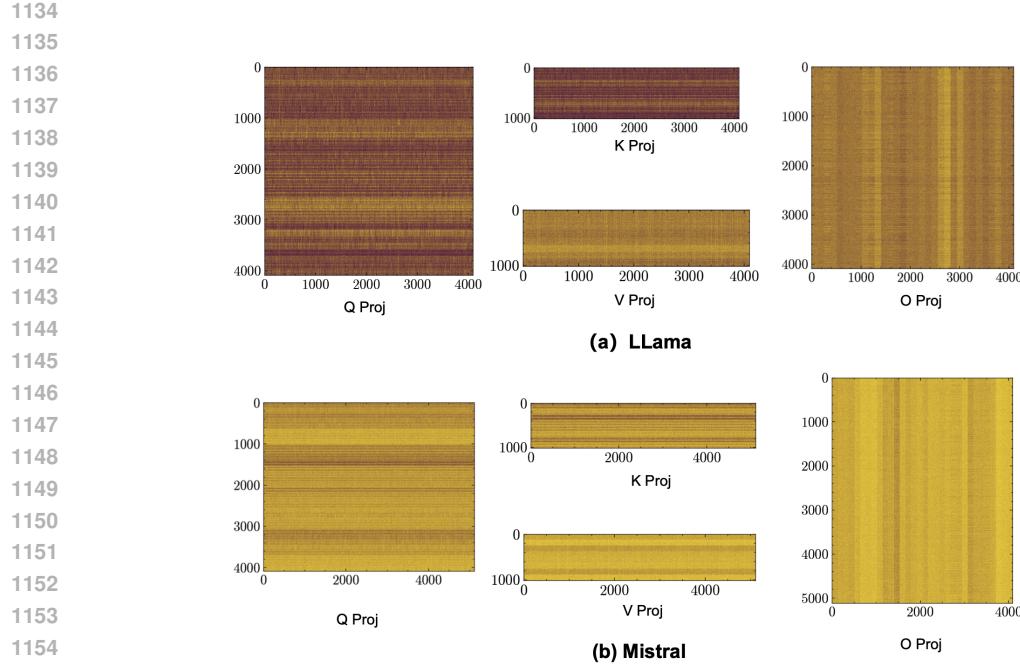


Figure 9: Structured Update observed on LLama and mistral models.

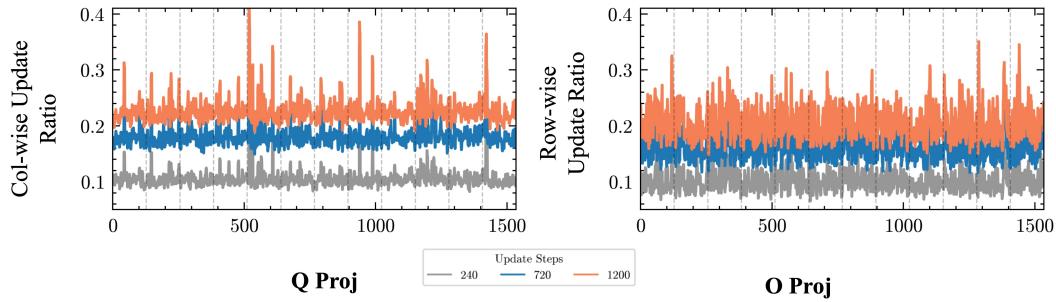
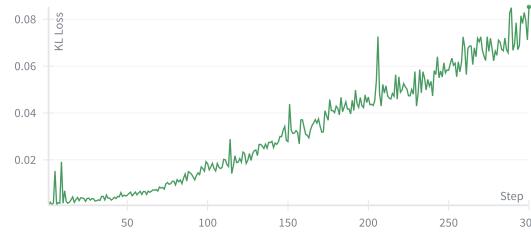
Figure 10: Temporal emergence of the optimization bias with row and column-wise update ratios for the 13th attention block across gradient update steps ($t \in \{240, 720, 1200\}$), smoothed with a 3-step window. The column-wise (Q) and row-wise (O) update ratios show a much weaker bias.

Figure 11: Token-wise KL loss. We show the token-wise KL loss during a DAPO run without a KL loss penalty, which shows a steadily increasing KL loss instead of being unconstrained.

1188
1189
1190
1191
1192
1193

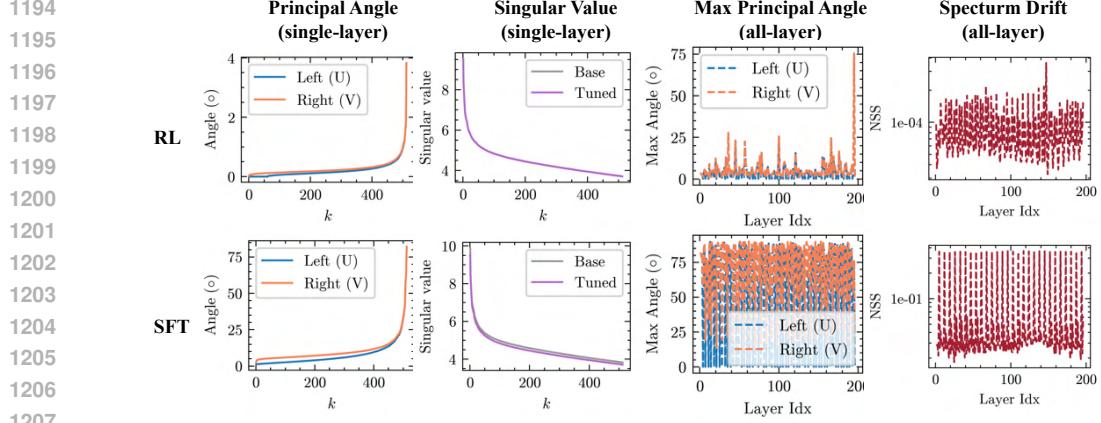


Figure 12: The spectrum probe results on the RL and SFT version on the DS-Distill-Qwen-1.5B Liu et al. (2025a). RLVR shows surprisingly stable top-k spectrum with minimal subspace rotation and top-k eigenvalue changes.

1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221

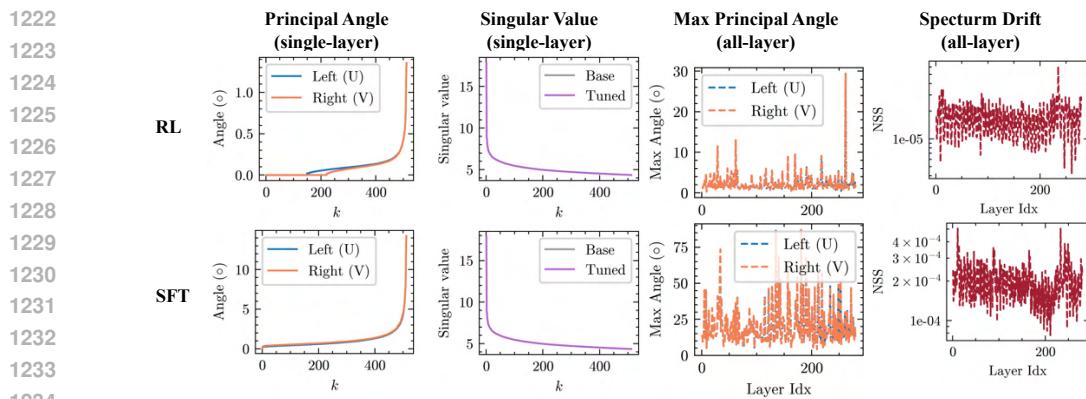


Figure 13: The spectrum probe results on the RL and SFT version on the Qwen3-14B Huan et al. (2025). RLVR shows surprisingly stable top-k spectrum with minimal subspace rotation and top-k eigenvalue changes.

1234
1235
1236
1237
1238
1239
1240
1241

1242 **G ADDED RESULTS DURING REBUTTAL**
12431244 **G.1 CONSENSUS RATIO AND ALGORITHM ROBUSTNESS**1245 To further validate our findings on the "Implicit Compass," we extended our analysis to **Llama-3.2-3b-Instruct** across five distinct reinforcement learning configurations. We varied the algorithm
1246 (Majority-Voting, Self-Certainty, Co-rewarding, GRPO) and the dataset (MATH vs. DAPO-14k).
1247 The resulting sparsity levels are detailed in Table 4.1249
1250 **Table 4: Additional Rebuttal Runs (Llama-3.2-3b-Instruct).** All from Co-reward Comparison of
1251 update sparsity across different RL algorithms and datasets.

Dataset	Method	Sparsity
MATH	Majority-Voting	71.21%
MATH	Self-Certainty	83.24%
MATH	Co-rewarding	71.86%
MATH	GRPO	71.28%
DAPO-14k	GRPO	66.00%

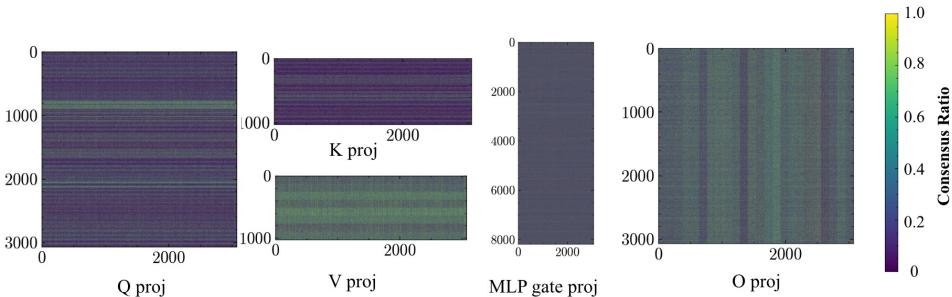
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
22210
22211
22212
22213
22214
22215
22216
22217
22218
22219
22220
22221
22222
22223
22224
22225
22226
22227
22228
22229
222210
222211
222212
222213
222214
222215
222216
222217
222218
222219
222220
222221
222222
222223
222224
222225
222226
222227
222228
222229
2222210
2222211
2222212
2222213
2222214
2222215
2222216
2222217
2222218
2222219
2222220
2222221
2222222
2222223
2222224
2222225
2222226
2222227
2222228
2222229
22222210
22222211
22222212
22222213
22222214
22222215
22222216
22222217
22222218
22222219
22222220
22222221
22222222
22222223
22222224
22222225
22222226
22222227
22222228
22222229
222222210
222222211
222222212
222222213
222222214
222222215
222222216
222222217
222222218
222222219
222222220
222222221
222222222
222222223
222222224
222222225
222222226
222222227
222222228
222222229
2222222210
2222222211
2222222212
2222222213
2222222214
2222222215
2222222216
2222222217
2222222218
2222222219
2222222220
2222222221
2222222222
2222222223
2222222224
2222222225
2222222226
2222222227
2222222228
2222222229
22222222210
22222222211
22222222212
22222222213
22222222214
22222222215
22222222216
22222222217
22222222218
22222222219
22222222220
22222222221
22222222222
22222222223
22222222224
22222222225
22222222226
22222222227
22222222228
22222222229
222222222210
222222222211
222222222212
222222222213
222222222214
222222222215
222222222216
222222222217
222222222218
222222222219
222222222220
222222222221
222222222222
222222222223
222222222224
222222222225
222222222226
222222222227
222222222228
222222222229
2222222222210
2222222222211
2222222222212
2222222222213
2222222222214
2222222222215
2222222222216
2222222222217
2222222222218
2222222222219
2222222222220
2222222222221
2222222222222
2222222222223
2222222222224
2222222222225
2222222222226
2222222222227
2222222222228
2222222222229
22222222222210
22222222222211
22222222222212
22222222222213
22222222222214
22222222222215
22222222222216
22222222222217
22222222222218
22222222222219
22222222222220
22222222222221
22222222222222
22222222222223
22222222222224
22222222222225
22222222222226
22222222222227
22222222222228
22222222222229
222222222222210
222222222222211
222222222222212
222222222222213
222222222222214
222222222222215
222222222222216
222222222222217
222222222222218
222222222222219
222222222222220
222222222222221
222222222222222
222222222222223
222222222222224
222222222222225
222222222222226
222222222222227
222222222222228
222222222222229
2222222222222210
2222222222222211
2222222222222212
2222222222222213
2222222222222214
2222222222222215
2222222222222216
2222222222222217
2222222222222218
2222222222222219
2222222222222220
2222222222222221
2222222222222222
2222222222222223
2222222222222224
2222222222222225
2222222222222226
2222222222222227
2222222222222228
2222222222222229
22222222222222210
22222222222222211
22222222222222212
22222222222222213
22222222222222214
22222222222222215
22222222222222216
22222222222222217
22222222222222218
22222222222222219
22222222222222220
22222222222222221
22222222222222222
22222222222222223
22222222222222224
22222222222222225
22222222222222226
22222222222222227
22222222222222228
22222222222222229
222222222222222210
222222222222222211
222222222222222212
222222222222222213
222222222222222214
222222222222222215
222222222222222216
222222222222222217
222222222222222218
222222222222222219
222222222222222220
222222222222222221
222222222222222222
222222222222222223
222222222222222224
222222222222222225
222222222222222226
222222222222222227
222222222222222228
222222222222222229
2222222222222222210
2222222222222222211
2222222222222222212
2222222222222222213
2222222222222222214
2222222222222222215
2222222222222222216
2222222222222222217
2222222222222222218
2222222222222222219
2222222222222222220
2222222222222222221
2222222222222222222
2222222222222222223
2222222222222222224
2222222222222222225
2222222222222222226
2222222222222222227
2222222222222222228
2222222222222222229
22222222222222222210
22222222222222222211
22222222222222222212
22222222222222222213
22222222222222222214
22222222222222222215
22222222222222222216
22222222222222222217
22222222222222222218
22222222222222222219
22222222222222222220
22222222222222222221
22222222222222222222
22222222222222222223
22222222222222222224
222222222222222

Table 5: Model List for analyzed checkpoints for agentic and embodied AI (robotics manipulation) tasks and RLHF algorithms.

Category	Base Model	FT Model	Algorithm	Data	Sparsity
Agent	Qwen3-8B	SkyRL-Agent-WebResearch-8B	GRPO	WebResearch	40.56%
	Qwen3-8B	VT-deepsearch-8B	GRPO	Deepsearch	89.67%
	Qwen3-8B	VT-SWE-8B	GRPO	SWE	84.32%
	Qwen2.5-7B-Instruct	agentflow-planner-7b	Flow-GRPO	Planning	80.99%
	Qwen2.5-7B-Instruct	GiGPO-Qwen2.5-7B-Instruct-WebShop	GiGPO	WebShop	51.7%
	Qwen2.5-7B-Instruct	GiGPO-Qwen2.5-7B-Instruct-ALFWORLD	GiGPO	ALFWORLD	62.08%
RLHF	Meta-Llama-3-8B-Instruct	Llama-3-Instruct-8B-DPO	DPO	instruction-following	82.38%
	Meta-Llama-3-8B-Instruct	Llama-3-Instruct-8B-SimPO	SimPO	instruction-following	71.00%
Embodied AI	openvla/openvla-7b	Openvla-oft-SFT-libero10-trajall	SFT	Robotic manipulation Liu et al. (2023)	3.46%
	openvla/oft-SFT-libero10-traj1	openvla-oft-libero10-traj1-rl	GRPO	Robotic manipulation Liu et al. (2023)	35.04%
	Qwen2.5-VL-3B-Instruct	Embodied-R1-3B-v1	GRPO	Robotic Manipulation	44.28%

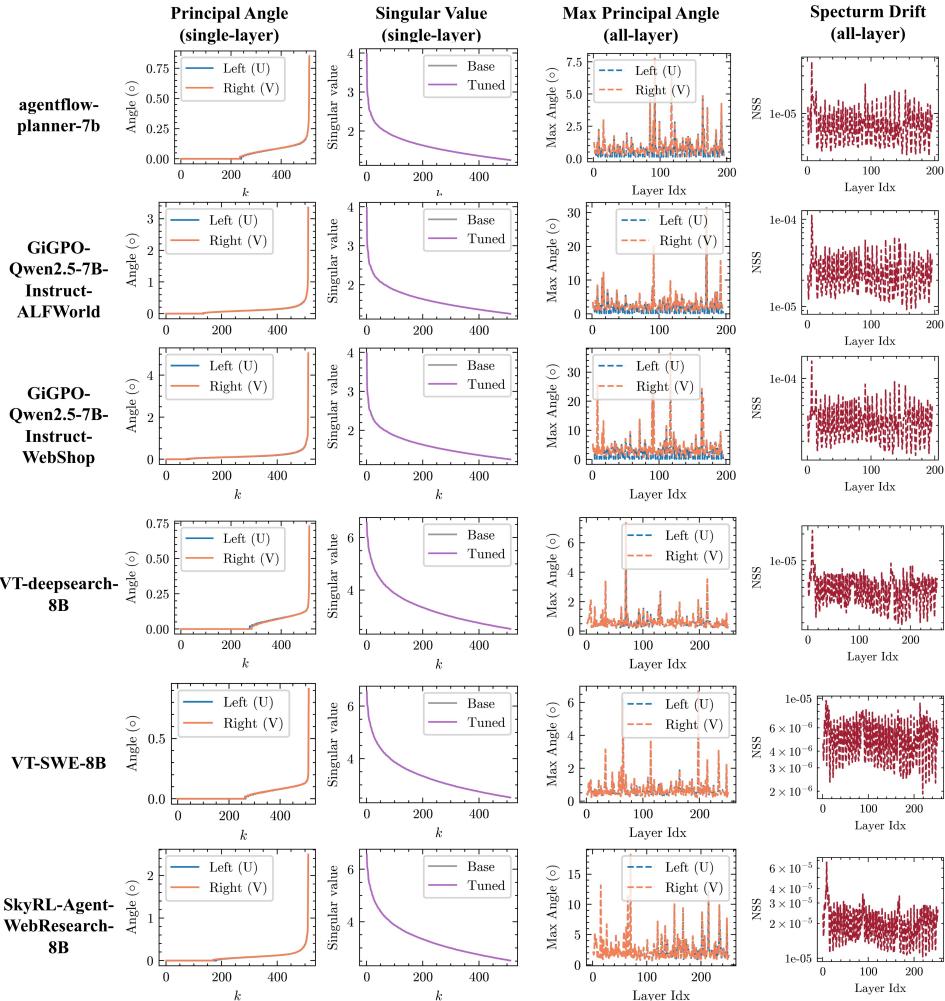


Figure 15: **Spectrum under RL in agent tasks.** In agent settings, including multi-turn interactions and tool use, RL leaves layer singular-value spectra nearly unchanged and induces only small rotations of the top- k singular subspaces, consistent with the spectrum-preserving, off-principal RLVR regime. Results for RL with human feedback (RLHF), which exhibit the same optimization signature, appear in Fig. 16. For consistency, we use the *second block* O -projection layer as an exemplar single-layer readout.

Off-Principal Routing. Finally, we verify the spatial distribution of the updates. Figure 18 visualizes the update masks (M_ℓ) against the principal masks ($P_\ell^{(k)}$) for representative agent and embodied models. Consistent with the "Implicit Compass" hypothesis, the updates strictly avoid the principal parts.

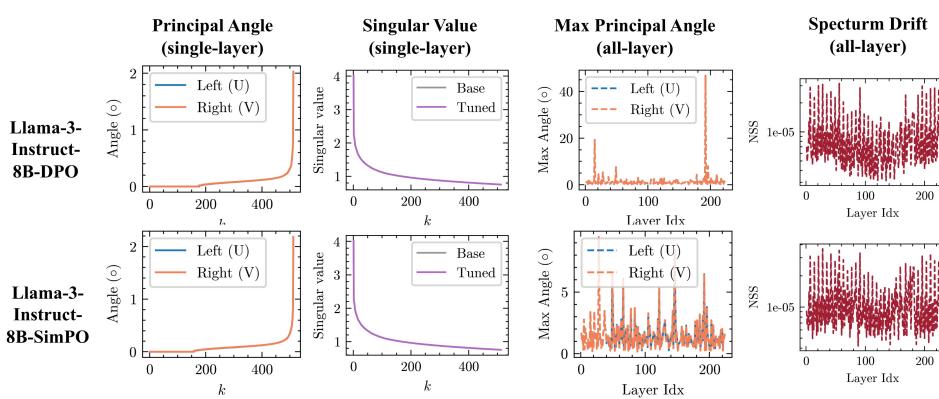


Figure 16: **Spectral geometry under RLHF setting Meng et al. (2024b).** Across RLHF checkpoints, RL training preserves layer spectra and induces only minor rotation of the top- k subspaces, consistent with the RLVR regime.

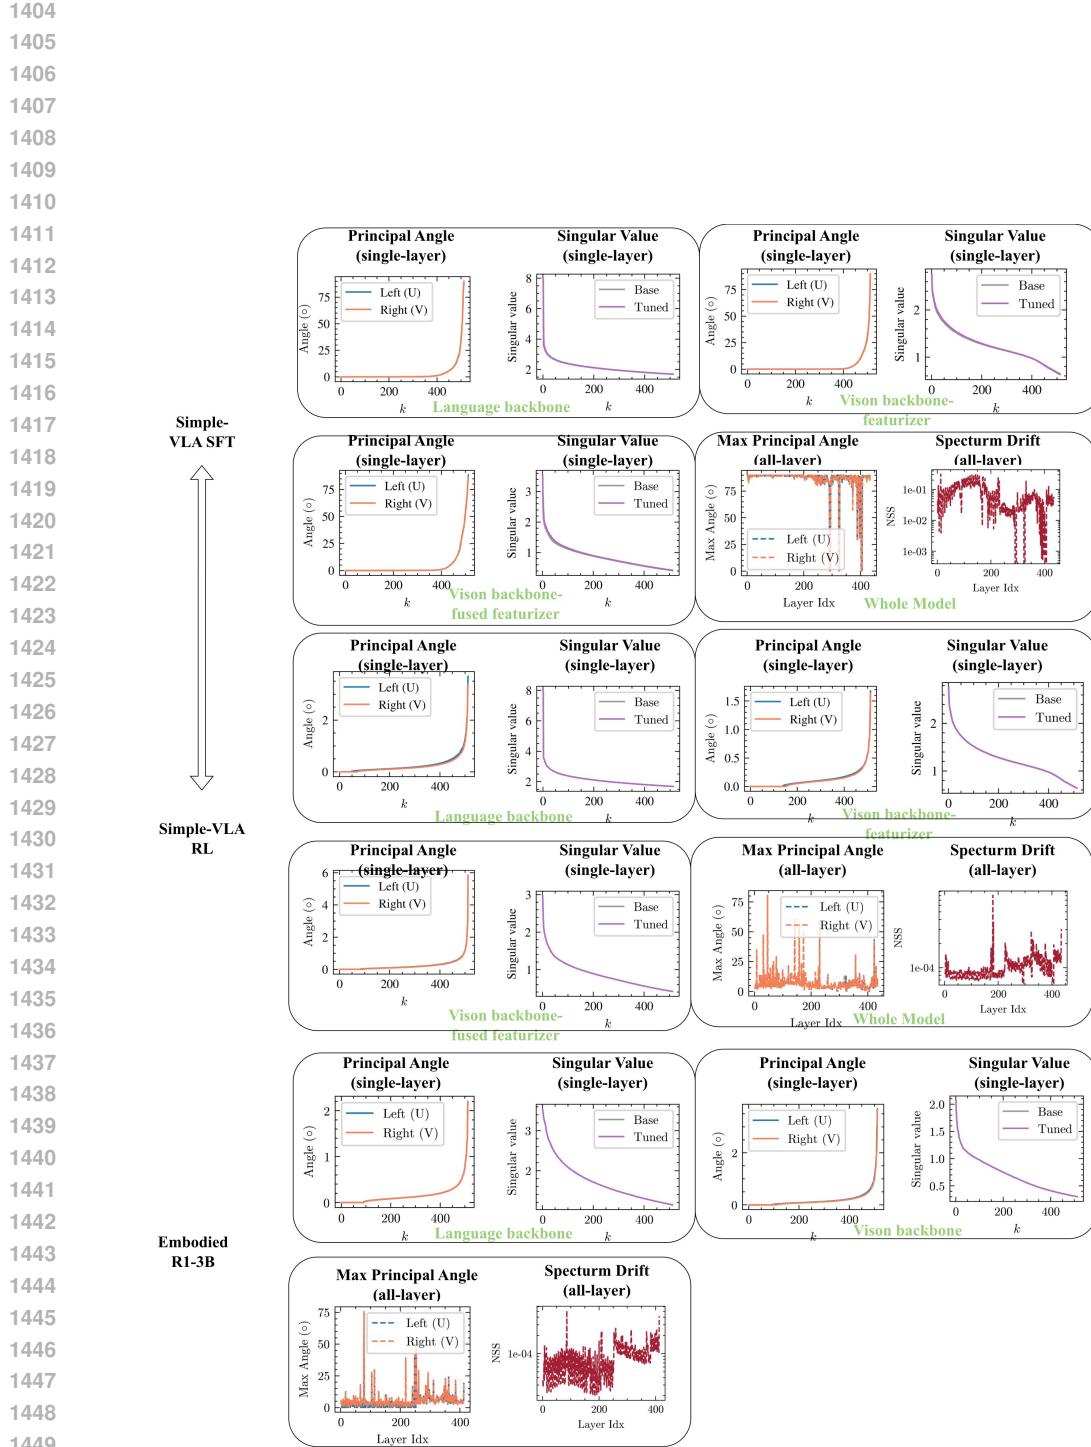
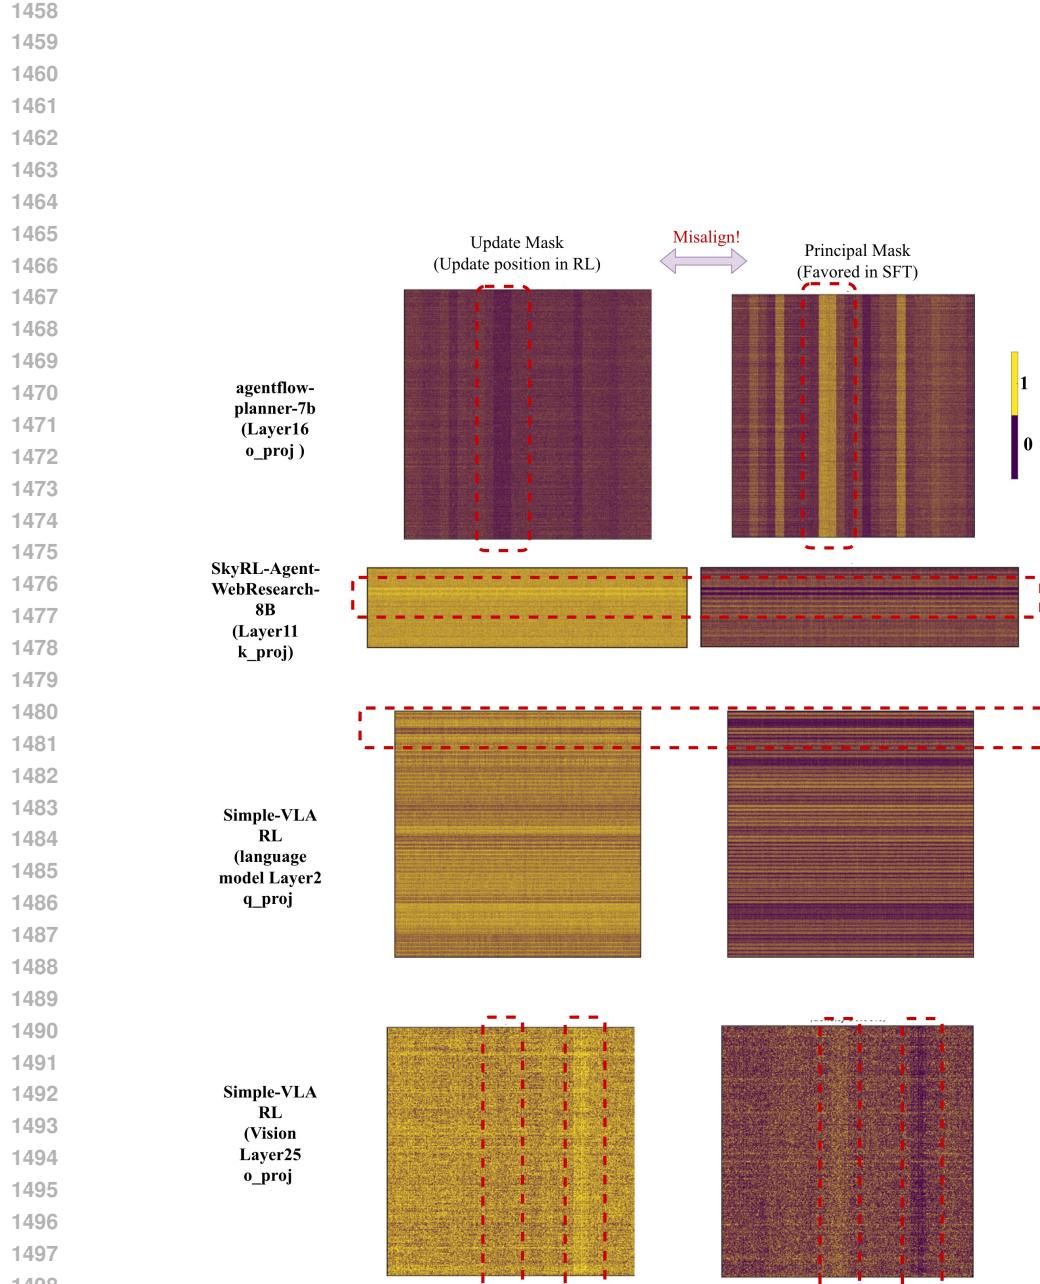


Figure 17: **Spectral geometry under embodied AI tasks: SimpleVLA-RL (Li et al., 2025a) and Embodied R1 (Yuan et al., 2025).** Across language and vision backbone, RL training preserves layer spectra and induces only minor rotation of the top- k subspaces.



1499 Figure 18: **Update–principal misalignment in RL-trained agents and embodied AI models.** The figure
 1500 visualizes the bf16-aware *update mask* M_ℓ (left, showing locations changed under RL) versus the *principal*
 1501 $M_\ell^{(k)}$ (right, showing top- k singular-subspace support) for representative layers across different domains.
 1502 Top row: AGENTFLOW-PLANNER-7B (Layer 16, o_{proj}). Second row: SKYRL-AGENT-WEBRESEARCH-8B
 1503 (Layer 11, k_{proj}). Third row: SIMPLE-VLA-RL (language model Layer 2, q_{proj}). Bottom row: SIMPLE-
 1504 VLA-RL (Vision Layer 25, o_{proj}). Dashed blue boxes highlight regions where RL updates concentrate *outside*
 1505 principal-weight bands. This consistently shows that RL updates are misaligned with the principal subspace,
 1506 indicating a robust off-principal routing mechanism across agent, tool-use, and embodied AI settings.

1507
 1508
 1509
 1510
 1511

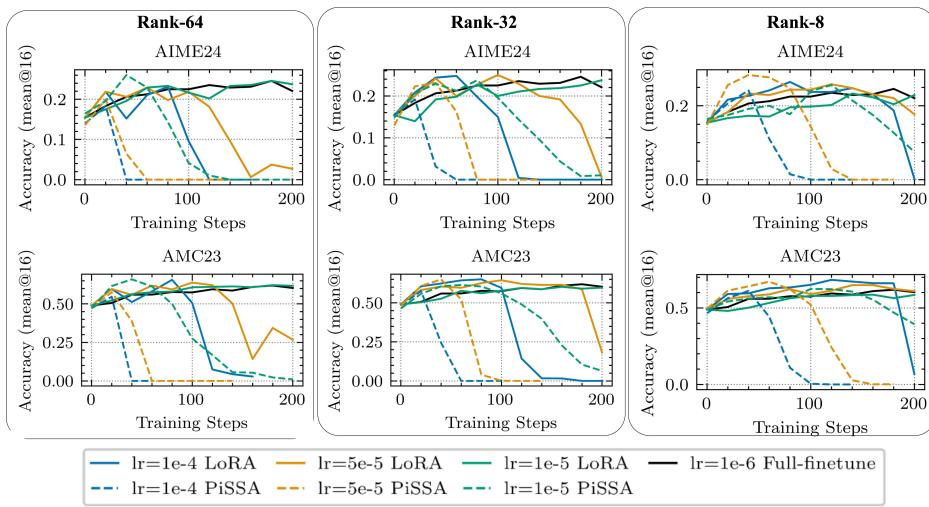


Figure 19: **LoRA vs. PiSSA on DS-Qwen-1.5B (DeepMath-103K).** We sweep ranks $\{8, 32, 64\}$ and learning rates $\{1 \times 10^{-4}, 5 \times 10^{-5}, 1 \times 10^{-5}\}$ for 200 steps, reporting pass@1 (avg@16) on AIME24 (top) and AMC23 (bottom). Across settings, PiSSA (principal-targeted) provides *no additional gains* over LoRA and, at higher learning rates that force principal-direction updates, *often collapses early*; LoRA remains more stable. This supports our geometric account: forcing updates into principal directions (favored in SFT) is misaligned with RL, offering no obvious gain and leading to training collapse when scaling up learning rates.

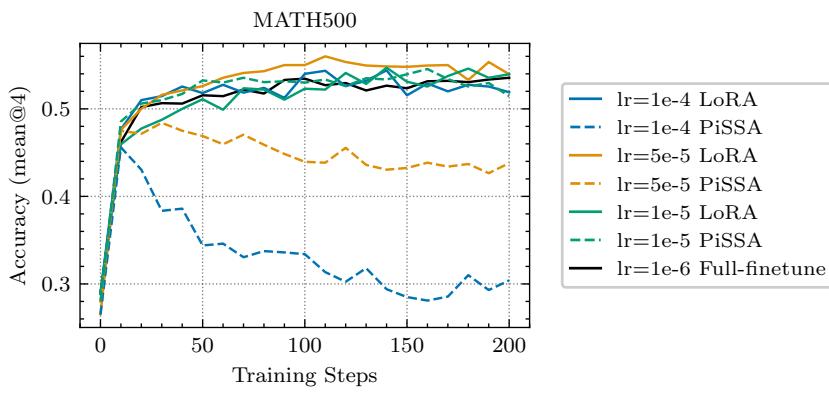


Figure 20: **LoRA vs. PiSSA on LLaMA-3.2-3B.** We sweep learning rates $\{1 \times 10^{-4}, 5 \times 10^{-5}, 1 \times 10^{-5}\}$ with a fixed rank of 64 for 200 steps, reporting pass@1 (mean@4) on MATH500. Consistent with the DS-Qwen-1.5B results in Fig. 19, PiSSA provides *no additional gain* over LoRA and, under higher learning rates that emphasize principal-direction updates, *often collapses early*.