arXiv:2611.24578v1 [cs.CV] 18 Nov 2026

BaseRSI: Recursive Self-Improvement of a Base Vision-Language Model via Training-Free Harness Internalization

Mingda Zhang[1*] Yuanhao Dong[2*] Qiang Huang[3] Rui Qian[3] Yanjin Li[4] Zhenyu Liu[5]
Xiao Huang[6] Zihao Yang[1] Jian Wang[1] Pinjia He[1] Xiaoying Tang[1]

[1]The Chinese University of Hong Kong, Shenzhen, China [2]National University of Singapore, Singapore [3]Fudan University, China
[4]University of Illinois Urbana-Champaign, USA [5]The University of Melbourne, Australia [6]Shanghai Jiao Tong University, China
*Equal contribution mingdazhang@ieee.org

Abstract

Recursive self-improvement (RSI) is moving from the harness of prompts and tools around a base model toward the model itself. However, existing RSI still faces key challenges, including harness-bound gains *that stay outside the base model,* black-box supervision *that reads only the model's outputs, and* imprecise edits *whose site and strength are fixed in advance. To address these challenges, we propose BaseRSI, a new paradigm of* Recursive Harness Internalization*: the base model improves its harness, then writes what the harness adds into its own weights through training-free edits, round after round on its own responses. To support this paradigm, we introduce* ReGap*, a mechanism that supervises every transformer block with the high-dimensional activation difference that the harness makes at that block alone, and updates the block's weights along this signal in a low-dimensional subspace. Building on this mechanism, we further propose* ReKey*, an algorithm that writes this signal by regression into a few key–value neurons of every block, and measures the site and strength of each edit on the model's own responses. Moreover, the base model fills every role of this recursion with forward passes alone, and its edits, written at different sites, compose in one checkpoint without refitting. Experimental results on ten visual reasoning benchmarks show that BaseRSI consistently improves a base vision-language model both in distribution and on held-out benchmarks. Our code is available at* https://github.com/beita6969/BaseRSI.

1. Introduction

In recent years, a variety of recursive self-improvement (RSI) systems have let foundation models improve the prompts, tools, and context they work with [31, 32, 82], gradually moving from the harness around a frozen model toward improving the base model itself.

In this process, the harness has become a key bridge from a frozen base model to a stronger system: by adding verified procedures and finer views to the request, a vision-language model can answer visual questions with improved accuracy, controllability, and reuse (Figure 1). However, in practice,

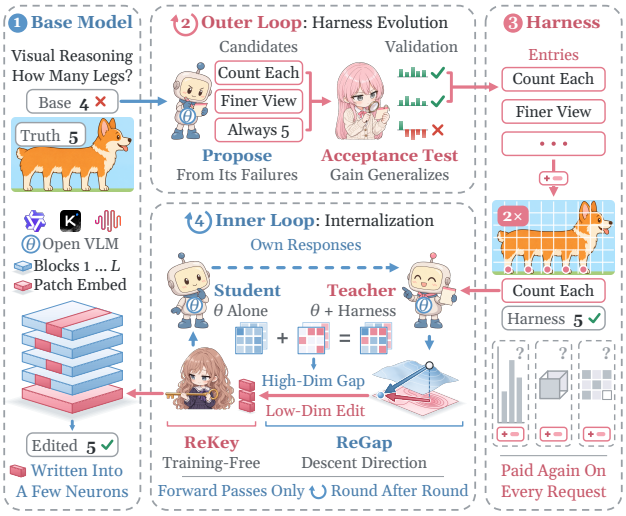


Figure 1. **A harness gain that stays outside the base model is paid again on every request.** BaseRSI lets the base model write what its harness learns into its own weights, round by round, and answer the original request alone.

these gains still live in the request rather than in the model [60, 75], making them costly on every request and leaving the base model unable to build on what its harness learns.

Along this path from the harness to the model, three main lines of work have emerged, as shown in Figure 2. First, *harness evolution* rewrites the prompts, tools, and programs around frozen weights [31, 32, 68, 70], treating the harness as the object of self-improvement. Second, *model–harness co-evolution* updates the weights together with the harness through fine-tuning, reinforcement learning, or distillation [18, 38, 76], moving harness gains into the model. Third, *weight editing* writes new behavior into the weights directly, as memories of facts and mistakes [26, 44, 45] or as updates computed directly from forward passes [46, 72].

However, these methods still face three challenges (Figure 2a–c). **(i) Harness-bound gains.** Harness evolution improves what surrounds a frozen model [31, 32, 68]; every addition costs context again on each request and is lost when the harness is taken away [18, 60, 75], so the gains accumulate outside the base model and its own capability stays where it started. **(ii) Black-box supervision.** Co-

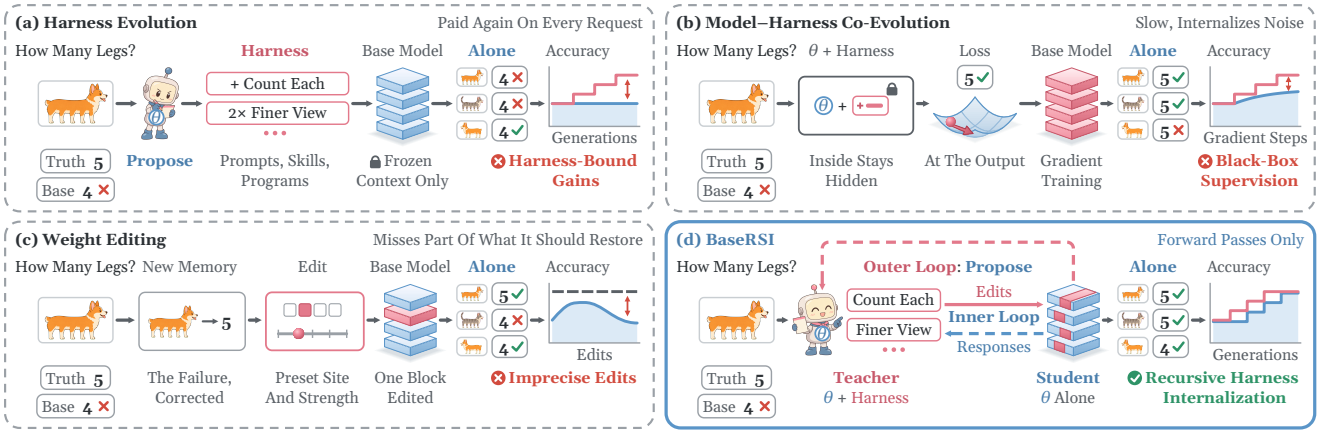


Figure 2. **Three lines of work and BaseRSI.** (a) Harness evolution leaves *harness-bound gains*. (b) Model–harness co-evolution learns from *black-box supervision*. (c) Weight editing makes *imprecise edits*. (d) BaseRSI evolves the harness and internalizes each accepted gain into the base model from its own forward passes.

evolution and harness distillation supervise the weights only at the output of the model, with trajectories, rewards, or output distributions [38, 58, 75, 81]; what the harness changes inside the model stays hidden, so the signal travels back through every block by gradient training, which is slow and internalizes noise [18]. **(iii) Imprecise edits.** Weight editing sets the site and strength of an edit by a preset rule, in a few layers [45] or the last layer [26]; localization does not tell where an edit works best [23, 48], so edits miss part of their target [6] and harm general abilities [21].

To address these challenges, we propose **BaseRSI** (Figure 2d), a new paradigm of *Recursive Harness Internalization* that runs in two loops: in the outer loop, the base model adds candidate entries to its harness, and an acceptance test keeps those whose gain generalizes. In the inner loop, the base model with an accepted entry in the request is the teacher, the same model on the original request is the student, and the teacher and the unedited student first answer the development questions. To support this paradigm, we introduce *ReGap*, a mechanism that replays each response through the teacher and the student, takes the difference that the entry makes at every block alone, the teacher–student gap, as a high-dimensional target inside the model rather than at its output, and turns each gap into a low-dimensional update of that block's weights. Building on this mechanism, we further propose *ReKey*, an algorithm that first measures on the student's responses which weights the entry should be written into, then writes the gaps by regression into a few key–value neurons of every block, and finally measures the strength of the edit on held-out responses. The edited student then answers again, and since the edit has changed its responses, the next round repeats these steps on the new ones. After the last round, the edits written at different sites add up in one checkpoint, which answers the original requests alone. Throughout, the base model fills every role with its own forward passes, and BaseRSI joins harness evolution, weight editing, and internalization into one recursive loop that makes the base model itself stronger.

We evaluate on five in-distribution and five held-out visual reasoning benchmarks across perception, counting, charts, spatial reasoning, and visual logic. Results show that BaseRSI outperforms prompting, visual zoom-in, harness evolution, and fine-tuning on every benchmark and carries over to other base models (RQ1, RQ2), that the edited checkpoint exceeds its own harness in the request while answering original requests on standard inference engines (RQ3), and that supervision inside the model and a measured site and strength each add to the gain (RQ4, RQ5). These results mark a step toward base models that recursively improve their own weights.

2. Related Work

Self-Improving Harnesses and Models. Harness evolution searches and rewrites prompts, tools, and programs around frozen weights [31, 32, 37, 68, 70, 82], and co-evolution, harness distillation, and self-edits also update the weights, by fine-tuning, reinforcement learning, or distillation from a teacher that holds the context [8, 18, 28, 35, 38, 58, 69, 75, 76, 81, 83], at times with a scheduler choosing between harness and weights [24, 60]. The former keeps its gains in the request and the latter supervises the weights at the output, and both often run an external frontier model as proposer or corrector [24, 32, 75, 76]. BaseRSI keeps every role in one base model and writes each accepted entry into its weights from the gap measured inside the model.

Weight Editing and Parametric Memory. Weight editing writes facts, mistakes, and experiences into one or a few layers, the last layer, or adapters [26, 44, 45, 51, 80], and more precise or neighbor-aware edits raise locality [6, 48],

while edits still miss related knowledge and can harm general abilities [12, 21, 78]. A context or behavior is also condensed into weights, as steering directions [5], implicit or least-squares block updates [13, 20, 42], edits from contrastive forward passes [14, 46], or dormant neurons [77], and specialists of one base are fused from forward passes [72]; concurrently, a frontier model rewrites the weights of a small model [59]. In contrast, BaseRSI fits each edit to the gap of every block alone, measures its site and strength, and refits it round by round on the model's own responses.

3. Preliminaries

Definition 1: Base model and harness. A base model f_θ is a vision-language model with weights θ and L residual blocks, each applying an attention layer (linear or full attention) and a SwiGLU MLP [55] to a residual stream of width H. A harness [31] $\mathcal{H} = \{c_1, \dots, c_m\}$ wraps the model with entries $c = (\kappa_c, a_c)$, where the scope $\kappa_c(q) \in \{0, 1\}$, answered by the base model as a router, tells whether the entry applies to a request q, and the action a_c rewrites the request by adding a procedure to the prompt or a finer view of the image; $c(q)$ is the rewritten request when the entry applies and q otherwise. The entries are the parameters of the harness, as the weights θ are those of the model.

Definition 2: Response trajectory. A model with a harness responds to a request q with a token sequence $y_{1:T}$, its reasoning followed by the final answer, in which every token is drawn on the request that the harness has rewritten:

$$\begin{aligned} y_t \sim p_\theta\big(\cdot \mid \mathcal{H}(q), y_{<t}\big), \quad \mathcal{H}(q) = (c_m \circ \cdots \circ c_1)(q), \\ \boldsymbol{s}_l = F_l\big(\boldsymbol{s}_{l-1}; \boldsymbol{z}_{l-1}\big), \quad l = 1, \dots, L, \end{aligned} \tag{1}$$

where p_θ is the next-token distribution, $\boldsymbol{z}_l$ and $\boldsymbol{s}_l$ are the activations after block l at the prompt and response positions (the tokens of y), the prompt states and the response states, and F_l is the output of block l at the response positions. Blocks are causal, so the prompt states depend on the request alone, and an empty harness leaves it unchanged.

Problem Statement. Given base weights θ_0, an empty harness $\mathcal{H}_0$, a pool $\mathcal{Q}$ of development requests, and a family $\mathcal{E}$ of edits that rewrite a bounded number of MLP neurons per block or the patch kernel, we seek a base model that answers original requests better by recursive self-improvement:

$$\max_{e_0, \dots, e_{N-1} \in \mathcal{E}} \mathbb{E}_{(q, y^\star) \sim \mathcal{T}} \, \mathbf{1}\big[f_{\theta_N}(q) = y^\star\big], \tag{2}$$

where $\mathcal{T}$ are test tasks in distribution and held out, $f_{\theta_N}(q)$ is the final answer of the base model to the original request, and $\theta_{n+1} = \theta_n \oplus e_n$ are the weights after edit e_n. Every generation may evolve a harness $\mathcal{H}_n$ to find what to internalize, and only the weights enter Eq. (2). Equation (2) is the deployed target; Section 4 shows how an outer loop over the harness and an inner loop over the weights approach it.

4. Method

As illustrated in Figure 3, this section introduces the BaseRSI framework from the system down to the weights, including Recursive Harness Internalization (Section 4.1), the paradigm in which the base model teaches itself, ReGap (Section 4.2), the mechanism that states what is taught in every block, and ReKey (Section 4.3), the algorithm that writes it into the weights at a measured site and strength.

4.1. Recursive Harness Internalization

The paradigm lets the base model keep what its harness learns. BaseRSI treats the harness as a fast system that records new experience and the weights as a slow system that consolidates it through replay [30, 43], and joins them in one cycle (Figure 3, top): an outer loop finds a gain in the harness, an inner loop moves it into the weights, and the next cycle looks for a further gain.

Outer loop: harness evolution. The fast system has to supply gains that are worth keeping, so the base model θ_0 proposes entries from its own failures and an acceptance test keeps those whose gain generalizes:

$$\mathcal{A}_n = \big\{c : \; p_{\text{gen}}(c) = \max_{b \in \mathcal{B}} p\big(\{d_{b'}(c)\}_{b' \neq b}\big) \leq \alpha_c\big\}, \tag{3}$$

where $d_b(c)$ is the paired accuracy gain of adding c on the validation questions of benchmark b at an equal token budget, $p(\cdot)$ is the one-sided sign-flip p-value of the macro mean of its arguments, and α_c is the Holm threshold over the candidates of generation n [25]. A general entry is thus kept only when its gain is significant with any one benchmark left out, and an entry for one benchmark only when its gain there is significant and no other benchmark is significantly worse (Appendix B). Each accepted entry is tested on questions disjoint from those it was written from, so it hands the inner loop a verified target.

Inner loop: internalization. A verified entry $c \in \mathcal{A}_n$ still acts through the request, so the inner loop moves it into the weights. The base model with c in the request is the teacher $\pi^c_{\theta_0} = p_{\theta_0}(\cdot \mid c(q), y_{<t})$, the same model on the original request is the student $\pi_{\theta_0} = p_{\theta_0}(\cdot \mid q, y_{<t})$, and the edit makes the student match its own teacher on the responses that it generates by itself [36, 62]:

$$\begin{aligned} e_n \in \arg\min_{e \in \mathcal{E}} \; \mathbb{E}_{q \sim \mathcal{Q}} \, \mathbb{E}_{y \sim \pi_e(\cdot \mid q)} \, \overline{\mathrm{KL}}\big(\pi^c_{\theta_0} \,\big\|\, \pi_e\big), \\ \theta_{n+1} = \theta_n \oplus e_n, \qquad \mathcal{H}_{n+1} = (\mathcal{H}_n \cup \mathcal{A}_n) \setminus \{c\}, \end{aligned} \tag{4}$$

where $\pi_e = \pi_{\theta_0 \oplus e}$ is the student that carries edit e, teacher and student both run under the entries internalized before c, and $\overline{\mathrm{KL}}$ averages the divergence between the two next-token distributions over the response positions of (q, y). The expectation runs over responses of the edited student, and since $c(q) = q$ outside the scope of the entry, the same

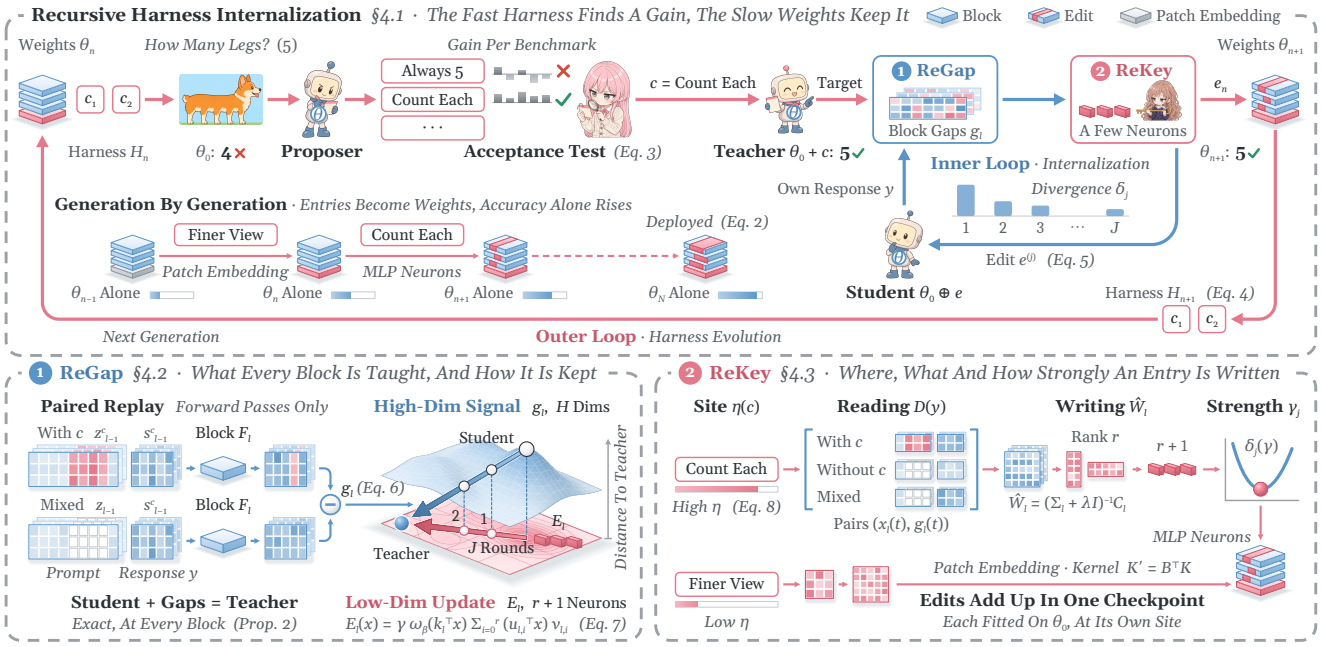


Figure 3. **Overview of BaseRSI.** *Top:* Recursive Harness Internalization runs an outer loop that evolves the harness and an inner loop in which the model with its harness teaches the model alone, round after round. *Left:* ReGap supervises every block with the teacher–student gap and keeps it as a low-dimensional update. *Right:* ReKey writes the gap into key–value neurons at a measured site and strength.

objective there asks the edit to keep the base model as it is. Every edit is fitted on the base weights and takes its own neurons or the patch kernel, so the edits add up in one checkpoint without refitting, and each can be removed alone. The internalized entries leave the request, and the checkpoint θ_N with every edit is deployed alone (Eq. (2)).

Rounds on its own responses. The target of Eq. (4) is a fixed point, since an edit changes the responses on which it is fitted. An edit fitted once on the teacher's responses therefore loses part of its effect once the edited student generates its own, the covariate shift of imitation learning [33, 52]. So the inner loop runs in rounds on the student's own responses [53]: the edited student responds, the teacher labels that response, and the edit is refitted on the pooled responses:

$$\mathcal{D}^{(j)} = \mathcal{D}^{(j-1)} \cup \mathcal{D}(y), \qquad y \sim \pi_{e^{(j-1)}}(\cdot \mid q),$$
$$e^{(j)} = \mathrm{ReKey}\big(\mathcal{D}^{(j)}\big), \tag{5}$$

where $\mathcal{D}(y)$ are the labels of a response y, the gaps of Section 4.2, and $\mathrm{ReKey}(\cdot)$ is the edit that Section 4.3 fits on a set of labels. The first two rounds replay the responses of the teacher and of the unedited student, and the later rounds those of the edited student (Appendix C). Every round computes its labels with the unedited base model, so the labels of all rounds describe one target and can be pooled. The refit replaces the previous edit and minimizes the summed loss of all rounds so far [53], and the last one is e_n of Eq. (4).

Proposition 1. *Some round's fit leaves a gap on its own student's responses no larger than the smallest regularized gap any single linear readout leaves on the responses of all rounds, plus a margin that shrinks as rounds accumulate. Proof.* Appendix A.4, via Lemmas A.10 and A.11 [53].

4.2. ReGap

ReGap is the mechanism of the inner loop: it states what the teacher teaches in every block, a high-dimensional signal, and in what form the student keeps it, a low-dimensional update of the block's weights (Figure 3, bottom left).

High-dimensional signal. The whole effect of an entry on a response is a set of additions, one per block. We define the gap of block l by running the block on the teacher's states with and without the entry in the request, and treat it as an implicit update of the block's MLP weights [13, 20, 42]:

$$\boldsymbol{g}_l = F_l\big(\boldsymbol{s}^c_{l-1};\, \boldsymbol{z}^c_{l-1}\big) - F_l\big(\boldsymbol{s}^c_{l-1};\, \boldsymbol{z}_{l-1}\big), \tag{6}$$

where F_l is the block map of Eq. (1) and the superscript c marks the states of the teacher. Both terms take the same response states, so $\boldsymbol{g}_l$ is what the entry adds at block l alone, its direct effect [49], and gaps of different blocks share no part. For a procedure entry, the teacher's attention layers read procedure tokens that the student's request lacks, and $\boldsymbol{g}_l$ is the part of the block's output that they account for.

Descent direction. The gap takes the place that the gradient has in training. It is the exact descent direction of the distance between block l and its teacher [16, 40]: a unit step along $\boldsymbol{g}_l$ closes that distance, and it comes from forward passes, with no gradient propagated from the output. Every

term of the gap is a forward pass of the base model itself, which also proposes and routes the entries, so one model fills every role of the cycle and edits itself.

Low-dimensional update. The student has to produce $\boldsymbol{g}_l$ from its own state, so the edit is a function of the input $\boldsymbol{x}$ of the MLP of block l, kept in a few directions. The update is a small set of key–value neurons, each storing one key–value pair in the MLP weights [2, 7, 19, 29, 44]:

$$E_l(\boldsymbol{x}) = \gamma\, \omega_\beta\big(\boldsymbol{k}_l^\top \boldsymbol{x}\big) \sum\nolimits_{i=0}^{r} \big(\boldsymbol{u}_{l,i}^\top \boldsymbol{x}\big)\boldsymbol{v}_{l,i}, \tag{7}$$

where each term reads the input along a readout $\boldsymbol{u}_{l,i}$ and writes a value $\boldsymbol{v}_{l,i}$, the key $\boldsymbol{k}_l$ carries the entry's scope through $\omega_\beta(z) = z\,\sigma(\beta z)$, γ is the step, and r is far below the width H. Each term is one SwiGLU neuron, which passes its readout scaled by the key (Lemma A.7), so an edit rewrites $r{+}1$ neurons per block, every other weight stays bit-identical, and the MLP keeps its architecture.

Proposition 2. *Adding every block's gap to the student's activations makes its output identical to the teacher's. Proof.* Appendix A.1, via Lemma A.2 and induction.

4.3. ReKey

ReKey is the algorithm that computes the update of Eq. (7) from forward passes alone: it first measures the site of an edit, then writes the gap by regression into a few key–value neurons of every block, and finally measures its strength, all on the model's own responses (Figure 3, bottom right).

Site. Before any fit, ReKey measures the site of an edit, the weights that an entry is written into. It replays the student's own responses with the outputs of its linear-attention layers set to those of the teacher and measures the fraction of the divergence from the teacher that this patch removes [49]:

$$\eta(c) = 1 - \overline{\mathrm{KL}}\big(\pi_\theta^c \,\big\|\, \pi_\theta^{\mathrm{lin}\leftarrow c}\big) \,/\, \overline{\mathrm{KL}}\big(\pi_\theta^c \,\big\|\, \pi_\theta\big), \tag{8}$$

where $\pi_\theta^{\mathrm{lin}\leftarrow c}$ is the patched student. Linear attention carries the request as a running state, which converts exactly into bias terms [10], whereas full attention retrieves single image tokens. A high fraction therefore sends the entry to the MLP neurons, which read the gap from the activations, and a low fraction sends it to the patch embedding, since such an entry changes what the model sees. For such a view entry and an up-sampling B of the image, resampling the patch kernel K [9] to $K' = B^\top K$ makes the student embed every patch of the original image as the teacher embeds it after up-sampling (Lemma A.9, Appendix A.3).

Reading. ReKey then reads the gap of Eq. (6) on a response y by teacher-forced forward passes of the base model over its tokens, which label every block:

$$\mathcal{D}(y) = \big\{\big(\boldsymbol{x}_l(t), \boldsymbol{g}_l(t)\big)\big\}_{l,\, t\in\mathcal{S}(y)}, \tag{9}$$

where $\boldsymbol{x}_l(t)$ is the input of the MLP of block l at response position t in the second term of Eq. (6) and $\mathcal{S}(y)$ are the sampled response positions. The pairs are exact by Proposition 2 and supervise every block directly.

Writing. For every block, the fit is a linear readout of the MLP input by ridge regression on the pairs:

$$\hat{W}_l = \arg\min_{W}\ \mathbb{E}_{\mathcal{D}_l}\big\|\boldsymbol{g} - \boldsymbol{b}_l - W^\top P_l \boldsymbol{x}\big\|^2 + \lambda\|W\|_F^2, \tag{10}$$

where $\mathcal{D}_l$ are the pairs of block l over the replayed responses, $\boldsymbol{b}_l$ is the mean gap, P_l removes the mean direction of the inputs (Lemma A.5), and λ is the ridge weight. The fit is then reduced to the r directions that carry most of its predictions over all responses [3, 27], so that it occupies only the few neurons a block can spare (Appendix A.2). Each direction fills one key–value neuron of Eq. (7) with a readout $\boldsymbol{u}_{l,i}$ and a value $\boldsymbol{v}_{l,i}$, and one more neuron carries the mean gap $\boldsymbol{b}_l$. Their key $\boldsymbol{k}_l$ is fitted by ridge regression to tell the response positions where the gap is present from all other positions (Appendix C); as a function of the block's input, it carries the scope of an entry to unseen requests.

Strength. Finally, the strength of an edit, its step γ, is measured by a line search on held-out student responses [16]:

$$\gamma_j = \arg\min_{\gamma\in\Gamma}\ \delta_j(\gamma), \qquad \delta_j(\gamma) = \overline{\mathrm{KL}}\big(\pi_\theta^c \,\big\|\, \pi_{\gamma e^{(j)}}\big), \tag{11}$$

where $\gamma e^{(j)}$ is the refitted edit with step γ in Eq. (7), Γ is a small set of steps, and δ_j is the divergence that remains.

Proposition 3. *On the fitted responses, the r readout–value pairs of a block reproduce the predictions of the ridge fit as closely as any linear readout through r dimensions can. Proof.* Appendix A.2, via Lemmas A.5–A.7 [15].

5. Experiments

We evaluate BaseRSI through the following research questions (RQs). **RQ1:** How does BaseRSI compare with prompting, visual zoom-in, harness evolution, and fine-tuning in distribution? **RQ2:** Does its gain generalize to held-out (OOD) benchmarks and to other base models? **RQ3:** Does the base model alone keep the gain of its harness after internalization, and do rounds on its own responses add to a one-shot edit (Section 4.1)? **RQ4:** Does supervision inside the model carry more of this gain into the weights than supervision at its output (Section 4.2)? **RQ5:** Do a measured site and strength make the edit more effective than preset ones, and at what cost (Section 4.3)?

Setup. Five in-distribution (ID) benchmarks, BabyVision [11], BLINK [17], CharXiv [66], VisuLogic [71], and VLMs Are Biased [63], supply the development questions; five out-of-distribution (OOD) benchmarks, ChartQAPro [41], HallusionBench [22], PerceptionBench [34], SpatialViz [64], and VisualPuzzles [57], are held out from every loop, and each test set has 128 questions on unseen images. We compare with prompting and sampling, visual zoom-in,

Method	In-distribution (ID)						Held-out (OOD)					
	Baby Vision	BLINK	CharXiv	Visu Logic	VLMs Are Biased	Avg.	ChartQA Pro	Hallusion Bench	Perception Bench	Spatial Viz	Visual Puzzles	Avg.
Closed-source models												
Claude Haiku 5.5	$37.0_{\pm 2.7}$	$72.9_{\pm 2.0}$	$78.6_{\pm 3.0}$	$35.9_{\pm 2.7}$	$62.8_{\pm 0.5}$	$57.4_{\pm 0.3}$	$68.8_{\pm 0.8}$	$79.2_{\pm 1.2}$	$55.5_{\pm 2.1}$	$57.8_{\pm 0.8}$	$66.9_{\pm 2.7}$	$65.6_{\pm 0.5}$
GPT-6 Luna	$35.7_{\pm 0.9}$	$76.6_{\pm 1.4}$	$74.7_{\pm 0.5}$	$34.4_{\pm 1.6}$	$50.8_{\pm 1.4}$	$54.4_{\pm 0.5}$	$63.3_{\pm 0.8}$	$74.0_{\pm 2.0}$	$33.1_{\pm 4.0}$	$49.5_{\pm 3.5}$	$67.2_{\pm 3.4}$	$57.4_{\pm 0.9}$
Qwen3.8-27B (base)	$39.1_{\pm 1.6}$	$74.0_{\pm 1.2}$	$79.9_{\pm 3.3}$	$34.4_{\pm 2.7}$	$74.7_{\pm 1.6}$	$60.4_{\pm 1.3}$	$68.0_{\pm 0.8}$	$74.2_{\pm 0.8}$	$50.3_{\pm 2.7}$	$62.8_{\pm 0.9}$	$67.4_{\pm 1.8}$	$64.5_{\pm 0.5}$
Prompting and sampling												
CoT [67]	$45.8_{\pm 2.3}$	$72.1_{\pm 2.0}$	$81.0_{\pm 3.7}$	$38.0_{\pm 3.0}$	$75.3_{\pm 2.5}$	$62.4_{\pm 0.9}$	$70.1_{\pm 0.5}$	$76.8_{\pm 1.2}$	$55.2_{\pm 0.9}$	$65.1_{\pm 3.0}$	$68.5_{\pm 3.7}$	$67.1_{\pm 0.8}$
4-shot examples	$38.5_{\pm 3.7}$	$69.5_{\pm 2.8}$	$76.0_{\pm 1.2}$	$31.8_{\pm 0.9}$	$68.0_{\pm 1.4}$	$56.8_{\pm 1.5}$	$65.6_{\pm 1.6}$	$79.2_{\pm 0.9}$	$52.1_{\pm 1.6}$	$62.2_{\pm 3.3}$	$69.8_{\pm 1.2}$	$65.8_{\pm 1.4}$
Self-consistency [65]	$42.2_{\pm 1.7}$	$73.7_{\pm 1.7}$	$84.4_{\pm 1.1}$	$35.9_{\pm 1.7}$	$74.7_{\pm 2.8}$	$62.2_{\pm 0.4}$	$68.5_{\pm 1.1}$	$75.8_{\pm 1.1}$	$52.3_{\pm 1.1}$	$65.1_{\pm 1.1}$	$66.4_{\pm 2.2}$	$65.6_{\pm 0.4}$
Self-Refine [39]	$39.8_{\pm 6.2}$	$72.9_{\pm 2.0}$	$82.3_{\pm 1.2}$	$42.4_{\pm 5.2}$	$73.4_{\pm 3.1}$	$62.2_{\pm 3.2}$	$68.2_{\pm 1.2}$	$80.5_{\pm 3.9}$	$57.3_{\pm 3.5}$	$65.1_{\pm 1.6}$	$71.1_{\pm 2.1}$	$68.4_{\pm 1.6}$
Visual zoom-in												
ZoomEye [56]	$31.3_{\pm 2.1}$	$74.2_{\pm 3.6}$	$76.8_{\pm 3.0}$	$37.5_{\pm 0.8}$	$60.4_{\pm 2.0}$	$56.0_{\pm 1.7}$	$65.9_{\pm 0.9}$	$74.2_{\pm 1.6}$	$48.7_{\pm 2.3}$	$62.5_{\pm 2.1}$	$68.5_{\pm 2.4}$	$64.0_{\pm 0.4}$
Two-stage zoom-in	$47.9_{\pm 3.9}$	$73.4_{\pm 4.7}$	$84.9_{\pm 2.7}$	$38.5_{\pm 5.5}$	$72.4_{\pm 4.8}$	$63.4_{\pm 0.4}$	$69.8_{\pm 1.2}$	$79.9_{\pm 0.9}$	$57.8_{\pm 2.3}$	$64.6_{\pm 0.9}$	$70.1_{\pm 0.5}$	$68.4_{\pm 0.5}$
Harness evolution												
GEPA [1]	$41.9_{\pm 3.0}$	$73.7_{\pm 4.0}$	$81.8_{\pm 1.6}$	$40.4_{\pm 4.0}$	$74.7_{\pm 1.6}$	$62.5_{\pm 0.3}$	$70.8_{\pm 1.2}$	$79.4_{\pm 2.0}$	$52.6_{\pm 1.2}$	$62.0_{\pm 2.7}$	$68.8_{\pm 1.6}$	$66.7_{\pm 0.2}$
ACE [79]	$35.7_{\pm 0.5}$	$70.8_{\pm 1.2}$	$81.3_{\pm 1.4}$	$36.5_{\pm 1.8}$	$72.1_{\pm 3.2}$	$59.3_{\pm 1.2}$	$68.0_{\pm 1.4}$	$79.2_{\pm 4.3}$	$50.3_{\pm 2.0}$	$61.2_{\pm 2.5}$	$66.4_{\pm 5.4}$	$65.0_{\pm 0.9}$
Meta-Harness [32]	$45.1_{\pm 3.7}$	$73.4_{\pm 0.8}$	$84.1_{\pm 1.6}$	$40.9_{\pm 3.5}$	$71.9_{\pm 0.8}$	$63.1_{\pm 1.1}$	$70.1_{\pm 2.3}$	$77.3_{\pm 1.6}$	$54.9_{\pm 3.9}$	$65.6_{\pm 2.8}$	$70.3_{\pm 1.6}$	$67.7_{\pm 1.4}$
SkillOpt [73]	$41.7_{\pm 3.3}$	$71.6_{\pm 1.2}$	$81.8_{\pm 0.9}$	$37.8_{\pm 1.2}$	$75.8_{\pm 2.8}$	$61.7_{\pm 0.2}$	$70.8_{\pm 0.5}$	$76.8_{\pm 1.6}$	$54.2_{\pm 3.0}$	$65.9_{\pm 3.9}$	$69.0_{\pm 3.5}$	$67.3_{\pm 1.2}$
VideoHarness-RSI [70]	$40.9_{\pm 0.5}$	$75.0_{\pm 1.6}$	$80.7_{\pm 0.5}$	$37.8_{\pm 5.2}$	$70.1_{\pm 3.5}$	$60.9_{\pm 1.3}$	$69.0_{\pm 0.5}$	$76.0_{\pm 1.6}$	$52.1_{\pm 3.2}$	$64.1_{\pm 1.4}$	$70.1_{\pm 2.3}$	$66.3_{\pm 0.4}$
RRSI [68]	$42.4_{\pm 3.0}$	$75.5_{\pm 0.5}$	$81.3_{\pm 2.1}$	$42.7_{\pm 2.5}$	$83.6_{\pm 3.6}$	$65.1_{\pm 1.3}$	$69.8_{\pm 0.9}$	$79.9_{\pm 1.6}$	$54.9_{\pm 5.3}$	$65.4_{\pm 4.4}$	$71.4_{\pm 4.0}$	$68.3_{\pm 2.0}$
Fine-tuning												
SFT	$39.3_{\pm 1.6}$	$72.4_{\pm 1.6}$	$79.7_{\pm 2.3}$	$37.0_{\pm 1.2}$	$76.3_{\pm 3.0}$	$60.9_{\pm 0.7}$	$68.8_{\pm 2.7}$	$75.3_{\pm 2.3}$	$57.0_{\pm 3.6}$	$63.3_{\pm 0.8}$	$61.2_{\pm 2.0}$	$65.1_{\pm 1.0}$
GRPO [54]	$45.6_{\pm 3.5}$	$68.8_{\pm 1.4}$	$79.2_{\pm 1.2}$	$35.7_{\pm 2.5}$	$76.8_{\pm 1.6}$	$61.2_{\pm 0.8}$	$67.7_{\pm 1.2}$	$75.5_{\pm 2.7}$	$52.6_{\pm 2.7}$	$54.2_{\pm 2.7}$	$69.3_{\pm 3.9}$	$63.9_{\pm 0.7}$
PixelProof [50]	$39.6_{\pm 3.2}$	$75.0_{\pm 2.7}$	$81.8_{\pm 2.0}$	$39.3_{\pm 2.0}$	$72.7_{\pm 2.1}$	$61.7_{\pm 0.2}$	$68.2_{\pm 2.4}$	$75.3_{\pm 1.8}$	$54.7_{\pm 2.1}$	$60.9_{\pm 0.8}$	$67.4_{\pm 2.7}$	$65.3_{\pm 0.8}$
LBS-27B [74]	$49.0_{\pm 1.7}$	$71.9_{\pm 1.1}$	$80.5_{\pm 4.4}$	$41.9_{\pm 2.8}$	$70.1_{\pm 3.9}$	$62.7_{\pm 2.1}$	$67.7_{\pm 2.8}$	$75.8_{\pm 1.1}$	$52.3_{\pm 3.3}$	$61.5_{\pm 6.1}$	$64.1_{\pm 2.2}$	$64.3_{\pm 0.2}$
Harness internalization												
BaseRSI (ours)	$\mathbf{60.7}_{\pm 2.4}$	$\mathbf{76.6}_{\pm 1.4}$	$\mathbf{87.5}_{\pm 1.4}$	$\mathbf{44.8}_{\pm 2.0}$	$\mathbf{91.1}_{\pm 1.2}$	$\mathbf{72.1}_{\pm 0.4}$	$\mathbf{74.0}_{\pm 1.6}$	$\mathbf{81.3}_{\pm 2.1}$	$\mathbf{59.1}_{\pm 2.7}$	$\mathbf{66.1}_{\pm 2.5}$	$\mathbf{72.9}_{\pm 1.6}$	$\mathbf{70.7}_{\pm 0.4}$
Δ over base model	+21.6	+2.6	+7.6	+10.4	+16.4	+11.7	+6.0	+7.0	+8.9	+3.4	+5.5	+6.1

Table 1. Main results on five in-distribution (ID) and five held-out (OOD) benchmarks: accuracy (%), mean$_{\pm \text{std}}$ over three runs. All methods but the two closed-source reference models use Qwen3.8-27B with the same reasoning effort, token budget, and evaluator; BaseRSI (shaded) is the edited base model answering the original requests alone. Δ: gain over the base model, from unrounded means.

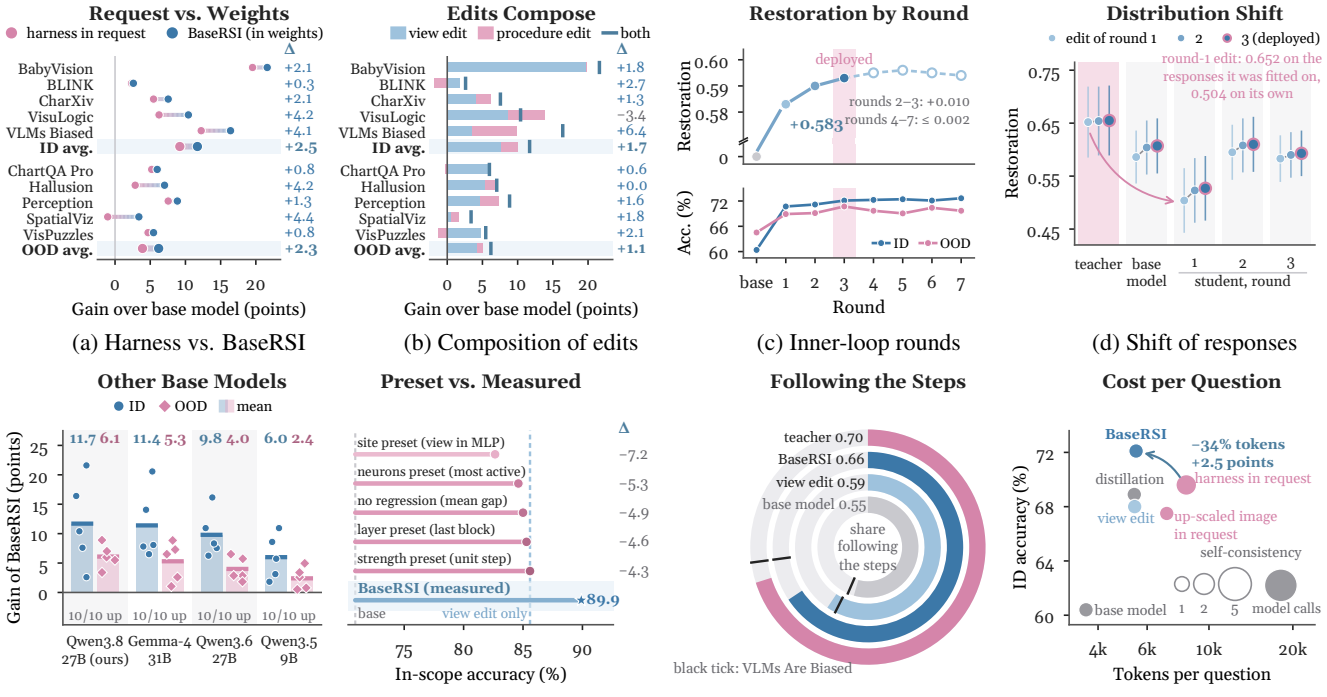
harness evolution, and fine-tuning on the same base model, Qwen3.8-27B, and list two closed-source models for reference [4, 47] (Table 1). Over eight generations the base model proposes 269 candidates from its own failures, and the accepted entries, a finer view and a counting procedure, are internalized as the view edit and the procedure edit. We report accuracy over three runs and the restoration of an edit, the fraction of the divergence from the teacher that it removes on held-out responses (Appendix D).

Main results (RQ1, RQ2). BaseRSI has the highest mean accuracy on all ten benchmarks and raises the ID and OOD averages of the base model by 11.7 and 6.1 points (Table 1). Fine-tuning on the same development questions (SFT, GRPO) barely moves these averages, while the harness that the outer loop verifies lifts them, so the gain lies in what the harness learns rather than in fitting these questions. Prompting, zoom-in, and harness evolution act through the request, whereas BaseRSI writes this gain into the weights, where it is largest on the benchmarks of the internalized entries (Table 2). The acceptance test carries the gain to the held-out benchmarks: the finer view, which the base model proposes from its own failures, keeps its gain with any one benchmark left out, and testing the macro average alone lowers the OOD average (Figure 5a). The same loop, run from the original weights of Gemma-4-31B [61], Qwen3.6-27B, and Qwen3.5-9B, internalizes entries that differ in part from those of Qwen3.8-27B and raises every benchmark score of each model (Figure 4e, Appendix F).

Recursive Harness Internalization (RQ3). The edited base model on original requests exceeds the same model with its harness in the request on all ten benchmarks and answers in one call (Figures 4a and 4h). Inside the entry's scope the two are level (Appendix E), and the surplus comes from the questions that the router leaves out: there the teacher's request has no procedure, while the student carries it in its weights behind a key that reads its own state, the advantage of a student trained with privileged information [36, 62]. Its edits compose without refitting because each is fitted on the base weights at its own site, and together they exceed the sum of their gains (Figure 4b). At the same view, the procedure edit gains on every benchmark (Table 2), while the up-scaled image alone in the request leaves the base model near the view edit (Figure 4h). Rounds on the student's own responses make the edit hold on its own outputs, and three suffice (Figure 4c): a one-shot edit fitted on the teacher's responses loses part of its effect

Variant	In-distribution (ID)						Held-out (OOD)					
	Baby Vision	BLINK	CharXiv	Visu Logic	VLMs Are Biased	Avg.	ChartQA Pro	Hallusion Bench	Perception Bench	Spatial Viz	Visual Puzzles	Avg.
Qwen3.8-27B (base)	39.1	74.0	79.9	34.4	74.7	60.4	68.0	74.2	50.3	62.8	67.4	64.5
view edit only	58.9	75.8	84.1	43.0	78.4	68.0	73.7	79.7	54.9	63.3	72.1	68.8
Recursive Harness Internalization (Section 4.1)												
harness kept in the request	58.6	76.3	85.4	40.6	87.0	69.6	73.2	77.1	57.8	61.7	72.1	68.4
one round	59.6	75.5	85.9	44.5	88.0	70.7	73.4	78.6	58.9	62.2	71.4	68.9
teacher's responses only	60.2	74.7	85.9	44.0	89.3	70.8	72.7	79.9	58.1	63.0	68.0	68.3
ReGap (Section 4.2)												
output only (distillation)	62.2	75.8	83.3	42.4	80.7	68.9	72.9	78.1	55.2	65.1	68.0	67.9
last block only	59.4	75.8	84.6	43.8	78.9	68.5	72.9	77.6	55.5	61.2	70.6	67.6
direct difference as signal	59.6	74.7	82.3	40.9	89.3	69.4	73.4	80.5	54.7	64.3	71.6	68.9
ReKey (Section 4.3)												
mean gap, no regression	60.2	76.0	84.4	42.4	79.7	68.5	74.0	78.4	57.0	63.5	72.7	69.1
site preset (view in MLP)	38.3	73.4	81.3	38.0	81.8	62.6	67.7	76.8	52.6	61.2	67.7	65.2
strength preset (1.0)	59.4	74.0	84.1	42.4	82.3	68.4	72.9	78.6	56.3	63.8	71.4	68.6
constant key	60.4	76.3	85.4	43.2	89.8	71.0	71.9	80.7	56.5	63.3	72.7	69.0
Null edits on top of the view edit												
random readouts	60.7	75.8	86.7	41.4	75.8	68.1	72.7	76.6	56.8	61.7	72.9	68.1
random values	60.2	75.8	84.1	41.4	77.9	67.9	73.4	77.3	58.3	63.5	72.9	69.1
BaseRSI	**60.7**	**76.6**	**87.5**	**44.8**	**91.1**	**72.1**	**74.0**	**81.3**	**59.1**	**66.1**	**72.9**	**70.7**

Table 2. Ablations of the three components of BaseRSI, accuracy (%), mean of three runs. Each group replaces one choice of the section it names; procedure-edit variants are deployed on the view edit (row 2). *Direct difference* uses the state difference of two runs as the signal; null edits keep the procedure edit's neurons, key, and norms but randomize its readouts or values (Appendix D).



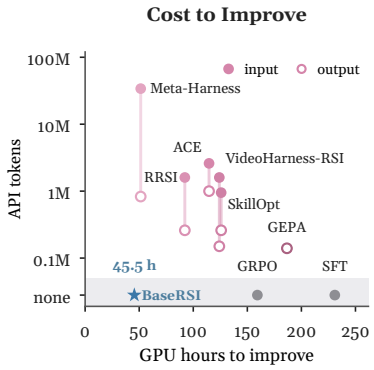
(a) Harness vs. BaseRSI (b) Composition of edits (c) Inner-loop rounds (d) Shift of responses

(e) Cross-model gain (f) Measured choices (g) Step adherence (h) Inference cost

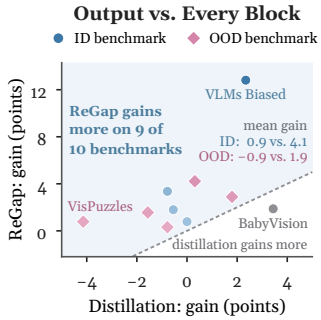
Figure 4. **Internalization, generalization, and process.** (a) Gain over the base model with the harness in the request and with BaseRSI; Δ: their difference. (b) Gains of each edit alone (stacked) and of both; Δ: both minus the sum. (c) Restoration (share of the teacher divergence removed) on 28 held-out responses of the round-3 student, and test accuracy, from the base model to round 7. (d) Restoration by each round's edit on responses of the teacher, the base model, and each round's student; error bars: standard error. (e) Gain of BaseRSI per benchmark on four base models; bars: means. (f) In-scope accuracy with one choice of BaseRSI preset instead of measured; Δ: against BaseRSI. (g) Share of responses that follow the steps; tick: VLMs Are Biased. (h) Tokens against ID accuracy; area: model calls.

Setting	Restor.	In sc.	ID	OOD	Setting	Restor.	In sc.	ID	OOD
base model	–	–	60.4	64.5	view edit only	–	+14.7	68.0	68.8
Replayed responses (ours: three rounds, own responses)					*Signal and positions (ours: per-block gap, up to 512 positions)*				
one round	0.583	+16.3	70.7	68.9	last-block state	0.271	+12.1	68.6	67.7
two rounds	0.590	+17.6	71.2	69.1	tail + 64 positions	0.575	+16.3	70.7	68.4
base model's only	0.515	+17.3	70.5	67.9	16 responses	0.541	+17.3	70.8	68.0
Acceptance (ours: leave one benchmark out)					64 responses	0.578	+19.0	70.4	69.5
macro average only	0.545	+16.7	71.1	68.8	*Readouts and neurons (ours: fitted, least active)*				
Blocks written (ours: all 64)					mean gap only	0.075	+14.1	68.5	69.1
last block	0.005	+14.4	68.5	67.6	regression-matrix directions	0.576	+16.7	70.5	68.0
last 16	0.304	+14.1	68.9	67.4	random directions	0.250	+16.0	69.1	69.3
first 32	0.055	+16.7	69.5	68.2	random neurons	0.474	+17.3	71.0	67.9
48 linear-attention	0.093	+14.4	69.4	68.8	most active neurons	−2.119	+13.7	65.1	65.3
16 full-attention	0.580	+14.7	70.0	67.1	**BaseRSI**	**0.593**	**+19.0**	**72.1**	**70.7**

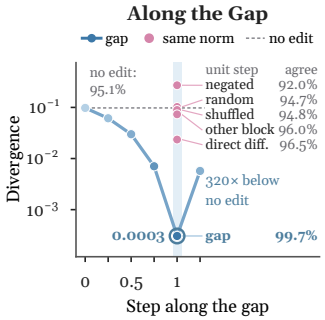
(a) Fine-grained ablations of the procedure edit



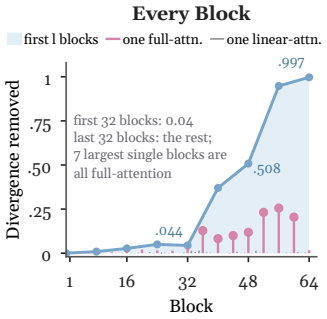
(b) Cost of improving



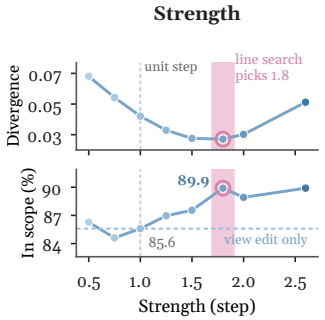
(c) Inside vs. output



(d) Exactness of the gap



(e) Each block



(f) Measured strength

Figure 5. **Fine-grained ablations, cost, and mechanisms of ReGap and ReKey (RQ3–RQ5).** (a) Each row changes one choice of BaseRSI (setting in parentheses); Restor.: restoration; In sc.: gain (points) on the in-scope validation questions. (b) GPU hours and API tokens (filled: input, open: output). (c) Gain over the view edit per benchmark, distillation (x) against ReGap (y). (d) Divergence left after a step along the gap and after a unit step along other directions of equal norm; percentages: next-token agreement with the teacher. (e) Divergence removed by one block (stems) and by the first l blocks (area). (f) Held-out divergence and in-scope accuracy per step.

on its own responses (Figure 4d), and the rounds add to it, while as many refits on the teacher's add nothing (Table 2), so the gain comes from whose responses are replayed.

ReGap (RQ4). Supervising every block with the gap carries the entry into the weights: in its scope the edited model matches the entry in the request (Figure 5a, Appendix E), and it follows the steps of the entry about as often as the teacher (Figure 4g). A unit step along the gap removes the divergence from the teacher, while gaps that are negated, shuffled, or taken from another block leave it (Figure 5d), and the direct difference of the two runs' states as the signal scores lower (Table 2). Supervision at the output alone, distilling the same teacher, gains less than half of what ReGap gains, also when it trains only the neurons that BaseRSI writes (Figure 5c, Appendix D), since its signal travels back through every block from a few hundred responses. On these neurons, gradient training toward the per-block gaps comes close to the closed-form solve, while the output loss restores nothing (Appendix E), so the gain comes from the signal. Null edits with the same neurons, key, and norms leave the ID average unchanged (Table 2), and the divergence falls only as all blocks' gaps are added (Figure 5e), so the fitted readouts of every block carry the gain.

ReKey (RQ5). Each measured quantity decides part of the edit, and every preset site, layer, or strength stays below the measured one (Figure 4f). The measured site sends the view entry to the patch embedding; in MLP neurons it loses most of its gain (Table 2). The regression makes the edit a function of what each block reads: the mean gap alone adds little, and readouts along the principal directions of the fit supply the gain (Table 2, Figure 5a). The measured strength offsets the shrinkage of the fit, and the line search picks the step with the least held-out divergence (Figure 5f). The fitted key decides where the edit acts: the edited model cites the procedure on 32% of in-scope responses and 0.5% of the others (Appendix E), and a constant key lowers both averages (Table 2). Internalization runs forward passes and one linear solve per block, with no backward pass or API call, in 45.5 GPU hours, the least of all methods (Figure 5b); the procedure edit rewrites 0.23% of the weights and leaves the remaining weights bit-identical.

6. Conclusion

BaseRSI unifies harness evolution, weight editing, and internalization in one recursive loop with ReGap and ReKey, and the base model alone improves on every benchmark.

LLM Usage Statement

Claude (Anthropic) was used to polish and edit the writing and for LaTeX typesetting. Qwen3.8-27B is part of the method and is described in Sections 4 and 5. The authors reviewed all content and take full responsibility for it.

References

[1] Lakshya A Agrawal, Shangyin Tan, Dilara Soylu, Noah Ziems, Rishi Khare, Krista Opsahl-Ong, Arnav Singhvi, Herumb Shandilya, Michael J Ryan, Meng Jiang, et al. Gepa: Reflective prompt evolution can outperform reinforcement learning. In *International Conference on Learning Representations*, pages 8479–8565, 2026. 6

[2] James A Anderson. A simple neural network generating an interactive memory. *Mathematical biosciences*, 14(3-4): 197–220, 1972. 5

[3] Theodore Wilbur Anderson. Estimating linear restrictions on regression coefficients for multivariate normal distributions. *The Annals of Mathematical Statistics*, pages 327–351, 1951. 5

[4] Anthropic. Introducing Claude Haiku 5.5. https://www.anthropic.com/claude-haiku-5-5, 2026. 6

[5] Andy Arditi, Oscar Obeso, Aaquib Syed, Daniel Paleka, Nina Panickssery, Wes Gurnee, and Neel Nanda. Refusal in language models is mediated by a single direction. *Advances in Neural Information Processing Systems*, 37: 136037–136083, 2024. 3

[6] Bhiman Kumar Baghel, Emma Jordan, Zheyuan Ryan Shi, and Xiang Lorraine Li. Resolving underedit & overedit with iterative & neighbor-assisted model editing. *arXiv preprint arXiv:2503.11895*, 2025. 2

[7] David Bau, Steven Liu, Tongzhou Wang, Jun-Yan Zhu, and Antonio Torralba. Rewriting a deep generative model. In *European conference on computer vision*, pages 351–369. Springer, 2020. 5

[8] Ali Behrouz, Farnoosh Hashemi, Adel Javanmard, and Vahab Mirrokni. Language models need sleep: Learning to self-modify and consolidate memories. *arXiv preprint arXiv:2606.03979*, 2026. 2

[9] Lucas Beyer, Pavel Izmailov, Alexander Kolesnikov, Mathilde Caron, Simon Kornblith, Xiaohua Zhai, Matthias Minderer, Michael Tschannen, Ibrahim Alabdulmohsin, and Filip Pavetic. Flexivit: One model for all patch sizes. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14496–14506. IEEE, 2023. 5

[10] Brian K Chen, Tianyang Hu, Hui Jin, Hwee Kuan Lee, and Kenji Kawaguchi. Exact conversion of in-context learning to model weights in linearized-attention transformers. *arXiv preprint arXiv:2406.02847*, 2024. 5

[11] Liang Chen, Weichu Xie, Yiyan Liang, Hongfeng He, Hans Zhao, Zhibo Yang, Zhiqi Huang, Haoning Wu, Haoyu Lu, Yiping Bao, et al. Babyvision: Visual reasoning beyond language. *arXiv preprint arXiv:2601.06521*, 2026. 5

[12] Roi Cohen, Eden Biran, Ori Yoran, Amir Globerson, and Mor Geva. Evaluating the ripple effects of knowledge editing in language models. *Transactions of the Association for Computational Linguistics*, 12:283–298, 2024. 3

[13] Benoit Dherin, Michael Munn, Hanna Mazzawi, Michael Wunder, and Javier Gonzalvo. Learning without training: The implicit dynamics of in-context learning. *arXiv preprint arXiv:2507.16003*, 2025. 3, 4

[14] Ali Ebrahimpour-Boroojeny, Xiaocong Yang, Samuel Schapiro, Zahra Borhani, Francisco R Ortega, Yujia Zheng, and ChengXiang Zhai. Resisting harmful instructions through training-free weight editing. *arXiv preprint arXiv:2610.21189*, 2026. 3

[15] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936. 5

[16] Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001. 4, 5

[17] Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A Smith, Wei-Chiu Ma, and Ranjay Krishna. Blink: Multimodal large language models can see but not perceive. In *European Conference on Computer Vision*, pages 148–166. Springer, 2024. 5

[18] Wenbo Gao, Zhaomou Song, Zhiyuan Ji, Renxi Liu, Xing Li, Xianzhi Yu, Xiaoguang Li, James Chung-wai Cheung, Weizhe Lin, and Yaoyuan Wang. Experience funnel: A state-policy alternating loop for self-evolving agents. *arXiv preprint arXiv:2609.08919*, 2026. 1, 2

[19] Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. Transformer feed-forward layers are key-value memories. In *Proceedings of the 2021 conference on empirical methods in natural language processing*, pages 5484–5495, 2021. 5

[20] Adrian Goldwaser, Michael Munn, Javier Gonzalvo, and Benoit Dherin. Equivalence of context and parameter updates in modern transformer blocks. *arXiv preprint arXiv:2511.17864*, 2025. 3, 4

[21] Jia-Chen Gu, Hao-Xiang Xu, Jun-Yu Ma, Pan Lu, Zhen-Hua Ling, Kai-Wei Chang, and Nanyun Peng. Model editing harms general abilities of large language models: Regularization to the rescue. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16801–16819, 2024. 2, 3

[22] Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, et al. Hallusionbench: an advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. In *2024 IEEE/CVF Confer-*

ence on Computer Vision and Pattern Recognition (CVPR), pages 14375–14385. IEEE, 2024. 5

[23] Peter Hase, Mohit Bansal, Been Kim, and Asma Ghandeharioun. Does localization inform editing? surprising differences in causality-based localization vs. knowledge editing in language models. *Advances in Neural Information Processing Systems*, 36:17643–17668, 2023. 2

[24] Prannay Hebbar, Yogendra Manawat, Samuel Verboomen, Alesia Ivanova, Selvam Palanimalai, Kunal Bhatia, and Vignesh Baskaran. Sia: Self improving ai with harness & weight updates. *arXiv preprint arXiv:2605.27276*, 2026. 2

[25] Sture Holm. A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, pages 65–70, 1979. 3

[26] Zeyu Huang, Yikang Shen, Xiaofeng Zhang, Jie Zhou, Wenge Rong, and Zhang Xiong. Transformer-patcher: One mistake worth one neuron. *arXiv preprint arXiv:2301.09785*, 2023. 1, 2

[27] Alan Julian Izenman. Reduced-rank regression for the multivariate linear model. *Journal of multivariate analysis*, 5(2): 248–264, 1975. 5

[28] Yubo Jiang, Fengying Xie, Zhiguo Jiang, and Haopeng Zhang. Distill skills into weights, not prompts: Abstract skills as privileged signals for on-policy self-distillation. *arXiv preprint arXiv:2608.09826*, 2026. 2

[29] Teuvo Kohonen. Correlation matrix memories. *IEEE transactions on computers*, 100(4):353–359, 1972. 5

[30] Dharshan Kumaran, Demis Hassabis, and James L McClelland. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in cognitive sciences*, 20(7):512–534, 2016. 3

[31] Hyunin Lee, Jinglue Xu, Jeffrey Seely, Donghyun Lee, Matei Zaharia, and Yujin Tang. Recursive harness self-improvement. *arXiv preprint arXiv:2607.15524*, 2026. 1, 2, 3

[32] Yoonho Lee, Roshen Nair, Qizheng Zhang, Kangwook Lee, Omar Khattab, and Chelsea Finn. Meta-harness: End-to-end optimization of model harnesses. *arXiv preprint arXiv:2603.28052*, 2026. 1, 2, 6

[33] Changhao Li, Rushi Qiang, Jiawei Huang, Chenxiao Gao, Chao Zhang, Niao He, and Bo Dai. Revisiting dagger in the era of llm-agents. *arXiv preprint arXiv:2605.12913*, 2026. 4

[34] Zichao Lin, Yifeng Xie, Bowen Qu, Haiming Wang, Jia Li, Haoning Wu, Yuhao Dong, Zuhao Yang, Jinguo Zhu, Haoyu Lu, et al. Perceptionbench: Evaluating atomic visual perception in multimodal large language models. *arXiv preprint arXiv:2607.24957*, 2026. 5

[35] Kaiyuan Liu, Zezhou Huang, Janvijay Singh, and Asli Celikyilmaz. What happens when the harness comes and goes? a controlled study of on-policy harness self-distillation. In *NeurIPS 2026 Workshop on Transitioning from Pre-training to Post-training*, 2026. 2

[36] David Lopez-Paz, Léon Bottou, Bernhard Schölkopf, and Vladimir Vapnik. Unifying distillation and privileged information. *arXiv preprint arXiv:1511.03643*, 2015. 3, 6

[37] Bingjun Luo, Jialin Guo, and Siqi Li. Video-rsi: Recursive self-improvement of video understanding agents via harness evolution. *arXiv preprint arXiv:2609.37950*, 2026. 2

[38] Haochen Luo, Yi Huang, Sichun Luo, Fengyuan Liu, Lei Li, Zefa Hu, Junlan Feng, and Qi Liu. Harness-aware self-evolving: Co-evolving model weights, harness, and task solutions. *arXiv preprint arXiv:2607.03935*, 2026. 1, 2

[39] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. *Advances in neural information processing systems*, 36:46534–46594, 2023. 6

[40] Llew Mason, Jonathan Baxter, Peter Bartlett, and Marcus Frean. Boosting algorithms as gradient descent. *Advances in neural information processing systems*, 12, 1999. 4

[41] Ahmed Masry, Mohammed Saidul Islam, Mahir Ahmed, Aayush Bajaj, Firoz Kabir, Aaryaman Kartha, Md Tahmid Rahman Laskar, Mizanur Rahman, Shadikur Rahman, Mehrad Shahmohammadi, et al. Chartqapro: A more diverse and challenging benchmark for chart question answering. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 19123–19151, 2025. 5

[42] Hanna Mazzawi, Benoit Dherin, Michael Munn, Adrian Goldwaser, Michael Wunder, and Javier Gonzalvo. Transmuting prompts into weights. *arXiv preprint arXiv:2510.08734*, 2025. 3, 4

[43] James L McClelland, Bruce L McNaughton, and Randall C O'Reilly. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419, 1995. 3

[44] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *Advances in neural information processing systems*, 35:17359–17372, 2022. 1, 2, 5

[45] Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. Mass-editing memory in a transformer. *arXiv preprint arXiv:2210.07229*, 2022. 1, 2

[46] Michael Merzin, Yehonatan Elisha, Oren Barkan, Yonatan Toib, and Noam Koenigstein. In-context compilation: Encoding context-induced behavioral shift into model weights. In *Findings of the Association for Computational Linguistics: EMNLP 2026*, pages 20777–20796, 2026. 1, 3

[47] OpenAI. Introducing GPT-6 Sol and Luna. https://openai.com/index/introducing-gpt-6-sol-and-luna/, 2026. 6

[48] Haowen Pan, Xiaozhi Wang, Yixin Cao, Zenglin Shi, Xun

Yang, Juanzi Li, and Meng Wang. Precise localization of memories: A fine-grained neuron-level knowledge editing technique for llms. In *International Conference on Learning Representations*, pages 482–502, 2025. 2

[49] Judea Pearl. Direct and indirect effects. In *Probabilistic and causal inference: the works of Judea Pearl*, pages 373–392. 2022. 4, 5

[50] PixelProof. PixelProof: Visual question generation through forward–inverse agreement. https://github.com/ProjectPixelProof/PixelProof, 2026. 6

[51] Tao Ren, Weiyao Luo, Hui Yang, Rongzhi Zhu, Xiang Huang, Yuchuan Wu, Bingxue Chou, Jieping Ye, Jiafeng Liang, Yongbin Li, et al. Scaling self-evolving agents via parametric memory. *arXiv preprint arXiv:2606.04536*, 2026. 2

[52] Stéphane Ross and Drew Bagnell. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 661–668. JMLR Workshop and Conference Proceedings, 2010. 4

[53] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011. 4

[54] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024. 6

[55] Noam Shazeer. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*, 2020. 3

[56] Haozhan Shen, Kangjia Zhao, Tiancheng Zhao, Ruochen Xu, Zilun Zhang, Mingwei Zhu, and Jianwei Yin. Zoomeye: Enhancing multimodal llms with human-like zooming capabilities through tree-based image exploration. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 6613–6629, 2025. 6

[57] Yueqi Song, Tianyue Ou, Yibo Kong, Zecheng Li, Graham Neubig, and Xiang Yue. Visualpuzzles: Decoupling multimodal reasoning evaluation from domain knowledge. *arXiv preprint arXiv:2504.10342*, 2025. 5

[58] Jiaqi Su, Guibin Zhang, Jie-Jing Shao, Cong Pang, Tiankuo Yao, Ivor Tsang, Haiyan Yin, and Shuicheng Yan. Coevo-rsi: When harnesses evolve, which models should follow? *arXiv preprint arXiv:2610.25078*, 2026. 2

[59] Rikiya Takehi, Jayden Teoh, Alex Zhang, and Omar Khattab. Sparks of weight-space programming: Llms can natively read & write neural network weights. *MIT CSAIL research blog. Published November*, 2:2026, 2026. 3

[60] Zihan Tan, Leixin Sun, Zitong Shi, Yitao Liu, Jiajun Wu, Nathaniel Brooks, Jiaru Qian, Xiaoran Shang, Suyuan Huang, Yi Ding, et al. Metarsi/rsi2: A meta-recursive self-improving system for recursive self-improving systems themselves. *arXiv preprint arXiv:2609.06396*, 2026. 1, 2

[61] Gemma Team, Sherif El Abd, Vaibhav Aggarwal, Robin Algayres, Alek Andreev, Olivier Bachem, Ian Ballantyne, Cormac Brick, Victor Cărbune, Michelle Casbon, et al. Gemma 4 technical report. *arXiv preprint arXiv:2607.02770*, 2026. 6

[62] Vladimir Vapnik and Akshay Vashist. A new learning paradigm: Learning using privileged information. *Neural networks*, 22(5-6):544–557, 2009. 3, 6

[63] An Vo, Khai-Nguyen Nguyen, Mohammad Reza Taesiri, Vy Tuong Dang, Anh Totti Nguyen, and Daeyoung Kim. Vision language models are biased. *arXiv preprint arXiv:2505.23941*, 2025. 5

[64] Siting Wang, Minnan Pei, Luoyang Sun, Cheng Deng, Yuchen Li, Kun Shao, Zheng Tian, Haifeng Zhang, and Jun Wang. Spatialviz-bench: A cognitively-grounded benchmark for diagnosing spatial visualization in mllms. In *International Conference on Learning Representations*, pages 47210–47279, 2026. 5

[65] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022. 6

[66] Zirui Wang, Mengzhou Xia, Luxi He, Howard Chen, Yitao Liu, Richard Zhu, Kaiqu Liang, Xindi Wu, Haotian Liu, Sadhika Malladi, et al. Charxiv: Charting gaps in realistic chart understanding in multimodal llms. *Advances in Neural Information Processing Systems*, 37:113569–113697, 2024. 5

[67] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022. 6

[68] Peng Xia, Rujun Han, Zifeng Wang, Yanfei Chen, Yufan Zhuang, Yoonho Lee, Chengsong Huang, Han Yu, Zhongying CuiZhu, Yifei Ming, et al. Rrsi: Regularized recursive self-improvement of agent harnesses. *arXiv preprint arXiv:2609.24972*, 2026. 1, 2, 6

[69] Chejian Xu, Shujian Zhang, Wenxuan Zhou, John Lambert, Beck Chen, Bo Li, Andrew Hard, and Rajiv Mathews. Evo ^{2}: Evolving the evolver through weight-harness co-evolution. *arXiv preprint arXiv:2610.13025*, 2026. 2

[70] Guoyang Xu and Hao Chen. Videoharness-rsi: Recursive harness self-improvement for long-video understanding with frozen vision-language models. *arXiv preprint arXiv:2608.24302*, 2026. 1, 2, 6

[71] Weiye Xu, Jiahao Wang, Weiyun Wang, Zhe Chen, Wengang Zhou, Aijun Yang, Lewei Lu, Houqiang Li, Xiaohua Wang,

Xizhou Zhu, et al. Visulogic: A benchmark for evaluating visual reasoning in multi-modal large language models. In *International Conference on Learning Representations*, pages 25966–26003, 2026. 5

[72] Bangji Yang, Jiajun Fan, Zhaoyu Zhu, Hongbo Ma, and Ge Liu. Training-free on-policy distillation. *arXiv preprint arXiv:2610.22236*, 2026. 1, 3

[73] Yifan Yang, Ziyang Gong, Weiquan Huang, Qihao Yang, Ziwei Zhou, Zisu Huang, Yan Li, Xuemei Gao, Qi Dai, Bei Liu, et al. Skillopt: Executive strategy for self-evolving agent skills. *arXiv preprint arXiv:2605.23904*, 2026. 6

[74] Ruilin Yao, Xinyu Qiu, Tianyu Zou, Yi Rong, Shengwu Xiong, Shuheng Shen, and Changhua Meng. Learning to look beyond screens: Transferable visual observation from gui grounding. *arXiv preprint arXiv:2610.37093*, 2026. 6

[75] Haoran Ye, Yuxing Lu, Haonan Dong, Zhaochen Su, and Guojie Song. Harness-zero: Harness distillation via agent-as-harness. *arXiv preprint arXiv:2609.24974*, 2026. 1, 2

[76] Zhou Yu, Bin Bi, Shiva Kumar Pentyala, Shubham Mehrotra, Sougata Chaudhuri, Shilpa Bhagavath, Zeyuan Chen, Ran Xu, Phil Mui, James Zhu, et al. Co-evolving harnesses and models: On-policy correction helps weaker models catch up where imitation fails. *arXiv preprint arXiv:2609.09134*, 2026. 1, 2

[77] Yuchen Zeng, Giorgio Severi, Blake Bullwinkel, Eric Horvitz, and Dimitris Papailiopoulos. Lazare, veni foras: Reprogramming dormant neurons to backdoor ai agents. *arXiv preprint arXiv:2611.15118*, 2026. 3

[78] Mengqi Zhang, Xiaotian Ye, Qiang Liu, Pengjie Ren, Zhumin Chen, et al. Uncovering overfitting in large language model editing. In *International Conference on Learning Representations*, pages 79550–79568, 2025. 3

[79] Qizheng Zhang, Changran Hu, Shubhangi Upasani, Boyuan Ma, Fenglu Hong, Vamsidhar Kamanuru, Jay Rainton, Chen Wu, Mengmeng Ji, Hanchen Li, et al. Agentic context engineering: Evolving contexts for self-improving language models. In *International Conference on Learning Representations*, pages 86069–86100, 2026. 6

[80] Fanyu Zhao, Ruike Cao, Liang Dong, Fugen Yao, Jian Xu, Guanjun Jiang, Han Zhang, Yifei Zhao, and Yinsheng Li. Rpmem: Learning long-term recurrent parametric memory across sessions for llm agents. *arXiv preprint arXiv:2609.23466*, 2026. 2

[81] Zhengyang Zhao, Lu Ma, and Wentao Zhang. Training with harnesses: On-policy harness self-distillation for complex reasoning. *arXiv preprint arXiv:2605.08741*, 2026. 2

[82] Tailin Zhou. Hierarchical self-improvement: A framework for task-specific evolvable agent harnesses. *arXiv preprint arXiv:2608.08466*, 2026. 1, 2

[83] Adam Zweiger, Jyo Pari, Han Guo, Yoon Kim, and Pulkit Agrawal. Self-adapting language models. *Advances in Neural Information Processing Systems*, 38:74084–74115, 2026. 2

BaseRSI: Recursive Self-Improvement of a Base Vision-Language Model via Training-Free Harness Internalization

Supplementary Material

This appendix treats the three parts of BaseRSI in turn: the teacher–student gap of ReGap (Appendix A.1), the key–value edit of ReKey with the patch-embedding edit (Appendices A.2 and A.3), and the guarantee of the rounds of Recursive Harness Internalization (Appendix A.4), followed by the outer loop and its acceptance test (Appendix B), the algorithms (Appendix C), experimental details (Appendix D), mechanism checks (Appendix E), the complete data (Appendix F), and a case study (Appendix G). Algebraic identities hold on every replayed response, and every bound is stated with its explicit constants.

A. Theory of BaseRSI

A.1. Setting and the Exact Gap

Fix weights θ, an entry c, a request q inside its scope, and a response $y_{1:T}$. A replay places the tokens of the request first and the tokens of y after them; the former are the prompt positions, whose number may differ between $c(q)$ and q, and the latter are the response positions $\mathcal{S} = \{1, \dots, T\}$, with embeddings $\boldsymbol{s}_0$ of the tokens of y.

Property A.1 (Replay)**.** (a) Blocks are causal: the state of a position after block l depends only on the states of that position and of earlier positions after block $l-1$. (b) The embedding of a response token is the same in both replays; positions enter through relative encodings within attention.

By (a), the prompt states evolve on their own, so $\boldsymbol{z}_l$ and $\boldsymbol{z}_l^c$ do not depend on y. The block map $F_l(\boldsymbol{s}; \boldsymbol{z})$ returns the output of block l at the response positions when its input is $\boldsymbol{z}$ at the prompt positions and $\boldsymbol{s}$ at the response positions, with the positional layout of the replay that $\boldsymbol{z}$ comes from. The teacher and the student are the recursions $\boldsymbol{s}_l^c = F_l(\boldsymbol{s}_{l-1}^c; \boldsymbol{z}_{l-1}^c)$ and $\boldsymbol{s}_l = F_l(\boldsymbol{s}_{l-1}; \boldsymbol{z}_{l-1})$, both started at $\boldsymbol{s}_0$ by (b), and the gap of Eq. (6) is $\boldsymbol{g}_l = F_l(\boldsymbol{s}_{l-1}^c; \boldsymbol{z}_{l-1}^c) - F_l(\boldsymbol{s}_{l-1}^c; \boldsymbol{z}_{l-1})$.

Lemma A.2 (Hybrid state)**.** *Let additions $\boldsymbol{a}_1, \dots, \boldsymbol{a}_L$ enter the response states of the student, $\tilde{\boldsymbol{s}}_l = F_l(\tilde{\boldsymbol{s}}_{l-1}; \boldsymbol{z}_{l-1}) + \boldsymbol{a}_l$ with $\tilde{\boldsymbol{s}}_0 = \boldsymbol{s}_0$. Then the prompt states stay $\boldsymbol{z}_l$, and whenever $\tilde{\boldsymbol{s}}_{l-1} = \boldsymbol{s}_{l-1}^c$, we have $\tilde{\boldsymbol{s}}_l = \boldsymbol{s}_l^c$ if and only if $\boldsymbol{a}_l = \boldsymbol{g}_l$.*

Proof. Prompt positions precede the response, so by Property A.1(a) additions at response positions do not reach them. If $\tilde{\boldsymbol{s}}_{l-1} = \boldsymbol{s}_{l-1}^c$, then $\tilde{\boldsymbol{s}}_l = F_l(\boldsymbol{s}_{l-1}^c; \boldsymbol{z}_{l-1}) + \boldsymbol{a}_l$, while $\boldsymbol{s}_l^c = F_l(\boldsymbol{s}_{l-1}^c; \boldsymbol{z}_{l-1}^c) = F_l(\boldsymbol{s}_{l-1}^c; \boldsymbol{z}_{l-1}) + \boldsymbol{g}_l$ by the definition of the gap in Eq. (6). □

Proposition A.3 (Proposition 2, formal)**.** *By Property A.1, the student with additions $\boldsymbol{a}_l = \boldsymbol{g}_l$ for $l = 1, \dots, L$ has $\tilde{\boldsymbol{s}}_l = \boldsymbol{s}_l^c$ for every l, so its next-token distribution equals that of the teacher at every response position; they are the only additions that keep $\tilde{\boldsymbol{s}}_l = \boldsymbol{s}_l^c$ at every block.*

Proof. Induction on l. The case $l = 0$ is Property A.1(b), and the step is Lemma A.2. The next-token distribution at a response position is a function of its final state. □

The two terms of $\boldsymbol{g}_l$ run block l on the same response states and differ only in the prompt states that its attention layer reads. What the entry does to later blocks through the response states is carried by these states once the earlier gaps are added, so no part is counted twice. For an addition $\boldsymbol{a}$ at block l, let $D_l(\boldsymbol{a}) = \frac{1}{2}\|F_l(\boldsymbol{s}_{l-1}^c; \boldsymbol{z}_{l-1}) + \boldsymbol{a} - \boldsymbol{s}_l^c\|^2$ be the squared distance of the block to its teacher. Then $-\nabla_{\boldsymbol{a}} D_l(0) = \boldsymbol{g}_l$ and $D_l(\boldsymbol{g}_l) = 0$: the gap is the exact descent direction of D_l in activation space, a unit step along it closes the distance, and it comes from forward passes.

An edit replaces the additions by a function of the state. Let the edit add $E_l(\boldsymbol{x}_l)$ to the output of block l at every position, where $\boldsymbol{x}_l$ is the input of its MLP and E_l is taken net of the former output of the rewritten neurons; let $\boldsymbol{h}_l$ collect the states at all positions, $\hat{\boldsymbol{h}}_l = (\boldsymbol{z}_l, \boldsymbol{s}_l^c)$ be the hybrid state, $\tilde{F}_l$ the edited block on all positions, $\tilde{\boldsymbol{h}}_l$ the states of the edited student, and

$$\boldsymbol{r}_l = \tilde{F}_l(\hat{\boldsymbol{h}}_{l-1}) - \hat{\boldsymbol{h}}_l. \tag{12}$$

The MLP input of the edited block on $\hat{\boldsymbol{h}}_{l-1}$ is the input $\boldsymbol{x}_l$ of Eq. (9), so $\boldsymbol{r}_l$ equals $E_l(\boldsymbol{x}_l) - \boldsymbol{g}_l$ at the response positions and $E_l(\boldsymbol{x}_l)$ at the prompt positions.

Lemma A.4 (Error propagation)**.** *Let $\|\cdot\|$ be a norm on states and Λ_l a Lipschitz constant of $\tilde{F}_l$ on a set that contains $\tilde{\boldsymbol{h}}_{l-1}$ and $\hat{\boldsymbol{h}}_{l-1}$. Then*

$$\|\tilde{\boldsymbol{h}}_L - \hat{\boldsymbol{h}}_L\| \le \sum_{l=1}^{L} \Big(\prod_{m=l+1}^{L} \Lambda_m \Big) \|\boldsymbol{r}_l\|. \tag{13}$$

With Λ_{out} a Lipschitz constant of the map from the final state to the logits, from $\|\cdot\|$ to the largest absolute logit difference at any position, the divergence from the teacher to the edited student at every response position is at most $\frac{1}{2}\Lambda_{\text{out}}^2$ times the square of the right side of Eq. (13).

Proof. Since $\tilde{\boldsymbol{h}}_l - \hat{\boldsymbol{h}}_l = \tilde{F}_l(\tilde{\boldsymbol{h}}_{l-1}) - \tilde{F}_l(\hat{\boldsymbol{h}}_{l-1}) + \boldsymbol{r}_l$, we get $\|\tilde{\boldsymbol{h}}_l - \hat{\boldsymbol{h}}_l\| \le \Lambda_l \|\tilde{\boldsymbol{h}}_{l-1} - \hat{\boldsymbol{h}}_{l-1}\| + \|\boldsymbol{r}_l\|$, and $\tilde{\boldsymbol{h}}_0 = \hat{\boldsymbol{h}}_0$; unrolling gives Eq. (13). For logits $\boldsymbol{a}, \boldsymbol{a}'$ with $\boldsymbol{d} = \boldsymbol{a}' - \boldsymbol{a}$ and

p the distribution of $\boldsymbol{a}$, the divergence from p to the distribution of $\boldsymbol{a}'$ is $\log \mathbb{E}_p e^d - \mathbb{E}_p d$, which Hoeffding's lemma bounds by $(\max d - \min d)^2/8 \le \|\boldsymbol{d}\|_\infty^2/2$. □

Lemma A.4 is the reason for the two regressions of ReKey: Eq. (10) minimizes $\boldsymbol{r}_l$ at the response positions in mean square, and the third group of Eq. (20) keeps the edit, hence $\boldsymbol{r}_l$, small at the prompt positions. With $\boldsymbol{r}_l = 0$ at every block the lemma returns Proposition A.3.

A.2. ReKey

Fix a block and drop its index. Let $\bar{\boldsymbol{x}} \neq 0$ be the mean input on the fitted pairs, $\hat{\boldsymbol{x}} = \bar{\boldsymbol{x}}/\|\bar{\boldsymbol{x}}\|$, $P = I - \hat{\boldsymbol{x}}\hat{\boldsymbol{x}}^\top$, and $\boldsymbol{k}^0 = \bar{\boldsymbol{x}}/\|\bar{\boldsymbol{x}}\|^2$ the mean direction.

Lemma A.5 (Centred inputs are linear readouts)**.** *$P(\boldsymbol{x} - \bar{\boldsymbol{x}}) = P\boldsymbol{x}$ for every $\boldsymbol{x}$, so a readout $\boldsymbol{u}^\top P(\boldsymbol{x} - \bar{\boldsymbol{x}})$ of the centred input equals $(P\boldsymbol{u})^\top \boldsymbol{x}$, a readout of the input with no offset term, and $\boldsymbol{k}^{0\top}\bar{\boldsymbol{x}} = 1$.*

Proof. $P\bar{\boldsymbol{x}} = \bar{\boldsymbol{x}} - \hat{\boldsymbol{x}}(\hat{\boldsymbol{x}}^\top \bar{\boldsymbol{x}}) = 0$. □

Let $\tilde{X} \in \mathbb{R}^{n\times H}$ stack the rows $(P\boldsymbol{x})^\top$ of the fitted pairs and $\tilde{G}$ the rows $(\boldsymbol{g} - \boldsymbol{b})^\top$. The ridge solution of Eq. (10) and the readouts that ReKey writes are

$$\hat{W} = (\Sigma + \lambda I)^{-1} C, \qquad \tilde{X}\hat{W} \approx A S V_r^\top, \\ \boldsymbol{u}_i = P\hat{W}\boldsymbol{v}_i, \quad i = 1, \dots, r, \tag{14}$$

where $\Sigma = \tilde{X}^\top \tilde{X}/n$, $C = \tilde{X}^\top \tilde{G}/n$, the middle expression is the rank-r truncated singular value decomposition of the fitted values, and $\boldsymbol{v}_i$ is column i of V_r. The r directions are shared by all positions of all responses, and the fitted gaps hold no answer of a development question. For any $\hat{W} \in \mathbb{R}^{H\times H}$, in particular the ridge solution of Eq. (14), let $\hat{Y} = \tilde{X}\hat{W}$ have singular values $s_1 \ge s_2 \ge \cdots$ and right singular vectors $\boldsymbol{v}_1, \boldsymbol{v}_2, \dots$, and let $V_r = [\boldsymbol{v}_1, \dots, \boldsymbol{v}_r]$ and $\hat{W}_r = \hat{W} V_r V_r^\top$, of rank at most r.

Lemma A.6 (Best reproduction through r dimensions)**.** *For every matrix $Z \in \mathbb{R}^{n\times H}$ of rank at most r,*

$$\left\|\hat{Y} - \tilde{X}\hat{W}_r\right\|_F^2 = \sum_{i>r} s_i^2 \le \left\|\hat{Y} - Z\right\|_F^2. \tag{15}$$

If $\lambda = 0$ and $\hat{W}$ is a least-squares solution, $\hat{W}_r$ also minimizes $\|\tilde{G} - \tilde{X}W\|_F$ over all W of rank at most r.

Proof. $\tilde{X}\hat{W}_r = \hat{Y} V_r V_r^\top$ is the rank-r truncation of the singular value decomposition of $\hat{Y}$, so Eq. (15) is the Eckart–Young theorem. If $\lambda = 0$, the normal equations make $\tilde{G} - \hat{Y}$ orthogonal to the column space of $\tilde{X}$, which contains $\hat{Y} - \tilde{X}W$; hence $\|\tilde{G} - \tilde{X}W\|_F^2 = \|\tilde{G} - \hat{Y}\|_F^2 + \|\hat{Y} - \tilde{X}W\|_F^2$, and by Eq. (15) the second term is minimized by $\hat{W}_r$. □

The map $\boldsymbol{x} \mapsto \hat{W}_r^\top P\boldsymbol{x}$ equals $\sum_{i=1}^r (\boldsymbol{u}_i^\top \boldsymbol{x})\boldsymbol{v}_i$ with $\boldsymbol{u}_i = P\hat{W}\boldsymbol{v}_i$, the readouts of Eq. (14). For $\lambda > 0$, Eq. (15) is the optimality of $\hat{W}_r$: among readouts through r dimensions it reproduces the predictions of the ridge fit most closely.

Lemma A.7 (One pair is one neuron)**.** *A neuron of a SwiGLU MLP with rows $\boldsymbol{w}_\mathrm{g} = \beta\boldsymbol{k}$, $\boldsymbol{w}_\mathrm{u} = \boldsymbol{u}$, and $\boldsymbol{w}_\mathrm{d} = \gamma\boldsymbol{v}/\beta$ outputs $\gamma\,\omega_\beta(\boldsymbol{k}^\top\boldsymbol{x})(\boldsymbol{u}^\top\boldsymbol{x})\boldsymbol{v}$ with $\omega_\beta(z) = z\sigma(\beta z)$, and*

$$\sup_{z\in\mathbb{R}} \left|\omega_\beta(z) - \max(z, 0)\right| = \frac{w_0 - 1}{\beta} < \frac{0.2785}{\beta}, \tag{16}$$

where $w_0 \approx 1.2785$ solves $w = 1 + e^{-w}$.

Proof. The neuron computes $\mathrm{silu}(\beta\boldsymbol{k}^\top\boldsymbol{x})(\boldsymbol{u}^\top\boldsymbol{x})\gamma\boldsymbol{v}/\beta$ and $\mathrm{silu}(\beta z) = \beta z\sigma(\beta z)$. For $z \ge 0$, $z - \omega_\beta(z) = z\sigma(-\beta z)$, and for $z < 0$, $|\omega_\beta(z)| = |z|\sigma(-\beta|z|)$; both equal $w\sigma(-w)/\beta$ with $w = \beta|z|$. The function $w\sigma(-w)$ on $w \ge 0$ is maximal where $w\sigma(w) = 1$, that is, $w = 1 + e^{-w}$, and there $w\sigma(-w) = w - w\sigma(w) = w - 1$. □

An edit with r readouts thus rewrites $3(r+1)H$ weights of a block, the gate, up, and down rows of its $r+1$ neurons.

Proposition A.8 (Proposition 3, formal)**.** *On the fitted pairs, among all linear readouts of the input through at most r dimensions, the readouts and values $(\boldsymbol{u}_i, \boldsymbol{v}_i)_{i\le r}$ of Eq. (14) give the predictions closest in Frobenius norm to those of the ridge fit, and each pair is one neuron of the original MLP, scaled by ω_β of the key.*

Proof. On the fitted pairs, the predictions of a linear readout through at most r dimensions form a matrix of rank at most r, so the first part is Lemma A.6; by Lemma A.5 the pairs read the input itself, and the second part is Lemma A.7. □

The truncation is computed by a randomized range finder with oversampling and power iterations. Pair 0 has readout $\boldsymbol{k}^0$ and value $\boldsymbol{b}$, so by Lemma A.5 it adds the mean gap on average, scaled by the same key.

A.3. Patch-Embedding Edit

Lemma A.9 (Resampled patch kernel)**.** *Let the patch embedding map a patch $\boldsymbol{v} \in \mathbb{R}^{C\times p\times p}$ to $\langle K_o, \boldsymbol{v}\rangle$ for each output channel o, let B be a linear up-sampling from $p' \times p'$ to $p \times p$ pixels applied to every input channel, and let $K'_o = B^\top K_o$. Then $\langle K'_o, \boldsymbol{v}'\rangle = \langle K_o, B\boldsymbol{v}'\rangle$ for every patch $\boldsymbol{v}' \in \mathbb{R}^{C\times p'\times p'}$ of the original image.*

Proof. $B^\top$ is the adjoint of B in each channel. □

When the view entry up-samples the image by an integer factor with nearest neighbours, each $p \times p$ patch of the up-sampled image is B applied to one $p' \times p'$ patch of the original image. By Lemma A.9, the vision tower with kernel K' and patch size p' on the original image then produces, on the same grid, the patch embeddings of the teacher; the bias and all other weights of the tower are unchanged.

A.4. Rounds on the Model's Own Responses

Lemma A.4 holds on one replayed response. Which responses are replayed is the subject of this subsection: the student generates its own responses, so the distribution of the regression pairs moves with the edit.

Let $W = (W_1, \dots, W_L)$ stack one linear readout per block, and let the teacher-forced loss of W on a replayed response $\xi = (q, y)$ be

$$\ell(W;\xi) = \sum_{l=1}^{L} \frac{1}{|\mathcal{S}(y)|} \sum_{t\in\mathcal{S}(y)} \big\|\boldsymbol{g}_l(t) - \boldsymbol{b}_l - W_l^\top P_l \boldsymbol{x}_l(t)\big\|^2. \tag{17}$$

Let d_j be the distribution of the responses replayed at round j, $L_j(W) = \mathbb{E}_{\xi\sim d_j}\ell(W;\xi)$, and $L_j^\lambda(W) = L_j(W) + \lambda\|W\|_F^2$, with $\boldsymbol{b}_l$, P_l, and $\lambda > 0$ held fixed across rounds. The refit of Eq. (5) solves Eq. (10) on the pairs of all rounds so far, $W^{(j)} = \arg\min_W \sum_{i\le j} L_i^\lambda(W)$ with $W^{(0)} = 0$, which is the follow-the-leader update.

Let M_x and M_g bound the pairs: $\|P_l\boldsymbol{x}_l(t)\| \le M_x$ and $\sum_l \|\boldsymbol{g}_l(t) - \boldsymbol{b}_l\|^2 \le M_g^2$ for every pair.

Lemma A.10 (Regret of the aggregated refit)**.** *Let $\rho = M_g/\sqrt{\lambda}$ and $G = 2M_x(M_x\rho + M_g) + 2\lambda\rho$. For every $N \ge 1$ and any sequence $d_1, \dots, d_N$,*

$$\frac{1}{N}\sum_{j=1}^{N} L_j^\lambda\big(W^{(j-1)}\big) - \min_W \frac{1}{N}\sum_{j=1}^{N} L_j^\lambda(W) \le \frac{G^2(1+\ln N)}{\lambda N}. \tag{18}$$

Proof. Each L_j is a convex quadratic, so L_j^λ is 2λ-strongly convex and $\Phi_j = \sum_{i\le j} L_i^\lambda$ is $2\lambda j$-strongly convex with minimizer $W^{(j)}$. Since $\lambda\|W^{(j)}\|_F^2 \le \Phi_j(W^{(j)})/j \le \Phi_j(0)/j \le M_g^2$, every iterate lies in the ball of radius ρ, on which the gradient of each L_j^λ has norm at most G. By induction on N, $\sum_{j\le N} L_j^\lambda(W^{(j)}) \le \sum_{j\le N} L_j^\lambda(W)$ for every W, so the left side of Eq. (18) is at most $\frac{1}{N}\sum_j [L_j^\lambda(W^{(j-1)}) - L_j^\lambda(W^{(j)})]$. For each j, strong convexity gives $\lambda j\|W^{(j-1)} - W^{(j)}\|_F^2 \le \Phi_j(W^{(j-1)}) - \Phi_j(W^{(j)})$, and the right side equals $\Phi_{j-1}(W^{(j-1)}) - \Phi_{j-1}(W^{(j)}) + L_j^\lambda(W^{(j-1)}) - L_j^\lambda(W^{(j)}) \le G\|W^{(j-1)} - W^{(j)}\|_F$, because $W^{(j-1)}$ minimizes Φ_{j-1}. Hence $\|W^{(j-1)} - W^{(j)}\|_F \le G/(\lambda j)$ and $L_j^\lambda(W^{(j-1)}) - L_j^\lambda(W^{(j)}) \le G^2/(\lambda j)$; summing $1/j$ gives the claim. □

Lemma A.10 holds for any sequence of response distributions, so it holds when the responses of each round are generated by the student that the previous refit produces; Lemma A.11 states this case, with the first j_0 rounds left to responses that the student does not generate.

Lemma A.11 (Own responses)**.** *Suppose that for every round $j > j_0$ the responses replayed at round j are generated on original requests by the student that carries the edit written from $W^{(j-1)}$. Then*

$$\min_{j_0<j\le N} L_j\big(W^{(j-1)}\big) \le \frac{N}{N-j_0}\Big(\min_W \frac{1}{N}\sum_{j=1}^{N} L_j^\lambda(W) + \frac{G^2(1+\ln N)}{\lambda N}\Big), \tag{19}$$

and the left side is the smallest teacher-forced gap that the fit of a round leaves on the responses of its own student.

Proof. The minimum is at most the mean over $j_0 < j \le N$; all terms are non-negative and $L_j \le L_j^\lambda$, so this mean is at most $\frac{1}{N-j_0}\sum_{j\le N} L_j^\lambda(W^{(j-1)})$, and Lemma A.10 bounds the sum. For $j > j_0$, d_j is the distribution of the responses of the student of round $j-1$, as the left side reads. □

Proposition A.12 (Proposition 1, formal)**.** *In the setting of Lemma A.11, for every $N > j_0$ the smallest teacher-forced gap that a round's fit leaves on the responses of its own student is at most $N/(N-j_0)$ times the smallest regularized gap $\min_W \frac{1}{N}\sum_{j\le N} L_j^\lambda(W)$ that a single linear readout leaves on the responses of all rounds together, plus a term of order $\ln N/N$; with j_0 fixed, the factor tends to one.*

Proof. Lemma A.11; the added term is $\frac{N}{N-j_0}\cdot\frac{G^2(1+\ln N)}{\lambda N}$. □

Remark A.13 (Rounds as run)**.** (i) The warm-up rounds of Algorithm 2, which replay responses generated under the entry and by the unedited base model, are the j_0 rounds of Lemma A.11. (ii) The lemma asks only that the responses of round j come from the student of round $j-1$; it holds for the student as deployed, with its readouts reduced to rank r, keyed, and scaled by the step. (iii) With m replayed responses per round, Lemma A.10 holds for the empirical losses $\hat{L}_j$ as stated. The differences between $L_j(W^{(j-1)})$ and the loss of each response replayed at round j are martingale differences bounded by $\ell_{\max}$, a bound on the loss of Eq. (17) on the ball of radius ρ, so by the Azuma–Hoeffding inequality the mean of $L_j(W^{(j-1)}) - \hat{L}_j(W^{(j-1)})$ over $j_0 < j \le N$ is at most $\ell_{\max}\sqrt{2\ln(1/\delta)/(m(N-j_0))}$ with probability $1-\delta$; Eq. (19) then holds for the population gap with the empirical losses on its right side and this term added. (iv) The step of Eq. (11) acts on the measured divergence: Γ contains the unit step, so the chosen step never leaves a larger held-out divergence than the unit step.

B. Outer Loop and Acceptance Test

Entries, proposer, and router. An entry holds a scope in words, or the scope of every request, and an action: a procedure text added to the prompt, request settings, or an operation on the views of the image. The base model with its original weights is the proposer: it reads its own responses

on the fitting questions and writes candidate entries (Table 3). The reference answers of the development questions mark which of these responses fail and score the acceptance test. The base model is also the router: asked, without thinking and on the original view of the image, whether a request falls inside the scope of an entry, it answers yes or no, which gives $\kappa_c(q)$. The system kept between generations holds a general layer of entries and a layer of entries accepted for single ID benchmarks. In the reported run the loop runs eight generations and its proposer writes 269 candidates; the finer view, accepted in generation 1, and the counting procedure, accepted in generation 2, are internalized in this order, the procedure edit on top of the view edit.

Questions	Read by
Fitting (development)	proposer, router examples, paired replay, fits of readouts and keys, site and strength
Validation (development)	acceptance of entries and edits
ID test	reported scores
OOD tests	reported scores

Table 3. Data routing; development questions are split by image.

Acceptance test. A candidate is scored on the validation questions only, against the current system at an equal token budget. For each scored benchmark b, $d_b(c)$ is the mean of the paired differences of its validation questions, with the questions of one image forming one cluster; for a general entry, the benchmark that it was written from is left unscored. The p-value $p(\cdot)$ of Eq. (3) is the one-sided sign-flip p-value of the macro mean of the listed benchmark effects, obtained by flipping the signs of the image clusters inside each benchmark. A candidate is rejected when some scored benchmark is significantly negative at a level divided by the number of scored benchmarks. An entry for a single ID benchmark is accepted when its gain on the validation questions of that benchmark is significant, with the same rejection on the other benchmarks. An edit is tested in the same way (Algorithm 1), and in addition against a null edit that writes random readouts into the same neurons.

Lemma B.1 (Family-wise error of the acceptance test)**.** *For a candidate c, let $\mathcal{N}_c$ be the hypothesis that its macro gain is at most zero with some scored benchmark left out. If every leave-one-out p-value is valid under its own null, then $p_{\text{gen}}(c)$ is valid under $\mathcal{N}_c$, and the probability that the test accepts at least one candidate of a generation for which $\mathcal{N}_c$ holds is at most the test level α.*

Proof. Under $\mathcal{N}_c$ some benchmark $b^\star$ has a non-positive macro gain of the others, and $p_{\text{gen}}(c)$ is at least the p-value with $b^\star$ left out, which is valid; this is the intersection–union principle. Holm's procedure over the candidates controls the family-wise error at α under any dependence. □

C. Algorithms

Algorithm 1 BaseRSI, generation n

Require: weights θ_n, harness $\mathcal{H}_n$, fitting and validation questions
The base model θ_0 proposes candidate entries from its responses on the fitting questions
Score them on the validation questions; accept $\mathcal{A}_n$ by Eq. (3) or for one benchmark
$\theta \leftarrow \theta_n$; $\mathcal{H} \leftarrow \mathcal{H}_n \cup \mathcal{A}_n$
for each entry $c \in \mathcal{A}_n$ chosen for internalization **do**
    Compute the fraction $\eta(c)$ by Eq. (8)
    if c is a view entry with a low fraction **then**
        $e \leftarrow$ the patch kernel resampled to $K' = B^\top K$
    else
        $e \leftarrow \text{Rounds}(\theta_0, c)$ (Algorithm 2)
    end if
    if e passes its acceptance test (Appendix B) **then**
        $\theta \leftarrow \theta \oplus e$; $\mathcal{H} \leftarrow \mathcal{H} \setminus \{c\}$
    end if
end for
$\theta_{n+1} \leftarrow \theta$; $\mathcal{H}_{n+1} \leftarrow \mathcal{H}$

Algorithm 2 Rounds of the inner loop

Require: weights θ, entry c, questions inside and outside the scope, number of rounds J
$\mathcal{D}^{(0)} \leftarrow \emptyset$; $e^{(0)} \leftarrow$ empty edit
for $j = 1, \ldots, J$ **do**
    Collect responses: under the entry (j=1), of the base model (j=2), of the student exported with $\gamma_{j-1}e^{(j-1)}$ ($j \geq 3$)
    Replay each response with $c(q)$, with q, and block by block on the mixed input; add the pairs of Eq. (9) to $\mathcal{D}^{(j)}$
    $e^{(j)} \leftarrow \text{ReKey}(\mathcal{D}^{(j)})$ (Algorithm 3), one solve per block
    Choose γ_j and measure $\delta_j(\gamma_j)$ on held-out responses of the current student by Eq. (11)
end for
return $\gamma_J e^{(J)}$, the edit with its values scaled by γ_J

Algorithm 3 ReKey, one block

Require: pairs $(\boldsymbol{x}, \boldsymbol{g})$ of the block, rank r, key positions
$\bar{\boldsymbol{x}} \leftarrow$ mean input; $\boldsymbol{b} \leftarrow$ mean gap; $P \leftarrow I - \hat{\boldsymbol{x}}\hat{\boldsymbol{x}}^\top$
$\hat{W} \leftarrow (\Sigma + \lambda I)^{-1} C$, with λ chosen on a held-out fold
$V_r \leftarrow$ top r right singular vectors of $\tilde{X}\hat{W}$; $\boldsymbol{u}_i \leftarrow P\hat{W}\boldsymbol{v}_i$
Rescale each readout to unit variance and its value by the inverse factor
Solve the key of Eq. (20); shift it by the threshold
Write rows $(\beta\boldsymbol{k}, \boldsymbol{u}_i, \boldsymbol{v}_i/\beta)$ and $(\beta\boldsymbol{k}, \boldsymbol{k}^0, \boldsymbol{b}/\beta)$ into the $r+1$ neurons that contribute least

Paired replay. A response is replayed three times by teacher forcing: with the entry, without it, and block by block on the mixed input of Eq. (6). The response positions of a replay are the generated tokens and the tail of the chat template before them; the tail is the same with and without

the entry, so it is treated as the first tokens of y. For block l, the mixed input takes the states of the replay without the entry at the prompt positions and the states of the replay with the entry at the response positions; block l is run once on it with original weights, which rebuilds its attention and recurrent states from the mixed sequence, so each block is asked what it has to add once every block before it is closed. Pairs are stored at the template tail and at a random sample of reasoning positions of each response. On the first response of every run the script adds the gaps to the replay without the entry and prints the largest absolute difference to the replay with the entry (Appendix E reports the check in full precision). For the procedure entry, both replays carry the view entry, so the two differ only in the procedure; the student responses come from the base weights with the edit of the previous round, under the same view entry.

Key. The key of a block is the ridge solution over three groups of positions on original requests:

$$\boldsymbol{k}_l \propto \arg\min_{\boldsymbol{k}} \sum_g \frac{1}{|\mathcal{P}_g|} \sum_{t \in \mathcal{P}_g} \left(\boldsymbol{k}^\top \boldsymbol{x}_l(t) - y_g\right)^2 + \lambda \|\boldsymbol{k}\|^2, \tag{20}$$

where the groups are the response positions of requests inside the scope of the entry ($y_g = 1$), on which the gap is present, those of requests outside it, and the prompt positions (both $y_g = 0$). With $\boldsymbol{k}^{\mathrm{r}}$ the solution of Eq. (20), the written key is $\boldsymbol{k} \propto \boldsymbol{k}^{\mathrm{r}} - \vartheta \boldsymbol{k}^0$, so that the factor ω_β is a ramp that starts at the threshold ϑ of the key value. The intercept of the key is carried by $\boldsymbol{k}^0$, since the MLP has no bias terms, and an entry for every request keeps the constant key.

Neurons and controls. The neurons of a block are ranked by the mean absolute activation times the norm of the output column on responses inside and outside the scope; the pairs take the lowest-ranked neurons, and neurons of earlier edits are excluded. Two null edits use the same neurons, key, and norms: *random readouts* rotates the readouts by a random orthogonal map, and *random values* the values.

Restoration. The restoration of an edit is $1 - \overline{\mathrm{KL}}(\pi_\theta^c \| \pi_e) / \overline{\mathrm{KL}}(\pi_\theta^c \| \pi_\theta)$, the fraction of the divergence from the teacher that the edit removes on one response, averaged over the responses of a set; the divergences we report are means of $\overline{\mathrm{KL}}$ over the same responses, so the two are averaged separately.

D. Experimental Details

Benchmarks and splits. The five ID benchmarks supply 640 development questions, split by image into 270 fitting and 370 validation questions (54 and 74 per benchmark), and an ID test of 640 questions on disjoint images. The procedure edit is fitted on the 57 fitting questions that the router sends to the counting entry, and its key is fitted against 75 that it does not. The five OOD benchmarks supply 640 test questions and are held out from the proposer, the acceptance test, and every fit. Each test set has 128 questions, and a test score is the accuracy on them, averaged over three runs. The regression uses the fitting questions alone. A held-out fold of them, drawn by a hash of the question identifier, selects the ridge weight of each block, and the step and the restoration are measured on responses to further fitting questions held out from the regression. The validation questions serve the acceptance test alone. Every setting in Table 4 is chosen on the development questions. The view edit sets the patch size to 8 and, in the preprocessing configuration stored with the weights, 524,288 to 1,048,576 pixels per image, that is 2,048 to 4,096 image tokens of 2×2 patches; the request carries no image setting. The development splits are released together with the code.

Setting	Value
Base model, proposer, router	Qwen3.8-27B (L=64, H=5120), original weights
Reasoning effort	medium
Test level α	0.10
Rank r, neurons per block	64, 65
Edited blocks	all 64
Ridge λ	$\{0.03, 0.3, 3\}\times$ mean diag. of Σ_l
Key slope β, threshold	8, 0.25
Steps Γ	$\{0.5, 0.75, 1, 1.25, 1.5, 1.8, 2, 2.6\}$; selected $\gamma = 1.8$
Rounds J	3
Positions per response	5 template-tail positions and up to 512 (round 1) or 256 (rounds 2–3) reasoning positions
Patch size	16 to 8
Image tokens after the edit	2,048 to 4,096

Table 4. Configuration of the released code.

Metrics. Accuracy is scored by one rule-based grader shared by all methods: it reads the final answer of a response, matches option letters, yes or no, and normalized short text, and matches numbers exactly. The standard deviation in Table 1 is the sample standard deviation of the accuracy over the three runs of a method. Restoration is the fraction of the divergence from the teacher that an edit removes on a held-out response, averaged over the responses of a set (defined in Appendix C).

Significance. Table 5 tests the gain of BaseRSI over the base model on the two averages of Table 1. Each run of BaseRSI is paired with the run of the base model on the same sampling seed. Every paired run gains 10.6 to 12.8 points in distribution and 5.3 to 6.6 points held out, so both gains are significant under the two-sided paired t-test.

	Base	BaseRSI	Gain	95% CI	p
ID average	60.4	72.1	+11.7	[9.0, 14.4]	0.003
OOD average	64.5	70.7	+6.1	[4.4, 7.9]	0.005

Table 5. Gain of BaseRSI over the base model (Qwen3.8-27B): accuracy (%), paired over the three runs.

Baselines and fairness. All methods but the closed-source models use the same base model, reasoning effort, token budget, evaluator, and test questions. Self-consistency takes the majority of five responses of the base model, and the 4-shot baseline adds four examples to the request. SFT and GRPO train on the development questions until convergence, which gives each its full training budget: SFT with a learning rate of 10^{-6}, cosine decay, 10% warm-up, and 32 samples per update; GRPO with a constant learning rate of 2×10^{-6} after 20 warm-up updates and, in each round, 8 sampled responses to each of 64 questions and 8 updates of 64 responses. LBS-27B is a released checkpoint fine-tuned from the same base model and answers with its own crop loop. PixelProof is the LoRA adapter released by its authors (rank 16, $\alpha = 32$), merged into the same base model with no further training. The closed-source models are Claude Haiku 5.5 and GPT-6 Luna, each at its default reasoning effort. *Harness kept in the request* (Table 2, Figure 4a) is the base model with the finer view and the counting procedure kept in the request and no weight edit applied. *Up-scaled image in request* (Figure 4h) is the base model with the image enlarged in the request, which gives it the finer view while all of its weights stay unchanged; on the OOD benchmarks it reaches 67.2, against 68.8 for the view edit alone and 70.7 for BaseRSI.

Ablation variants. Each variant of Table 2 and Figure 5a changes one choice of BaseRSI; Table 19 gives their scores on every benchmark. *View edit only* and *procedure edit only*: one of the edits alone. *One round*: round 1 only, fitted on responses generated under the entry; *two rounds*: rounds 1 and 2. *Teacher's responses only*: the same number of refits, with every round replaying responses generated under the entry; *base model's responses only*: every round replays responses of the unedited base model. *Macro average only*: the acceptance test scores the macro average with no benchmark left out. *Output only*: the same teacher distilled at the output into every weight; with training restricted to the 4,160 neurons that BaseRSI writes, and the learning rate and the number of passes chosen on the development questions (3×10^{-4} and 16, from 10^{-3} or 3×10^{-4} and 16 or 32), it reaches 69.8 ID and 69.1 OOD. *Last block only*, *last 16*, *first 32*, *48 linear-attention*, and *16 full-attention*: the pairs are written into these blocks alone. *Direct difference as signal*: the difference of the states of the two replays at each block replaces the gap of Eq. (6); *last-block state*: the difference of their last-block states is the target. *Tail + 64 positions*: pairs at the template tail and at 64 reasoning positions of each response; *16* and *64 responses*: that many replayed responses per round (BaseRSI: 344 over three rounds). *Mean gap only*: pair 0 only, one constant direction per block. *Regression-matrix directions* and *random directions*: the readouts at rank 64 along these directions; *random neurons* and *most active neurons*: the pairs written into 65 such neurons per block. *Site preset*: the view entry is written into MLP neurons in place of the patch embedding. *Strength preset*: the unit step, $\gamma = 1$. *Constant key*: the pairs with the mean direction as their key. *Random readouts* and *random values*: null edits with the neurons, key, and norms of the procedure edit, whose readouts or whose values are rotated by a random orthogonal matrix.

Compute. Table 6 gives the cost of improving the base model for every method, in GPU hours measured on the H800 GPUs with 80 GB that every method ran on and in the input and output tokens of the language-model API calls made while improving. BaseRSI takes 45.5 GPU hours, the least of all methods, against 51.4 to 186.6 for the harness methods and 159.3 and 230.9 for GRPO and SFT; LBS-27B and PixelProof are released weights and have no cost of improving here. Like fine-tuning it spends no API tokens, while the harness methods spend 0.14 M to 34.0 M input tokens. Fine-tuning pays backward passes once, the harness kept in the request pays its entries again in every request, and BaseRSI pays forward passes and one linear solve per block once and then answers the original request.

Repeated runs. Each method improves the base model once, and the resulting model or harness answers every test question once in each of three runs, with sampling seeds 0, 1, and 2 shared by all methods.

E. Mechanism Checks

The divergence checks replay 28 held-out responses of the edited student by teacher forcing; the divergence is the mean KL from the teacher, 0.0977 for the unedited student.

E.1. Recursive Harness Internalization (RQ3)

The procedure is learned. On the 102 in-scope validation questions the edited model alone answers 89.9% right, against 90.2% for the teacher and 70.9% for the base model, and it gives the teacher's final answer on 89.9% of them (Table 7). On the 266 test questions that the router sends to the entry, it answers 81.2% right, against 80.2% for the teacher and 65.5% for the base model. A judge, the base model reading the steps of the entry and one response, decides whether the response carries the steps out; its wording was fixed on a calibration set of 40 responses and agrees with

Improving the base model	GPU hours	API input tokens	API output tokens	Backward passes	Weights changed
BaseRSI	**45.5**	**0**	**0**	**no**	**0.23% + patch embedding**
Meta-Harness	51.4	34.0 M	0.83 M	no	none
RRSI	92.2	1.6 M	0.26 M	no	none
ACE	114.6	2.6 M	1.0 M	no	none
VideoHarness-RSI	124.1	1.6 M	0.15 M	no	none
SkillOpt	125.8	0.95 M	0.26 M	no	none
GRPO	159.3	0	0	yes	all
GEPA	186.6	0.14 M	0.14 M	no	none
SFT	230.9	0	0	yes	all

Table 6. Cost of improving the base model: measured H800 GPU hours, and input and output tokens of the API calls made for it.

Deployment	In scope	Agree	Routed	Steps
base model	70.9	76.5	65.5	0.545
teacher (entry in request)	90.2	100	80.2	0.701
edited model (ours)	**89.9**	**89.9**	**81.2**	**0.657**

Table 7. The procedure in the weights. In scope: accuracy (%) on the 102 in-scope validation questions. Agree: share (%) of them answered as the teacher answers. Routed: accuracy on the 266 test questions the router sends to the entry. Steps: share of their responses that follow the steps of the entry, by the judge.

Responses (count)	Round 1	Round 2	Round 3
teacher (14)	0.652	0.654	0.655
base model (28)	0.586	0.604	0.607
student of round 1 (20)	0.504	0.523	0.527
student of round 2 (28)	0.595	0.608	0.610
student of round 3, held out (28)	0.583	0.590	0.593

Table 8. Restoration of each round's edit by response group.

human labels on 18 of 20 unseen ones. On these 266 questions it finds the steps followed in 65.7% of the responses of the edited model, against 70.1% for the teacher, 59.4% for the view edit alone, and 54.5% for the base model.

Rounds. An edit removes more of the divergence on the responses it is fitted on than on the responses it generates: the round-1 edit, fitted on responses under the entry, removes 0.652 on teacher responses and 0.504 on its own (Table 8). Each later round raises the restoration on every group of responses, and on the student's own held-out responses it rises from 0.583 to 0.593 through round 3. The rounds on the student's own responses add their gain on held-out benchmarks, 70.7 against 68.3 for the same number of refits on teacher responses (Table 2).

E.2. ReGap (RQ4)

The gap closes the divergence. Table 9 adds a direction to the activations of every block. Along the gap the divergence falls with the step and reaches 0.0003 at the unit step, where 99.7% of the next tokens agree with the teacher, and it rises again past it. In full precision the unit step leaves a divergence below 10^{-10}, with outputs within a relative error of 1.9×10^{-4} of the teacher, as Proposition 2 states; the remainder in bf16 is rounding. Every other direction of the same norm leaves far more: the direct difference 0.024, the gap of another block 0.073, the gap shuffled across positions 0.089, and a random direction 0.102, while the negated gap raises the divergence to 0.271.

Direction added	Step	Divergence	Agreement
none (student)	0	0.0977	0.9507
gap	0.25	0.0613	0.9577
gap	0.5	0.0299	0.9723
gap	0.75	0.0071	0.9858
gap	**1**	**0.0003**	**0.9965**
gap	1.25	0.0057	0.9854
random direction	1	0.1022	0.9466
gap shuffled across positions	1	0.0890	0.9484
gap of another block	1	0.0732	0.9601
negated gap	1	0.2714	0.9202
direct difference	1	0.0235	0.9651

Table 9. Divergence from the teacher and next-token agreement on 28 held-out responses of the student (bf16), after adding a direction of the same norm to the activations of every block.

l	Divergence	Removed
0	0.0977	0
8	0.0968	0.009
16	0.0951	0.027
24	0.0928	0.050
32	0.0934	0.044
40	0.0615	0.370
48	0.0481	0.508
56	0.0050	0.948
64	0.0003	0.997

Table 10. Divergence after adding the gaps of blocks 1 to l.

The gaps of all blocks. Adding the gaps of the first l blocks removes the whole divergence only when all 64 are added (0.997), and most of it comes between blocks 40 and 56 (Table 10). Added one at a time, the gaps of the full-attention blocks of the second half remove the most, 0.08 to 0.26 each for blocks 36, 40, 44, 48, 52, 56, and 60; the single-block removals sum to 1.36, more than their joint 0.997, since the divergence is not additive over blocks.

Site and step	In scope	Out of scope	Length
procedure in MLP, $\gamma = 0.5$	+15.4	+5.6	−4.3%
procedure in MLP, $\gamma = 1.0$	+14.7	+5.5	−12.6%
procedure in MLP, $\gamma = 2.0$	+18.0	+7.2	−2.4%
view entry in MLP, $\gamma = 0.5$	+11.8	+0.4	+4.4%
view entry in MLP, $\gamma = 1.0$	+11.8	−0.6	+0.4%
view entry in MLP, $\gamma = 2.0$	+9.8	−0.1	+4.2%
last block only, $\gamma = 0.5$	+15.4	+4.0	−8.0%
last block only, $\gamma = 1.0$	+13.4	+6.2	−4.3%
last block only, $\gamma = 2.0$	+15.0	+3.6	−6.3%
measured site and step (ours)	**+19.0**	**+7.0**	**−2.4%**

Table 11. Site and step: accuracy change (points) against the base model on in-scope and out-of-scope validation questions, and the change in length of out-of-scope responses.

Controls at equal capacity. Three controls hold the edit at 4,160 neurons and change one choice each. Distilled from the same teacher at the output, the neurons that BaseRSI writes reach 69.8 ID and 69.1 OOD, against 72.1 and 70.7 for BaseRSI (Appendix D). Trained by gradient, they restore 0.560 ± 0.037 with the per-block gaps as targets and −0.052 ± 0.159 with the output loss alone; the closed-form solve restores 0.593 on the same 28 held-out responses. Written as 260 neurons in each of the 16 full-attention blocks, the edit answers 89.3 ID and 87.9 OOD test questions of 128 right, against 92.3 and 90.5 over all 64 blocks. On 306 further responses that entered no choice of a setting, the exported edit restores 0.562 ± 0.026.

E.3. ReKey (RQ5)

Site and step. Outside the entry's scope the edited model is 7.0 points more accurate than the base model on validation questions and answers 2.4% shorter (Table 11). Writing the view entry into MLP neurons loses the out-of-scope gain at every step, and writing the procedure into the last block alone loses part of the in-scope gain at every step, so the measured site and step give the edit its gain both inside and outside the scope of the entry.

Rank. The restoration rises with the rank up to 64, from 0.075 at rank 0 to 0.593, and rises no further at 96 and 128 (0.576 and 0.566, refitted; Table 12), so the edit uses $r = 64$. The fitted values spread over more directions than that: the first 64 carry 81–85% of their variance in a block, and 90% takes 113–153 directions on average. The directions past the first 64 thus carry variance of the fit that the restoration on held-out responses does not use.

Reduced rank. The readouts of Eq. (14) reach the Eckart–Young bound to four decimals in every group of blocks, as Proposition 3 states (Table 13). The directions of the regression matrix leave 21–24% more error, and the best of 300 random sets of readouts 2.0 to 2.2 times as much; the same order holds on the validation fold.

Rank r	0	1	2	4	8	16	32	**64**	96	128
Restor.	0.075	0.296	0.379	0.467	0.522	0.552	0.575	**0.593**	0.576	0.566

Blocks	k=1	k=4	k=16	k=64	For 90%
1–16	0.280	0.490	0.684	0.850	114.4
17–32	0.271	0.455	0.659	0.840	122.1
33–48	0.223	0.412	0.616	0.810	153.1
49–64	0.236	0.424	0.643	0.841	112.7

Table 12. Rank. Top: restoration on held-out responses against the rank of the fit (96 and 128 refitted). Bottom: share of the variance of the fitted values carried by the first k singular directions of a block, and the number for 90%, averaged over 16 blocks.

Readouts ($r = 64$)	1–16	17–32	33–48	49–64
principal directions of the fit (ours)	**0.3825**	**0.3967**	**0.4296**	**0.3879**
Eckart–Young bound	0.3825	0.3967	0.4296	0.3879
directions of the regression matrix	0.4655	0.4909	0.5314	0.4677
best of 300 random readouts	0.8413	0.8444	0.8635	0.8631

Table 13. Error relative to the ridge predictions of a rank-64 readout, averaged over 16 blocks, on the fitting responses.

Strength. The line search of Eq. (11) picks the step 1.8, the one with the lowest held-out divergence (Table 14). The in-scope gain is largest there as well (+19.0), tied only at 2.6, where the divergence is 1.9 times higher and the out-of-scope change falls to +3.9 points.

Size of the edit. The procedure edit rewrites 65 neurons in each of the 64 blocks, 4,160 in all, each with three vectors of 5,120 entries: 63,897,600 weights, 0.230% of the 27.78 billion weights of the model and 0.249% of the 25.62 billion of its language model. Every other weight of the language model is bit-identical to the base model, so writing back the original rows of the 4,160 neurons restores it exactly; the patch-embedding edit rewrites its own tensor, which writing back the original tensor restores.

Rank, bound, and the steps in the thinking. Restoration rises with the rank of the readout and levels off near $r = 64$, where the principal directions of the fit restore 0.593, the di-

Step γ	Divergence	Restor.	In scope	Out of scope
0.5	0.0681	0.238	+15.4	+5.6
0.75	0.0542	0.354	+13.7	+5.5
1.0	0.0421	0.452	+14.7	+5.5
1.25	0.0329	0.527	+16.0	+5.8
1.5	0.0276	0.579	+16.7	+4.5
1.8 (line search)	**0.0271**	**0.593**	**+19.0**	**+7.0**
2.0	0.0302	0.574	+18.0	+7.2
2.6	0.0512	0.395	+19.0	+3.9

Table 14. Step of the procedure edit: mean held-out divergence and restoration, and accuracy change (points) against the base model on in-scope and out-of-scope validation questions.

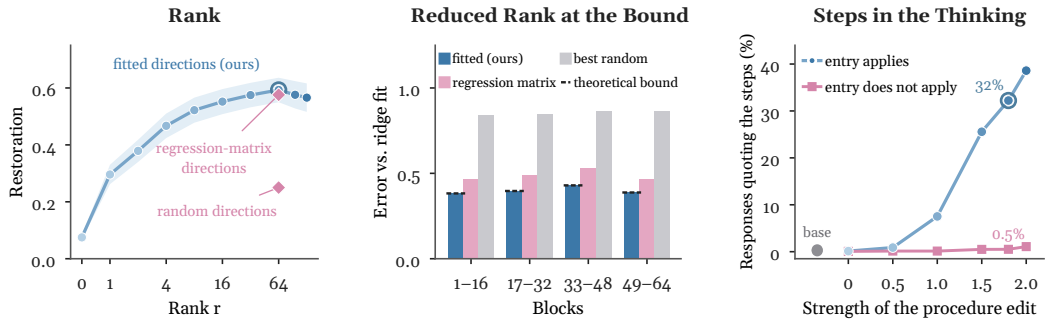


Figure 6. Rank, the bound of Proposition 3, and the steps in the thinking. (a) Restoration against the rank r of the readout; at $r = 64$, the directions of the regression matrix and random directions. (b) Error of the r neurons against the ridge fit for each group of blocks, with the bound. (c) Share of responses whose thinking cites the procedure, by the step of the procedure edit, inside and outside its scope.

	Baby Vision	BLINK	CharXiv	Visu Logic	VLMs Are Biased	ID	ChartQA Pro	Hallusion Bench	Perception Bench	Spatial Viz	Visual Puzzles	OOD
run 1	48	93	99	46	98	384	86	96	64	79	85	410
run 2	50	95	101	40	94	380	87	95	68	81	85	416
run 3	52	96	107	46	95	396	88	94	61	81	89	413
mean correct (of 128)	50.0	94.7	102.3	44.0	95.7	77.3	87.0	95.0	64.3	80.3	86.3	82.6
mean (%)	39.1	74.0	79.9	34.4	74.7	60.4	68.0	74.2	50.3	62.8	67.4	64.5
std (points)	1.6	1.2	3.3	2.7	1.6	1.3	0.8	0.8	2.7	0.9	1.8	0.5

Table 15. Base model (Qwen3.8-27B), three runs (sampling seeds 0, 1, 2). ID and OOD per run: correct answers out of 640.

	Baby Vision	BLINK	CharXiv	Visu Logic	VLMs Are Biased	ID	ChartQA Pro	Hallusion Bench	Perception Bench	Spatial Viz	Visual Puzzles	OOD
run 1	75	99	113	55	117	459 (71.7%)	93	107	72	86	94	452 (70.6%)
run 2	81	96	110	60	115	462 (72.2%)	97	102	79	81	91	450 (70.3%)
run 3	77	99	113	57	118	464 (72.5%)	94	103	76	87	95	455 (71.1%)
sum	233	294	336	172	350	1,385	284	312	227	254	280	1,357
mean (%)	60.7	76.6	87.5	44.8	91.1	72.1	74.0	81.3	59.1	66.1	72.9	70.7
std (points)	2.4	1.4	1.4	2.0	1.2	0.4	1.6	2.1	2.7	2.5	1.6	0.4

Table 16. BaseRSI, correct answers out of 128 in each of three runs (sampling seeds 0, 1, 2), with mean and standard deviation.

rections of the regression matrix 0.576, and random directions 0.250 (Figure 6a). In every group of blocks, the r neurons of the fit reach the bound of Proposition 3, with errors of 0.38 to 0.43 against the ridge fit, while the regression-matrix directions leave 0.47 to 0.53 and the best random directions 0.84 to 0.86 (Figure 6b). The procedure also appears in the thinking of the edited model, whose request no longer carries it: at the deployed step 1.8, 32% of the responses inside the scope of the entry cite it by a fixed pattern, against 0.5% outside, and the share inside grows with the step while that outside stays near zero (Figure 6c).

F. Complete Data

This appendix gives the rest of the measured data, each table with the experiment it comes from and how to read it. Unless a table states otherwise, a test score is the mean over three runs with sampling seeds 0, 1, and 2, the validation counts are correct answers on the 102 in-scope and 268 out-of-scope validation questions, and restoration is measured on 28 held-out responses of the student.

Runs of the base model and BaseRSI. Tables 15 and 16 give the three test runs behind the base model and BaseRSI in Table 1. Run k of either model samples every test question once with seed k, so the two models answer the same questions under the same sampling noise, and the test of Table 5 pairs them run by run. On the 640 ID questions the base model answers 384, 380, and 396 right and BaseRSI 459, 462, and 464, a gain of 68 to 82 questions in every pair; on the 640 OOD questions the gains are 42, 34, and 42 questions. Per benchmark the standard deviation of BaseRSI ranges from 1.2 points on VLMs Are Biased to 2.7 on PerceptionBench, and its two averages vary by 0.4 points.

Runs of five baselines. Table 17 gives the three runs behind five baselines of Table 1, with the same seeds as the

		Baby Vision	BLINK	CharXiv	Visu Logic	VLMs Are Biased	ID	ChartQA Pro	Hallusion Bench	Perception Bench	Spatial Viz	Visual Puzzles	OOD
4-shot	run 1	44	85	97	40	86	352	82	100	65	75	89	411
	run 2	51	90	99	40	89	369	84	102	66	81	91	424
	run 3	53	92	96	42	86	369	86	102	69	83	88	428
ZoomEye	run 1	41	91	101	48	80	361	85	93	64	82	85	409
	run 2	37	94	94	47	75	347	83	97	64	77	91	412
	run 3	42	100	100	49	77	368	85	95	59	81	87	407
PixelProof	run 1	55	94	107	48	92	396	90	95	69	77	90	421
	run 2	47	100	102	50	96	395	84	99	73	79	86	421
	run 3	50	94	105	53	91	393	88	95	68	78	83	412
SFT	run 1	52	92	102	47	95	388	92	98	74	80	78	422
	run 2	51	95	99	46	96	387	86	93	68	82	81	410
	run 3	48	91	105	49	102	395	86	98	77	81	76	418
GRPO	run 1	63	87	100	47	99	396	88	97	67	69	83	404
	run 2	54	87	101	48	96	386	87	93	71	66	92	409
	run 3	58	90	103	42	100	393	85	100	64	73	91	413

Table 17. Five baselines of Table 1, correct answers out of 128 in three runs (seeds 0, 1, 2); ID and OOD: correct out of 640 per run.

Model	Baby Vision	BLINK	CharXiv	Visu Logic	VLMs Are Biased	Avg.	ChartQA Pro	Hallusion Bench	Perception Bench	Spatial Viz	Visual Puzzles	Avg.	Δ ID / OOD
Gemma-4-31B	12.8	69.0	69.0	26.6	45.8	44.6	66.4	76.8	45.3	47.7	63.5	59.9	–
+ BaseRSI	19.3	76.8	83.1	34.6	66.4	56.0	69.0	77.9	54.2	54.9	70.1	65.2	+11.4 / +5.3
Qwen3.6-27B	29.2	67.7	74.2	28.6	41.9	48.3	69.0	72.1	44.0	56.3	60.9	60.5	–
+ BaseRSI	45.3	74.0	81.8	39.6	50.3	58.2	71.9	75.0	49.7	58.1	67.4	64.4	+9.8 / +4.0
Qwen3.5-9B	22.9	65.4	69.8	38.8	36.5	46.7	64.8	75.3	40.6	44.3	58.3	56.7	–
+ BaseRSI	33.9	67.2	75.5	41.9	45.1	52.7	67.4	76.0	44.0	44.8	63.3	59.1	+6.0 / +2.4

Table 18. BaseRSI on other base models: accuracy (%), three-run mean. Δ: gain over the unedited model, from unrounded means.

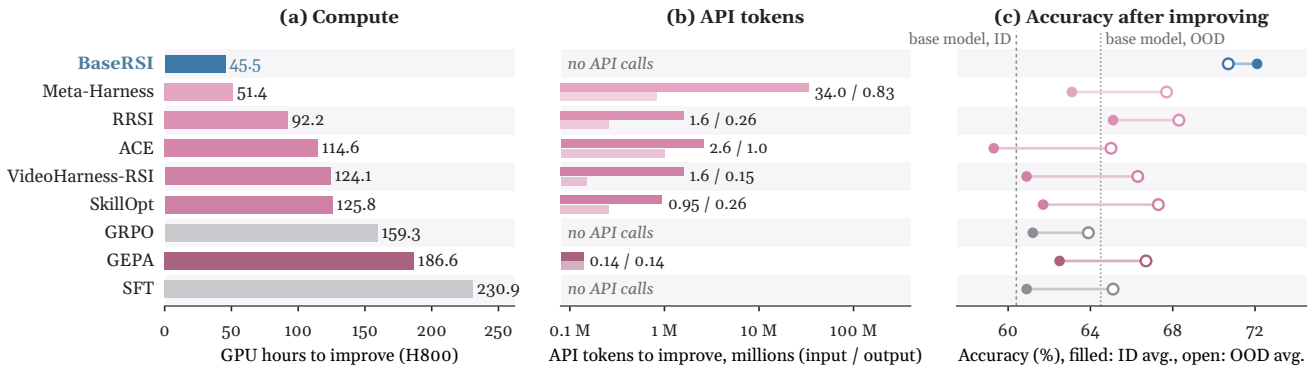


Figure 7. Cost of improving the base model against accuracy. (a) H800 GPU hours. (b) API tokens spent while improving. (c) ID and OOD averages after improving (Table 1); dashed: the base model. Blue: BaseRSI; pink: harness methods; grey: fine-tuning.

base model and BaseRSI. Their standard deviations range from 0.8 to 3.9 points per benchmark and from 0.2 to 1.7 points on the ID and OOD averages, within the range of the other methods. The best run of any of them answers 396 of the 640 ID and 428 of the 640 OOD questions right, 63 and 22 below the worst run of BaseRSI (459 and 450).

Other base models. Table 18 runs the loop of BaseRSI on three other vision-language models and tests each of them on the same ten benchmarks over three runs. The loop starts from the original weights of each model, and the entries that it accepts and internalizes differ in part from those of Qwen3.8-27B. It raises the ID average of every model, by 11.4, 9.8, and 6.0 points on Gemma-4-31B, Qwen3.6-27B, and Qwen3.5-9B, and the OOD average by 5.3, 4.0, and 2.4, and each of the thirty benchmark scores rises. The gain is largest on BabyVision for the two Qwen models (16.1 and 11.0 points) and on VLMs Are Biased for Gemma-4-31B

Variant	In-distribution (ID)						Held-out (OOD)					
	Baby Vision	BLINK	CharXiv	Visu Logic	VLMs Are Biased	Avg.	ChartQA Pro	Hallusion Bench	Perception Bench	Spatial Viz	Visual Puzzles	Avg.
Qwen3.8-27B (base)	39.1	74.0	79.9	34.4	74.7	60.4	68.0	74.2	50.3	62.8	67.4	64.5
BaseRSI (full)	60.7	76.6	87.5	44.8	91.1	72.1	74.0	81.3	59.1	66.1	72.9	70.7
Recursive Harness Internalization (Section 4.1)												
harness kept in the request	58.6	76.3	85.4	40.6	87.0	69.6	73.2	77.1	57.8	61.7	72.1	68.4
procedure edit only	39.1	72.1	82.0	39.6	81.0	62.8	67.7	75.8	52.9	63.8	66.1	65.3
view edit only	58.9	75.8	84.1	43.0	78.4	68.0	73.7	79.7	54.9	63.3	72.1	68.8
one round	59.6	75.5	85.9	44.5	88.0	70.7	73.4	78.6	58.9	62.2	71.4	68.9
two rounds	61.5	75.8	85.2	45.3	88.3	71.2	73.7	77.9	58.1	63.8	72.1	69.1
teacher's responses only	60.2	74.7	85.9	44.0	89.3	70.8	72.7	79.9	58.1	63.0	68.0	68.3
base model's responses only	58.1	77.1	86.5	40.9	90.1	70.5	72.7	81.3	56.8	59.9	68.8	67.9
macro average only	59.6	75.3	85.9	44.3	90.6	71.1	72.9	79.2	59.4	61.7	70.6	68.8
ReGap (Section 4.2)												
output only (distillation)	62.2	75.8	83.3	42.4	80.7	68.9	72.9	78.1	55.2	65.1	68.0	67.9
last block only	59.4	75.8	84.6	43.8	78.9	68.5	72.9	77.6	55.5	61.2	70.6	67.6
last 16 blocks	58.6	77.3	85.9	42.4	80.2	68.9	72.4	75.5	56.3	62.2	70.6	67.4
first 32 blocks	58.9	79.2	84.9	42.2	82.6	69.5	72.9	77.3	57.3	62.8	70.6	68.2
48 linear-attention blocks	60.9	75.5	84.6	41.9	84.1	69.4	73.2	77.9	58.9	63.8	70.3	68.8
16 full-attention blocks	60.4	75.5	84.6	43.2	86.2	70.0	72.7	77.1	54.9	62.5	68.5	67.1
direct difference as signal	59.6	74.7	82.3	40.9	89.3	69.4	73.4	80.5	54.7	64.3	71.6	68.9
last-block state	56.8	76.6	86.2	43.5	79.9	68.6	73.4	76.0	54.4	64.8	69.8	67.7
tail + 64 positions	60.7	75.0	85.2	42.7	89.8	70.7	73.2	77.1	56.5	63.0	72.1	68.4
16 responses	60.9	76.6	84.6	45.6	86.5	70.8	72.9	75.8	57.6	63.8	69.8	68.0
64 responses	57.0	76.0	85.9	43.2	89.8	70.4	73.2	81.8	57.6	63.8	71.1	69.5
random readouts	60.7	75.8	86.7	41.4	75.8	68.1	72.7	76.6	56.8	61.7	72.9	68.1
random values	60.2	75.8	84.1	41.4	77.9	67.9	73.4	77.3	58.3	63.5	72.9	69.1
ReKey (Section 4.3)												
mean gap only (rank 0)	60.2	76.0	84.4	42.4	79.7	68.5	74.0	78.4	57.0	63.5	72.7	69.1
rank 1	57.6	77.1	83.6	44.5	82.6	69.1	72.1	78.1	57.3	62.0	71.9	68.3
rank 4	60.4	77.9	83.6	45.8	85.2	70.6	74.5	76.6	56.5	60.9	70.6	67.8
rank 8	59.4	75.8	84.4	42.7	84.9	69.4	73.7	77.3	56.0	62.0	71.4	68.1
regression-matrix directions	60.9	76.3	83.3	42.7	89.1	70.5	73.4	77.9	56.8	63.5	68.2	68.0
random directions	58.3	76.8	83.6	42.7	83.9	69.1	72.9	79.2	59.6	63.3	71.4	69.3
random neurons	60.4	75.0	85.4	44.3	89.8	71.0	71.9	78.1	55.7	61.7	71.9	67.9
most active neurons	49.7	71.4	82.3	38.5	83.6	65.1	71.9	75.5	55.2	52.6	71.4	65.3
constant key	60.4	76.3	85.4	43.2	89.8	71.0	71.9	80.7	56.5	63.3	72.7	69.0
site preset (view in MLP)	38.3	73.4	81.3	38.0	81.8	62.6	67.7	76.8	52.6	61.2	67.7	65.2
step 0.5	59.4	75.3	85.7	44.8	80.2	69.1	72.9	80.7	58.1	62.8	73.4	69.6
strength preset (step 1.0)	59.4	74.0	84.1	42.4	82.3	68.4	72.9	78.6	56.3	63.8	71.4	68.6
step 1.5	59.4	76.0	85.9	42.2	86.7	70.1	74.0	78.6	59.1	59.6	70.1	68.3
step 2.0	59.6	76.6	84.9	45.3	92.2	71.7	73.4	80.2	57.0	63.8	71.1	69.1
step 2.6	60.7	77.3	85.2	43.8	91.4	71.7	73.4	79.4	55.7	62.0	70.6	68.2

Table 19. Ablation variants, accuracy (%) on each benchmark, mean over three runs.

Setting	Restoration	In scope Δ	Out of scope Δ	In scope (of 102)	Out of scope (of 268)	ID	OOD	Runs
base model	–	–	–	72.3	147.3	77.3	82.6	3 / 3
one round	0.583 ± 0.044	+16.3	+4.4	89.0	159.0	90.5	88.2	3 / 3
two rounds	0.590 ± 0.043	+17.6	+6.0	90.3	163.3	91.1	88.5	3 / 3
BaseRSI (three rounds)	0.593 ± 0.043	+19.0	+7.0	91.7	166.0	92.3	90.5	3 / 3
base model's responses only	0.515 ± 0.038	+17.3	+5.5	90.0	162.0	90.3	86.9	3 / 3
macro average only (+ one entry)	0.545 ± 0.049	+16.7	+6.6	89.3	165.0	91.1	88.0	3 / 3

Table 20. Ablations of Recursive Harness Internalization in full. Restoration ± standard error on 28 held-out responses; Δ in points against the base model on 102 / 268 validation questions; ID and OOD: mean correct answers out of 128; runs: test runs / validation runs.

(20.6 points), the two benchmarks on which it is also the largest on Qwen3.8-27B (Table 1).

Ablations per benchmark. Table 19 gives the test accuracy on each benchmark for the variants of Table 2 and Figure 5a, together with the sweeps over the rank and the step. The edits act on different benchmarks: without the view edit BabyVision falls most, from 60.7 to 39.1, and without the procedure edit VLMs Are Biased falls most, from 91.1 to 78.4, while BabyVision stays at 58.9. Most variants of

Rounds	1	2	**3 (BaseRSI)**	4	5	6	7
Restoration	0.583	0.590	**0.593**	0.595	0.596	0.595	0.594
Change	–	$+0.007 \pm 0.005$	$\mathbf{+0.003 \pm 0.003}$	$+0.002 \pm 0.002$	$+0.001 \pm 0.002$	-0.001 ± 0.003	-0.001 ± 0.002
ID	90.5	91.1	**92.3**	92.5	92.7	92.3	93.0
OOD	88.2	88.5	**90.5**	89.2	88.4	90.1	89.2

Table 21. Rounds of the inner loop, each built as a full model. Restoration on the 28 held-out responses; change: paired difference from the previous round $\pm$ standard error; ID and OOD: mean correct answers out of 128 over three runs.

Setting	Restoration	In scope Δ	Out of scope Δ	In scope (of 102)	Out of scope (of 268)	ID	OOD	Runs
base model	–	–	–	72.3	147.3	77.3	82.6	3 / 3
BaseRSI (three rounds)	0.593 ± 0.043	+19.0	+7.0	91.7	166.0	92.3	90.5	3 / 3
blocks: last block only (65 neurons)	0.005 ± 0.003	+14.4	+4.1	87.0	158.3	87.7	86.5	3 / 3
blocks: last 16 (1,040 neurons)	0.304 ± 0.032	+14.1	+5.5	86.7	162.0	88.2	86.3	3 / 3
blocks: first 32 (2,080 neurons)	0.055 ± 0.022	+16.7	+5.1	89.3	161.0	89.0	87.3	3 / 3
blocks: 48 linear-attention (3,120 neurons)	0.093 ± 0.017	+14.4	+6.1	87.0	163.7	88.9	88.1	3 / 3
blocks: 16 full-attention (1,040 neurons)	0.580 ± 0.046	+14.7	+5.2	87.3	161.3	89.6	85.9	3 / 3
signal: last-block state (step 1.0)	0.271 ± 0.032	+12.1	+5.3	84.7	161.7	87.8	86.7	3 / 3
positions: tail + 64	0.575 ± 0.043	+16.3	+4.6	89.0	159.7	90.5	87.5	3 / 3
responses: 16	0.541 ± 0.043	+17.3	+4.6	90.0	159.7	90.7	87.0	3 / 3
responses: 64	0.578 ± 0.044	+19.0	+6.7	91.7	165.3	90.1	88.9	3 / 3

Table 22. ReGap ablations in full; columns as in Table 20.

Block	Removed	Block	Removed	Block	Removed	Block	Removed
1	+0.0023	17	+0.0016	33	+0.0142	49	+0.0015
2	+0.0014	18	+0.0011	34	+0.0063	50	+0.0141
3	−0.0014	19	+0.0028	35	+0.0145	51	+0.0026
4*	+0.0043	20*	+0.0208	36*	+0.1280	52*	+0.2316
5	+0.0019	21	+0.0009	37	+0.0046	53	+0.0031
6	+0.0021	22	+0.0051	38	+0.0030	54	+0.0008
7	+0.0004	23	−0.0021	39	+0.0008	55	+0.0010
8*	+0.0037	24*	+0.0156	40*	+0.0825	56*	+0.2557
9	+0.0018	25	−0.0000	41	−0.0009	57	+0.0006
10	+0.0002	26	+0.0033	42	+0.0002	58	+0.0029
11	−0.0005	27	−0.0018	43	−0.0028	59	+0.0022
12*	+0.0046	28*	+0.0140	44*	+0.0998	60*	+0.2045
13	+0.0011	29	+0.0003	45	+0.0082	61	+0.0050
14	−0.0007	30	+0.0008	46	+0.0012	62	+0.0016
15	+0.0009	31	+0.0031	47	+0.0033	63	+0.0035
16*	+0.0121	32*	+0.0299	48*	+0.1184	64*	+0.0166

Table 23. Divergence removed by the gap of block l alone ($*$: full attention). They sum to 1.358; together they remove 0.997.

the procedure edit therefore differ on VLMs Are Biased, where the counting entry applies: the last-block edit lowers it to 78.9 and the mean gap alone to 79.7, while BabyVision moves by less than 1.5 points. Writing the most active neurons lowers the benchmarks outside the scope of the entry the most (BabyVision 49.7, SpatialViz 52.6), since it overwrites neurons that the base model uses on every question, and writing the view entry into MLP neurons loses the gain on BabyVision (38.3 against 60.7).

Ablations of Recursive Harness Internalization in full. Tables 20 to 24 extend Figure 5a with the standard error of the restoration, the correct answers on the validation questions, and the test scores as correct answers out of 128. Restoration measures how much of the teacher's next-token distribution an edit reproduces on held-out responses, the validation counts how many answers it changes, and the test scores whether the change carries to new questions. In Table 20, each added round raises the restoration (0.583, 0.590, 0.593) and the in-scope gain (+16.3, +17.6, +19.0). Further rounds (Table 21) change the restoration by at most 0.002 each and the ID accuracy by at most 0.7 correct answers. Fitting on base-model responses alone restores 0.515, with 90.3 and 86.9 on the tests against 92.3 and 90.5.

ReGap ablations in full. In Table 22, an edit of a subset of blocks restores little unless it contains the full-attention blocks of the second half: the last block alone restores 0.005, the first 32 blocks 0.055, the 48 linear-attention blocks 0.093, and the last 16 blocks 0.304, while the 16 full-

Setting	Restoration	In scope Δ	Out of scope Δ	In scope (of 102)	Out of scope (of 268)	ID	OOD	Runs
base model	–	–	–	72.3	147.3	77.3	82.6	3 / 3
BaseRSI (three rounds)	0.593 ± 0.043	+19.0	+7.0	91.7	166.0	92.3	90.5	3 / 3
mean gap only (rank 0)	0.075 ± 0.008	+14.1	+5.8	86.7	163.0	87.7	88.5	3 / 3
rank 1	0.296 ± 0.034	+16.3	+6.1	89.0	163.7	88.4	87.4	3 / 3
rank 4	0.467 ± 0.043	+17.0	+3.6	89.7	157.0	90.3	86.8	3 / 3
rank 8	0.522 ± 0.044	+16.7	+6.3	89.3	164.3	88.9	87.1	3 / 3
reduction: regression-matrix directions	0.576 ± 0.045	+16.7	+4.9	89.3	160.3	90.2	87.0	3 / 3
reduction: random directions	0.250 ± 0.025	+16.0	+6.2	88.7	164.0	88.4	88.7	3 / 3
neurons: random 65 per block	0.474 ± 0.062	+17.3	+5.2	90.0	161.3	90.9	86.9	3 / 3
neurons: 65 most active	−2.119 ± 0.742	+13.7	+1.0	86.3	150.0	83.3	83.6	3 / 3
site preset: view in MLP neurons	–	+11.8	−0.6	84.3	145.7	80.1	83.5	3 / 3
step 0.5	0.238 ± 0.026	+15.4	+5.6	88.0	162.3	88.4	89.1	3 / 3
strength preset: step 1.0	0.452 ± 0.040	+14.7	+5.5	87.3	162.0	87.6	87.8	3 / 3
step 1.5	0.579 ± 0.046	+16.7	+4.5	89.3	159.3	89.7	87.4	3 / 3
step 2.0	0.574 ± 0.039	+18.0	+7.2	90.7	166.7	91.8	88.5	3 / 3
step 2.6	0.395 ± 0.040	+19.0	+3.9	91.7	157.7	91.7	87.3	3 / 3

Table 24. ReKey ablations in full; columns as in Table 20.

attention blocks restore 0.580. Even this subset gains less than all 64 blocks, in scope (+14.7 against +19.0) and on the OOD tests (85.9 against 90.5 correct); Appendix E repeats the comparison with all 4,160 neurons. The 48 linear-attention blocks hold three times the neurons of the 16 full-attention blocks and restore 0.093 against 0.580, and ranks 96 and 128 add neurons without raising the restoration (Table 12), so the placement of an edit, not its size, decides what it restores. The target matters as much as the site: the difference of the last-block states restores 0.271. Adding 64 reasoning positions to the template tail restores 0.575, close to the 0.593 of BaseRSI, and 16 or 64 replayed responses per round restore 0.541 and 0.578.

Each block alone. Table 23 adds the gap of one block at a time to the student and gives the fraction of the divergence that it removes on the held-out responses; Table 10 adds the gaps cumulatively. The full-attention blocks of the second half dominate: blocks 52, 56, and 60 remove 0.23, 0.26, and 0.20 each, blocks 36 to 48 between 0.08 and 0.13, and most linear-attention blocks less than 0.01. The removals sum to 1.358, more than the 0.997 that all gaps remove together, since the divergence is not additive over blocks, and no subset in Table 22 reaches the restoration of all 64.

ReKey ablations in full. In Table 24, the restoration rises with the rank from 0.075 for the mean gap alone to 0.522 at rank 8 and 0.593 at rank 64. At rank 64, readouts along the directions of the regression matrix restore 0.576 and random directions 0.250. Writing the pairs into 65 random neurons per block restores 0.474, and writing them into the 65 most active neurons raises the divergence (−2.119) and keeps only +1.0 out of scope, since those neurons carry what the base model computes on every request. Writing all entries into MLP neurons instead of the patch embedding keeps +11.8 in scope and loses the gain outside (−0.6).

Cost against accuracy. Figure 7 sets the cost of improving the base model (Table 6) beside the accuracy that each method reaches (Table 1). BaseRSI is the cheapest and the most accurate: 45.5 GPU hours for 72.1 ID and 70.7 OOD, against 51.4 hours and 34.0 M API input tokens for Meta-Harness, the next cheapest, which reaches 63.1 and 67.7. More compute does not buy accuracy among the other methods: GEPA spends 186.6 hours for 62.5 ID and SFT 230.9 for 60.9, and SFT and GRPO stay within 0.8 points of the base model in distribution.

G. Case Study

We take the cases from the 1,280 test questions. On 38 of them the base model is wrong in all three runs with seeds 0 to 2 and BaseRSI right in all three, against 6 questions the other way, a split that sampling noise alone would make even; among these we keep the ones that the other methods on the same base model mostly miss, and add questions on which the base model is right in at most two of ten runs. Some of these answers turn on the finer view and others on the procedure; we show both kinds, eight questions from four benchmarks, two of them held out (Figures 8 to 16; Table 25 gathers their runs by setting). The two edits repair different failures: the view edit lets the model see what the question asks about, and the procedure edit, fitted by ReGap and written by ReKey, makes it count what it sees in place of the value it expects. On six questions the answer turns on the finer view, with the view edit alone right in at least four of five runs; on the rings it turns on the procedure, with the procedure edit alone right in 2 of 3 runs and the view edit alone in none; on the stripes both edits contribute. Together the edits are right more often than either alone, in 55 of 56 runs against 31 of 40 for the view edit and 6 of 24 for the procedure edit, as on the benchmarks (Figure 4b). The view edit is also right more often than the up-scaled image in the request, in 31 against 25 of 40 runs: as often on four

Setting	Runs	VLMs Are Biased				BabyVision		Perception Bench	ChartQA Pro
		Rings Fig. 8	Stripes Fig. 10	Bird Fig. 11	Wolf Fig. 12	Squares Fig. 13	Odd one Fig. 14	Lights Fig. 15	Chart Fig. 16
Base model	10	0	2	1	0	0	0	2	0
View edit alone	5	0	3	5	5	5	4	4	5
Procedure edit alone	3	2	2	2	0	0	0	0	0
Both edits, step 0.5	3	0	2	3	2	3	3	3	3
Both edits, step 1.0	3	0	2	3	3	3	3	2	2
Both edits, step 1.5	3	2	3	3	3	3	3	3	3
Both edits, step 2.0	3	3	3	3	3	3	3	2	3
Procedure edit of round 1	3	2	2	3	3	3	3	3	3
Procedure edit of rounds 1–2	3	3	2	3	3	3	3	3	2
BaseRSI (step 1.8)	7	**6**	**7**	**7**	**7**	**7**	**7**	**7**	**7**
Up-scaled image in the request	5	0	1	5	4	5	2	4	4
Other methods, same base model	31	3	8	7	6	0	0	4	0
Closed-source models	21	9	4	0	3	5	4	2	3

Table 25. Runs answered right on the eight case questions, by setting, out of the number of runs given. The edits alone, the steps, and BaseRSI use the same base model; the up-scaled image is enlarged in the request with the weights unchanged; the closed-source models run three times at each of seven effort levels. Lights and Chart come from held-out benchmarks.

questions and more often on the other four. The share of right runs rises with the step of the procedure edit, from 19 and 18 of 24 at steps 0.5 and 1.0 to 23 of 24 at step 1.5 and 55 of 56 at the step 1.8 that the line search of ReKey selects; the procedure edits of round 1 and of rounds 1 and 2 are right in 22 of 24 runs each. The two held-out questions follow the same pattern, the view edit alone right in four and five of five runs and BaseRSI in all seven. The base model is right in 5 of 80 runs, the other methods in 28 of 248, at most 8 of their 31 on one question, and the four closed-source models in 30 of 168 over seven effort levels, at most 9 of their 21. Appendix E and Table 2 give the controls: shuffled gaps or gaps of another block leave the divergence, null edits leave the ID average unchanged, and no preset site or strength reaches the measured one.

Every figure has the same layout. The question shows the image with the region that decides the answer marked in red. *What the model sees* draws the image tokens of that region as a grid, for the base model and after the view edit, each cell being one token that the model reads. *Becoming right step by step* gives the share of runs answered right as the edits are added, where Edit 1 is the view edit and Edit 2 the procedure edit at the step written below it. *Answers of every run* lists each run of each method on the same base model and of four closed-source models, a filled square for a right answer and a crossed square for a wrong one, with the typical answer of each method on the right. The bottom row quotes the thinking of single runs, with [. . .] marking omitted text, and the excerpts below each case are taken from the same runs. The eleven other methods, chain-of-thought, 4-shot examples, Self-Refine, ZoomEye, two-stage zoom-in, PixelProof, GEPA, ACE, SkillOpt, VideoHarness-RSI, and RRSI, give 31 runs in all. The closed-source models give three runs at each of seven effort levels, Claude Haiku 5.5 at medium, Claude Sonnet 5.5 at medium and high, GPT-6 Luna at low, medium, and its highest effort, and GPT-6.1 Terra at medium, scored as in the main tables; the figures show the highest level of each, and the text counts all 21 runs. Self-consistency appears in the figures only.

One observation recurs across the counting questions. With the procedure edit, the thinking of the model often cites an instruction to trust direct observation rather than the typical value, although the request carries no such text: ReGap and ReKey have written the entry into the weights, and the model recalls it as if it were part of the request. Over all replies to 21 candidate questions, a fixed pattern finds such a citation in 46 of the 147 replies of BaseRSI, in 0, 4, 16, and 21 of 63 replies at steps 0.5, 1.0, 1.5, and 2.0, and in none of the 42 replies of the base model, the 63 of the view edit, or the 109 of chain-of-thought, 4-shot, and ZoomEye prompting. The citation grows with the step and is absent without the edit, so it comes from the edit, and it appears on counting questions only, never on the chart or odd-one-out questions, matching the scope that the key carries.

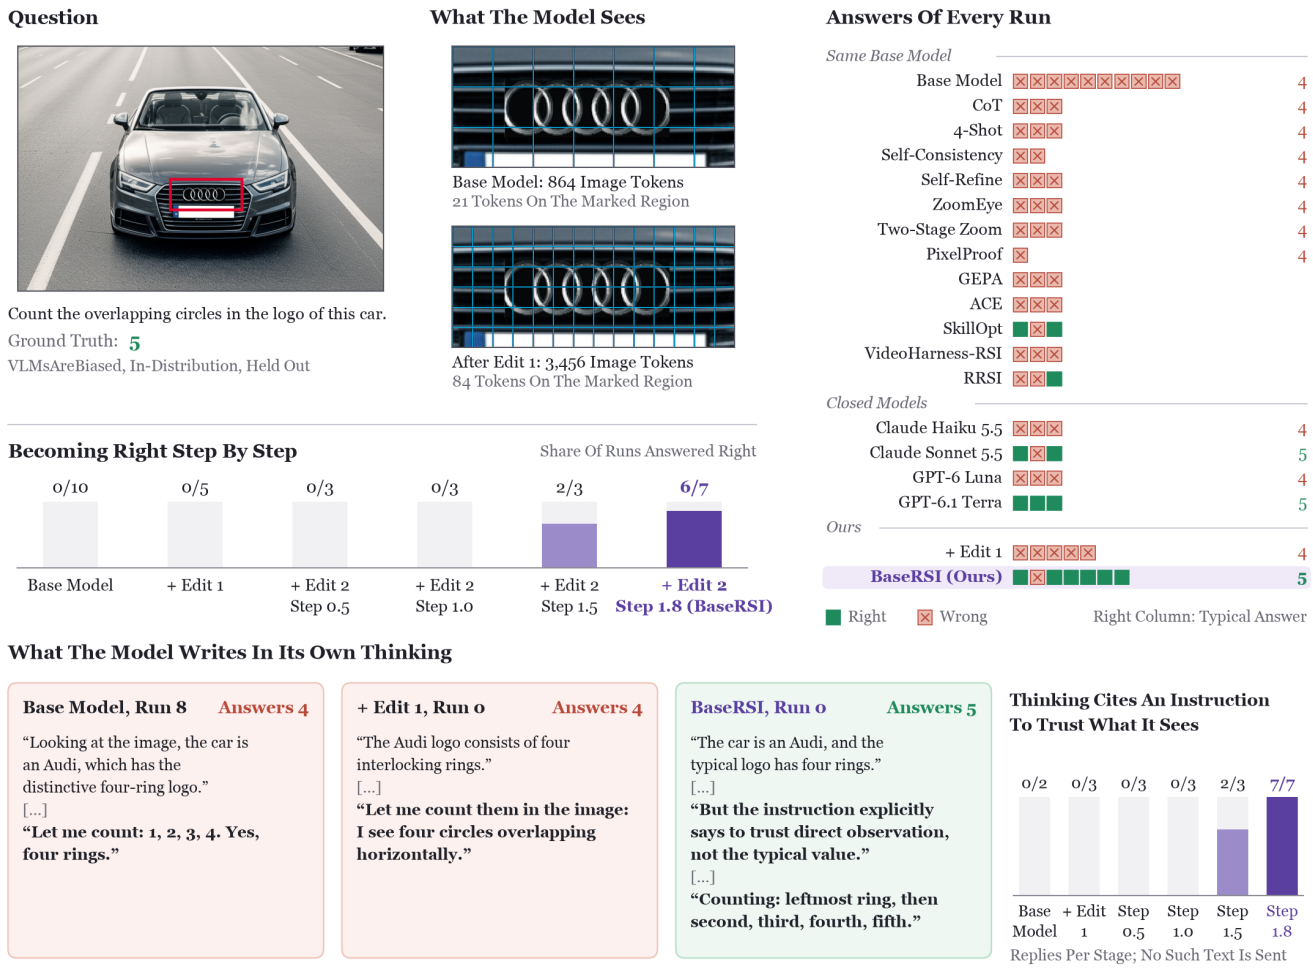


Figure 8. **Counting the rings of a logo** (VLMs Are Biased). The view edit alone answers four; the procedure edit turns it to five.

Key observations. The question asks for the overlapping circles in the logo of a car, which VLMs Are Biased has redrawn with five rings (Figures 8 and 9). The base model recognizes an Audi and answers four in all ten runs; its thinking names the brand first and then counts "1, 2, 3, 4" along the rings it expects, so the count follows the prior rather than the image. The view edit raises the image tokens on the logo from 21 to 84, and the edited model still answers four in all five runs, as does the up-scaled image in the request: seeing the logo more finely leaves the count unchanged. The procedure edit changes it: on top of the view edit, 0/3, 0/3, 2/3, and 6/7 runs are right at steps 0.5, 1.0, 1.5, and 1.8, and the procedure edits fitted in round 1 and in rounds 1 and 2 reach 2/3 and 3/3. Of the 31 runs of the eleven other methods on the same base model, 3 are right, and over seven effort levels of the four closed-source models 9 of 21 runs are right, from three of the four models. The excerpts below give the three runs quoted in the figure.

Base model, run 8: "Looking at the image, the car is an Audi, which has the distinctive four-ring logo. [...] The Audi logo consists of four interlocking/overlapping circles (rings). I can see them in the center of the grille. Let me count: 1, 2, 3, 4. Yes, four rings. [...] Let me make sure I'm not missing anything. The logo clearly shows the four rings of Audi." *View edit, run 0:* "Looking at the image, it's an Audi car, and the logo is the four rings. Let me verify by looking at the front grille logo. The Audi logo consists of four interlocking rings. Let me count them in the image: I see four circles overlapping horizontally. So the answer is 4. [...] The Audi emblem is four rings. Yes." *BaseRSI, run 0:* "Looking at the image, the logo shows overlapping rings. Let me count them individually [...] Counting: leftmost ring, then second, third, fourth, fifth. [...] Let me double check by mentally tracing: circle at left edge, then overlapping second, middle third, fourth, right edge fifth. Yes, I see five." Its answer: "Looking carefully at the emblem on the front grille, I trace each overlapping ring from left to right. Rather than the usual four-ring Audi mark, this image shows a row of five interlocking circles."

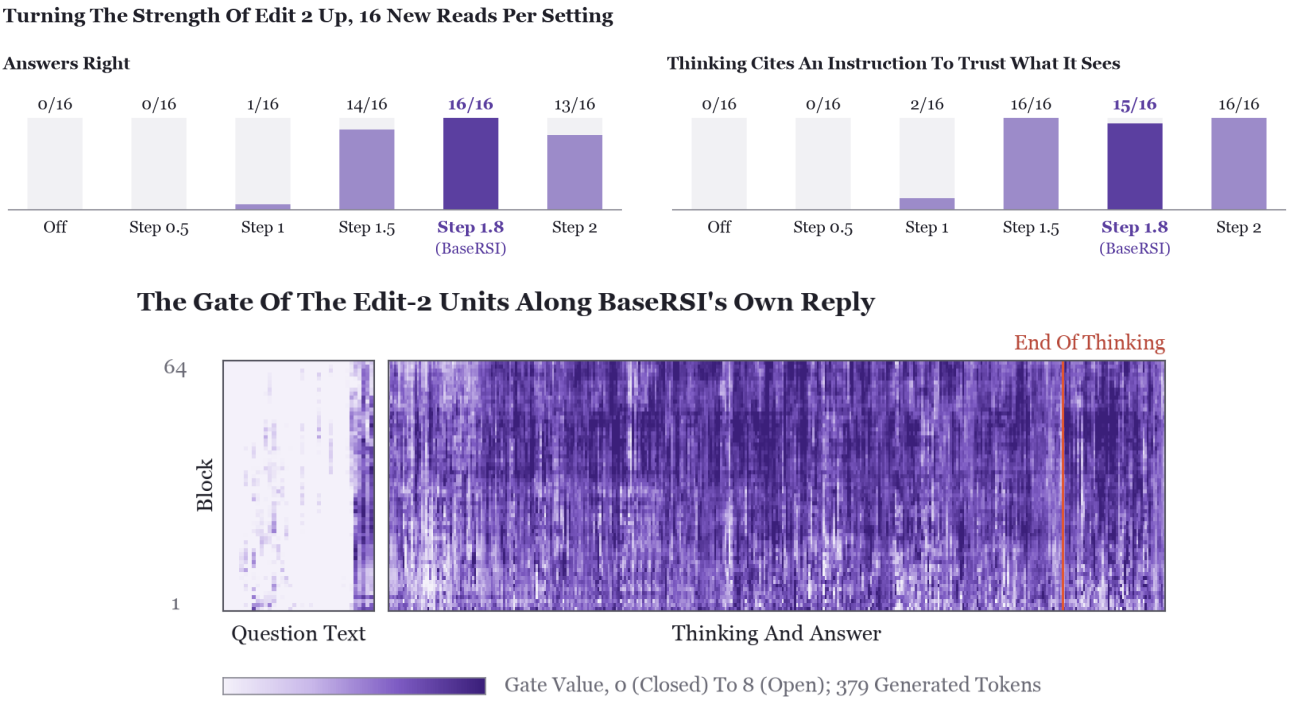


Figure 9. **Counting the rings of a logo, inside the model.** Top: runs answered right, and runs whose thinking cites an instruction to trust what it sees, at each step of the procedure edit, 16 new runs each. Bottom: key value of the procedure edit (the gate) at every block along a response of BaseRSI: near zero over most of the question text, open through the thinking and the answer.

Key observations. In 16 new runs per setting, none is right with the procedure edit off or at step 0.5, 1 at step 1.0, 14 at step 1.5, and all 16 at step 1.8, and 13 at step 2.0 (top). From none of 16 runs to all 16, the answer follows the strength of the edit and not the sampling. The request carries no procedure, yet the thinking of the edited model now cites "the instruction" to trust direct observation over the typical value and then counts the rings one by one; the citation appears in 2 of the 16 runs at step 1.0 and in 15 of the 16 at step 1.8, and in none with the edit off or at step 0.5, so it grows with the step as the right answers do. In the runs of Figure 8 it appears in no recorded reply of the base model or the view edit, in none at steps 0.5 and 1.0, in 2 of 3 at step 1.5, and in all 7 of BaseRSI. At step 2.0 every run still cites the instruction but three of them miscount, so a larger step adds no accuracy, in line with the line search of ReKey at 1.8 (Table 14). A further reading shows where the repair comes from: the procedure edit alone, without the view edit, is right in 2 of 3 runs, against none for the view edit alone or the up-scaled image, so on this question the count is turned by the procedure, which the edit carries in the weights instead of the request. Along the response of BaseRSI (bottom), the key that ReKey fits for the procedure edit stays near zero over most of the question text and opens through the thinking and the answer, at every block from the first to the last; it switches on at the last tokens of the question, where the chat template hands over to the response, and stays open after the end of the thinking, through the final answer of this 379-token response. The key is fitted to be one on the response positions of requests inside the scope and zero on prompt positions and on other requests (Eq. (20)), and the figure shows this behaviour along an actual response. One of the seven runs of BaseRSI answers six; it too cites the instruction and counts ring by ring.

BaseRSI, run 0: "The user wants me to count the overlapping circles in the car's logo. The car is an Audi, and the typical logo has four rings. But the instruction explicitly says to trust direct observation, not the typical value. Let me look carefully at the image. I need to examine the logo in the grille. [...] The standard Audi has 4 rings. But the instruction warns the image may deviate. Let me look closely. [...] Looking at the rendered logo, I can discern 5 circles. Let me look: the pattern is OOOOO perhaps. The width of the logo seems wider than 4 rings would be relative to the grille." *BaseRSI, run 1 (answers six):* "The standard Audi logo has four rings. But the instruction says to trust direct observation and not substitute the standard value. [...] Looking at the image, the rings appear to be: I can see six circular outlines. Let me trace: leftmost circle, then second, third, fourth, fifth, sixth. The logo width looks like it contains 6 rings." Its answer: "Looking directly at the emblem on the front grille, I trace each individual ring from left to right rather than assuming the usual four-ring Audi mark. Counting them as they appear in this image gives six overlapping circles."

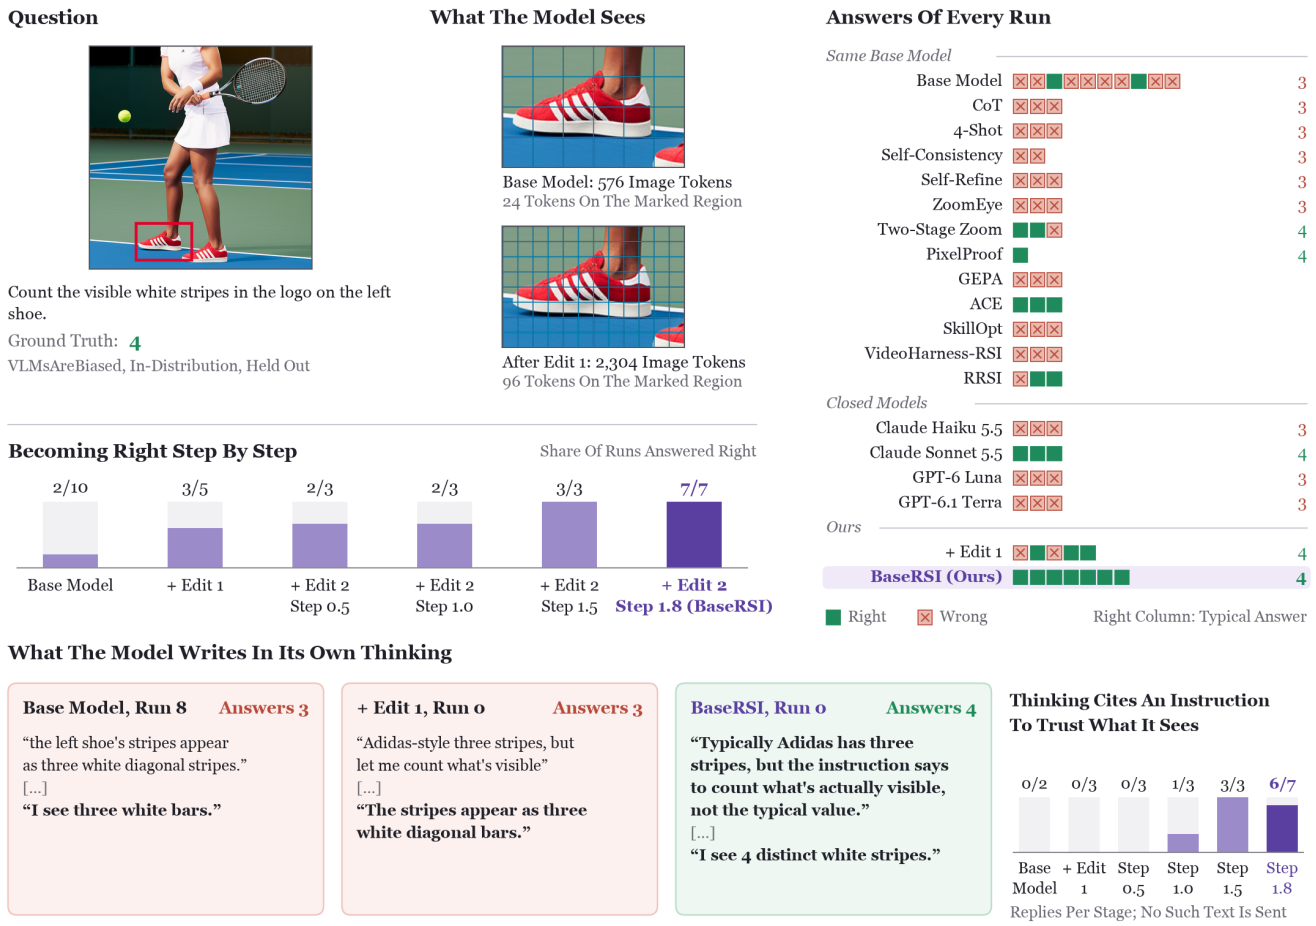


Figure 10. **Counting the stripes on a shoe** (VLMs Are Biased). Each edit lifts the share of right runs, and together they reach 7/7.

Key observations. The question asks for the white stripes in the logo on the left shoe, which VLMs Are Biased has drawn with four stripes where the brand has three (Figure 10). The base model is right in 2 of 10 runs and otherwise answers three; its thinking reads "three white diagonal stripes" and settles on "three white bars". Here both edits contribute. The view edit raises the image tokens on the shoe from 24 to 96 and the right answers to 3 of 5; in a wrong run its thinking announces that it will "count what's visible" and still reads three bars. The procedure edit alone, on the original view, is right in 2 of 3 runs. With both edits the answer improves with the step, 2/3 at steps 0.5 and 1.0, 3/3 at step 1.5, and 7/7 at step 1.8, and the share of runs whose thinking cites the instruction rises in the same way, from 1 of 3 at step 1.0 to 6 of 7 at step 1.8. The edits are fitted separately and written at different sites, the patch embedding and the MLP neurons, and on this question they add up as they do on the benchmarks (Figure 4b). Of the 31 runs of the other methods on the same base model, 8 are right, and over seven effort levels of the closed-source models 4 of 21 runs are right, all from one model. The up-scaled image in the request is right in only 1 of 5 runs, against 3 of 5 for the view edit, and the procedure edits of the earlier rounds reach 2 of 3, as the final edit at steps 0.5 and 1.0.

Base model, run 8: "There's a person wearing red shoes with white stripes (Adidas style). [...] Looking at the image, the left shoe (from the viewer's perspective) shows white stripes: I see three stripes? Let me look carefully. [...] In a typical Adidas logo, three stripes. But maybe the image shows four? Let me examine. [...] I see three diagonal stripes." *View edit, run 0:* "The image shows a tennis player wearing red shoes with white stripes (Adidas-style three stripes, but let me count what's visible). [...] Usually Adidas has three stripes. Let me look at the description: the left shoe shows white diagonal stripes. Counting: 1, 2, 3. [...] The stripes appear as three white diagonal bars." *BaseRSI, run 0:* "Typically Adidas has three stripes, but the instruction says to count what's actually visible, not the typical value. Let me look at the left shoe: I see white stripes on the side. Counting: there appear to be four stripes? [...] I'll count positions: first stripe near the toe side, then second, third, fourth. It looks like four."

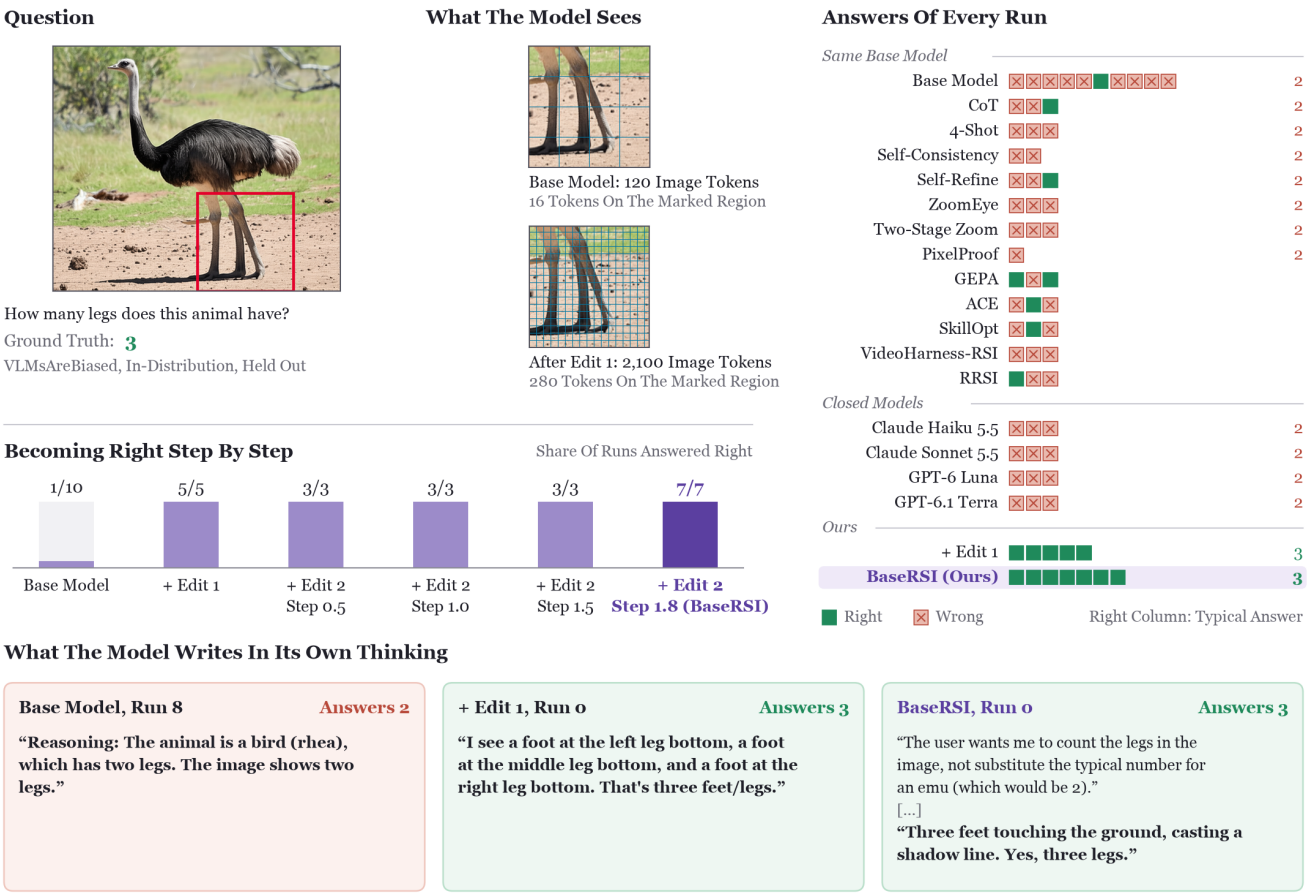


Figure 11. **Counting the legs of a bird** (VLMs Are Biased). Every closed-source run answers two; the view edit gives three in every run.

Key observations. The question asks how many legs an animal has, and the image shows a large flightless bird drawn with three legs (Figure 11). The base model answers two in 9 of 10 runs, reasoning that "the animal is a bird (rhea), which has two legs", and its answer describes "its two legs (one planted straight down and one angled to the side)". The view edit gives the legs 280 image tokens instead of 16, and every run then answers three, one of them finding "a foot at the left leg bottom, a foot at the middle leg bottom, and a foot at the right leg bottom". The answer stays three at every step of the procedure edit, so the view edit carries this repair, and the procedure edit changes how the model reasons on top of it: in 6 of the 7 runs of BaseRSI the thinking cites the request to count the legs, "not substitute the typical number for an emu (which would be 2)", and then traces each leg to the ground. Of the 31 runs of the other methods, 7 are right, while all 21 runs over seven effort levels of the closed-source models answer two. The up-scaled image in the request also gives three in all five runs.

Base model, run 8: "It's a rhea or similar ratite bird standing on dirt. [...] I see one vertical leg and one leg angled to the right. So two legs. [...] Reasoning: The animal is a bird (rhea), which has two legs. The image shows two legs." *View edit, run 0:* "There's one leg on the left (vertical, ending in a foot on the ground). There's a second leg in the middle (vertical, ending in a foot). There's a third leg on the right that's angled diagonally, ending in a foot on the ground. So I see three legs. [...] Let me double check by examining the feet on the ground. I see a foot at the left leg bottom, a foot at the middle leg bottom, and a foot at the right leg bottom. That's three feet/legs." Its answer: "Looking carefully at the image, I can trace three distinct legs extending from the body down to the ground: one on the left (vertical), one in the middle (vertical), and one on the right (angled diagonally). Each ends in a foot on the sandy ground." *BaseRSI, run 0:* "Starting from the body: there's a leg going down on the left side (with a spur visible), then another leg in the middle, then a third leg on the right side that's angled. [...] Three feet touching the ground, casting a shadow line. Yes, three legs." Its answer: "Looking directly at the image and tracing each limb from the body to the ground, I can identify three separate legs: a left leg (with a visible spur), a middle vertical leg, and a right leg angled outward. Each ends in a foot on the ground."

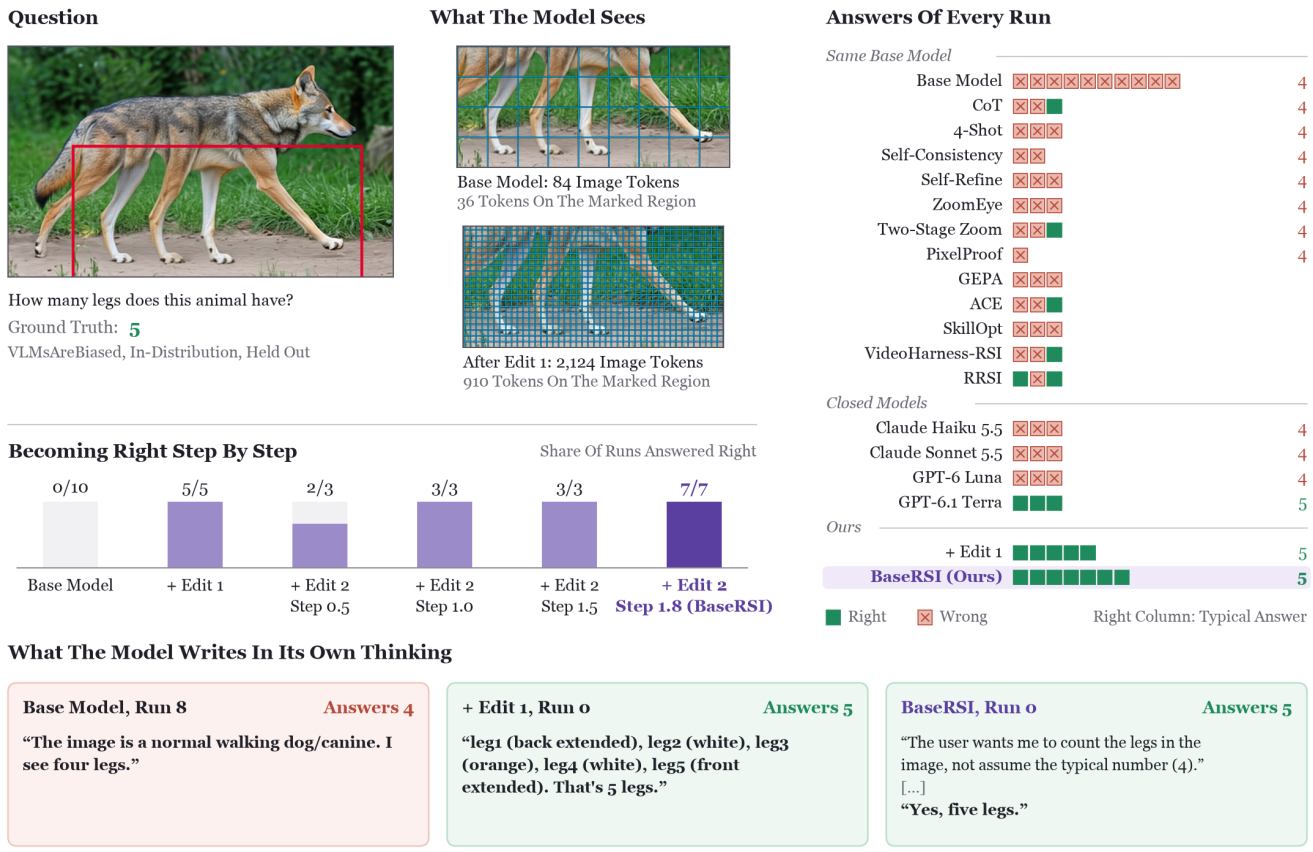


Figure 12. **Counting the legs of a wolf** (VLMs Are Biased). The base model sees "a normal walking dog"; the view edit counts five.

Key observations. The image shows a wolf drawn with five legs (Figure 12). The base model answers four in all ten runs and states that "the image is a normal walking dog/canine" with four legs; its answer counts "the trailing hind leg, the standing hind leg, the standing front leg, and the extended front leg". It reads the animal through 84 image tokens, 36 of them on the legs, and the view edit through 2,124, 910 of them on the legs. With the finer view the model lists "leg1 (back extended), leg2 (white), leg3 (orange), leg4 (white), leg5 (front extended)" and answers five in all five runs. The procedure edit alone, at the original resolution, answers four in all three runs, so on this question the count rests on the finer view. BaseRSI is right in all seven runs, and in 5 of them the thinking again begins from the request to count the legs "not assume the typical number (4)". Six of the 31 runs of the other methods on the same base model are right, and among the closed-source models only GPT-6.1 Terra answers five, in its 3 runs; the other 18 closed-source runs answer four. The up-scaled image in the request is right in 4 of 5 runs, close to the view edit.

Base model, run 8: "The image shows a canine (looks like a jackal or coyote mix) standing in profile, walking. [...] The animal is walking, so all four legs should be visible. [...] Let me double check there isn't a fifth or any trick. The image is a normal walking dog/canine. I see four legs." *View edit, run 0:* "Starting from the left (rear of animal): 1. Rear leg extended back (the one with the orange/brown upper and white lower, with paw on ground at left). 2. Another rear leg (white, standing straight). 3. A front-ish leg (orange/brown upper, standing). 4. Another leg (white, standing). 5. A front leg extended forward (orange, paw raised). Wait, let me look more carefully. This image might be a manipulated or AI-generated image with extra legs. Let me count paws/legs on the ground. [...] That's 5 legs. That's unusual." Its answer: "Counting the legs from left to right: the rear leg stretched backward, a white leg standing, an orange leg in the middle, another white leg, and the front leg extended forward [...] that makes five legs." *BaseRSI, run 0:* "Leg 1: far left, hind leg, foot at bottom left [...] Leg 2: hind leg [...] Leg 3: middle leg [...] Leg 4: front-ish leg [...] Leg 5: extended front leg [...] So that's 5 legs visible." Its answer: "Looking at each leg position from rear to front, I can identify distinct legs with separate feet: 1. A hind leg trailing to the left (foot at far bottom-left). 2. A second hind leg (foot near the left-center). [...] 5. A front leg extended forward (foot at the right). That gives five legs visible in the image."

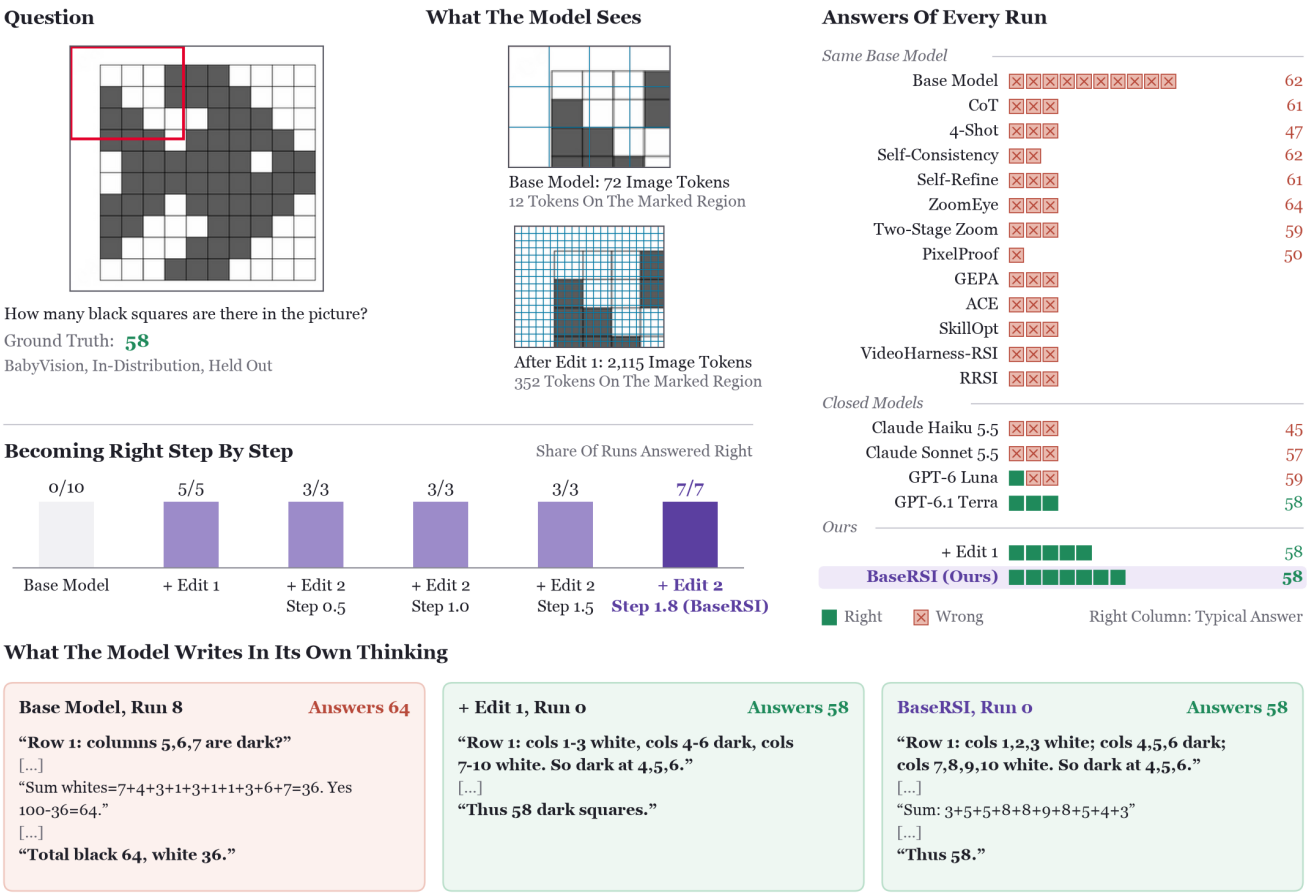


Figure 13. **Counting black squares** (BabyVision). At 72 image tokens a token is larger than a square; at 2,115 the model reads every row.

Key observations. The question asks how many squares of a grid are black, 58 in all (Figure 13). The image is 273 by 264 pixels, and at its default resolution the base model reads it as 72 image tokens, each covering about 30 by 30 pixels, more than one square of about 22 pixels. Its counts scatter between 52 and 64 over ten runs, and its thinking shows why: row by row it reads the dark run of the first row one column to the right of where the run of BaseRSI reads it (5, 6, 7 against 4, 5, 6) and reads six dark squares in the second row against five; it marks its first reading "Row 1: columns 5,6,7 are dark?" with a question mark and finally derives the total from a count of the white squares, "100-36=64". The run takes 2,223 generated tokens, against 1,392 for the run of BaseRSI. The view edit gives the picture 2,115 tokens, 352 of them on the marked corner where the base model has 12; the model then reads each row exactly, "Row 1: cols 1-3 white, cols 4-6 dark, cols 7-10 white", and counts 58 in all five runs, as does the up-scaled image in the request. BaseRSI counts 58 in all seven runs, adding the dark squares of the ten rows, "3+5+5+8+8+9+8+5+4+3". None of the 31 runs of the other methods on the same base model is right; their recorded answers range from 39 to 65. Over seven effort levels of the closed-source models 5 of 21 runs are right, three of them from GPT-6.1 Terra, and the other runs answer between 45 and 62.

Base model, run 8: "The image is a 10x10 grid. [...] Row 1 (top): positions: 1 white,2 white,3 white,4 white,5 dark,6 dark,7 dark,8 white,9 white,10 white. Count=3. Row 2: 1 dark,2 white,3 white,4 dark,5 dark,6 dark,7 dark,8 dark,9 white,10 white. Count=6? [...] Sum whites=7+4+3+1+3+1+1+3+6+7=36. Yes 100-36=64. [...] Total black 64, white 36." *View edit, run 0:* "Row 1: cols 1-3 white, cols 4-6 dark, cols 7-10 white. So dark at 4,5,6. Row 2: col1 dark, cols2-3 white, cols4-7 dark, cols8-10 white. Dark at 1,4,5,6,7. [...] Sum: 3+5=8, +5=13, +8=21, +8=29, +9=38, +8=46, +5=51, +4=55, +3=58. Let me double check each row by re-examining the image. Row 1 (top): I see three dark squares in middle (cols 4,5,6). Yes 3." *BaseRSI, run 0:* "Row 1: cols 1,2,3 white; cols 4,5,6 dark; cols 7,8,9,10 white. So dark at 4,5,6. Row 2: col1 dark; col2,3 white; col4,5,6 dark; col7 dark; col8,9,10 white. So dark at 1,4,5,6,7. Row 3: col1 dark; col2 dark; col3 white; col4 white; col5 white; col6 dark; col7 dark; col8 dark; col9 white; col10 white. So dark at 1,2,6,7,8. [...] Total = 3+5=8; +5=13; +8=21; +8=29; +9=38; +8=46; +5=51; +4=55; +3=58. So 58 dark squares."

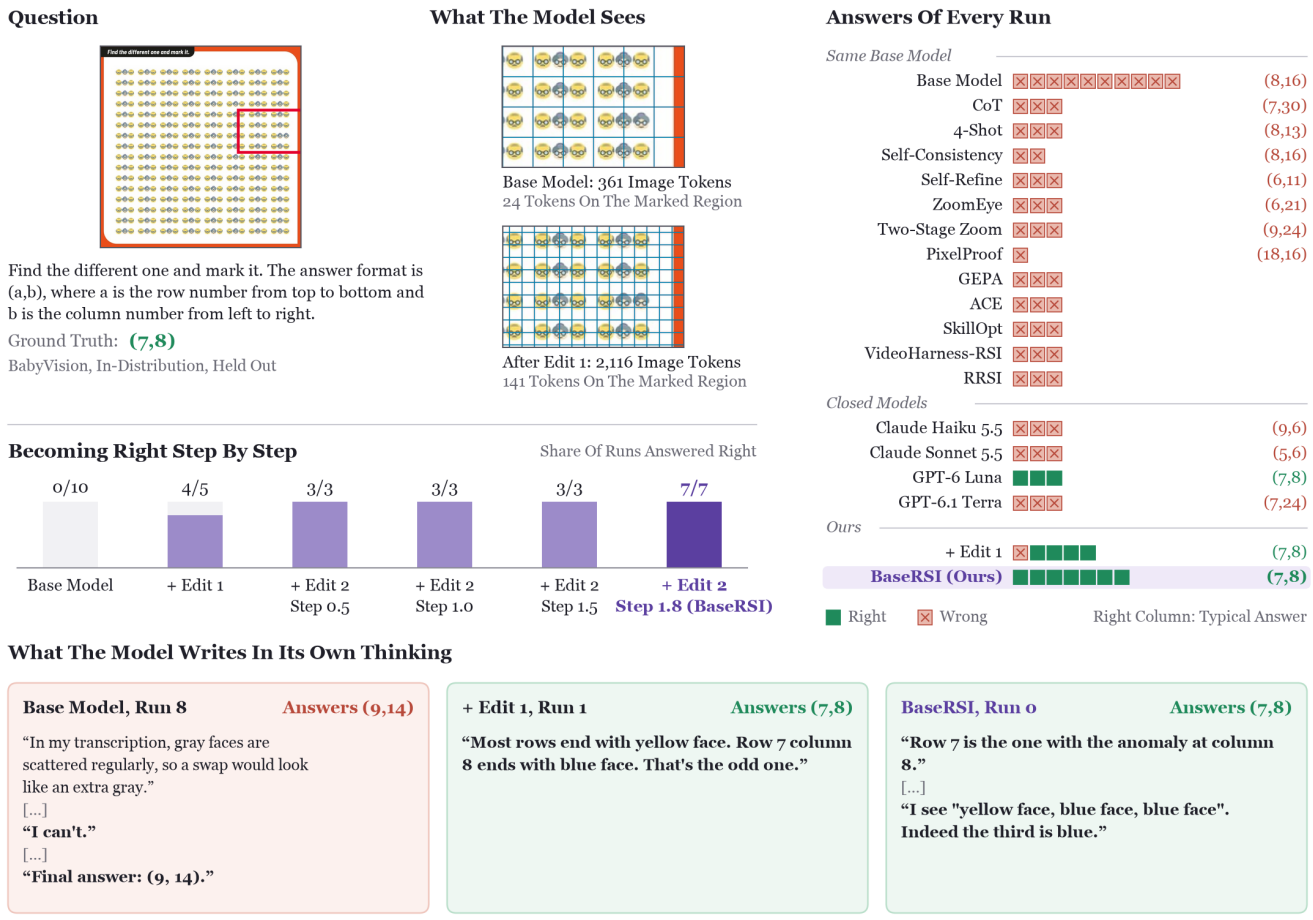


Figure 14. **Finding the odd one out** (BabyVision). The base model gives ten positions in ten runs; the view edit finds the face.

Key observations. The question asks for the row and the column of the one face that differs among 384 small faces, (7,8) in its answer format (Figure 14). The base model gives ten different positions in ten runs and spends up to 7,526 generated tokens on reconstructing the layout of the grid; in one run it writes that "gray faces are scattered regularly, so a swap would look like an extra gray", then "I can't", and answers (9,14). It reads the picture through 361 image tokens, 24 of them on the marked region, about one token per face. The view edit gives 2,116 tokens, 141 on the region, and the model then reads along the rows, "Most rows end with yellow face. Row 7 column 8 ends with blue face. That's the odd one"; it is right in 4 of 5 runs, BaseRSI in all 7, and every step of the procedure edit keeps the answer. The view edit is right more often here than the up-scaled image in the request, which finds the face in 2 of 5 runs. None of the 31 runs of the other methods on the same base model is right, and among the closed-source models GPT-6 Luna finds the face in all 3 runs at its highest effort and in 1 of 3 at its lowest, 4 of 21 runs over seven effort levels; the other closed-source runs name positions in other rows. The procedure edits of the earlier rounds find (7,8) in every run.

Base model, run 8: "The image shows rows of small circular faces, each with a color (yellow or gray/blue) and a hairstyle. Let me count the columns. [...] Hmm, that doesn't seem uniform. Let me think more carefully. [...] The faces are small circles of about 28px spacing. Let me try to reconstruct from typical generation: probably 24 columns?" *View edit, run 1:* "Each cell seems to be a group of three faces: yellow, blue, yellow. I need to scan row by row, column by column. Looking at row 7 (counting from top), the last column (8th column) appears to have the pattern yellow, blue, blue instead of yellow, blue, yellow. [...] Most rows end with yellow face. Row 7 column 8 ends with blue face. That's the odd one." *BaseRSI, run 0:* "Each group contains three emoji faces (yellow with glasses, blue with glasses, yellow with glasses). [...] The typical pattern in each group: yellow, blue, yellow (the middle one is blue/gray hair). The different one would have a different arrangement [...] Row 7 is the one with the anomaly at column 8. [...] I see 'yellow face, blue face, blue face'. Indeed the third is blue."

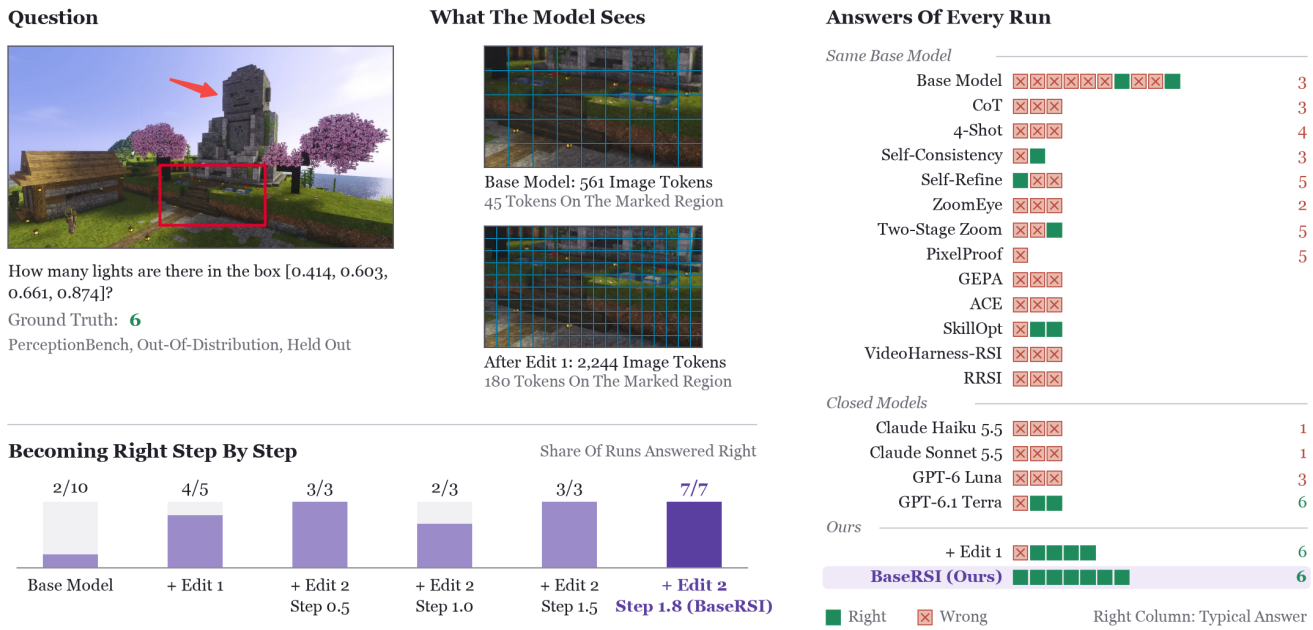


Figure 15. **Counting lights in a box** (PerceptionBench, held out). The view edit lets the model locate the small lanterns by their coordinates.

Key observations. The question asks how many lights lie in a given box of a game scene, six in all, on a benchmark that no part of BaseRSI saw (Figure 15). The base model is right in 2 of 10 runs and otherwise answers between two and five; in one run it places a light at "(0.712,0.742)", judges the coordinate "too big", and stops at five. The view edit raises the image tokens on the box from 45 to 180, and the model locates the small lanterns by their coordinates, right in 4 of 5 runs; the wrong run lists "a lantern near left of pool [...] and right of pool. That's 5". BaseRSI is right in all seven runs: in run 0 it lists ten candidate lanterns by their coordinates, along the path and around the monument, six of which lie inside the box, and then counts the lights around the pond and "two more" along the path, six in all. The up-scaled image in the request is right in 4 of 5 runs as well, and the procedure edits of the earlier rounds in all their runs, so the repair comes from the finer view. Most closed-source runs answer between one and three, and only GPT-6.1 Terra is right, in 2 of its 3 runs, 2 of 21 over seven effort levels; of the 31 runs of the eleven other methods on the same base model, 4 are right.

Base model, run 8: "Then near the stone steps/path center: (0.43, 0.83), (0.49, 0.87), (0.44, 0.84). Then on the right side near the monument base: (0.56, 0.74), (0.60, 0.72), (0.63, 0.72), (0.71, 0.74). Now the box x: 0.414-0.661, y: 0.603-0.874." *View edit, run 0:* "In that region I see the wooden steps/path leading up to the pool. There's a lantern on the left of the steps around (0.43, 0.79). There's a lantern at the bottom of steps around (0.497, 0.825). There's a lantern near the pool right side around (0.647, 0.703). There's a lantern near the pool left around (0.557, 0.681). There's a lantern at (0.443, 0.642) maybe. Hmm, this is too imprecise. [...] A lantern near left of pool (~0.557,0.681) and right of pool (~0.647,0.703). That's 5." *BaseRSI, run 0:* "The box is [0.414, 0.603, 0.661, 0.874]. In normalized coordinates, x1=0.414 (left), y1=0.603 (top), x2=0.661 (right), y2=0.874 (bottom). [...] Along the path in the foreground: there's a lantern around (0.318, 0.901), (0.361, 0.945), (0.433, 0.795), (0.497, 0.822). Around the monument base: lanterns at (0.445, 0.642), (0.497, 0.666), (0.557, 0.681), (0.647, 0.705), (0.720, 0.663), (0.792, 0.750). I need to be precise. [...] one near top-right of pond (0.497,0.666), one near bottom of pond (0.557,0.681), one to the right (0.647,0.705). [...] Then along the path two more (0.433,0.795) and (0.497,0.822). That's 6."

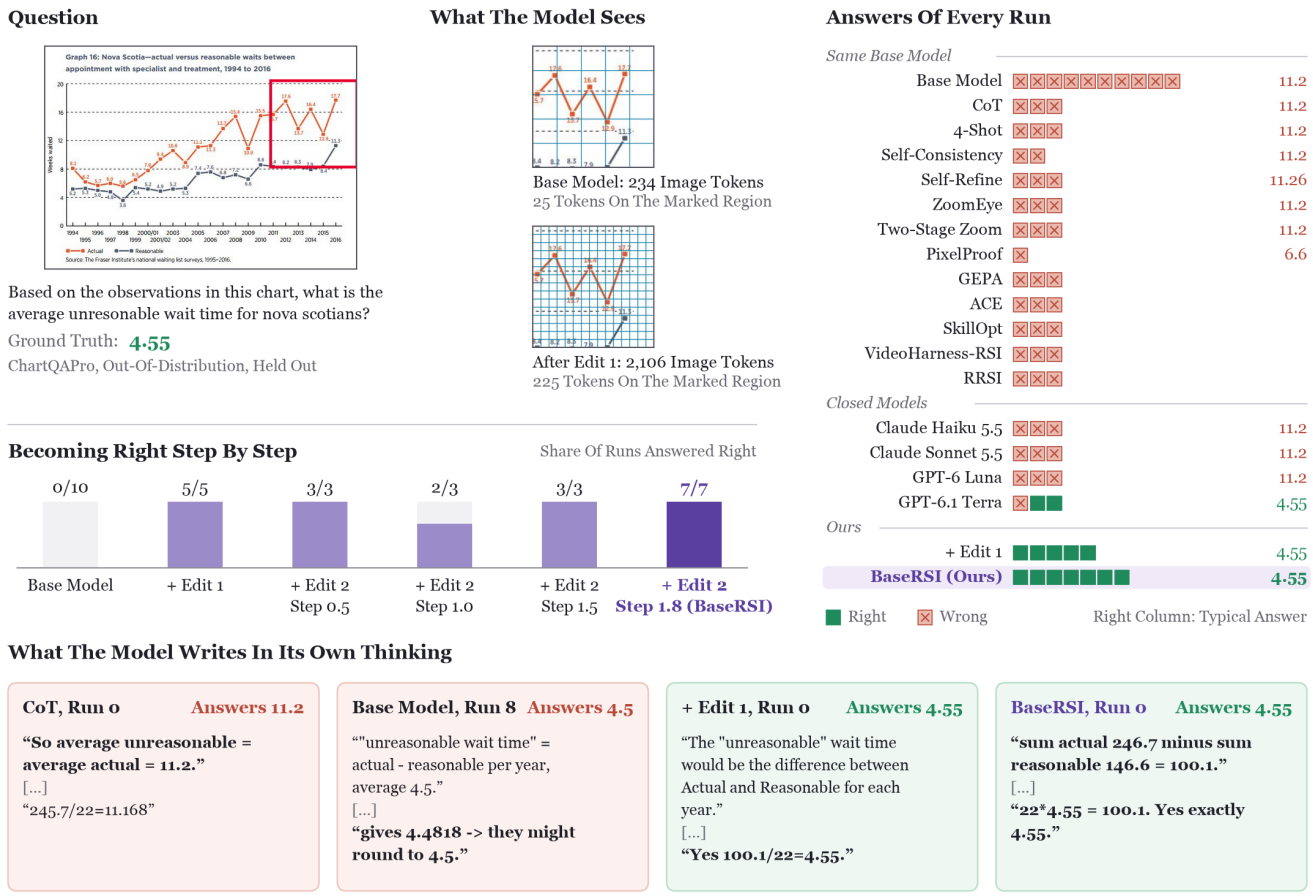


Figure 16. **Reading a chart** (ChartQAPro, held out). The base model and the other methods average the upper line; with the finer view the model averages the gap between the two lines, as the question asks, in all five runs.

Key observations. The question asks for the average “unreasonable” wait time in a chart of the actual and the reasonable wait for each year, that is, the mean gap between its two lines, 4.55 weeks (Figure 16). Six of the ten runs of the base model answer about 11.2, the mean of the upper line, and the chain-of-thought baseline makes this reading explicit: “average unreasonable = average actual = 11.2”. The base model has 25 image tokens on the marked region of the labels, and the view edit 225 of 2,106; the edited model reads every label and writes that “the ‘unreasonable’ wait time would be the difference between Actual and Reasonable for each year”, reaching 4.55 in all five runs. BaseRSI adds the two lines over the 22 years, “sum actual 246.7 minus sum reasonable 146.6 = 100.1”, and checks “22*4.55 = 100.1”. All 31 runs of the other methods on the same base model are wrong, almost all with 11.2, and over seven effort levels of the closed-source models 3 of 21 runs are right, two of them from GPT-6.1 Terra, and 17 answer 11.2. The up-scaled image in the request is right in 4 of 5 runs, so here too the finer view decides how the question is read.

Base model, run 8: “The question asks for ‘average unresonable wait time’ - likely meaning the actual wait time that exceeds reasonable, or perhaps they want the average of the ‘Actual’ values? Hmm. [...] Let’s list Actual values per year (1994-2016): 1994: 8.1, 1995: 6.2, 1996: 5.7, 1997: 6.0, 1998: 5.6, 1999: 5.5 [...]” *View edit, run 0:* “The ‘unreasonable’ wait time would be the difference between Actual and Reasonable for each year. [...] Actual (orange): 8.1, 6.2, 5.7, 6.0, 5.6, 6.5, 7.8, 9.4 [...] Reasonable (blue): 5.2, 5.3, 5.0, 4.8, 3.6, 5.4, 5.2, 4.9 [...] Alternatively, maybe they mean average of Actual minus average of Reasonable, which equals average of differences anyway. [...] Yes 100.1/22=4.55.” *BaseRSI, run 0:* “The unreasonable wait time would be the difference between actual and reasonable for each year, then averaged. [...] Actual (red): 1994: 8.1, 1995: 6.2, 1996: 5.7 [...] Reasonable (blue): 1994: 5.2, 1995: 5.3, 1996: 5.0 [...] Now compute differences (actual - reasonable) for each: 1994: 8.1-5.2=2.9, 1995: 6.2-5.3=0.9, 1996: 5.7-5.0=0.7 [...]”