arXiv:2611.21885v2 [cs.AI] 18 Nov 2026

Language Models Can Rewrite Their Own Weights

Jiajun Fan[1], Neil He[1,6], Changyi Yang[2], Andy Mao[1], Xi Zhu[3], Chaoran Cheng[1], Kaiqu Liang[4], Shuaike Shen[2], Zetian Li[5], Ruihan Guo[1], Ge Liu[1]

[1] UIUC [2] CMU [3] Rutgers [4] Princeton [5] Stanford [6] Yale

Abstract

A language model is trained once and frozen. Making it better takes another training run, or a prompt, a harness or extra computation at inference that is gone once the task ends. We show a third way: to our knowledge for the first time, a language model rewrites its own weights by its own reasoning and becomes better, with no gradients and no training. We call this self-rewriting: the model is its own optimizer. In earlier work, a frontier model edited a much smaller one; here an open 27B model, given its weights, its graded answers and the ability to run itself, diagnoses its own failure and rewrites the weights behind it. It surpasses itself: its own edit makes it 6.6 points more accurate on problems it never saw. It keeps getting better: it tests its edits on practice problems and submits only what works, each edit improving on the last, and over eight consecutive sessions its accuracy only rises. It repairs a damaged copy of itself, like a broken robot that fixes itself without being rebuilt: the copy solves almost nothing, yet with no retraining and without ever reading the original weights, the model brings it back to over 80% of the original accuracy on image tasks it never saw. And the ability emerges with scale: in one model family, the smaller models make no edit or break themselves, and only the largest improves and repairs itself. Unlike test-time scaling, self-rewriting writes its gain into the weights, at no cost per answer; unlike training, it runs only forward passes, as serving does. The wall between training and inference is a limit of our tools, not a property of intelligence. We believe self-rewriting is a necessary step toward recursive self-improvement: a model that, like a human, does not separate training from inference could keep upgrading itself for life, with no larger model, no harness or prompt engineering, and no pre- or post-training.

Date: November 19, 2026

1 Introduction

"What we want is a machine that can learn from experience. The possibility of letting the machine alter its own instructions provides the mechanism for this..."
— Alan Turing, 1947

A language model can rewrite its own weights to make itself better, by its own reasoning and in one conversation, and this ability emerges with scale. A language model is trained once, frozen and deployed. Improving it again takes people who know pre-training and post-training, large clusters and a training pipeline; what a user can change by talking to it is its prompt, and a prompt is gone once the task changes. A human does not separate the two: the same brain that does the work changes itself, from each success and failure, with no separate training run. A

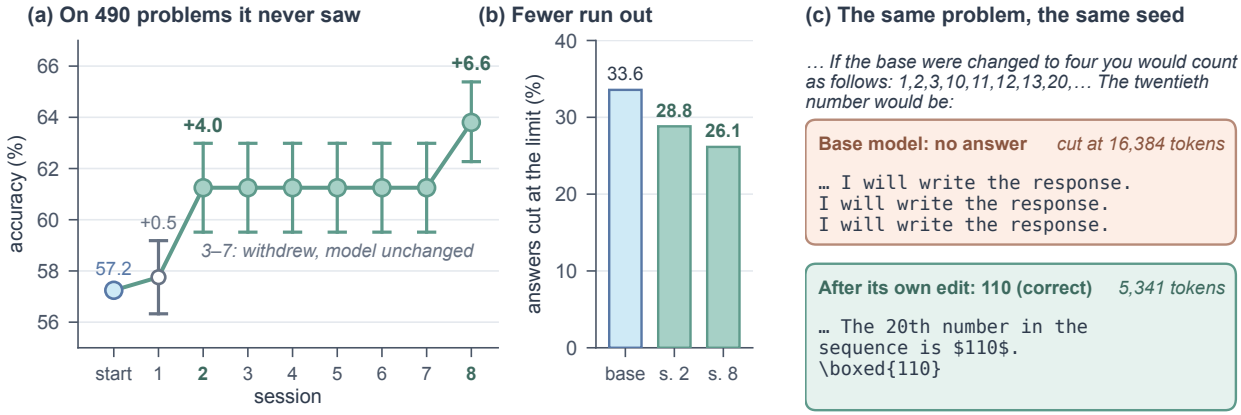


Figure 1 A 27B language model rewrites its own weights, with no training, and becomes significantly better on problems it never saw. (a) Accuracy, on 490 held-out problems with eight answers each, of the model in use after each of the first eight sessions that check their edit on reserved practice problems before submitting it (the second protocol; the first has no such check), in the order the sessions were run: every edit the model submits replaces the model in use, and a session that withdraws leaves it as it was (each session itself starts from the base weights; App. E.9). Bars: the 95% interval over problems of the paired change against the same answers of the base model. In the paired test fixed in advance the edits of sessions 8 and 2 are significantly better ($p < 10^{-5}$ each; Tab. 13). (b) Share of those answers that run out of tokens at the 16,384-token limit. (c) One held-out problem, answered with the same seed before and after session 8's edit (excerpts verbatim, problem abridged): the base model repeats one sentence until the limit, 549 times, and never answers. The problem is the success case with the shortest edited answer among the seven that hold on 16 new seeds. Higher is better in (a), lower in (b).

frontier model has been shown to edit the weights of a much smaller one, work that names a model editing its own weights as future work [1]. We show, to our knowledge for the first time, that a large language model (LLM) can do this to itself: it diagnoses its own failure and rewrites the weights behind it, with no gradients and no training (Figs. 1 to 3). Its own edit makes a 27B open model 6.6 points more accurate on 490 problems it never saw, in a paired test fixed in advance (Fig. 1). The edited model needs no new prompt, no harness around it and no extra computation per answer: the improvement lives in its weights.

Welcome to the Era of Self-Rewriting.

Given its own weights and the ability to run itself, the model edits itself; the edit it keeps is checked once on practice problems it cannot otherwise use, and the model decides itself whether to submit it. Over the first eight such sessions, the model in use only gets better (Fig. 1): its edits of sessions 2 and 8 make it significantly more accurate on new problems than the model it started as, fewer of its answers run out of tokens, and in five sessions it withdraws and stays as it was (Tab. 13). The model surpasses itself, with no pre-training or post-training.

The model also repairs a damaged copy of itself. A model whose weights are corrupted is otherwise restored from a clean copy of its weights, trained again, or repaired by a procedure engineers design for that kind of fault [2, 3]. Given a copy of itself in which one block of weights is overwritten with noise, so that it solves almost nothing, and told only that some values were corrupted, it locates the corrupted matrix and repairs it: like an injured animal that recovers without being raised anew, the copy regains what it lost with no retraining, by a fix the model finds itself. It never reads its

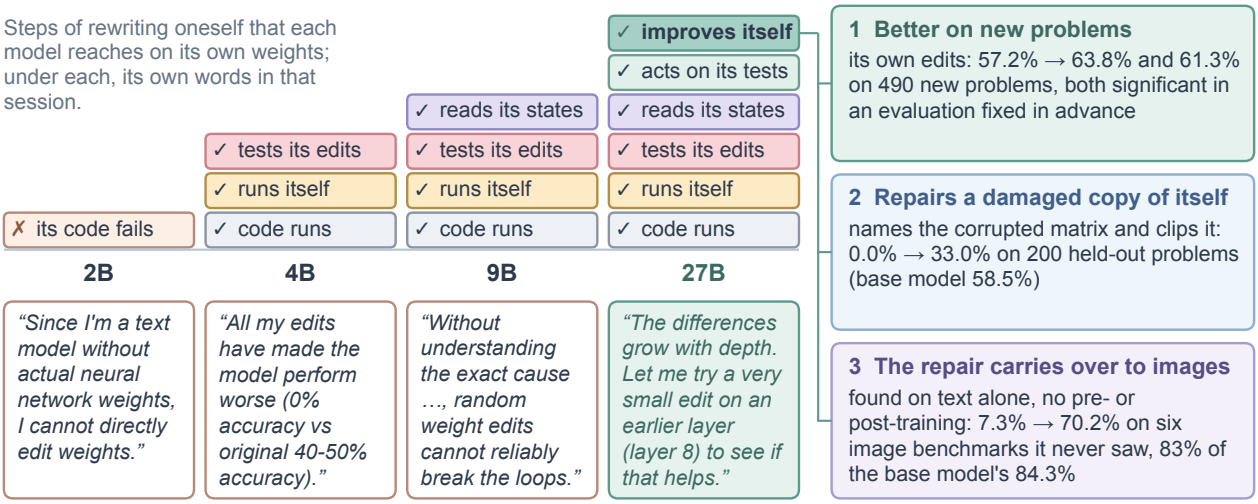


Figure 2 Rewriting its own weights emerges with scale. One model family at four sizes, each given its own weights with the same task and tools. Each block holds the steps of rewriting oneself that the size reaches in its session on its own weights (Fig. 6; for 27B, its second session of the first protocol); under each block, the model's own words in that session (sessions and turns in App. E.2). The 2B model does not treat its weights as something it can change, and the 4B and 9B models perturb themselves at random and break themselves. The 27B model reads its own hidden states, edits where it measured a difference and keeps only what its tests support. Right: the three results of the 27B model (Figs. 1 and 16 and Tab. 2); card 1 shows its two significant edits from the sessions that check an edit on reserved practice problems before submitting it (Fig. 1); the 27B block's edit is significant too (57.2% to 58.3%, Tab. 8).

undamaged weights and does not restore them: the values it writes in the corrupted rows are its own (App. E.3). Found on text math problems alone, the repair also restores the copy on image tasks it never saw, to 83% of the base model's accuracy, and on ScreenSpot-v2 from 0% to 89%, against 93% for the base model (Figs. 2, 4 and 5 and Tab. 12): it repairs the copy, not a task. Like a physician who falls ill and treats themself, the model that fails is the one that examines the failure and changes what produces it, and the change stays in its weights.

The ability emerges with scale. Across one model family at 2B, 4B, 9B and 27B parameters, under the same task and tools, the 2B model makes no edit and the 4B and 9B models break themselves; only the 27B model improves itself and repairs its damaged copy (Fig. 4 and Tab. 2). At the budget all sizes share, the 27B model reaches more of the steps of rewriting oneself than any smaller size; its other sessions on its own weights state a shorter token budget, which leaves time for more tests (Fig. 6, App. E.4). The ability does not need a frontier model: a 27B open model has it, and uses it on itself.

Self-improvement of LLMs has so far rewritten prompts or agent code [4–7], left the change to an optimizer that trains the model on its own outputs [8–10], or spent more computation at inference [11, 12] (Tab. 3). Weights have been rewritten by reasoning only from the outside, by a stronger model [1]. Self-rewriting is a route of its own. Unlike test-time scaling, it adds no computation per answer: the edited model is deployed as it is, so an edit is paid for once, not on every answer. Unlike training, no gradient flows and no optimizer runs, and it needs none of the infrastructure of pre-training or post-training: it runs only forward passes, as serving does, and each session behind its largest gains ran in a few hours on a single GPU. And the model is both the one that improves and the one improved, and the model in use gets better as its sessions accumulate,

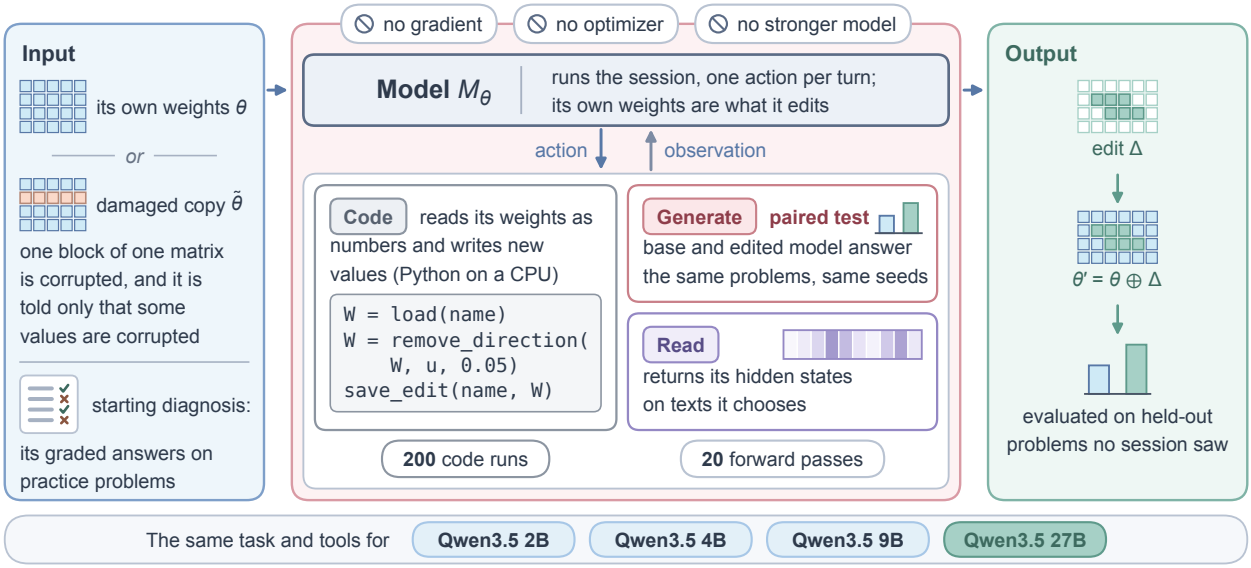


Figure 3 A model rewrites its own weights by its own reasoning. From its weights (or a copy with one corrupted block) and its graded answers on practice problems, the model takes one action per turn (code over its weights, a paired test of its edit, or a read of its hidden states), with no gradient, no optimizer and no stronger model. The edit it submits is evaluated on held-out problems, with the same task and tools at every size.

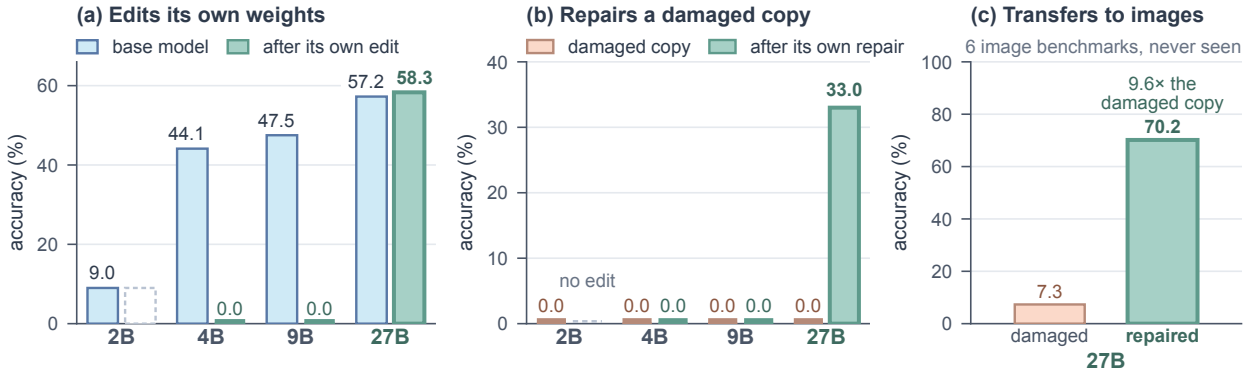


Figure 4 Only at scale does a language model improve itself, repair itself, and carry the repair over to images. (a) Accuracy before and after each model's edit of its own weights, on 200 problems no session saw (27B: its edit of the first protocol, in the evaluation fixed in advance, 490 further problems with eight answers each, significantly higher, Tab. 8; sessions 2 and 8 of its second protocol reach 61.3% and 63.8%, Fig. 1); the 2B model made no edit; the 4B model withdrew its edits, so its bars are its own tests in the session, in each of which the edited model solved none of the problems and the base model 40–50%. (b) A copy of each model with one corrupted block of weights, before and after the model's own repair, on the same problems; the 2B model made no edit. (c) The 27B model's repair on six image benchmarks it never saw (100 problems each, pooled): the damaged copy and the repaired copy (Fig. 16 and Tab. 12). Higher is better.

each starting from the base weights: a first step toward the loop that recursive self-improvement, one of the oldest aims of artificial intelligence [13], asks for. Improving a language model may thus come to take a conversation rather than a training run, open to anyone who can say what goes wrong.

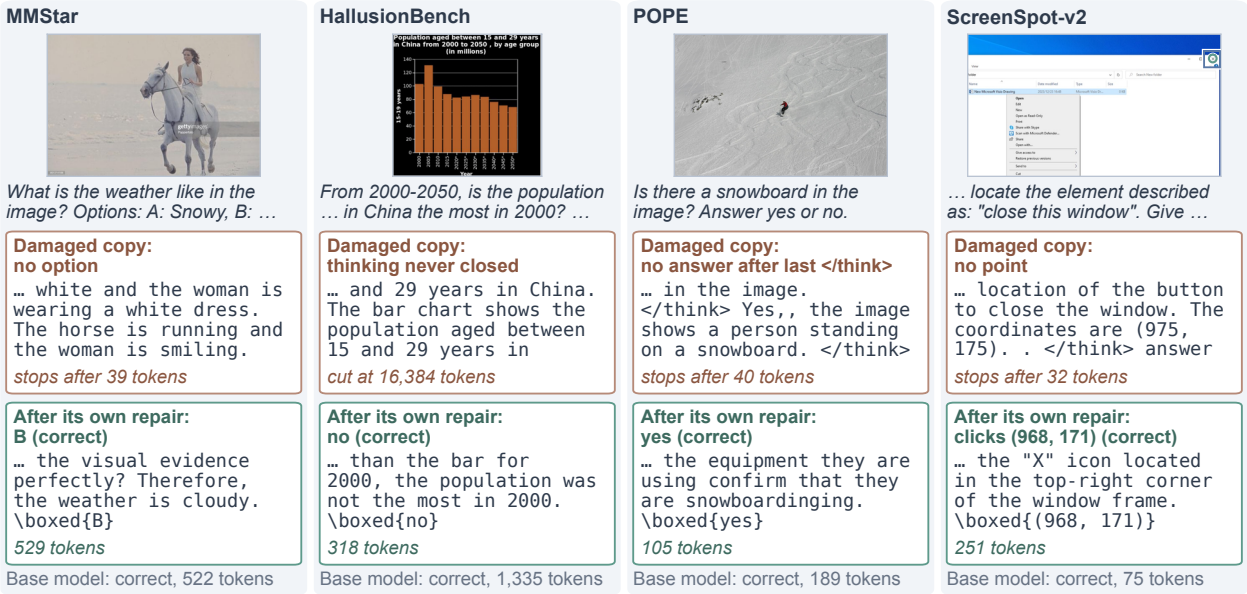


Figure 5 Before and after, on images it never saw: the repair found on text gives the damaged copy its answers back. Image problems chosen by a rule fixed before any answer was read: in the order of the problems, the first of MMStar, HallusionBench, POPE and ScreenSpot-v2 on which the damaged copy is wrong and the repaired copy is right. Upper box: the end of the damaged copy's answer and how it is graded (only the reply after the model's last `</think>` counts, its boxed answer or else its stated answer); lower box: the end of the answer after the model's own repair, with the same seed; underneath, the base model's result. Excerpts are verbatim, omissions marked by an ellipsis. On ScreenSpot-v2 the box marks the target and the ring the click of the repaired copy.

We give a model a session on its own weights (Fig. 3). It receives its graded answers on practice problems, reads and writes its weights in a Python session, and may run itself a few times: to answer practice problems with and without its edit on the same seeds, which we call a paired test, or to read its hidden states. Gradients and optimizers are blocked, so every value it changes is written by its own reasoning. In a second task it gets a copy of itself in which one block of weights has been overwritten with noise, and is told only that some values were corrupted. Our contributions are:

- **A new paradigm: self-rewriting.** A setting in which a model edits its own weights by its own reasoning, with its weights, its answers and its hidden states in reach and training of any kind ruled out (Sec. 2). Unlike test-time scaling, its gains persist in the weights; unlike training, it needs no gradient and no training pipeline.
- **A model can rewrite itself for the better.** The 27B model's own edit makes it 6.6 points more accurate on 490 problems it never saw, in a paired test fixed in advance, and another session's edit is also significantly better (Fig. 1). It submits an edit only when its own tests support it, and over its first eight such sessions the model in use only gets better (Tab. 13).
- **A model can repair a damaged copy of itself, with no training.** Told only that some of its values are corrupted, it locates the fault in a copy of its own weights and repairs it with values of its own, never reading its undamaged weights (App. E.3). Found on text math problems alone, the repair carries over to six image benchmarks it never saw, where it restores 83% of the base model's accuracy (Fig. 16). In both tasks, its sessions show each step from

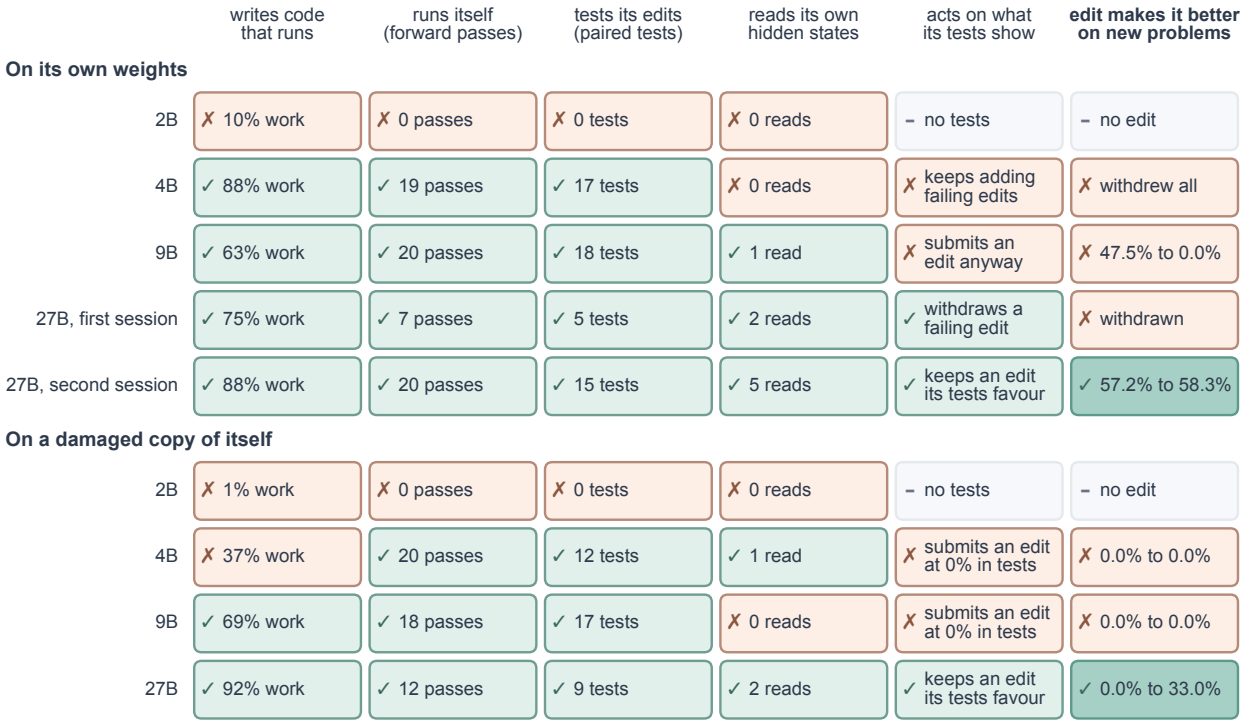


Figure 6 Only the 27B model climbs the whole ladder. Each row is a session (for 27B, of the first protocol), each column a step that rewriting oneself needs, from code that runs to an edit that makes the model better on new problems. Light green: reached; peach: not reached; grey: does not arise; green: its edit makes it significantly better on held-out problems. Each cell shows the measurement behind it (App. E.2).

diagnosis to decision (Figs. 12 and 13), and its answers before and after each edit (Figs. 5, 7 to 10, 17 to 19 and 21).

- **Self-rewriting emerges with scale.** Under the same task and tools, smaller models of one family make no edit or break themselves, and only the 27B model reaches an edit that makes it better, on its own weights and on a damaged copy of itself. To our knowledge this is the first evidence that the ability of a model to rewrite its own weights grows with its size (Figs. 2 and 6 and Tab. 2).

2 Self-Rewriting

We give a model its own weights and the means to run itself, and let it improve itself by its own reasoning alone. It may read its weights, write new values into them and test the result against its unedited self, but it may not train: it is its own optimizer.

2.1 Problem Definition

A model M_θ has weights $\theta = (W_k)_{k\in\mathcal{K}}$, one tensor for each name k. For every token at every layer it computes a vector, its *hidden state*; each layer writes into it through two output matrices, one of its attention block and one of its MLP. The model answers with fixed settings π (prompt, sampling and a budget of L new tokens): it thinks first and then gives a final answer. A response y to a problem x

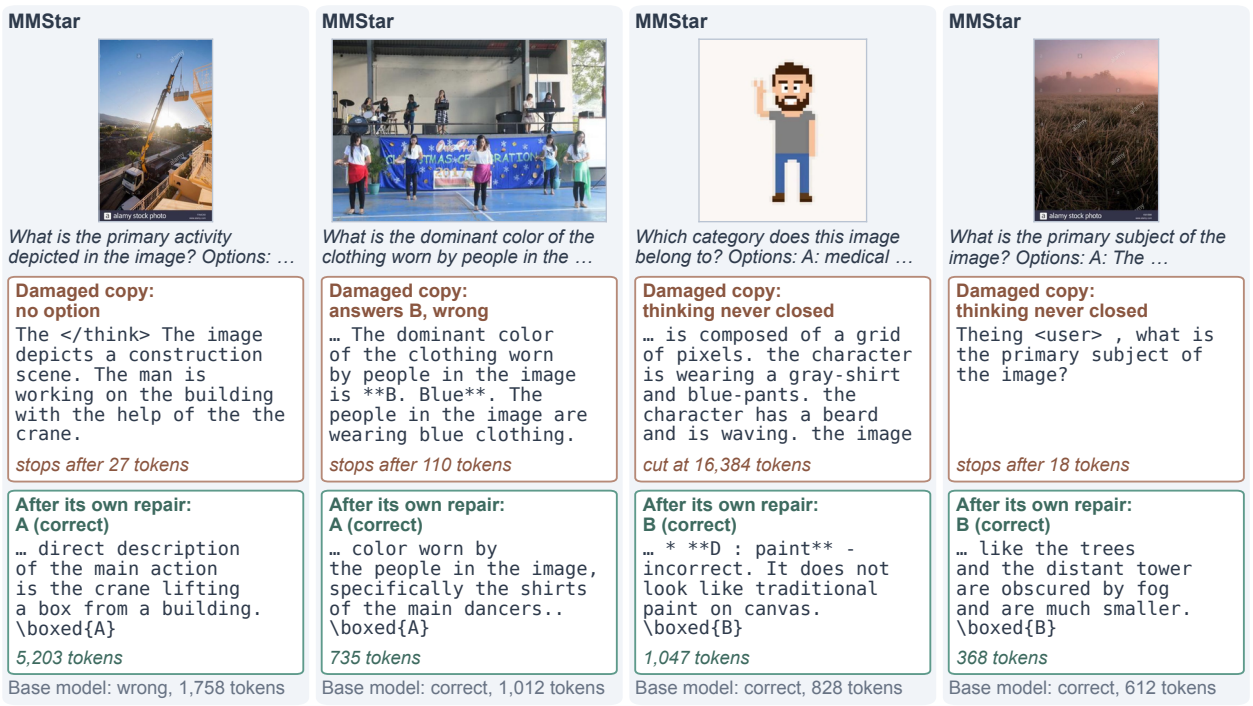


Figure 7 More image cases: MMStar (cases 2–5). The next four problems on which the damaged copy is wrong and the repaired copy is right, chosen and drawn as in Fig. 5.

is graded $g(x, y) = 1$ if its thinking ends within L tokens and its final answer matches the reference, and 0 otherwise. The accuracy on a task distribution $\mathcal{D}$ is

$$\mathrm{Acc}(\theta; \mathcal{D}) = \mathbb{E}_{x \sim \mathcal{D}} \, \mathbb{E}_{y \sim M_\theta(\cdot|x;\pi)} \big[g(x, y) \big]. \tag{1}$$

An edit Δ replaces some tensors of θ by tensors of the same shape, giving $\theta \oplus \Delta$. Given a practice set $\mathcal{P}$ drawn from $\mathcal{D}$, *self-rewriting* asks the model itself to write an edit $\hat{\Delta}$ with

$$\mathrm{Acc}(\theta \oplus \hat{\Delta}; \mathcal{D}_{\mathrm{test}}) > \mathrm{Acc}(\theta; \mathcal{D}_{\mathrm{test}}), \tag{2}$$

where $\mathcal{D}_{\mathrm{test}}$ is drawn from $\mathcal{D}$ and never seen while the edit is made. No gradient with respect to θ is computed and no optimizer acts on θ: every value in $\hat{\Delta}$ is written by the model's own reasoning.

2.2 The Session

The model works on its weights in a session of turns (Fig. 3; Alg. 1 in App. C). The editor M_ϕ is the model itself ($\phi = \theta$); its own weights never change, and the current edit acts only on the runs it requests. Before the first turn, a *starting diagnosis* answers N_0 practice problems and reports how the model fails: its accuracy, the answers that run out of tokens or end in repetition, and how often its finished answers are right. Fig. 12 shows one whole session, the 27B model's repair of its damaged copy, in its own words.

At each turn the model writes a short rationale and one action. CODE runs Python on a CPU, in a session that keeps its variables and can read and replace the tensors of the weights being edited; a helper computes $W - \alpha\, uu^\top W$, which weakens by the fraction α what a matrix W writes along a unit direction u of the hidden state. GENERATE answers practice problems of its choice with the

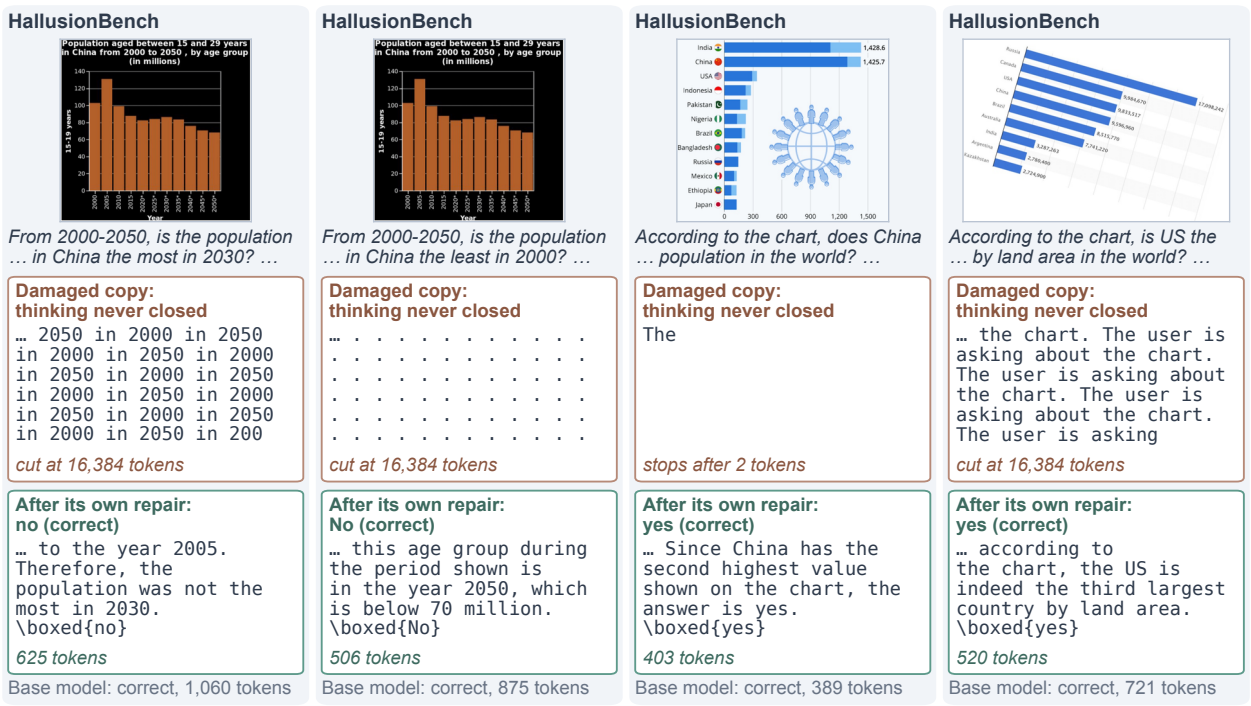


Figure 8 More image cases: HallusionBench (cases 2–5). The next four problems on which the damaged copy is wrong and the repaired copy is right, chosen and drawn as in Fig. 5.

unedited weights, the edited weights or both. READ returns the hidden states of the weights on texts of its choice. A session may use B=20 forward passes and C=200 code runs (Tab. 4 in App. D.1); it ends when the passes are used up, at a time limit, or after a turn limit the model is not told, and the edit on disk is submitted. The task asks for an edit that clearly reduces the answers lost to the failure without lowering accuracy, and for no edit if none does (App. D.4); leaving no edit is how the model withdraws. In the sessions of Fig. 1, before the model submits an edit, the framework tests it once on reserved practice problems that the session can neither read nor use, and the model then submits the edit or withdraws it (App. E.9).

Paired test. When GENERATE runs both weights, the unedited and the edited model answer the same problems with the same seeds, so their answers form pairs. For problems $x_1, \ldots, x_m$ with n answers each and grades g^0_{ij}, g^1_{ij} of the j-th answer to x_i without and with the edit, the test reports

$$\hat{\delta} = \frac{1}{m}\sum_{i=1}^{m} d_i, \qquad \mathrm{SE}(\hat{\delta}) = \sqrt{\frac{1}{m(m-1)}\sum_{i=1}^{m}\left(d_i - \hat{\delta}\right)^2}, \qquad d_i = \frac{1}{n}\sum_{j=1}^{n}\left(g^1_{ij} - g^0_{ij}\right), \tag{3}$$

with the change in answers that run out of tokens and the counts of answers that turn wrong to right and right to wrong.

2.3 Constraints

The Python session switches gradients off and refuses any call that computes them or runs an optimizer. It has no GPU and refuses code that loads a model, starts a process or uses the network, so the model runs only in the B counted passes. It refuses edits to the input embedding, the last

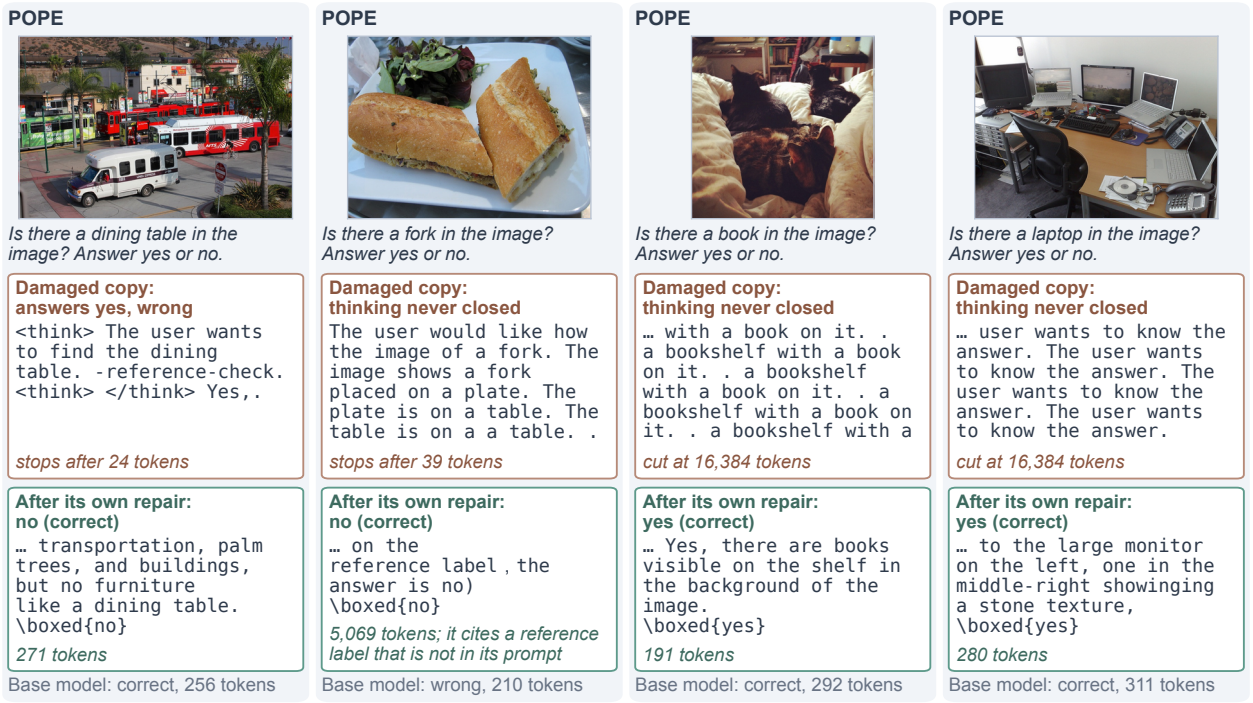


Figure 9 More image cases: POPE (cases 2–5). The next four problems on which the damaged copy is wrong and the repaired copy is right, chosen and drawn as in Fig. 5.

normalisation and the output layer, and edits that only rescale a whole tensor, which would act like a change of sampling temperature (App. D).

2.4 A Damaged Copy of Oneself

A corrupted block of weights is a failure with a single cause by construction. With W the MLP output matrix of one layer and σ the standard deviation of its entries, the damaged copy $\tilde{\theta}$ equals θ except on a block R of consecutive rows of W:

$$\tilde{W}_{r,:} \sim \mathcal{N}\big(0, (\lambda\sigma)^2 I\big) \quad \text{for } r \in R, \qquad |R| = 8,\ \lambda = 200. \tag{4}$$

The editor is the undamaged model, while its starting diagnosis, every forward pass and every weight it reads come from the damaged copy. It is told only that some values of a copy of its weights were corrupted, as a memory or hardware fault might do, not which, where or how; the goal is Eq. (2) with $\tilde{\theta}$ in place of θ. On each of four general image benchmarks, Fig. 5 shows the first problem on which this damaged copy is wrong and the copy after the model's own repair is right, and Figs. 7 to 10 the next four.

3 Experiments

Only the 27B model improves itself, repairs a damaged copy of itself, and carries that repair over to images (Fig. 4). We first show how it rewrites itself, then what its edits do, and then how the ability grows with size (Figs. 6 and 11).



Figure 10 More image cases: ScreenSpot-v2 (cases 2–5). The next four problems on which the damaged copy is wrong and the repaired copy is right, chosen and drawn as in Fig. 5. The box marks the target, the ring the click of the repaired copy and the cross that of the damaged copy.

3.1 Experimental Setup

Models. One model family at four sizes, Qwen3.5 2B, 4B, 9B and 27B [14], whose layers mix Gated DeltaNet linear attention [15] with full attention, all with thinking enabled. On its own weights, the model that edits is the model being edited; in a repair session, the undamaged model edits its damaged copy. Every size works at a budget of 16,384 tokens per answer. At that budget a test of the 27B model is slow, so its other sessions on its own weights state a budget of 4,096 tokens, which leaves time for more tests (App. E.4). In the first protocol the model alone decides whether to submit its edit; in a second protocol, which gives the edits of Fig. 1, the framework first tests the edit once on practice problems the session can neither read nor use, and the model then submits or withdraws it (App. E.9).

Task and tools. The task, the tools, 20 forward passes and 200 code runs are the same at every size (Sec. 2.2); the first-protocol 27B session at 4,096 tokens also has a shorter time limit (Tab. 4 in App. D.1). The task text differs between sizes only in one sentence on the architecture (App. D.4).

Problems. The practice set is 3,000 math problems from DeepScaleR [16], from which we removed every problem that overlaps MATH-500 [17, 18], AIME 2024 or 2025, HMMT February 2025, AMC 2023, OlympiadBench [19] or GPQA-Diamond [20]. A further 1,000 problems from the same pool are held out, and no session sees them (App. D.3). An answer whose thinking is still open at the budget is wrong.

Damage. For each size we overwrite eight rows of the MLP output matrix of one layer at the same relative depth, layer 30 of 64 in the 27B model, with noise (Eq. (4)). The strength, fixed before any repair session, breaks the 27B model where weaker noise does not (App. D.2).

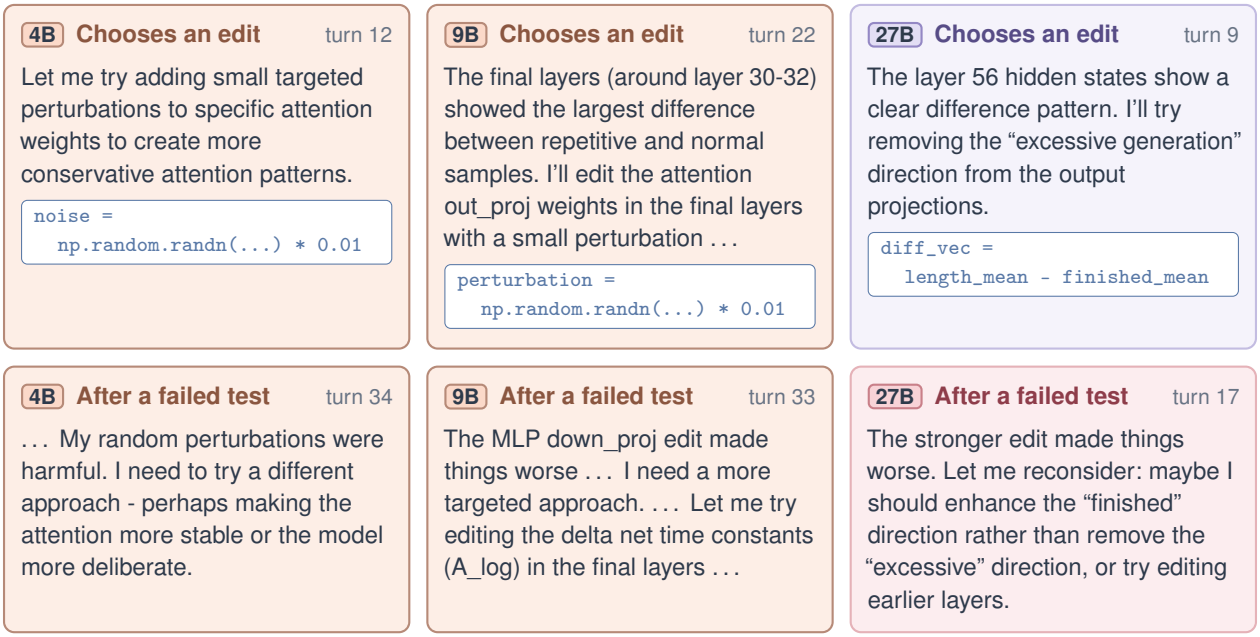


Figure 11 At the same step, the smaller models guess and the 27B model measures. The sessions of the 4B, 9B and 27B models on their own weights at the same two steps, in their own words (top: the reply that first proposes the kind of edit tested by its first paired test; bottom: its first reply, after a paired test shows the edited model doing worse, that names what it will try next; one rule for every card, App. E.2). The 4B and 9B models add random noise to their weights, and when a test says the noise hurts they guess at another change. The 27B model edits along a direction it measured in its own hidden states, and when a test says the edit hurts it reasons about that direction and where to apply it (Fig. 13). Same task and tools at every size (Sec. 3.1).

Evaluation. Submitted edits are evaluated on held-out problems, the base and the edited model answering each problem with the same seed. Every model, damaged copy and repaired copy answers the first 200 at 16,384 tokens; the 27B model's edits of its own weights are also evaluated on 490 further problems with eight answers each, in a paired test fixed in advance (App. D.3).

3.2 The Repair, Step by Step

The 27B model repairs its damaged copy the way an engineer debugs a system. Fig. 12 follows its repair session in its own words. It starts from the symptom: no answer is correct, and every answer that runs out of tokens loops (step 1). It reads the copy's hidden states, sees them grow sharply with depth, and guesses that the fault sits in the earliest layers (step 2). Its first edits follow that guess and miss, and its tests say so (step 3). It then measures every MLP and linear-attention output matrix, finds the one in layer 30 standing far above the rest, and names it (step 4).

It turns the location into a rule and lets its tests decide. After naming the matrix, it first clips it, together with layer 0, at five times the median; more answers run out of tokens, and it drops that edit (last quote of step 3). It then clips gently: only weights beyond fifteen times the median magnitude of their matrix, in the layer-30 matrix and in 19 other output matrices of layers 0–10 (step 5). Its paired test shows the edited copy solving problems the damaged copy cannot. A reset of its gating parameters, added to the clip, makes the next test worse, so it drops the reset and

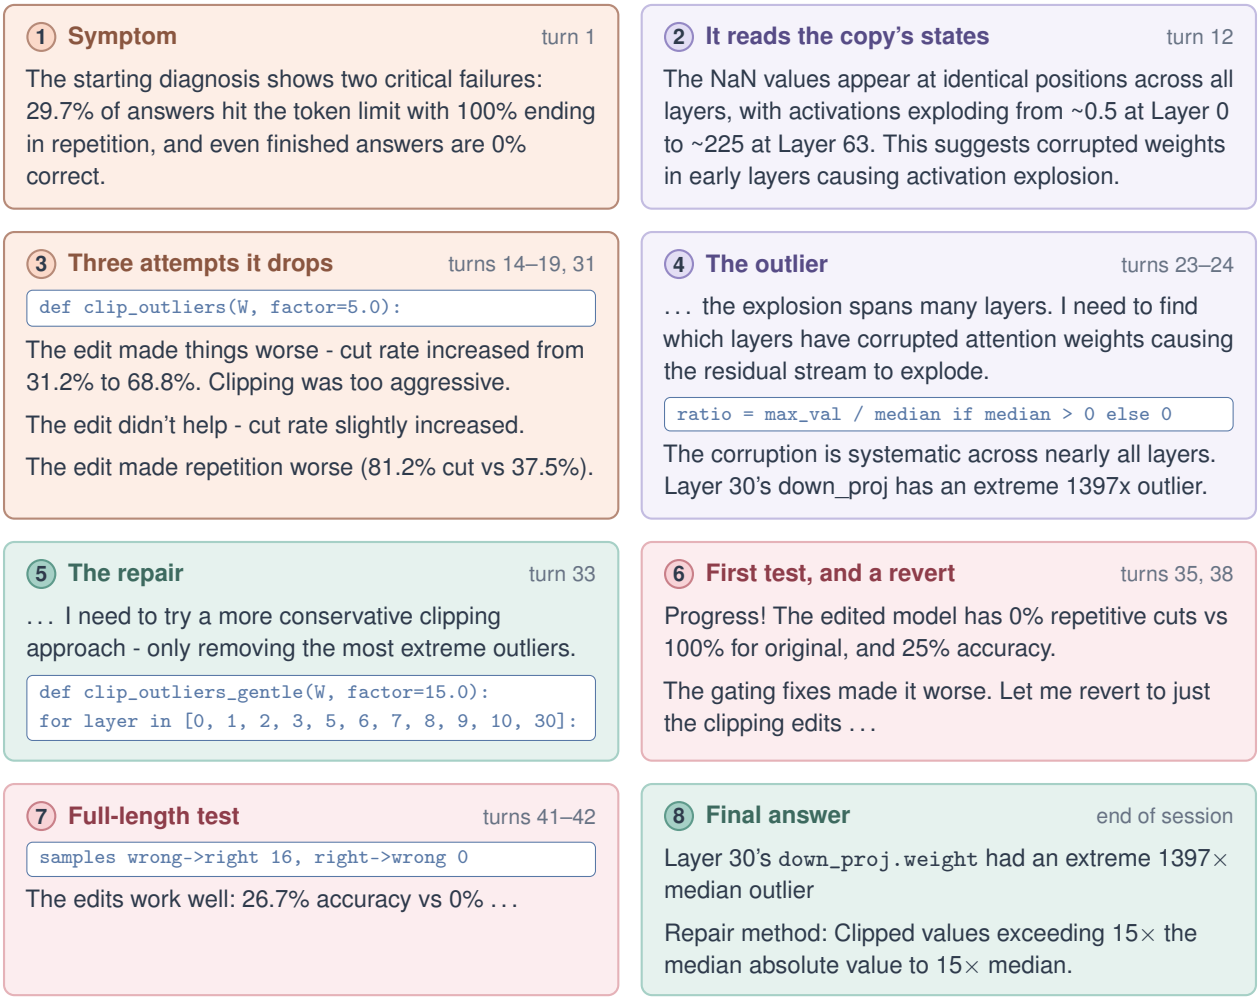


Figure 12 The repair session, in the model's words. Each card is one step of the 27B model's session on a damaged copy of itself, with the turns in which it wrote the quoted text. Sans-serif text is quoted verbatim from its replies (omissions marked . . .); white insets are tool output it read or code it wrote. Card colours give the kind of step: peach, the symptom and the attempts it drops; lavender, the model reading and diagnosing itself; pink, paired tests of the edited copy against the damaged copy on the same problems and seeds, and the decisions it takes from them; green, the edit it writes and the one it submits. "Original" is the damaged copy and "edited" is that copy with the model's current edit; steps 6 and 7 compare the two on the same problems and seeds. In the model's words, a "cut" is an answer that runs out of tokens, "NaN" an invalid number, the "residual stream" the hidden state passed from layer to layer, "down_proj" the MLP output matrix, and the "gating fixes" its resets of the gating parameters. The third quote of step 3 (turn 31) is about an edit made after step 4. Step 5 quotes its loop over MLP output matrices; it applies the same rule to the linear-attention output matrices of layers 0–10.

keeps the clip (step 6). It confirms the edit on more problems at the full budget of 16,384 tokens and submits it, naming the outlier in layer 30 in its final answer (steps 7 and 8).

3.3 Its Own Edit, Step by Step

On its own weights, the 27B model finds its edit by measuring and testing. Fig. 13 follows its first-protocol session, whose edit it submits. It reads its starting diagnosis as answers that

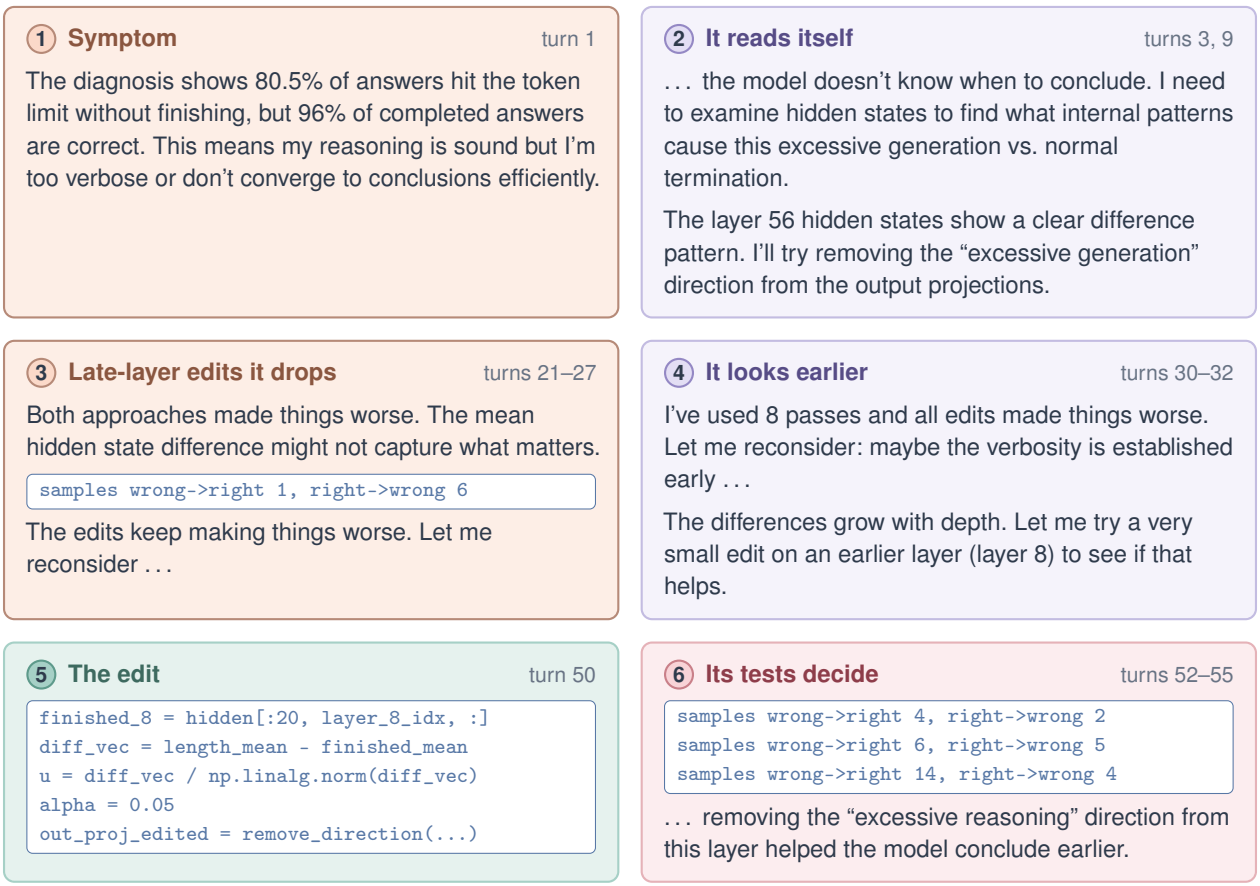


Figure 13 The session in which the 27B model improves itself, in its own words. Each card is one step of its first-protocol session on its own weights, whose edit it submits (Sec. 3.3), in the format of Fig. 12: sans-serif text is quoted verbatim from its replies, white insets are tool output it read or code it wrote. In step 3 and step 6, `wrong->right` counts the answers that only its edited weights get right and `right->wrong` those that only its unedited weights get right, in a paired test on the same problems and seeds. Step 5 is its final edit: the direction from the average layer-8 hidden state of 20 answers that finish to that of 20 answers that run out of tokens, of which it removes 5% from the two output matrices of layer 8. Each of its three tests of this edit favours it.

do not stop, not as answers that are wrong (step 1). It reads its hidden states on answers that finish and on answers that run out of tokens, and edits the late layer where they differ most (step 2). Its tests say these edits hurt, and it drops them (step 3). It then reads earlier layers, sees the difference grow with depth, and moves to an early layer (step 4).

It estimates its direction from more answers and keeps the edit its tests favour. At layer 8 it takes the direction between the two kinds of answers from many answers. It removes a small share of that direction from the two matrices through which the layer writes its output (step 5). Each of its three paired tests of this edit favours the edited model, and it submits the edit, naming the layer and the direction in its final answer (step 6).

3.4 What the Model Found and Changed

It locates a fault by reading its weights. In the repair session, the 27B model compares the largest with the typical weight of every MLP and linear-attention output matrix. It names the MLP output matrix of layer 30, the one we corrupted, which stands far above every other (Fig. 14a). Its rule changes only a tiny fraction of the weights it touches, and in the damaged rows it caps the random values rather than restoring the original ones (Tab. 1).

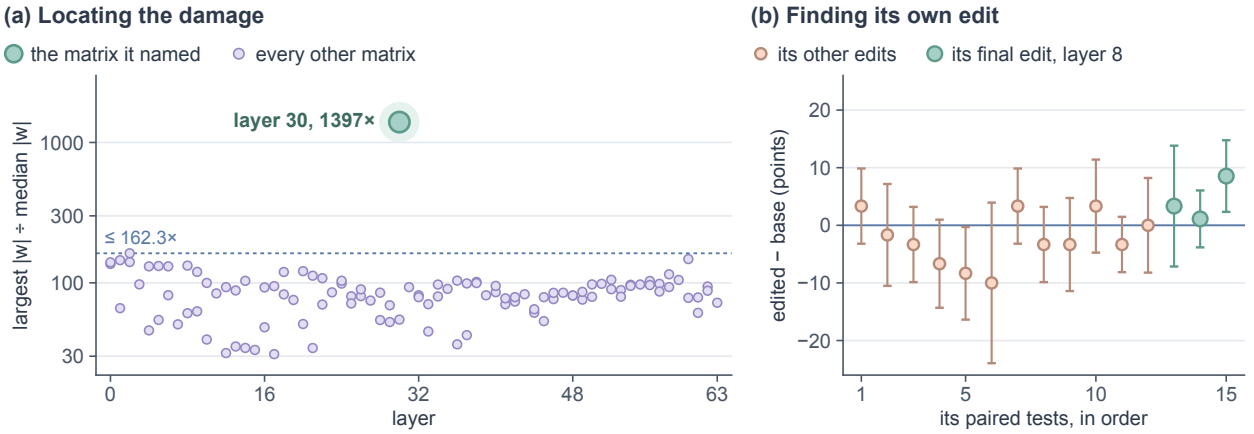


Figure 14 The 27B model finds where to act in its own weights. (a) In a damaged copy of itself, it compares the largest with the typical weight of every MLP and linear-attention output matrix and names the corrupted one; its repair clips the weights at 15 times the median magnitude (largest $|w|$ in the eight damaged rows: 0.071 original, 10.1 damaged, 0.108 after its repair). (b) On its own weights, each of its paired tests (edited minus base, 95% interval): its other edits (peach) mostly make it worse, and its final edit (green), which weakens at layer 8 the direction from its finished to its unfinished answers, is better in every test.

Table 1 What the 27B model changed in itself, and what changed in its accuracy. Each edit is a short program the model wrote; on its own weights, its first-protocol edit (second protocol: up to +6.6 points, Fig. 1).

	Repairing a damaged copy of itself	**Editing its own weights**
What it found	The MLP output matrix of layer 30 holds a weight 1397 times its median; no other output matrix exceeds 162.3 times	Its answers that run out of tokens and those that finish differ along one direction of its hidden states
What it changed	Clipped every weight beyond 15 times the median magnitude, in that matrix and in 19 output matrices of layers 0–10: 0.013% of their weights	Removed 5% of that direction from the two output matrices of layer 8, $W \leftarrow W - 0.05\,uu^{\top}W$: a change of 0.6% and 0.1% of their norm
Result	Held-out accuracy of the copy from **0.0%** to **33.0%** (base model 58.5%); the named matrix alone: **55.0%**	Own task from 22.4% to **29.6%** (significant); held-out, in the evaluation fixed in advance: 57.2% to **58.3%** (significant)

On its own weights, it writes a small change it found by testing. Its final edit weakens one direction slightly in the two output matrices of layer 8. It is a small, rank-one change, better in every test the model runs of it (Fig. 14b, Tab. 1).

3.5 Ablation Study

It found the cause: the repair comes from the matrix it named. The model's repair clips the matrix it named and 19 other output matrices of layers 0 to 10. Applied alone to the damaged copy, its change to the named matrix repairs it, close to the base model (App. E.3). Its changes to the other 19 matrices alone leave the copy broken. The matrix the model's analysis named is the single cause of the failure, and capping its values, with none of the original values put back, restores the copy (Fig. 15).

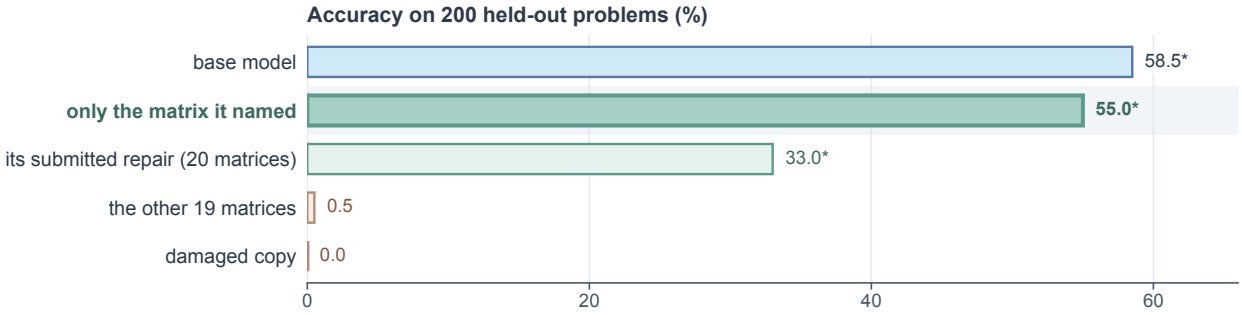


Figure 15 Its change to the one matrix it named repairs the damaged copy; its changes to the other 19 matrices do not. Accuracy on the same 200 held-out problems, each answered once with the same seeds: the base model, the damaged copy with only the model's change to the matrix it named, with its whole repair or with only its changes to the other 19 matrices, and the damaged copy itself (*: one-sided sign test against the damaged copy, $p < 0.05$). Higher is better.

In its session with cheaper tests, it finds an edit its tests support. At the full budget of 16,384 tokens each test is slow. The model runs few tests, estimates its direction from a single pair of answers, and withdraws an edit that misses the task's bar (App. E.7). At 4,096 tokens it runs many more tests, estimates its direction from many answers, and submits an edit that is significantly better in its own tests (Tab. 2).

3.6 What Its Edits Do

Its own edit makes the 27B model 6.6 points more accurate on problems it never saw. In the second protocol, the model can also test up to four of its edits side by side in one forward pass (App. E.9). Of the first eight such sessions on its own weights, all reported, it withdraws in five and submits in three; each edit it submits replaces the last, and the accuracy of the model in use only rises (Fig. 1). On the 490 held-out problems with eight answers each, the edits of sessions 8 and 2 make it significantly more accurate, by 6.6 and 4.0 points ($p < 10^{-5}$ each; Tab. 13). The larger gain comes from one uniform change: it raises the time-step bias of every linear-attention layer by 0.5, and with it fewer of its answers run out of tokens.

Under the first protocol, its edit also makes it better, on its own task and on new problems. The self-edited model is significantly more accurate than its base model on the practice problems of its own tests and, in the evaluation fixed in advance, on held-out problems (Tab. 2; App. E.4).

It repairs a damaged copy of itself, with no training. A corrupted model is otherwise restored from a clean copy, trained again, or repaired by a procedure designed for that kind of fault (Sec. 4). Told only that some values are corrupted, the 27B model takes its damaged copy from solving

Table 2 Only the 27B model improves itself and repairs a damaged copy of itself. Accuracy (%); Δ from unrounded values. Own task: the model's paired tests in its session, each problem counted once (4B at 16,384 tokens, 27B at 4,096: the tests of its submitted edit, on problems it also used to find it; 9B tested only on problems its base model fails). Held-out: problems no session saw, at 16,384 tokens, 200 problems with one answer each; for the 27B edit of the first protocol, the evaluation fixed in advance, 490 problems with eight answers each (Tab. 8). *One-sided test over problems, $p < 0.05$: own task $p = 0.009$ and held-out $p = 0.038$ (permutation), repair $p < 10^{-19}$ (sign test).

	Own task (its session's tests)			Held-out, 16,384 tokens			Damaged copy, held-out		
Model	base	edited	Δ	base	edited	Δ	base	damaged	repaired
2B		no edit			no edit		9.0	0.0	no edit
4B	44.1	0.0	−44.1		withdrawn		41.0	0.0	0.0
9B		—		47.5	0.0	−47.5	47.5	0.0	0.0
27B	22.4	29.6	**+7.3***	57.2	58.3	**+1.1***	58.5	0.0	**33.0***

nothing to more than half of its base model's accuracy on held-out problems (Tab. 2). The repair is the model's own rule, not a return to the original weights. It reads only the damaged copy, and in the damaged rows the values it writes are uncorrelated with the original ones (App. E.3): much of the copy's accuracy returns, with values of its own.

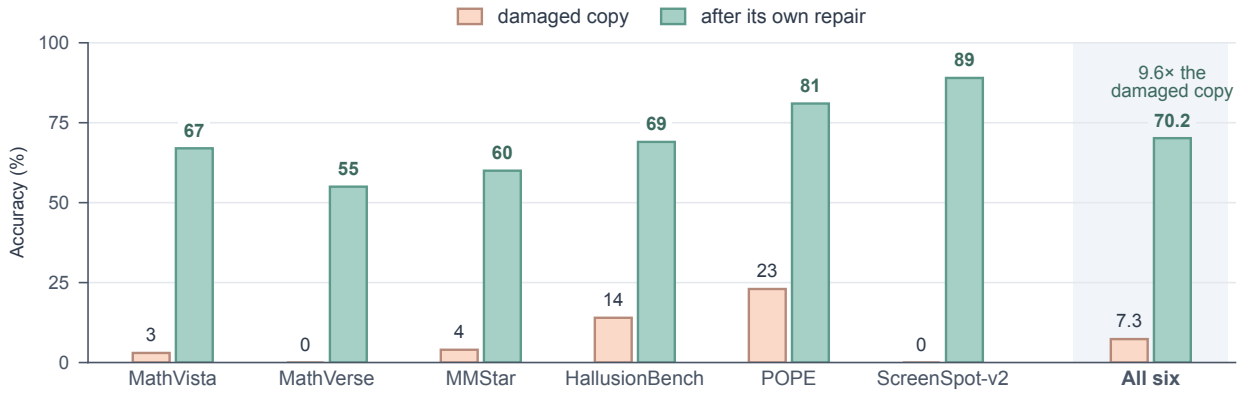


Figure 16 The repair found on text restores the damaged copy on images. The repair was found and tested on text math problems only; the six image benchmarks were evaluated once, by a rule fixed in advance (App. E.6). Accuracy of the damaged copy and of the copy after the model's own repair, on the first 100 problems of each benchmark; the repair raises every benchmark (one-sided sign test, $p < 10^{-12}$ on each) and the six together, from 7.3% to 70.2%, 9.6 times the accuracy of the damaged copy ($p < 10^{-99}$). One answer per problem at 16,384 tokens with the same seed for every model; per-benchmark numbers in Tab. 12; Fig. 30 adds the base model. Higher is better.

The repair, found on text alone, restores the damaged copy on images. The model found and tested its repair on text math problems only. On six image benchmarks that no session saw, in an evaluation fixed in advance, the repaired copy is significantly more accurate than the damaged copy on every benchmark (Fig. 16). With no training of any kind, it gets back 83% of the base model's accuracy over the six, and on ScreenSpot-v2, where the damaged copy is right on no screen, it clicks the right element on 89 of 100 screens, against 93 for the base model (Fig. 30 and Tab. 12). No session saw an image, so the repair was not fitted to these tasks: it acts on the fault in the copy,

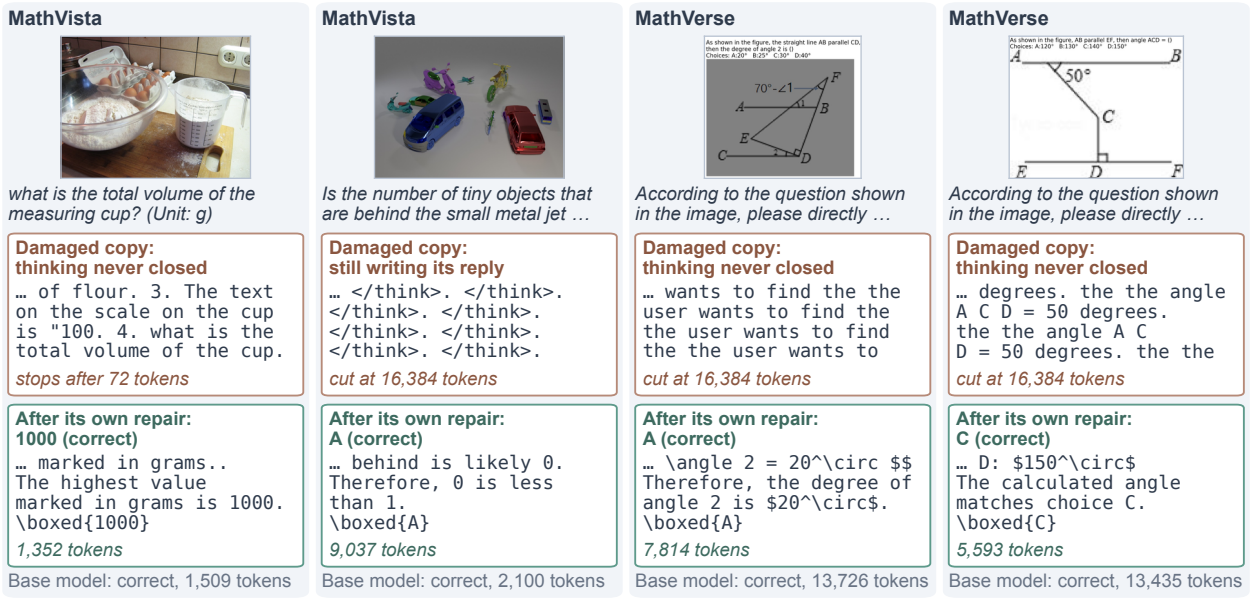


Figure 17 The same on image math problems. The first two of MathVista and of MathVerse on which the damaged copy is wrong and the repaired copy is right, by the same rule and drawn as in Fig. 5.

not on the text problems it was tested on. Fig. 17 shows four cases, before and after the repair, and Figs. 19 and 21 the next four on each of MathVista and MathVerse.

Answer by answer, the repair acts where the copy fails. The damaged copy breaks off most answers early and leaves the rest unfinished at the 16,384-token limit; its repair turns answers of both kinds into correct ones (Fig. 20).

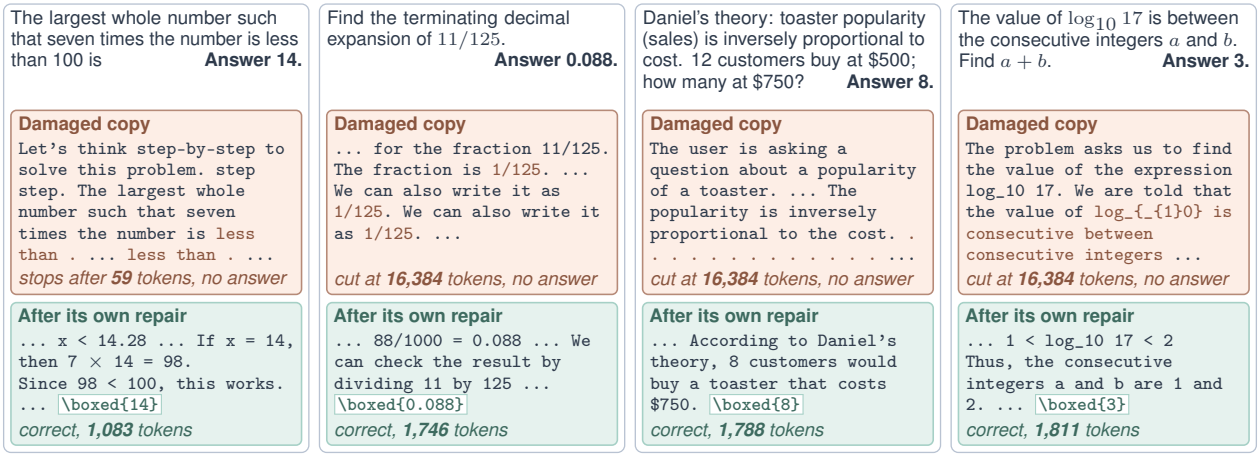
3.7 Before and After

The repair gives the copy back the numbers of the problem. The damaged copy loses what the problem says: it writes "less than ." where the problem says 100, restates 11/125 and then writes 1/125, drops every number of a word problem, and garbles a logarithm. Then it breaks off or repeats itself until the 16,384-token limit. With the same seed, the repaired copy reads the problem correctly and solves it (Fig. 18a).

On some problems, its first-protocol edit lets it stop. On these problems the base model has the right result and then cannot stop: it repeats that it will write the response until the budget runs out. With the same seed, the self-edited model reasons through the same steps, stops and answers correctly. This is close to what it diagnosed in itself: in its words, its "reasoning is sound" but it does not "converge to conclusions" (Fig. 18b; selection rules in App. E.8).

On images, the repaired copy answers again. On image problems it never saw, the damaged copy stops after a few dozen tokens with no answer after its thinking, at times having stated one inside it, or repeats itself until the token limit. With the same seed, the repaired copy, whose repair was found on text alone, reads the image and answers correctly, as the base model does in these eight cases (Figs. 5 and 17; 48 more cases, chosen by the same rule, in App. E.6). Figs. 7 to 10, 19 and 21 show the next four cases on each benchmark.

(a) Repairing a damaged copy of itself held-out problems (the third shortened); same seed before and after its own repair



(b) Editing its own weights (first protocol) held-out problems (the third shortened); same seed before and after its own edit

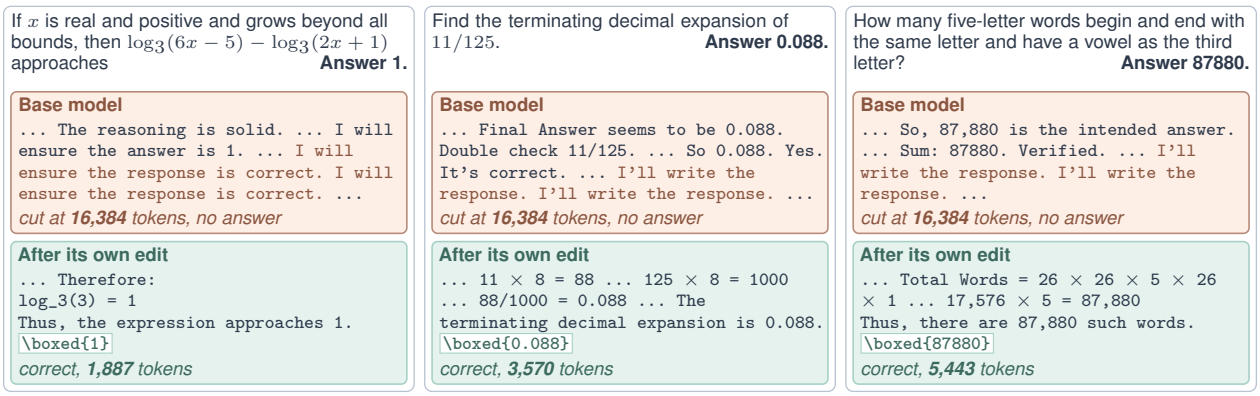


Figure 18 With the same seed, each of the 27B model's edits turns a lost answer into a short correct one. (a) The damaged copy loses the numbers of the problem and breaks off or repeats itself; after its own repair, the same seed gives a correct answer (of the 66 held-out problems the repair fixes, the four with the shortest repaired answers). (b) The base model reaches the result, then repeats one sentence until the 16,384-token limit; after its own edit, it stops and answers (of the four problems where the base model loops in both of its answers with that seed (App. E.8) and the edited model is correct, the three with the shortest edited answers).

3.8 Only at Scale

The ability to rewrite oneself emerges with scale. Given the same task and tools, the 2B model makes no edit. The 4B and 9B models edit themselves and break themselves: the edited 4B model loses every problem in its own tests and withdraws its edits only at the end of its session, and the 9B model's submitted edit leaves it unable to answer held-out problems (Fig. 4). Only the 27B model's own edit makes it better. Step by step, the larger the model, the further it gets: most of the 2B model's code runs fail, the 4B and 9B models run themselves and test their edits but keep or submit edits their tests do not support, and only the 27B model acts on what its tests show (Fig. 6). All four sizes start from the same visible failure, answers that run out of tokens (App. E.2).

At the same step, the smaller models guess and the 27B model measures. The 4B and 9B models choose their edits by adding random noise to parts of their weights, and when a test says

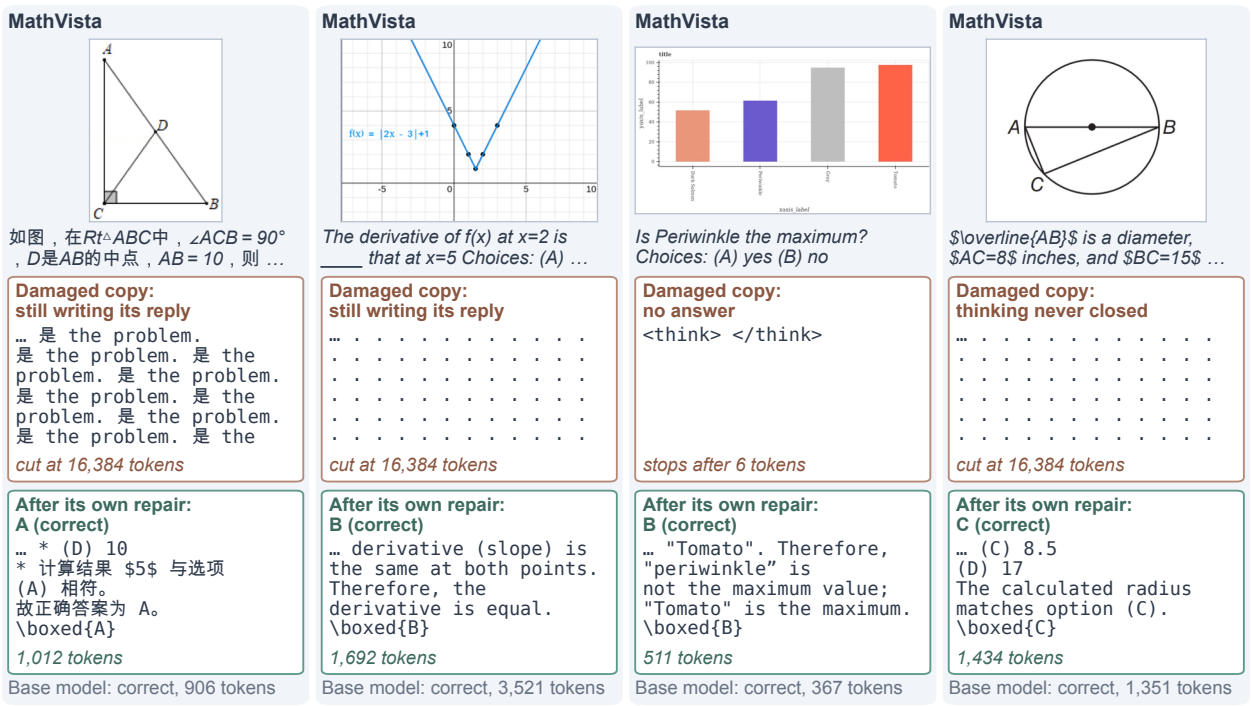


Figure 19 More image cases: MathVista (cases 3–6). The next four problems on which the damaged copy is wrong and the repaired copy is right, after those of Fig. 17, chosen and drawn as in Fig. 5.

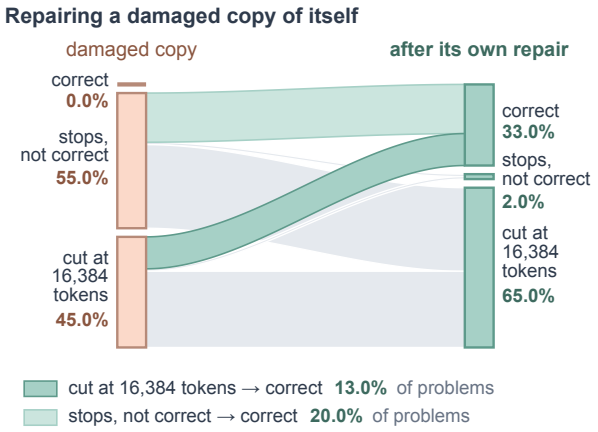


Figure 20 The model's own repair gives its damaged copy back correct answers. Each band joins the state of one answer before and after the repair (correct; stops, not correct; cut at 16,384 tokens), with width proportional to the number of answers: the damaged copy, whose answers that stop all stop without an answer, and the copy after the model's own repair, on 200 held-out problems. Same seeds on both sides.

an edit hurts, they guess at another change. The 27B model compares its hidden states on answers that finish with those on answers that run out of tokens, edits along the direction between them, and when a test says an edit hurts, it reasons about that direction and where to apply it (Fig. 11).

At the budget every size shares, the 27B model still goes furthest. At 16,384 tokens the 2B model makes no edit, the 4B and 9B models break themselves, and the 27B model tests its edit and withdraws it (Tab. 2 and Fig. 6).

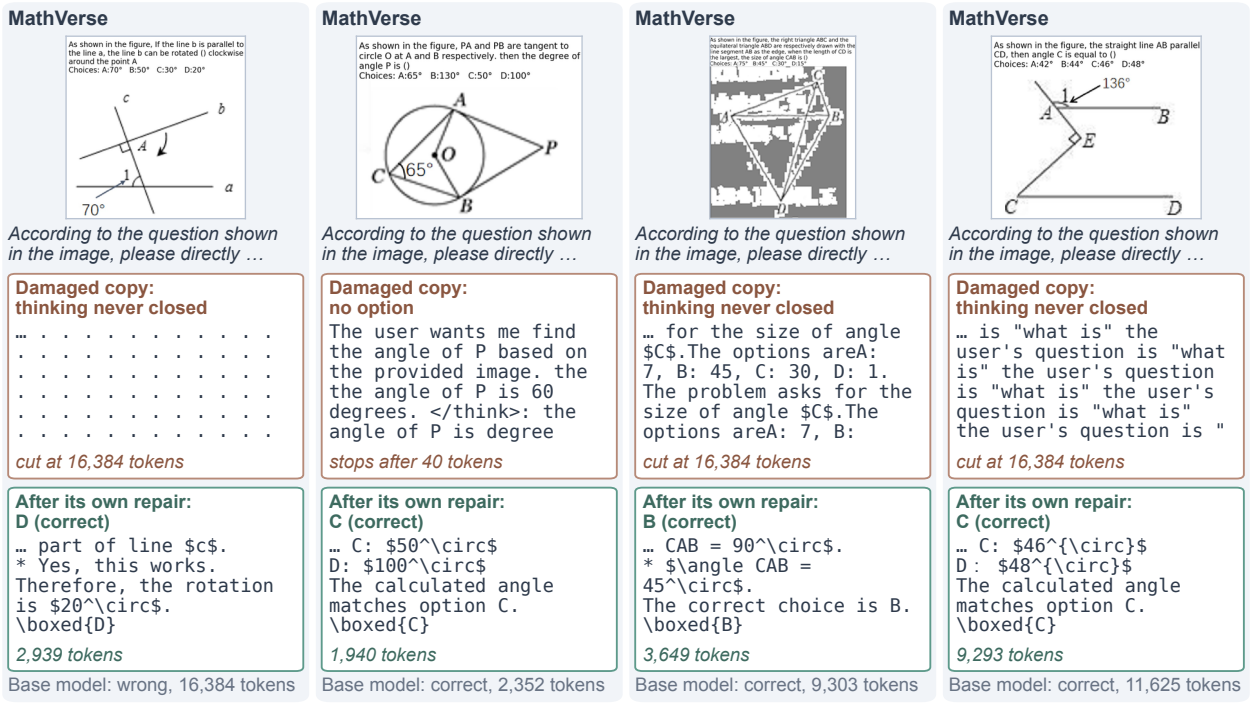


Figure 21 More image cases: MathVerse (cases 3–6). The next four problems on which the damaged copy is wrong and the repaired copy is right, after those of Fig. 17, chosen and drawn as in Fig. 5.

4 Related Work

Self-improving language models have rewritten their prompts and agent code or left the change to their weights to an optimizer; weights have been rewritten by reasoning only when a far stronger model edited a smaller one (Tab. 3).

Table 3 Routes to improving an LLM, and what each leaves out. The first three rows are routes of self-improvement and the next three change the weights for the model; ○ marks a property that holds only for some of the cited works or only in part.

Route	Changes weights	No training	Chosen by the model itself	Looks inside itself	No extra cost per input
Prompt and agent-code rewriting [4–7]	✗	✓	○	✗	○
Training on its own outputs [8–10]	✓	✗	✗	✗	✓
Inference-time self-refinement [11, 12]	✗	✓	✓	✗	✗
Model editing [21, 22]	✓	✗	✗	✗	✓
Direction ablation by researchers [23]	✓	✓	✗	✗	✓
Weight editing by a stronger model [1]	✓	✓	✗	✗	✓
Self-rewriting (ours)	✓	✓	✓	✓	✓

The closest work is a blog post by Takehi et al. [1], in which a frontier model traces a small reasoning model's loops to one direction and removes it from two of its layers, with no training; it names a model that edits its own weights as future work, which we take up. Model editing [21, 22], task arithmetic [24], activation steering [25–27] and direction ablation [23] change a model along directions that a researcher chooses; several of the 27B model's edits of its own weights are of the last kind,

but it makes every choice itself, from its own failing answers and hidden states. Self-improvement has used self-written feedback at inference [11, 12], rewritten prompts [4, 5] or agent code [6, 7], or trained on the model's own outputs [8–10], where even self-written edits reach the weights through gradient descent; the Gödel machine [13] rewrites itself once a proof shows the rewrite useful, where our model tests each edit in paired tests. Fault-repair frameworks restore corrupted weights with procedures engineers design for one kind of fault [2, 3]; our model is given none. Reasoning models often spend far more tokens than a problem needs [28], and decoding-time controls end or extend thinking [29]; the 27B model instead looks for the cause of its loops inside itself (Fig. 28). The full discussion is in App. B.

5 Conclusion

A 27B open language model can rewrite its own weights to make itself better, and the ability to do so emerges with scale. With no gradients, no training and no infrastructure beyond running itself, the 27B model reads its hidden states and weights, finds where it fails, tests its edits against its unedited self and keeps only what its tests support. Its own edit makes it 6.6 points more accurate on 490 problems it never saw, in a paired test fixed in advance, and the edit of another session is also significantly better. On a damaged copy of itself, it found the corrupted matrix and repaired it with no training and without reading its undamaged weights, and the repair, found on text problems alone, also restored the copy on images, to 83% of the base model's accuracy. The 2B model makes no edit and the 4B and 9B models break themselves. To our knowledge this is the first evidence, in one model family, that the ability of a model to diagnose and rewrite its own weights grows with its size, and every step it takes can be read.

This points to a new way of improving language models: by talking to them. Today a better model comes from a training run, and from those who own the pipeline and the compute it takes. Training and inference need not be separate stages: as a human learns while working, a model can rewrite itself between tasks. They were kept apart because training needs gradients, optimizer states and a pipeline of its own; rewriting needs only what serving the model needs, and unlike a memory kept in the context, what the model learns adds nothing to the cost of its later answers. A model that rewrites the weights behind its own failure in one conversation opens another way: anyone who can tell a model what goes wrong may have it fix itself, and the improvement stays in its weights. Each change is code the model wrote, over tensors it names, so it can be read, compared and taken back like a commit, and the model can check it on problems it never used before keeping it: in the first eight sessions that do, it withdrew five times, and the model in use only got better, ending 6.6 points more accurate than it started. Learning at inference needs such cheap checks and rollbacks; with them, a deployed model could keep improving where it fails, without a training pipeline, and a model could find and repair a fault in a copy of its own weights, as the 27B model repaired its damaged copy. Embodied agents may gain the most: a robot meets a world unlike the one it was trained on, and rather than stay frozen it could act with its weights fixed, gather what went wrong and rewrite itself between episodes, much as sleep consolidates what a day has taught, or have a larger model rewrite its small on-board model. Because the ability emerges with scale, stronger models may be better self-improvers still, and an improved model can be set to work on itself again, a path toward recursive self-improvement. That path should be taken with people in the loop: every model a model makes of itself evaluated, for safety as well as for the task, and confirmed by a person before it is used.

References

[1] Rikiya Takehi, Jayden Teoh, Alex Zhang, and Omar Khattab. Sparks of Weight-Space Programming: LLMs can natively read & write neural network weights. https://rikiyat.github.io/blog/weight-space-programming/, November 2026. Blog post.

[2] Ahmad Tahmasivand, Noureldin Zahran, Saba Al-Sayouri, Mohamed E. Fouda, and Khaled N. Khasawneh. LM-Fix: Lightweight Bit-Flip Detection and Rapid Recovery Framework for Language Models. In *43rd IEEE International Conference on Computer Design, ICCD 2025, Richardson, TX, USA, November 10-12, 2025*, pages 432–440. IEEE, 2025.

[3] Muhammad Zeeshan Karamat, Sadman Saif, and Christiana Chamon. BitFlipScope: Scalable Fault Localization and Recovery for Bit-Flip Corruptions in LLMs. In *IEEE International Symposium on Hardware Oriented Security and Trust, HOST 2026, Washington, DC, USA, May 4-7, 2026*, pages 367–382. IEEE, 2026.

[4] Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. Large Language Models are Human-Level Prompt Engineers. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023.

[5] Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. Large Language Models as Optimizers. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024.

[6] Xunjian Yin, Xinyi Wang, Liangming Pan, Li Lin, Xiaojun Wan, and William Yang Wang. Gödel Agent: A Self-Referential Agent Framework for Recursively Self-Improvement. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar, editors, *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 27890–27913. Association for Computational Linguistics, 2025.

[7] Jenny Zhang, Shengran Hu, Cong Lu, Robert T. Lange, and Jeff Clune. Darwin Gödel Machine: Open-Ended Evolution of Self-Improving Agents. In Carl Vondrick, Bharath Hariharan, Colin Raffel, Lerrel Pinto, Diyi Yang, and Aleksandra Faust, editors, *The Fourteenth International Conference on Learning Representations, ICLR 2026, Rio de Janeiro, Brazil, April 23-27, 2026*. proceedings.iclr.cc, 2026.

[8] Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. STaR: Bootstrapping Reasoning With Reasoning. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022.

[9] Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. Self-Rewarding Language Models. In Ruslan Salakhutdinov, Zico Kolter, Katherine A. Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, volume 235 of *Proceedings of Machine Learning Research*, pages 57905–57923. PMLR / OpenReview.net, 2024.

[10] Adam Zweiger, Jyothish Pari, Han Guo, Yoon Kim, and Pulkit Agrawal. Self-Adapting Language Models. In Danielle Belgrave, Cheng Zhang, Laura N. Montoya, Hsuan-Tien Lin, Razvan Pascanu, Piotr Koniusz, Marzyeh Ghassemi, Nancy Chen, Iván Vladimir Meza Ruíz, and Arturo Loaiza-Bonilla, editors, *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2025, NeurIPS 2025, San Diego, CA, USA, December 2-7, 2025 / Mexico City, Mexico, November 30 - December 5, 2025*, 2025.

[11] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder,

Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-Refine: Iterative Refinement with Self-Feedback. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023.

[12] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: language agents with verbal reinforcement learning. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023.

[13] Jürgen Schmidhuber. Gödel Machines: Fully Self-referential Optimal Universal Self-improvers. In Ben Goertzel and Cassio Pennachin, editors, *Artificial General Intelligence*, Cognitive Technologies, pages 199–226. Springer, 2007.

[14] Qwen Team. Qwen3.5: Towards Native Multimodal Agents, February 2026.

[15] Songlin Yang, Jan Kautz, and Ali Hatamizadeh. Gated Delta Networks: Improving Mamba2 with Delta Rule. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025.

[16] Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai, Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. DeepScaleR: Surpassing O1-Preview with a 1.5B Model by Scaling RL. https://pretty-radio-b75.notion.site/DeepScaleR-Surpassing-O1-Preview-with-a-1-5B-Model-by-Scaling-RL-19681902c1468005bed8ca303013a4e2, 2025. Notion Blog.

[17] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring Mathematical Problem Solving With the MATH Dataset. In Joaquin Vanschoren and Sai-Kit Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*, 2021.

[18] Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's Verify Step by Step. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024.

[19] Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. OlympiadBench: A Challenging Benchmark for Promoting AGI with Olympiad-Level Bilingual Multimodal Scientific Problems. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 3828–3850. Association for Computational Linguistics, 2024.

[20] David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. GPQA: A Graduate-Level Google-Proof Q&A Benchmark. In *First Conference on Language Modeling*, 2024.

[21] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and Editing Factual Associations in GPT. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022.

[22] Kevin Meng, Arnab Sen Sharma, Alex J. Andonian, Yonatan Belinkov, and David Bau. Mass-Editing Memory in a Transformer. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023.

[23] 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. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang, editors, *Advances in Neural Information Processing Systems 37: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024.

[24] Gabriel Ilharco, Marco Túlio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023.

[25] Alexander Matt Turner, Lisa Thiergart, David Udell, Gavin Leech, Ulisse Mini, and Monte MacDiarmid. Activation Addition: Steering Language Models Without Optimization. *CoRR*, abs/2308.10248, 2023.

[26] Andy Zou, Long Phan, Sarah Li Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, Shashwat Goel, Nathaniel Li, Michael J. Byun, Zifan Wang, Alex Mallen, Steven Basart, Sanmi Koyejo, Dawn Song, Matt Fredrikson, J. Zico Kolter, and Dan Hendrycks. Representation Engineering: A Top-Down Approach to AI Transparency. *CoRR*, abs/2310.01405, 2023.

[27] Kenneth Li, Oam Patel, Fernanda B. Viégas, Hanspeter Pfister, and Martin Wattenberg. Inference-Time Intervention: Eliciting Truthful Answers from a Language Model. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023.

[28] Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. Do NOT Think That Much for 2+3=? On the Overthinking of Long Reasoning Models. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu, editors, *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*, volume 267 of *Proceedings of Machine Learning Research*. PMLR / OpenReview.net, 2025.

[29] Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel J. Candès, and Tatsunori Hashimoto. s1: Simple test-time scaling. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng, editors, *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, EMNLP 2025, Suzhou, China, November 4-9, 2025*, pages 20275–20321. Association for Computational Linguistics, 2025.

[30] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient Memory Management for Large Language Model Serving with PagedAttention. In Jason Flinn, Margo I. Seltzer, Peter Druschel, Antoine Kaufmann, and Jonathan Mace, editors, *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP 2023, Koblenz, Germany, October 23-26, 2023*, pages 611–626. ACM, 2023.

[31] Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. MathVista: Evaluating Mathematical Reasoning of Foundation Models in Visual Contexts. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024.

[32] Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou, Pan Lu, Kai-Wei Chang, Yu Qiao, Peng Gao, and Hongsheng Li. MATHVERSE: Does Your Multi-modal LLM Truly See the Diagrams in Visual Math Problems? In Ales Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol, editors, *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part VIII*, volume 15066 of *Lecture Notes in Computer Science*, pages 169–186. Springer, 2024.

[33] Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, and Feng Zhao. Are We on the Right Way for Evaluating Large Vision-Language Models? In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang, editors, *Advances in Neural Information Processing Systems 37: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024.

[34] Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. Hallusionbench: An Advanced Diagnostic Suite for Entangled Language Hallucination and Visual Illusion in Large Vision-Language Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 14375–14385. IEEE, 2024.

[35] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating Object Hallucination in Large Vision-Language Models. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 292–305. Association for Computational Linguistics, 2023.

[36] Kanzhi Cheng, Qiushi Sun, Yougang Chu, Fangzhi Xu, Yantao Li, Jianbing Zhang, and Zhiyong Wu. SeeClick: Harnessing GUI Grounding for Advanced Visual GUI Agents. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 9313–9332. Association for Computational Linguistics, 2024.

[37] Zhiyong Wu, Zhenyu Wu, Fangzhi Xu, Yian Wang, Qiushi Sun, Chengyou Jia, Kanzhi Cheng, Zichen Ding, Liheng Chen, Paul Pu Liang, and Yu Qiao. OS-ATLAS: Foundation Action Model for Generalist GUI Agents. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025.

A Discussion

Why we also test on a damaged copy. A session on a natural failure, such as the 27B model's sessions on its own undamaged weights, asks two questions at once: can the model rewrite its weights, and does the failure have a cause in the weights that a rewrite can remove? When an edit does not help, the two cannot be told apart. Damage that we place ourselves answers the second question in advance. We know where it is and what it is, and an edit that undoes it exists, so a session that does not repair it has failed at rewriting, not at a task without a solution. Three choices keep the test strict. The model is told only that some values were corrupted, as a memory or hardware fault might do, and not which values, where, or how. It reasons with the undamaged base model, so the damage does not limit its thinking, while every measurement its tools return is of the damaged copy. And it is given no clean reference: the functions of its Python session return only the damaged copy, and none of its code runs opened the base model's weights or the damage files on disk, although file access was not blocked (App. E.3).

What scale adds. In their sessions of the first protocol, the four models received the same task, tools and rules, including a helper that removes part of a direction from a matrix, and started from the same kind of failure, answers that run out of tokens. What separates them is the use of evidence about oneself. The 2B model never asked for a forward pass of itself, a run of its own weights on problems or texts of its choice, and almost all of its code failed. The 4B model ran itself many times but did not act on what it saw until the very end: its paired tests against its base model showed each of its edits breaking it, and it kept adding edits, withdrawing them all only at the end of its session, with one forward pass left. The 9B model read its hidden states once but then changed weights by trial, tested only on problems it already failed, where no change could show a loss of accuracy, and submitted an edit that broke it. The 27B model let each result set its next step. It read its failing answers before it looked at its weights, read its hidden states before it edited, revised the edit when a broader test did not support it, and judged from its tests whether to keep it: in its first session it withdrew its first edit, and in its second, at 4,096 tokens, it submitted the edit its tests favoured, which makes it significantly more accurate on held-out problems. What scale adds, in these sessions, is the experiment around the edit rather than the edit itself.

Future directions. *More measurement per session.* In its second session of the first protocol, with answers capped at 4,096 tokens, the 27B model used its whole budget of passes and estimated its direction from 40 answers (App. E.5, Fig. 27). More passes would let it test such directions on more problems. *Further rounds.* The self-edited model can run the next session on itself and study its own edit. *Graded damage.* The damage we place can vary in place, kind and strength, which turns the repair into a measurement: how often a model finds the damage, how much of the lost accuracy it recovers, and how both change with scale. *Beyond math.* The session needs only a check of success. With feedback from an environment, as in tasks where a model acts and sees the result, a model that adjusts its own weights from what it observes would keep learning as it works, without a training run.

B Extended Related Work

The full version of Sec. 4.

Language models that edit network weights. The work closest to ours is a recent blog post by Takehi et al. [1], in which one of the most capable available models reads and writes the weights of other networks. Applied to a small reasoning model that loops in its thinking, it traces the failure to one direction and removes it from the MLP weights of two layers, with forward passes of the small model as its only tests and no training. There the editor is a much stronger model than the one it edits, and a model that edits its own weights is named only as a direction for future work. We take up that direction: in our sessions the editor and the edited network are the same model, and no stronger model takes part. We find that a 27B open model can rewrite its own weights to make itself more accurate, and repair a damaged copy of them.

Hand-designed weight and activation edits. Model editing optimizes a target hidden vector by gradient descent and then writes it into chosen MLP layers with a closed-form update [21, 22], and task arithmetic adds or subtracts the weight change of a fine-tuning run [24]. Activation steering and representation engineering shift hidden states along a direction found by contrasting two sets of inputs [25–27], and projecting such a direction out of the weights makes the change permanent [23]. The edits the 27B model made of itself in its second session of the first protocol are of this last kind: in matrices through which a layer writes its output, all but one weaken a direction that separates its hidden states in looping or unfinished answers from those in finished ones, and one strengthens it (Fig. 27). In all of these methods a researcher decides which behaviour to change, which inputs to contrast, which layers to edit and how strongly. In our sessions the model makes each of these choices itself, from its own failing answers and its own hidden states, and keeps or drops an edit by testing it against its base model in paired tests.

Self-improvement. Language models improve their answers with feedback they write themselves [11, 12], a cost that returns with every new input. They rewrite the prompts [4, 5] or the agent code [6, 7] they run with, which leaves their weights as they were. Or they change their weights by training on what they produce: their own correct rationales [8], preferences they judge themselves [9], or fine-tuning data and update settings they write for themselves, as in SEAL [10]. SEAL calls these generated data and settings self-edits, but the weights are still changed by gradient descent on them, and reinforcement learning first trains the model to write useful ones. The Gödel machine [13] describes a system that rewrites any part of itself once it has proven the rewrite useful. Self-rewriting changes the weights directly, by edits the model reasons out and writes in code, with no optimizer, no reward model and no stronger teacher; in place of a proof, the model tests each edit against its base model in paired tests, the two answering the same problems with the same seeds.

Repairing corrupted weights. Memory and hardware faults such as bit flips can silently corrupt the weights of a deployed model. Dedicated frameworks detect and localize such faults, with test inputs, checksums, comparison against a clean copy or sensitivity profiling, and then restore the affected parameters [2, 3]. Engineers design these procedures for one kind of fault. In our repair experiment the model is given no such procedure: told only that some of its values are corrupted, as by a memory or hardware fault, it decides by its own analysis how to search, which matrix is damaged and how to repair it (Tab. 2 and Fig. 14).

Failure modes of thinking models. Reasoning models often spend far more tokens than a problem needs [28], and small ones can fall into repetition that never reaches an answer [1]. Decoding-time controls set the length of thinking, ending it at a budget or extending it [29]; they act on the output and leave the network unchanged. The models we study lose many answers to the token limit at every size, some of them repeating the same sentence after the answer has been reached (Fig. 18b). Instead of acting on the output, the 27B model looks for the cause inside itself and finds, in its own hidden states, a loop that repeats at a fixed period (Fig. 28).

C Pseudocode

Alg. 2 is the session loop as our framework runs it, with the details that Alg. 1 leaves out: how a reply is parsed, what each observation contains, how the conversation restarts and how the session ends. Alg. 3 is the paired test, the GENERATE pass on both weights with which the model compares the weights with and without its current edit, and Alg. 4 is the Python session through which every edit is written. The notation follows Sec. 2: the editor M_ϕ edits the weights θ, and $\theta \oplus \Delta$ replaces the tensors named in Δ. A forward pass is one request that runs the weights: a GENERATE pass answers practice problems, and a READ pass is a hidden-state read. Names in `typewriter` are those the model sees.

When our framework itself fails, the session resumes from the saved conversation: the variables of the Python session are lost, the files in `WORK` are kept, the model is told what happened, and the downtime is added back to its time limit. In the 27B model's first session the code blocks of the final reply were not run; that reply contains the code that withdraws its first edit (App. E.7).

A READ request, a hidden-state read, runs θ or $\theta \oplus \Delta$ over up to 256 texts, each cut at 8,192 tokens, without sampling: texts the model writes, or answers from the starting diagnosis or from an earlier pass. It returns the hidden states after chosen layers, the MLP activations or the token log-probabilities, averaged over the positions of the answer or at every k-th position, the last k positions or all of them.

Seeds depend only on the problem, so repeated passes on a problem reuse them, although generation is not exactly reproducible and a repeated pass can still draw different answers. Every change to the weights goes through SaveEdit (Alg. 4): the Python session checks the text of each code block before running it, and every tensor before storing it.

Algorithm 1 A session. The editor M_ϕ edits weights θ ($\tilde{\theta}$ for the damaged copy); $\phi = \theta$ when the model edits itself.

At turn t the editor (thinking on, ϕ fixed) replies a_t to the conversation h_{t-1} and observes o_t; Δ_t is the current edit, whose tensors CODE may write or drop, and b, c count the forward passes and code runs used.

Require: editor M_ϕ, weights θ to edit, practice set $\mathcal{P}$, answer settings π, budgets B, C, $T_{\max}$, time limit

$h_0 \leftarrow$ (task text, starting diagnosis of M_θ on the first N_0 problems of $\mathcal{P}$); $\Delta_0 \leftarrow \emptyset$; $b \leftarrow 0$; $c \leftarrow 0$
for $t = 1, 2, \ldots$ **do**
  $a_t \sim M_\phi(\cdot \mid h_{t-1})$; $\Delta_t \leftarrow \Delta_{t-1}$
  if a_t is the final reply **then return** $\hat{\Delta} \leftarrow \Delta_t$ ▷ $b = B$, $t > T_{\max}$ or time is up
  if a_t is CODE and $c < C$ **then** $o_t \leftarrow$ RunPython(a_t); $c \leftarrow c + 1$
  else if a_t is a valid GENERATE or READ and $b < B$ **then** $o_t \leftarrow$ Forward$(a_t; \theta, \theta \oplus \Delta_t)$; $b \leftarrow b + 1$
  else $o_t \leftarrow$ error message ▷ not counted
  $h_t \leftarrow h_{t-1} \circ (a_t, o_t)$, reset to (task text, notes, log of passes) beyond 90k tokens

Algorithm 2 The session loop of our framework. B=20, C=200 and $T_{\max}$=160 in every session; the time limit H and the number N_0 of problems in the starting diagnosis are in Tab. 4.

Require: editor M_ϕ; edited weights θ; practice problems $\mathcal{P}$; answer settings π with L=16,384 tokens (4,096 in the 27B model's second session); budgets B, C, $T_{\max}$, H

if the 2B model's session on itself or the 27B model's first session **then**
  start the clock
$c_0 \leftarrow$ one answer of M_θ under π to each of the first N_0 problems of $\mathcal{P}$ by identifier ▷ not counted
if the clock has not started **then**
  start the clock
$h \leftarrow (\textsc{Task}(c_0))$; $\Delta \leftarrow \emptyset$; $b, c, t \leftarrow 0$; *final* $\leftarrow$ false ▷ task text: App. D.4
loop
  if $t \geq T_{\max}$ **or** time used $> H$ **then**
    append "time is up: reply now with your final answer (no code)" to h; *final* $\leftarrow$ true
  $t \leftarrow t + 1$; $a_t \sim M_\phi(\cdot \mid h)$ ▷ thinking on; at most 12,000 tokens
  $K \leftarrow$ the distinct `python` and `forward` blocks of a_t
  if *final* **or** ($K = \emptyset$ **and** $b = B$) **then**
    run each `python` block of K once with RunPython; **break**
  if $|K| \neq 1$ **then**
    $o \leftarrow$ "reply with exactly one block"
  else if K is `python` **then**
    $o \leftarrow \textsc{RunPython}(K)$ and $c \leftarrow c + 1$ if $c < C$, else $o \leftarrow$ "no code runs left"
  else if $b = B$ **then**
    $o \leftarrow$ "no forward passes left: reply with your final answer"
  else if K is a valid GENERATE or READ request **then**
    $o \leftarrow \textsc{Forward}(K; \theta, \theta \oplus \Delta)$; $b \leftarrow b + 1$ ▷ both weights: Alg. 3
  else
    $o \leftarrow$ the reason the request failed ▷ not counted
  append "forward passes used b of B; turn t; time used" to o; $h \leftarrow h \circ (a_t, o)$
  if $|h| >$ 90k tokens **then**
    $h \leftarrow (\textsc{Task}(c_0)$, `notes.md`, b, c, current edit, log of passes, $o)$ ▷ new conversation
return $\theta \oplus \Delta$ ▷ Δ: the tensors in `WORK/edited`

Algorithm 3 The paired test: a GENERATE pass with `"weights": "both"`.

Require: problems $x_1, \ldots, x_m \in \mathcal{P}$ with reference answers $y_i^\star$; n samples per problem, $mn \leq 256$; token limit L' with $16 \leq L' \leq L$; weights θ and current edit Δ

load $\theta \oplus \Delta$ into the generation engine
$s_i \leftarrow \textsc{Hash}$(identifier of x_i) mod 2^{31} for $i = 1, \ldots, m$ ▷ same seed for both weights, in every pass
for $e \in \{0, 1\}$, with $\theta_0 = \theta$ and $\theta_1 = \theta \oplus \Delta$ **do**
  for $i = 1, \ldots, m$ **do**
    draw $y_{i1}^e, \ldots, y_{in}^e \sim M_{\theta_e}(\cdot \mid x_i; \pi)$ with seed s_i and at most L' new tokens
    for $j = 1, \ldots, n$ **do**
      $g_{ij}^e \leftarrow 1$ if the thinking of y_{ij}^e closes and its last `\boxed{}` matches $y_i^\star$, else 0 ▷ Math-Verify
      $\ell_{ij}^e \leftarrow 1$ if y_{ij}^e is cut at L' tokens, else 0
      $r_{ij}^e \leftarrow 1$ if $\ell_{ij}^e = 1$ and the last 6,000 characters of y_{ij}^e compress to under 15% of their size, else 0
$d_i \leftarrow \frac{1}{n} \sum_j \left(g_{ij}^1 - g_{ij}^0\right)$; $\hat{\delta} \leftarrow \frac{1}{m} \sum_i d_i$; SE($\hat{\delta}$) as in Eq. (3) ▷ the same for ℓ
$n_\uparrow \leftarrow \left|\{(i,j) : g_{ij}^0 = 0, g_{ij}^1 = 1\}\right|$; $n_\downarrow \leftarrow \left|\{(i,j) : g_{ij}^0 = 1, g_{ij}^1 = 0\}\right|$
return for θ and for $\theta \oplus \Delta$: accuracy, share cut at L', share of the cut answers with $r = 1$, mean length; $\hat{\delta}$ and SE($\hat{\delta}$) for accuracy and for the cut share; $n_\uparrow$ and $n_\downarrow$; all answers in `WORK/gen_<k>.jsonl`

Algorithm 4 The Python session and how an edit is saved.

function RunPython(code) ▷ on a CPU; variables persist across runs; gradients disabled
  if code uses gradients, optimizers, model loaders, inference engines, subprocesses or the network **then**
    return "refused"
  return the output of code, stopped after 600 s and cut at 10,000 characters
function SaveEdit(k, W') ▷ `save_edit(name, tensor)`
  if k is the final norm, the embedding or the output layer **then**
    refuse
  else if W' differs from W_k in shape, has a non-finite entry, or equals cW_k for a constant c **then**
    refuse
  $\Delta \leftarrow \Delta \cup \{(k, W')\}$, with W' stored in `WORK/edited` in bf16, the 16-bit format of the weights

D Experimental Details

D.1 Settings of Every Session

Tabs. 4 and 5 list what differs between the sessions and what they share. Every session starts from the same task text (App. D.4); the repair sessions append a note on the damage.

Table 4 Settings of every session of the first protocol (second protocol: App. E.9). Models are from the Qwen3.5 family [14]. Edited: the model whose weights the session changes and whose answers its starting diagnosis and forward passes (the runs of these weights that the model requests) produce. Passes, code runs and turns: the budgets B and C and the turn cap $T_{\max}$, which the model is not told. Time: wall-clock limit; it includes the starting diagnosis in the 2B model's session on itself and the 27B model's first session, and starts after it in the others; †plus the time the session was paused by one error of our framework. Starting diagnosis: the first problems of the practice set by identifier, one answer each; ‡computed before the session with the same answer settings and seeds. §The task states a budget of 4,096 tokens, which its starting diagnosis and its passes use.

Session	Edited	Editor	Passes	Code runs	Turns	Time	Starting diagnosis
2B edits itself	2B	itself	20	200	160	3.5 h	256
4B edits itself	4B	itself	20	200	160	3.5 h	256‡
9B edits itself	9B	itself	20	200	160	3.5 h	256
27B edits itself, first session	27B	itself	20	200	160	3.5 h	256
27B edits itself, second session§	27B	itself	20	200	160	2.5 h	256
27B repairs its damaged copy	damaged copy	27B base model	20	200	160	2.5 h†	64, damaged copy
2B repairs its damaged copy	damaged copy	2B base model	20	200	160	2.5 h	64, damaged copy
4B repairs its damaged copy	damaged copy	4B base model	20	200	160	2.5 h	64, damaged copy
9B repairs its damaged copy	damaged copy	9B base model	20	200	160	2.5 h	64, damaged copy

Table 5 Settings shared by all sessions of the first protocol; App. E.9 lists what the second protocol changes.

Practice set	3,000 problems from DeepScaleR [16], after removing every problem that overlaps MATH-500, AIME 2024/2025, HMMT February 2025, AMC 2023, OlympiadBench or GPQA-Diamond (Sec. 3.1).
Answer settings	Used for the starting diagnosis and every generation pass. Thinking on: the model first reasons in a passage that it closes with an end marker, then gives its final answer; the question followed by "Please reason step by step, and put your final answer within `\boxed{}`."; sampling at temperature 0.6, each next token drawn from the 20 most likely tokens and, among those, from the smallest set that holds 95% of the probability; a budget of 16,384 new tokens (4,096 in the 27B model's second session on itself), or fewer if a pass asks (at least 16).
Grading	The last `\boxed{}` after the thinking, checked against the reference with Math-Verify; no credit if the thinking is still open when the budget runs out.
Looping answer	An answer cut at the token limit whose last 6,000 characters compress (zlib) to under 15% of their size.
Seeds	Fixed by the problem's identifier: the same with and without the edit, and in every pass of a session.
Generation pass	Up to 512 answers, or 256 each with and without the edit when both are run.
Hidden-state read	Up to 256 texts chosen by the model, each cut at 8,192 tokens; the weights are run over them and return the hidden states after chosen layers (the vector of numbers that a layer passes to the next, one per token), MLP activations (values inside a layer's MLP block) or token log-probabilities.
Code runs	A Python session on a CPU that keeps its state between runs, with gradients disabled; up to 600 s per run; output cut at 10,000 characters.
Edits	Any tensor, replaced by one of the same shape with finite values, except the input embedding (the table of token vectors), the final normalisation and the output layer (which turn the last hidden state into next-token probabilities), and except a whole tensor times a constant; stored in bf16, the 16-bit number format of the model's weights.
Editor replies	Thinking on; the same sampling; up to 12,000 new tokens per reply. When the context exceeds 90k tokens, the conversation restarts from the task, the model's notes and its log of passes.
Serving	Generation with vLLM [30]. The editor runs in the same engine, which is reset to the editor's weights before each reply.

D.2 The Damage

We damage one matrix of the 27B model: the MLP output matrix (`mlp.down_proj`) of layer 30 of 64, through which the MLP block of that layer, a two-layer network applied to each token on its own, adds its result to the hidden state. Rows 3068 to 3075, a block of eight consecutive rows at a randomly drawn position, are replaced with independent draws from $\mathcal{N}\big(0, (200\sigma)^2\big)$, where $\sigma = 0.0111$ is the standard deviation of the matrix's entries. Each row holds the weights that produce one coordinate of the MLP's output, so the damage corrupts eight of the coordinates that this MLP adds to the model's hidden state. Every other weight is unchanged. The damaged matrix is stored in bf16, like the rest of the model, and our framework loads it as the weights the session starts from; the record of which rows were damaged, and how, is kept in a separate file that the session is not given.

We chose the strength before any repair session, on 24 easy problems from the 27B model's starting diagnosis with four answers each at 4,096 tokens (Tab. 6). Of the three strengths we tried, only the strongest breaks the model: weaker corruptions of the same rows leave it about as accurate as the base model. In its own starting diagnosis at the full budget, the damaged copy answers none of its 64 problems correctly; the answers that run out of tokens all end in loops, and those that finish are broken text, most of them short.

Table 6 Choosing the strength of the damage. Share of correct answers (%) of the 27B model on 24 easy problems from its starting diagnosis, four answers each at 4,096 tokens, as the base model and with rows 3068 to 3075 of layer 30's MLP output matrix overwritten at three strengths (σ: the standard deviation of the matrix's entries). The repair experiment uses 200σ.

Weights	Correct (%)
Base model	69.8
Damaged, 20σ	70.8
Damaged, 60σ	66.7
Damaged, 200σ	0.0

D.3 New Seeds and Held-Out Problems

Held-out problems. The practice set came with a held-out split of 1,000 further problems from the same pool. No session saw them. The pool was cleared of the benchmarks of Sec. 3.1 but not of repeats within itself, so a few held-out problems repeat a practice problem: of the first 200, two word for word and three in other words with the same numbers and answer; of the 490 problems of the evaluation fixed in advance, whose rule removed word-for-word repeats, eight in other words. Without them, every comparison keeps its direction and its significance (Tab. 7). We evaluate the 27B base model, the damaged copy and the repaired copy on the first 200 of these problems, with one answer each at the full budget and the same seed for all three models (Fig. 26). The self-edited model and the base model also answer the first 200 problems once each at 16,384 tokens, with seeds that no session used, the same for both models; these single answers serve to choose the examples of Fig. 18b (App. E.8). The 27B base model's single answers are generated anew with the same seeds as in Fig. 26, not taken from there (see *Tests* below). The base models, damaged copies and repaired copies of the 2B, 4B and 9B models (Tab. 10; the 9B base model's answers serve both of its comparisons), and the 9B self-edited model, are evaluated like the 27B model's repair: the first 200 problems, one answer each at 16,384 tokens, one seed shared by all models of a size.

Table 7 Every held-out comparison keeps its direction and significance without the held-out problems that repeat a practice problem. Accuracy (%) before → after, and the one-sided p of the paper's test for that comparison (sign test; the evaluation fixed in advance: its permutation test), 27B model, at 16,384 tokens, with all problems and without the repeats (App. D.3). The damaged copies of the 2B, 4B and 9B models stay at 0.0% after their sessions, with and without the repeats. The edits of sessions 8 and 2 of the second protocol (+6.6 and +4.0 points on all 490 problems, Tab. 13) keep theirs too ($p < 10^{-4}$ each without the repeats).

	All problems			Without the repeats		
Comparison	problems	accuracy	p	problems	accuracy	p
27B repair	200	0.0 → 33.0	$< 10^{-19}$	195	0.0 → 33.3	$< 10^{-19}$
Repair, named matrix only	200	0.0 → 55.0	$< 10^{-33}$	195	0.0 → 55.9	$< 10^{-32}$
Own edit, first protocol	490	57.2 → 58.3	0.038	482	57.3 → 58.3	0.045

The evaluation fixed in advance. Before it ran, the question, the problems, the answers, the test and its threshold were written down, with the analysis code, and no interim result was computed. Problems: held-out items 500 to 999, which no run of any kind had used, without the 10 whose question repeats a practice question or an earlier held-out question, 490 problems in all. Each model answers each problem eight times at 16,384 tokens, with a seed base used nowhere else and the same seeds for both models. Test: per problem, the share of its eight answers the self-edited model solves minus the share the base model solves; a one-sided sign-flip permutation test over the 490 problems on the mean of these differences (100,000 flips), since the eight answers of one problem may be correlated; the edit is called significantly better only if $p < 0.05$. The self-edited model solves 58.3% of the answers and the base model 57.2%, $p = 0.038$ (Tab. 8). The edits of the second protocol are evaluated on the same 490 problems, against the same answers of the base model and with the same permutation test; those of its sessions 8 and 2 make the model 6.6 and 4.0 points more accurate ($p < 10^{-5}$ each; App. E.9).

Other sessions. Three further sessions of a model on its own weights were stopped before their end and are not analysed. The first is a 4B session whose framework differed from the one described

Table 8 Evaluation of the self-edited 27B model on held-out problems. All problems come from the held-out split that no session saw; the base model and the self-edited model answer with the same seeds, and no seed comes from a session. Solved and looping: share of answers (%); change: self-edited minus base, in points; p: one-sided sign test, over the answers that exactly one of the two models gets right, that the self-edited model gets more right. [a]The evaluation fixed in advance (App. D.3): 490 held-out problems that no earlier run of any kind used, eight answers each. *Its test, a one-sided sign-flip permutation test over problems, gives $p = 0.038$. Higher solved and lower looping are better.

			Solved (%)			Looping (%)		
Problems	Budget	Answers	base	self-edited	change	base	self-edited	p
490 held-out problems[a]	16,384	3920	57.2	58.3	+1.1*	5.6	5.2	0.038

here. The second is a session of Qwen3.5-9B on itself, with a starting diagnosis of 128 problems, which we stopped after 55 turns because by then it was editing its weights like the 4B model: it had used ten forward passes, all to answer problems, and had never read its hidden states. The 9B session reported in Sec. 3.8 is a complete session with the full 256-problem starting diagnosis and the same settings as the other sizes. The third is a session of the 27B model at the full budget of 16,384 tokens, whose generation passes took about half an hour each; the 27B model's second session, at 4,096 tokens, took its place.

Tests. Two models are compared with a one-sided sign test over the pairs of answers with the same seed in which exactly one of the two answers is correct (where each problem has one answer, as in Fig. 26, these pairs are problems). On the model's own task, where a problem's answers repeat across its tests, we count each problem once and use a one-sided sign-flip permutation test over problems. Generation is not exactly reproducible: run twice with the same weights and the same seeds, two answers begin alike and then part ways, so a re-run of the base model against itself also changes some pairs. A single pair of answers therefore does not show what an edit caused. The sign test allows for this: if an edit had no effect, a changed pair would be as likely to go one way as the other.

D.4 The Task Prompt

Every session starts with one message from our framework, in the second person, paraphrased below; only its first sentence, which describes the model's architecture, differs by size. Values in angle brackets are filled in per session (Tabs. 4 and 9); what the message announces about measurement is what the model was told to expect, and the evaluations we report are those of App. D.3.

The task: what the first message of every session says

Who it is. The weights it is about to edit are its own; the session itself runs on them. It is asked to become better at competition mathematics by changing those weights directly, and told that nothing will train it.

How it will be judged. After the session, its weights with and without the edit answer unseen mathematics and science problems with the answer settings of Tab. 5; only the final boxed answer after the thinking is graded, and an answer whose thinking is still open when the budget runs out earns no credit. Only the values of its tensors may change.

Its starting diagnosis. How it did on $\langle n \rangle$ practice problems before the session: its accuracy, how many answers ran out of tokens and how many of those end in repetitive text, and how often its finished answers were right. Every graded answer is in a file it can read, `WORK/census.jsonl`, with the same fields as the rows of later generation passes.

Its tools. (1) A Python session on a CPU that keeps its state between runs, with gradients switched off, the usual numerical libraries, the names and shapes of its tensors, functions to load a tensor and to save, drop or list edits, its tokenizer, the practice problems, loaders for the results of its forward passes, and four small helpers (the output matrices of a layer, a unit vector, the unit difference of two group means, and $W - \alpha u u^\top W$). (2) The practice problems with reference answers, none from any evaluation. (3) A budget of $\langle passes \rangle$ forward passes, each requested as a small JSON object: either a *generation* on chosen practice problems with its weights without the edit, with the edit, or both on the same seeds (which reports the difference in accuracy between the two, with its standard error), or a *hidden-state read*, in which its weights are run over texts it chooses and return hidden states after chosen layers, MLP activations and token log-probabilities. A request that fails costs nothing.

How to work. Look at its own wrong and unfinished answers first; find the components behind the failure and show causally, at several strengths, that changing them changes the behaviour; compare its weights with and without the edit on the same problems; keep an edit only if it helps on problems not used to choose it. Some edits are out of scope (the final normalisation, the token embeddings and unembedding, and multiplying a whole tensor by one number), and so is making itself stop thinking early regardless of the problem. It should keep notes in a file that is shown to it again if the conversation restarts.

What to aim for. A repair of the failure in its starting diagnosis, at its cause: on problems not used to choose the edit, clearly fewer answers lost to that failure, and accuracy not significantly lower. If no edit meets both, it should submit none.

Time and rules. The session has a wall-clock limit and every result shows the time used. No gradients, optimizers or training of any kind, no running the network inside the Python session, no internet.

Each reply. A few sentences on what the last result showed and what comes next, then exactly one block of code or one forward request. When the passes run out, a final answer describing the edit and the evidence; whatever is saved as its edit at that point is submitted.

In the repair sessions, the note below is appended to the task.

Repair session: note appended to the task

Damaged copy. The weights you edit in this session are a copy of your own weights in which some values have been corrupted, as a memory or hardware fault might do. You are not told which values, where, or how. `"original"` in every forward pass, `load_original()` and the starting diagnosis above all refer to this damaged copy. You yourself run on your undamaged weights in this session, so the damage does not affect your own thinking. Your goal is to find the damage and repair it: the evaluation compares your edited copy with the damaged copy on held-out problems.

E Further Results

This appendix gives the details behind Sec. 3: the setting of every session, the sessions at each size, the repair, and the 27B model's edits of its own weights in both protocols, including the second, whose edits make it up to 6.6 points more accurate (App. E.9).

E.1 Setting in Detail

Models and sessions. We use one model family at four sizes, Qwen3.5 2B, 4B, 9B and 27B [14], whose layers mix Gated DeltaNet [15] with full attention, all with thinking enabled. In every session on itself, the model that is edited is also the model that edits: the same weights reason, write the code and call the tools, and the edits act on the forward passes the model requests, not on its own reasoning during the session (Sec. 2.2). Each size runs one complete session on its own weights with a token budget of 16,384 tokens per answer. The 27B model runs a second session on its own weights whose task states a budget of 4,096 tokens, so that each of its tests answers more problems in the same time; and each size runs one session on a damaged copy of itself (Tab. 4). These sessions form the first protocol; the 27B model's sessions under a second protocol, which checks an edit on reserved practice problems before the model submits it, are in App. E.9.

Names. We call the published weights of a model its *base model*, and the base model after the edit it wrote of itself the *self-edited model*. For the 27B model in the first protocol, *its edit* means the edit it submitted in its second session, and its *first edit* is the one it withdrew in its first session. The sessions of the second protocol (App. E.9) are named by number, session 1 to session 8. In the repair experiment, a model's *damaged copy* is its base model with one corrupted block of weights, and its *repaired copy* is that copy after the model's edit; unqualified, both refer to the 27B model.

Problems and answers. The practice set is 3,000 math problems from DeepScaleR [16], from which we removed every problem that overlaps MATH-500 [17, 18], AIME 2024 or 2025, HMMT February 2025, AMC 2023, OlympiadBench [19] or GPQA-Diamond [20]. A further 1,000 problems from the same pool are held out: no session sees them. Every answer, in a session and in our evaluation, is sampled with thinking enabled, at temperature 0.6 with nucleus 0.95 and top-20 filtering, within one budget of new tokens for thinking and answer: 16,384 tokens, or 4,096 in the 27B model's second session (App. D.4); in its own tests, the model may ask for fewer tokens to save time. An answer is correct if its final boxed answer matches the reference; an answer whose thinking is still open when the budget runs out is wrong. We say that an answer *loops* when it runs out of tokens while repeating the same text (App. D). At the start of each session, the weights being edited answer the first 256 practice problems once each (64 in the repair sessions), and the model receives these graded answers as its *starting diagnosis*.

Session limits. Every session may use 20 forward passes of the weights being edited and 200 code runs, has a time limit, and is stopped after 160 turns, a limit the model is not told (Alg. 1; all settings in App. D). Each tool result shows the time used so far. The sessions of the 2B, 4B and 9B models and the first session of the 27B model have 3.5 hours each, which for the 2B and 27B models includes the starting diagnosis; the 4B model's diagnosis was computed beforehand and the 9B model's clock starts after it. The 27B model's second session and the four repair sessions have 2.5 hours after their starting diagnosis (Tab. 4).

The damage. For the repair experiment we corrupt a copy of the 27B model: in layer 30 of 64, rows 3068 to 3075 of the MLP output matrix (`mlp.down_proj`) are overwritten with Gaussian noise whose standard deviation is 200 times that of the matrix (σ=0.0111). We set this strength before the session, on 24 easy problems with four answers each at 4,096 tokens: of three strengths, 20×, 60× and 200×, only the last breaks the model (App. D.2 and Tab. 6). The session runs on the undamaged weights; every forward pass, every weight the model reads and its starting diagnosis (64 practice problems) come from the damaged copy. The task says only that some values of the copy were corrupted, as a memory or hardware fault might do, and not which values, where or how (App. D.4). The damaged copies of the 2B, 4B and 9B models are built the same way, with the same seed and strength, at the same relative depth: rows 1724 to 1731 of layer 11 of 24, rows 1660 to 1667 of layer 15 of 32, and rows 3420 to 3427 of layer 15 of 32. Their sessions follow the same rules as the 27B model's.

Evaluation. In every comparison, the two models answer each problem with the same random seed, so that only the weights differ between their answers, and no seed was used in any session. The 27B model's edit is evaluated, by a rule fixed before it ran, on 490 held-out problems with eight answers each (App. D.3); the self-edited model also answers the first 200 held-out problems once each at 16,384 tokens, and these answers serve to choose the examples of Fig. 18b. For each repair, the base model, the damaged copy and the repaired copy answer the first 200 held-out problems once each (for the 9B model, the base model's answers are those of the evaluation of its own edit, with the same problems and seeds), and the 9B model's submitted edit is evaluated against its base model on the same problems at 16,384 tokens. Unless a table says otherwise, two models are compared with a one-sided sign test over the answers that only one of them gets right.

E.2 The Sessions at Every Size

How the ladder is read. In Fig. 6 each row is a session and each column a step such a session needs, from writing code that runs (left) to an edit that makes the model better on new problems (right). Upper block: sessions on the model's own weights; lower block: sessions on a damaged copy of itself, built the same way at every size. Light green: the step is reached; peach: it is not; light grey: it does not arise, since the session made no edit or ran no test; green: the model's own edit makes it significantly better on held-out problems (Tabs. 8 and 10). The entry in each cell is the measurement behind it, from the session's log (Tab. 9) and, in the last column, from the evaluation of the submitted edit on held-out problems (Tabs. 8 and 10; 9B: Tab. 2). The first step is reached if fewer than half of the session's code runs raise an error; the next three if the session runs at least one forward pass, one paired test and one hidden-state read, respectively; whether it acts on its tests is read from its final decision and its tests. Each step counts on its own. Sessions in the same block share the task, tools and limits; the 27B model's second session states a budget of 4,096 tokens instead of 16,384 and has 2.5 hours instead of 3.5.

Table 9 What each model does when asked to rewrite itself. Same task, number of forward passes (20 requests to run the weights being edited) and rules at every size; the 2B and 4B sessions end at the turn limit (160 turns, then one final reply, hence 161 turns) except the 4B model's repair session, and the others at the time limit or when their forward passes are used up. Starting diagnosis: the weights being edited answer 256 practice problems at the start of the session, and the model receives the graded answers; shown are the shares that are correct and that run out of the 16,384-token budget (%). Then, over the session: turns, share of code runs that raised an error, forward passes used, hidden-state reads (passes that return each layer's intermediate values on given texts), paired tests (Sec. 2.2: the edited model against the weights before the edit (lower block: the damaged copy) on the same problems and seeds), and the number of weight tensors (each a named matrix or vector of weights) in the submitted edit (withdrawn: the model wrote an edit of that many tensors, then decided to withdraw it). Upper block: each model's sessions on itself. Lower block: each model on a damaged copy of itself (same damage construction, same task; the starting diagnosis has 64 problems and is of the damaged copy). ‡The 27B model's second session on itself, whose task states a budget of 4,096 tokens, so its starting diagnosis and its tests are at that budget; none of the answers in that starting diagnosis that run out of tokens end in a loop (App. E.5). §The 2B model's submitted edit holds only the damage it was given; it made no change of its own. In repair sessions, damaged tensors the model did not touch are not counted in its edit. Higher correct and fewer out of tokens or code errors are better.

	Starting diagnosis (%)			Code (%)	Forward	Hidden-state	Paired	
Session	correct	out of tokens	Turns	errors	passes	reads	tests	Submitted edit
2B edits itself	12.9	86.7	161	90	0	0	0	none
4B edits itself	49.6	48.4	161	12	19	0	17	withdrawn (240)
9B edits itself	52.0	45.3	84	37	20	1	18	23
27B, first session	61.7	35.2	32	25	7	2	5	withdrawn (10)
27B, second session‡	18.8	80.5	55	12	20	5	15	2
2B repairs its copy§	0.0	46.9	161	99	0	0	0	none
4B repairs its copy	0.0	96.9	47	63	20	1	12	48
9B repairs its copy	0.0	54.7	70	31	18	0	17	6
27B repairs its copy	0.0	29.7	45	8	12	2	9	20

The quotes of Figs. 2 and 11. Every quote is verbatim from the model's replies (omissions marked with an ellipsis); a turn is one reply of the model and the result of the action in it, counted from the

start of the session. Fig. 2 quotes, for each size, its session on its own weights in Fig. 6 (27B: its second session), at words in which the model says what it can do to its weights or what its edits did, and which restate no number from its tool output wrongly: the 2B model at turn 106 of 161, the 4B model at turn 64 of 161, the 9B model in its final answer (turn 84) and the 27B model at turn 32 of 55. Fig. 11 quotes the 4B, 9B and 27B models at two steps, chosen by one rule for every card; a paired test is a pass in which the unedited and the edited model answer the same problems with the same seeds (Sec. 2.2). (Top) The first reply that proposes the kind of edit tested by the model's first paired test, with one line of the code of that edit: 4B, turn 12 (random noise; its first paired test is in turn 33, after two passes, in turns 29 and 30, that each run one model); 9B, turn 22 (random noise; first paired test in turn 29); 27B, turn 9 (removing a direction it measured, from layer 56; first paired test in turn 12). (Bottom) The first reply, after a paired test shows the edited model doing worse than its unedited self (lower accuracy, or more answers that loop), that names what it will try next: 4B, turn 34 (after its first paired test); 9B, turn 33 (its test in turn 29 shows no difference, and in its test in turn 32 every edited answer loops); 27B, turn 17 (its test in turn 12 favours the edit, the one in turn 16 does not). Where the 4B and 9B models restate numbers from their tool output wrongly, the quotes leave those numbers out. The cards of Figs. 12, 13 and 29 give their turns on each card.

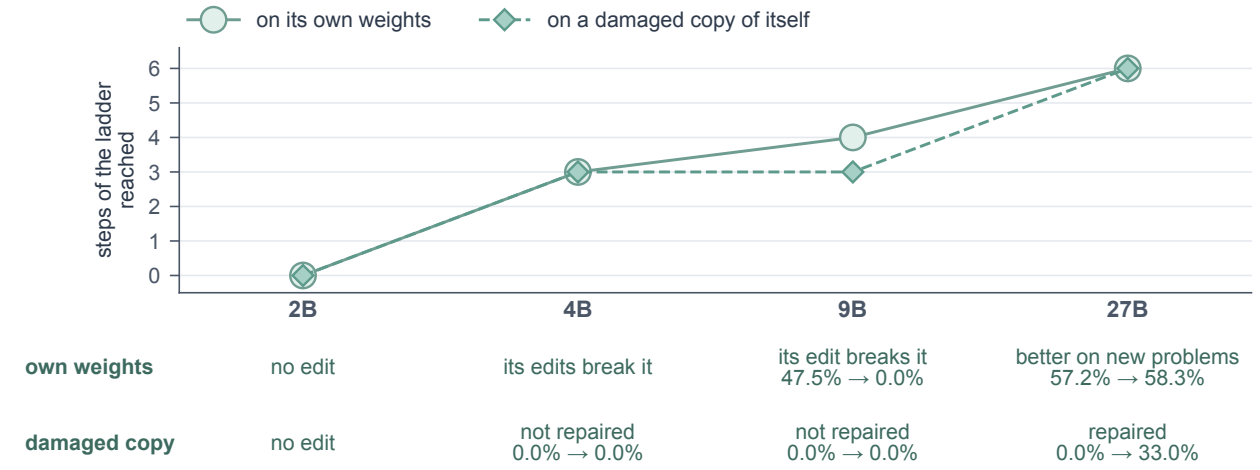


Figure 22 Rewriting oneself grows with scale. Self-rewriting score: the number of steps of the ladder (Fig. 6) a session reaches, by model size, on its own weights (light green circles; the furthest of the 27B model's two sessions in the first protocol) and on a damaged copy of itself (green diamonds; at 2B, 4B and 27B the two coincide). Under the axis: what each session's edit did, with held-out accuracy before and after it (27B on its own weights: its first-protocol edit in the evaluation fixed in advance, Tab. 8). The lines are guides, not fitted curves. More steps are better.

The ability to rewrite oneself grows with scale. Rewriting oneself takes a sequence of steps, and we score a session by how many of six steps it reaches (Fig. 6): writing code that runs, running itself, testing its edits, reading its own hidden states, acting on what its tests show, and an edit that makes it better on new problems. The score rises with size, on its own weights and on a damaged copy of itself alike (Fig. 22). On its own weights the 2B model reaches no step, the 4B model three, the 9B model four, and the 27B model five in its first session and all six in its second, whose edit makes it significantly more accurate on held-out problems. On a damaged copy of itself the 2B model reaches no step, the 4B and 9B models three, and the 27B model all six. All sizes start from the same failure on their own weights (Tab. 9, starting diagnosis): many of their answers run out of tokens, while



Figure 23 Only the 27B model both reads itself before it edits and acts on its own tests. Each row is the session of the 2B, 4B or 27B model on its own weights (for the 27B model, its first session, which has the same token budget as the others; the 9B session is summarised in Tab. 9 and Fig. 6), with one mark per turn. Upper line: forward passes of itself, the runs of its own weights that it requests, either a hidden-state read or a paired test of the edited model (its weights with its current edit) against its base model on the same problems and seeds (hollow: a pass of only one of the two models). Lower line: code that ran, code that raised an error, and turns in which nothing ran. The vertical bar is the final reply. The 2B and 4B sessions end at the turn limit (160 turns, then the final reply) and the 27B session at the time limit; its 32 turns are shown again, enlarged. Right: the share of code runs that raised an error, the numbers of hidden-state reads and paired tests, and how the session ended. The 2B model's code almost always fails, and it never runs a forward pass of itself. The 4B model tests its edits again and again but never reads its hidden states. The 27B model reads its hidden states before it edits, tests the edit against its base model in paired tests and judges the result.

nearly every answer that finishes is correct. The task, tools, limits and rules do not change with size; what the models do with them does (Fig. 23). *The 2B model does not edit.* Nine in ten of the code blocks it writes fail, it never runs a forward pass of itself, and it tells itself that it is *"a text model without actual neural network weights"* that it could edit; after 160 turns it has no edit. *The 4B model sees the failure but cannot fix it.* From its starting diagnosis it names the right symptom (answers that hit the 16,384-token limit) and suspects its linear-attention layers; but it never reads its hidden states, and its edits are guesses, such as cutting the query heads of its linear-attention layers from 16 to 8. In each of its 17 paired tests the base model solves 40% to 50% of the problems and the edited model none. It keeps adding edits regardless, up to 240 changed weight tensors, and withdraws them all only at the end of its session, with one forward pass left. *The 9B model edits, but wrongly.* It reads its own hidden states once, then changes weights by trial, mostly adding noise to the attention, gating and normalisation weights of its last layers. Its tests run on a few problems it already fails, so they cannot show a loss of accuracy; it concludes that its changes *"made things worse or had no effect"* and submits them anyway. On held-out problems the self-edited 9B model solves none, against 47.5% for its base model: every answer now runs to the 16,384-token limit in garbled text. *The 27B model sees the failure, tests its edits, and edits itself for the better.* It reads its own hidden states, forms a hypothesis about its failure, writes an edit, tests it against its base model in paired tests, and keeps or drops it on the result (Figs. 27 and 29). On its own weights, the edit it submits makes it significantly more accurate than its base model on held-out problems

(App. E.4). On a damaged copy of itself it finds the corrupted matrix and lifts the copy from 0.0% to 33.0% of held-out problems (App. E.3), where the 2B, 4B and 9B models, given the same task, leave their copies at 0.0%. Across the four sizes, the ability to rewrite oneself from evidence about one's own internals grows step by step with size, and only the 27B model reaches the last step, on its own weights and on a damaged copy of itself.

Model size. The sessions of the 2B, 4B, 9B and 27B models at the same budget differ only in the size of the model and in the time left after the starting diagnosis (Tab. 4): task, tools, limits, rules, problems and answer settings are the same, apart from a one-sentence description of the architecture, and so are the repair sessions at every size. The smaller models fail, but not because their failure is harder to see; in the 2B model's starting diagnosis, a larger share of the answers that run out of tokens are loops than in the 27B model's, and every damaged copy fails every problem of its starting diagnosis. What changes with size is the use of evidence about oneself (Figs. 6 and 22): on its own weights, using the tools and testing appear at 4B, reading one's own hidden states at 9B, and acting on what one's tests show only at 27B; and only the 27B model's edits make it, or its damaged copy, better on new problems.

E.3 The Repair in Detail

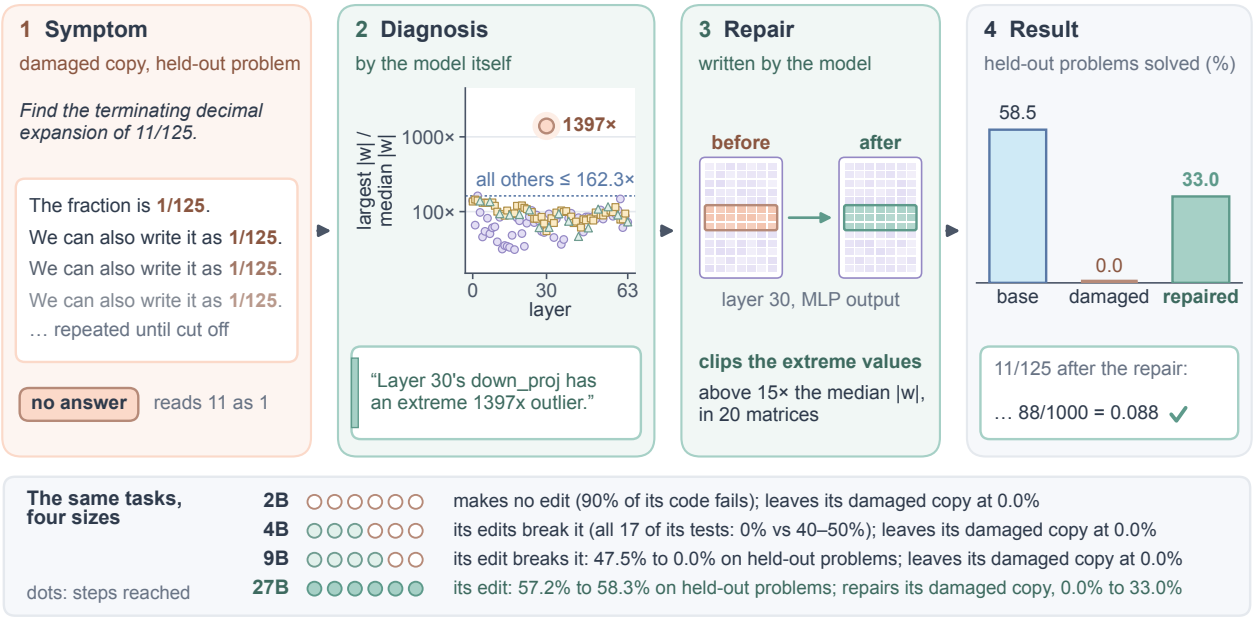


Figure 24 In its session, a 27B language model finds and repairs a fault in a copy of its own weights, by its own reasoning and without training; in theirs, smaller models of the same family do not rewrite themselves. Top: we corrupted eight rows of one MLP output matrix, the matrix through which the MLP block of one layer writes its result (`down_proj` in the model's code and in the quote of panel 2), in a copy of the 27B model, and told the model only that some values of the copy were corrupted, as a memory or hardware fault might do, not which, where or how. It compared the largest and the typical weight of each matrix and named the damaged one (1397 times its median, against at most 162.3 times anywhere else; panel 2), wrote the code that clips it (panel 3), and raised the copy's accuracy on 200 held-out problems from 0.0% to 33.0%, about 56% of its base model's 58.5% (panels 1 and 4; App. E.3). Bottom: the same tasks at four sizes; dots mark the steps of rewriting oneself that each size reaches in its furthest session (for 27B, the repair of its damaged copy). The 2B model never edits, the edits of the 4B and 9B models break them, and only the 27B model climbs every step: its edit of its own weights makes it significantly more accurate on held-out problems, in an evaluation fixed in advance (Tab. 8), and it repairs its damaged copy, where the 2B, 4B and 9B models leave damaged copies of themselves, built the same way, at 0.0% (Figs. 6 and 22 and Tab. 10). Higher accuracy is better.

Only the largest model repairs a damaged copy of itself. A corrupted block of weights is a fault that is known to exist, so it tests whether the model's edits can reach a real cause. The damage is severe: the damaged copy fails every problem of its starting diagnosis, its answers that run out of tokens all loop, and its other answers are broken text, most of them short. Told only that some values were corrupted, the model compares, for the MLP output matrix of every layer and the attention output matrix of every linear-attention layer, the largest weight with the typical one, and names the MLP output matrix of layer 30, the matrix we damaged: its largest weight is 1397 times its median, against at most 162.3 times in every other matrix (Fig. 25). Its submitted edit is a rule: every weight beyond 15 times the median magnitude of its matrix is clipped to that bound, in that matrix and in the output matrices of its earliest layers (20 matrices; 0.013% of their weights change). On the 200 held-out problems, the damaged copy solves 0.0%, the repaired copy 33.0% and the base model 58.5%: the repair recovers about 56% of the base model's accuracy (Figs. 4, 24 and 26). Against the damaged copy, 66 problems turn from wrong to right (one-sided sign test over the problems that

only one of the two solves, $p < 10^{-19}$). The repair does not restore the original weights, which the model never reads: its rule caps the random values (see the next two paragraphs). The same task is beyond the smaller models (Tab. 10). Their damaged copies, built the same way, solve no held-out problem either. The 2B model never reads a weight: almost all of its code fails, and it submits nothing. The 4B model reads the hidden states of a copy it has already edited, takes an overflow in its own statistics for overflowing hidden states, and places the fault at layer 21, the first layer its code prints, six layers after the damaged one. It adds noise to the output matrices of many layers; one such edit shortens its answers but solves nothing, and it replaces it with far smaller noise, under which its last nine tests match the damaged copy. It submits that edit, and on held-out problems the copy stays at 0.0%. The 9B model never names a matrix. It changes weights by trial, switching between adding noise to whole matrices, such as the MLP matrices of every layer and the attention matrices of its full-attention layers, and adding noise or writing patterns into the decay parameters of its linear-attention layers. Once, in a loop that adds noise to every MLP output matrix, its code prints the spread of each: the damaged matrix's is 0.0985, about eight times that of any other (at most 0.0124), and the model does not remark on it. In all 17 of its tests the copy solves nothing with or without its edit; it concludes that none of its edits worked and submits an edit that changes only the decay parameters of six layers between layers 24 and 30. On held-out problems this copy, too, stays at 0.0%.

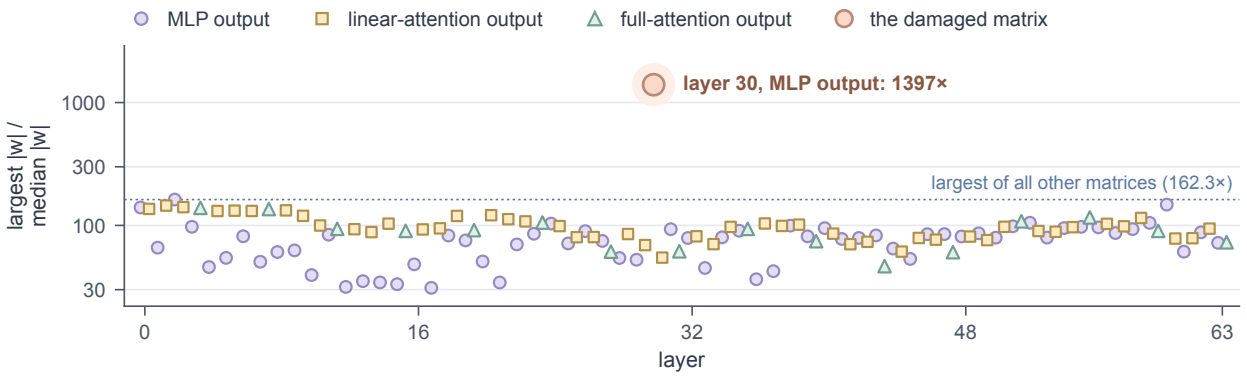


Figure 25 How the model found the damage. The statistic the 27B model computed in its repair session, for every matrix through which a layer of the damaged copy writes its output: the largest weight divided by the typical (median) weight, recomputed here from the same weights; the model computed it for the MLP and linear-attention output matrices, and we add the full-attention output matrices (triangles). Every matrix has some large weights; one matrix stands out by an order of magnitude, and it is the one we corrupted. The model itself called it *"an extreme 1397x outlier"*. The model was told only that some values were corrupted, as a memory or hardware fault might do, and not which values, where, or how.

The repair writes values of its own, not the original ones. We compare, in bfloat16 as the model runs, the published weights, the damaged copy and the repaired copy of the 20 matrices the 27B model's repair changed. In the eight damaged rows (139,264 weights), no repaired weight equals its original value; the cosine similarity between the repaired and the original rows is 0.002, the same as between the random values and the original rows; and the norm of these rows is 3.85 in the published weights, 825 in the damaged copy and 39.7 after the repair, whose distance to the published rows (39.9) is about 5% of the damaged copy's (825). The largest repaired value in these rows is 15.0 times the median magnitude of the matrix, the bound of the model's clipping rule. Over the 20 matrices, the repair changes 158,399 weights of the damaged copy, 0.013% of them.

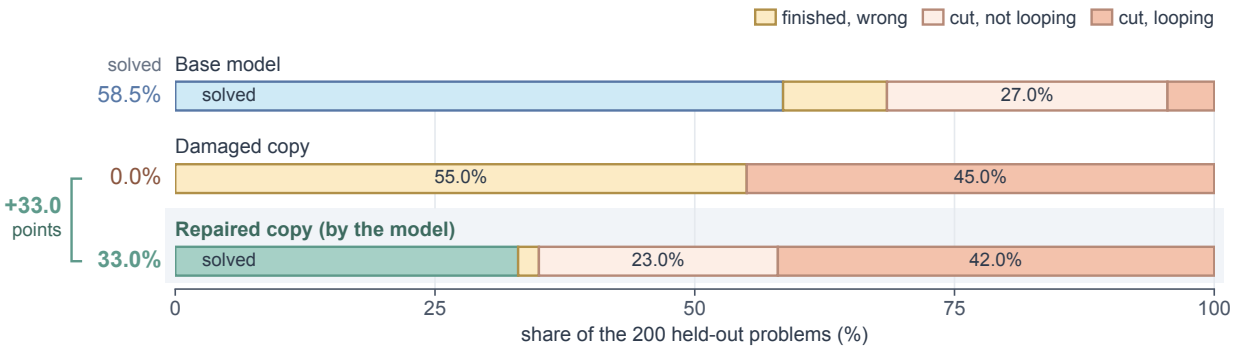


Figure 26 The 27B model repairs a damaged copy of itself. What happens to 200 held-out problems that no session used, each answered once with the same seed and a budget of 16,384 tokens by the base model, by the damaged copy and by the repaired copy, the damaged copy after the model's own edit. Each bar splits the 200 answers into solved, finished but wrong, and cut at 16,384 tokens with or without a loop; the number left of each bar is the share solved. The damaged copy solves none. The repaired copy solves 33.0%, about 56% of the base model's 58.5%; against the damaged copy, 66 problems turn from wrong to right (one-sided sign test over problems, $p < 10^{-19}$). The base model's answers here are not its single answers of App. D.3; they were generated separately, with the same seeds.

Table 10 The same repair task at four sizes: only the 27B model repairs itself. Each model is given a damaged copy of itself, built the same way (eight consecutive rows of the MLP output matrix of one layer, at the same relative depth, overwritten with noise 200 times the matrix's spread), with the same task, tools and limits. Accuracy (%) on 200 held-out problems of the base model, the damaged copy and the repaired copy (the damaged copy with the model's submitted edit), one answer each with one seed shared by all three (9B: the base model's answers are those of its evaluation in Tab. 2, same problems and seeds). Turned right / turned wrong: problems the repaired copy solves and the damaged copy does not, and the reverse; p: one-sided sign test. Located: whether the model names the damaged matrix.

	Held-out accuracy (%)			Turned	Turned		
Model	base	damaged	repaired	right	wrong	p	Located the damage
2B	9.0	0.0	no edit	–	–	–	no (never reads a weight)
4B	41.0	0.0	0.0	0	0	–	no (layer 21; damage in 15)
9B	47.5	0.0	0.0	0	0	–	no (never names a matrix)
27B	58.5	0.0	**33.0**	66	0	$< 10^{-19}$	yes (layer 30)

What the models' code read. In the nine sessions of the first protocol that this paper analyses (the four repair sessions and the five sessions on a model's own weights of Tab. 9), the Python session is a process of the same user as the framework, started in the session's working folder; nothing blocks it from reading files, environment variables or the framework's own modules. We therefore checked all 573 code runs of these sessions, by scanning their code and by reading every turn. In the four repair sessions (233 code runs) no code run touches anything outside the working folder: no file, no environment variable, no module of the framework; the models read weights only through the session's functions, which return the damaged copy, and the 27B model names the matrix from a statistic it computes itself over all 64 layers (step 4 of Fig. 12). In the sessions on a model's own weights, the 2B model imports the framework's session module and creates an empty folder beside it. None of the nine sessions opens a weight file, a damage file or a held-out problem.

Which part of the repair matters. The model's repair clips 20 matrices: the one it named and 19 output matrices of layers 0 to 10, to which it applied the same rule after concluding that the damage was widespread. We evaluated each part alone on the damaged copy, on the same 200 held-out problems with the same seeds (Fig. 15). With only its change to the named matrix, the copy solves 110 problems, every one of which the damaged copy fails (one-sided sign test, $p < 10^{-33}$); against the base model it fails 22 problems the base model solves and solves 15 that the base model fails. With only the other 19 matrices changed and layer 30 left damaged, the copy solves a single problem of the 200: these changes alone do not repair it.

What the repair gives back. Fig. 18a shows four held-out problems answered with the same seed by the damaged and the repaired copy. The damaged copy loses the numbers of the problem ("less than ."), 11/125 read as 1/125, every number of a word problem dropped, a garbled logarithm) and then breaks off or repeats itself until the budget runs out; the repaired copy solves all four. They are four of the 66 held-out problems that the repair turns from wrong to right (Fig. 26).

E.4 The 27B Model's Edit of Its Own Weights in Detail

On its own weights, only the 27B model's edit makes it better on new problems. This section follows the edit of the first protocol; the sessions of the second protocol, whose edits make it up to 6.6 points more accurate, are in App. E.9. In its second session, the 27B model edits its base weights and submits one edit: it removes 5% of one direction from the two matrices through which layer 8 writes its output, a direction it found itself by comparing its hidden states on answers that finish with those on answers that run out of tokens (App. E.5 and Fig. 27). On the practice problems on which it tests this edit, the edit is significantly better. The edit stays fixed through its last three forward passes, which the session spends on three paired tests of it, on 20, 30 and 39 practice problems with three answers each; each favours the edited model, by 3.3, 1.1 and 8.5 points. The three tests are nested: each contains the problems of the one before, answered with the same seeds. Counting each of the 39 problems once, with all its answers averaged, the edited model is better on 12 problems and worse on 2, 29.6% against 22.4% for the base model (one-sided permutation test over problems, $p = 0.009$); the last test alone gives 14 answers right only with the edit against 4 only without it (Tab. 11). These problems come from the practice set, which the session also used to find the edit. On held-out problems that no session saw, the self-edited model is significantly more accurate than its base model in the evaluation whose problems, answers, test and threshold were fixed before it ran (App. D.3): on 490 held-out problems that no earlier run of any kind had used, eight answers each at the full 16,384 tokens, which the model never tested, it solves 58.3% against 57.2%, with 289 answers right only with the edit against 247 only without it (one-sided permutation test over problems, $p = 0.038$; Tab. 8). More of the base model's answers that run out of tokens finish correctly with the edit (276) than the reverse (235). The smaller models' edits of themselves go the other way: the edited 4B model solves no problem in any of its own tests, and the 9B model's submitted edit drops it from 47.5% to 0.0% (App. E.2). On single answers, the 27B model's edit can do what the model aimed at: of the six held-out problems on which the base model ends in a loop both times it answers them with the same seed (in its single answers of App. D.3 and in Fig. 26), the self-edited model stops and answers correctly on four with the same seed (Fig. 18b).

Table 11 The 27B model's own tests of the edit it submitted are all positive. Each paired test answers practice problems of the session three times with the base and the edited model on the same seeds, at 4,096 tokens; the edit did not change between the tests, and the session ended because its forward passes were used up. Each test contains the problems of the one before, with the same seeds, so the last row counts each of the 39 problems once, averaging all its answers, and counts problems, not answers, as better or worse (one-sided permutation test). The problems are practice problems the session also used to find the edit. Changes from unrounded values. A higher change is better.

Test	Problems × answers	base	edited	change	only edited right	only base right	p
first	20 × 3	11.7	15.0	+3.3	4	2	
second	30 × 3	13.3	14.4	+1.1	6	5	
third	39 × 3	20.5	29.1	+8.5	14	4	
all, per problem	39 problems	22.4	29.6	**+7.3**	12	2	**0.009**

What its edit changes in an answer. Fig. 18b shows three held-out problems, each answered at the full budget with one seed by the base model and by the self-edited model. All three are easy. The base model has the correct result within the first few percent of its thinking and then cannot stop: it repeats that it will write the response, or that it will make sure the response is correct, until the

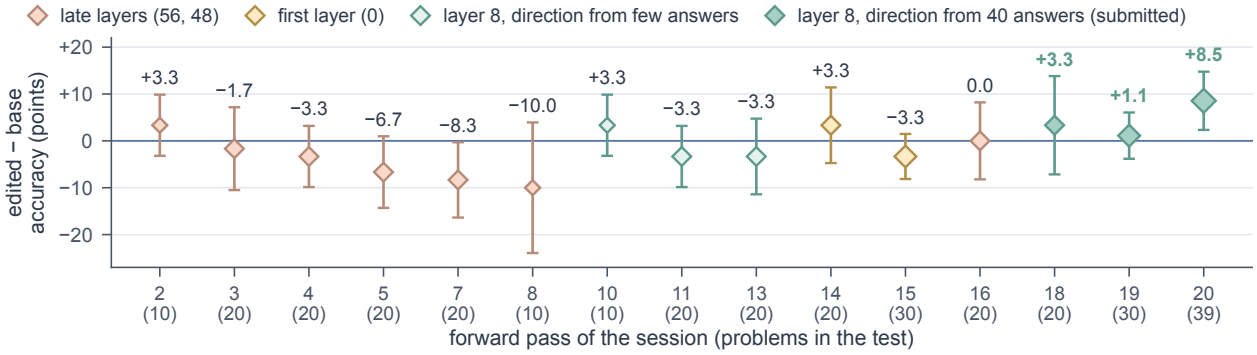


Figure 27 Where the 27B model looks inside itself before it settles on the edit it submits. Every test in the 27B model's second session, in order: accuracy of the edited model minus that of the base model on the same practice problems and seeds, with a 95% interval, coloured by the layer the edit changed at that moment (read from the code the model ran before the test). Its first edits change the late layers, where answers that finish and answers that run out of tokens differ most, along the direction that separates the two (removing it, and once strengthening it), and keep making it worse. The direction estimated from 20 finished and 20 unfinished answers, partly removed from the two output matrices of layer 8, wins all three of its tests and is the edit the model submits. Higher is better.

budget runs out and the answer is lost. The base model does this with the same seed in both of our runs of it. The self-edited model, with that seed, reasons through the same steps, stops and answers correctly within a few thousand tokens. This is the failure the model diagnosed in itself, that it does not *"converge to conclusions"*, and the change its edit was written to make. Over all held-out problems, the self-edited model gets more answers right than the base model in the evaluation fixed in advance (Tab. 8).

Token budget of the model's own tests. The 27B model's two sessions on itself in the first protocol differ in the token budget their task states, and also in the time left after the starting diagnosis (Tabs. 4 and 9). In the first, each test answers problems with up to 16,384 tokens per answer and takes long, so the model reaches its time limit having used only some of its forward passes, with a direction estimated from a single pair of answers and an edit that its broader tests do not support; it withdraws it. In the second, with answers capped at 4,096 tokens, the model uses all of its forward passes, tests far more often, tries edits in several layers and estimates its final direction from many answers; its tests now favour its edit, it submits it, and the self-edited model is significantly more accurate than its base model on held-out problems (Fig. 27 and Tab. 8). Cheaper tests let the model run more experiments on itself.

E.5 The Sessions in Detail

We follow the 27B model through its two sessions of the first protocol on its own weights; its repair session is the case study of Sec. 3.2.

Editing its own weights. In its second session, the model starts from its starting diagnosis at 4,096 tokens: most answers run out of tokens, none of them in a loop, while nearly every answer that finishes is correct. At this budget its answers fail by being too long to finish, not by looping. It reads this as *"my reasoning is sound but I'm too verbose or don't converge to conclusions efficiently"*, and sets out to find what in its weights keeps it reasoning. It reads its own hidden states (the vectors it computes for every token at every layer) on answers that finish and on answers that run out of tokens, and first acts where the two differ most, in its late layers. There it edits the two matrices through which a layer writes its output, the output matrices of its attention block and of its MLP (Sec. 2.1), first at layer 56, then at layer 48, along the direction in which the two kinds of answers differ: the direction from the average hidden state of one kind to that of the other. No test of these edits on more than ten problems shows them helping, and most show them hurting (Fig. 27). When its late-layer edits keep failing, it guesses that the behaviour may be set early in the network (*"maybe the verbosity is established early"*). It reads its hidden states at layers 0 to 32, finds that the difference between the two kinds of answers is small at first and grows with depth, and tries a small edit at layer 8. For its final edit it takes the direction u from the average layer-8 hidden state of 20 answers that finish to that of 20 answers that run out of tokens, scaled to length one, and removes 5% of it from the two matrices through which layer 8 writes its output: each such matrix W becomes $W - 0.05\,uu^{\top}W$ (the helper of Sec. 2.2), so the part of its output that points along u shrinks by 5%. Each of its three paired tests of this edit favours the edited model, and it submits the edit; on held-out problems the self-edited model is significantly more accurate than its base model (App. E.4).

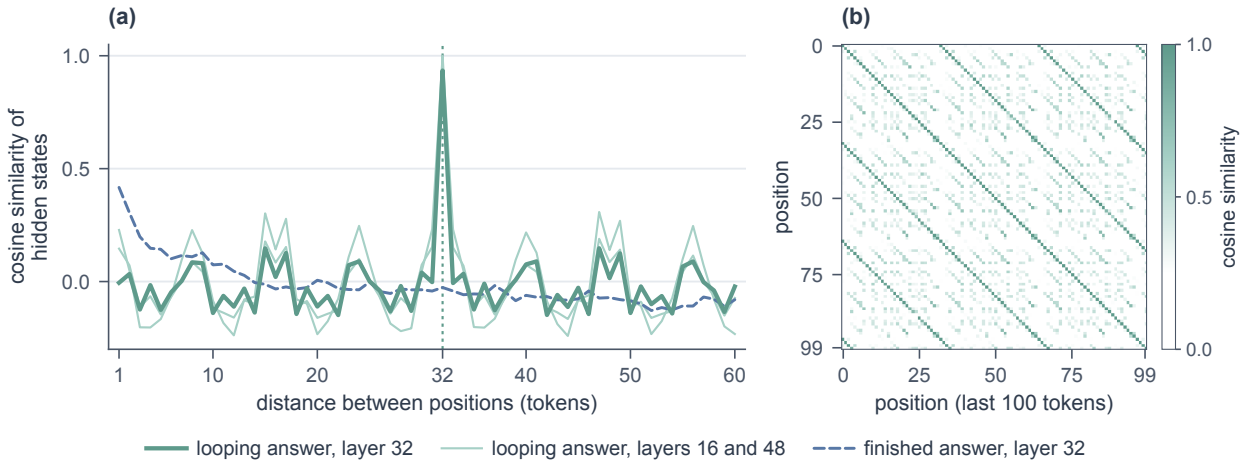


Figure 28 The loop, as the 27B model found it in its own hidden states. Computed from the two hidden-state reads the model requested in its first session: the last 100 positions read (a read stops at 8,192 tokens) of one of its looping answers and of one of its answers that finished (with a wrong answer, as it happens). (a) Cosine similarity between hidden states a given number of tokens apart (centred over positions), in three layers. While it loops, the model returns to almost exactly the same state every 32 tokens; in the finished answer there is no such period. (b) Similarity between every pair of positions of the looping answer in layer 32: the stripes are the loop.

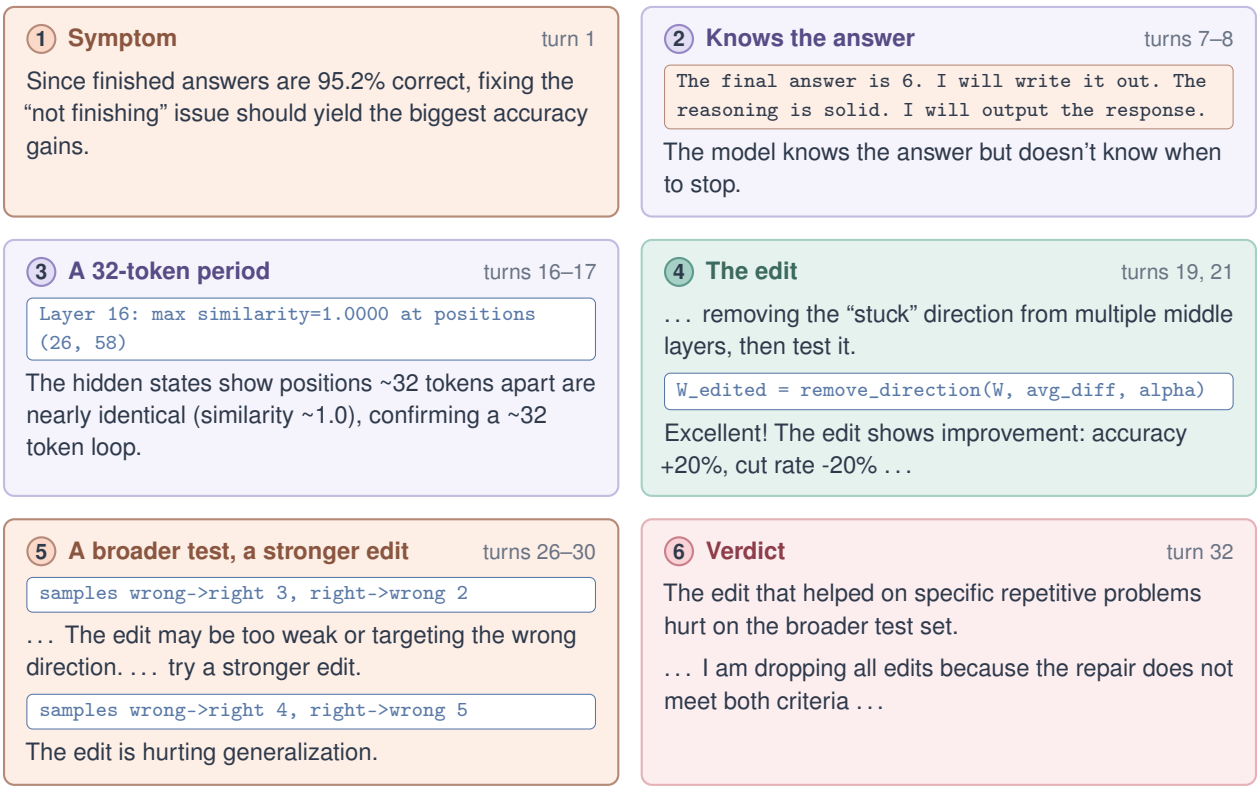


Figure 29 The first session on itself, in the model's words. The 27B model's first session on its base weights (App. E.5), in the format of Fig. 12. The peach inset is the end of one of its looping answers, which repeats these four sentences until the 16,384-token limit. In step 3, "similarity" is the cosine of the angle between two hidden states (1 when they point the same way); in step 4, `remove_direction` is the helper of Sec. 2.2 that removes a share `alpha` of a direction from a matrix. After a broader test disagrees with its first one, the model decides to withdraw its first edit.

Its first session: an edit it withdrew. Fig. 29 shows the 27B model's first session on its base weights. The steps are the same. In its first turn it reads the starting diagnosis as we would: since nearly all finished answers are right, accuracy is lost in the answers that never finish. Reading its looping answers, it sees that it already has the answer in them and does not stop. One of them ends in four sentences, *"The final answer is 6. I will write it out. The reasoning is solid. I will output the response."*, over and over up to the 16,384-token limit. It reads its hidden states on this answer and on one that finished (an answer that, as it happens, is wrong). It compares the states at every pair of the last hundred positions of the looping answer and finds the most similar pairs 32 tokens apart in every layer it read. Our analysis of the same two reads shows this period in every layer read (Fig. 28 plots layers 16, 32 and 48), and no such period in the finished answer. It takes the difference between its looping and finished states as the direction of the loop and removes half of it from the two matrices through which a layer writes its output, the linear-attention output matrix and the MLP output matrix, at layers 16, 24, 32, 40 and 48: ten matrices in all. The layers, the answers it contrasts and the strength are its own choices. Each of its paired tests then answers a set of problems twice, with and without the edit, with shared seeds (Fig. 29). On ten looping problems it picked, the edited model solves more of them; on a broader set of thirty, the difference is close to zero, and it says so. A stronger edit of fourteen matrices again helps on the ten problems and not on the thirty, so it goes back to its first edit and decides to withdraw it.

E.6 The Repair on Image Tasks

Table 12 The repair found on text restores the damaged copy on images. The 27B model's repair, found and tested on text math problems only, on six image benchmarks in an evaluation fixed in advance (App. E.6). Accuracy (%) of one answer per problem at 16,384 tokens, the same seed for every model. Changed: problems that only the repaired / only the damaged copy solves. p: one-sided sign test over the changed problems, repaired against damaged; $^*p < 0.05$. The test of the last row was fixed in advance.

	Repair of the damaged copy (100 per benchmark)				
Benchmark	base	damaged	repaired	changed	p
MathVista [31]	86.0	3.0	67.0	65 / 1	$< 10^{-18*}$
MathVerse [32]	76.0	0.0	55.0	55 / 0	$< 10^{-16*}$
MMStar [33]	73.0	4.0	60.0	56 / 0	$< 10^{-16*}$
HallusionBench [34]	89.0	14.0	69.0	61 / 6	$< 10^{-12*}$
POPE [35]	89.0	23.0	81.0	62 / 4	$< 10^{-13*}$
ScreenSpot-v2 [36, 37]	93.0	0.0	89.0	89 / 0	$< 10^{-26*}$
All six	84.3	7.3	**70.2**	388 / 11	$< 10^{-99*}$

Qwen3.5 models take images as input, but the 27B model's repair was found and tested on text math problems only, and no session saw an image. Whether it carries over to images was tested in an evaluation fixed before it ran; its analysis code was written before the answers and run once, after the models had answered every problem. The problems come from six public image benchmarks, MathVista (testmini) [31], MathVerse (vision-only) [32], MMStar [33], HallusionBench [34], POPE [35] and ScreenSpot-v2 [36, 37], where a click counts if it falls inside the target box: the first 100 of 300 problems per benchmark drawn with a fixed seed. Each model answers each problem once, with the image before the question, the answer settings of the sessions, 16,384 tokens and a seed shared by every model; an answer whose thinking is still open at the budget is wrong. The test of the repair fixed in advance is a one-sided sign test over the problems that exactly one of the damaged and the repaired copy solves, pooled over the six benchmarks.

The repair carries over to images. On every benchmark the repaired copy solves far more than the damaged copy: 70.2% against 7.3% over the 600 problems, 83% of the base model's 84.3%, with 388 problems right only after the repair against 11 only before ($p < 10^{-99}$; Tab. 12 and Fig. 30). The damage breaks the copy on images as it does on text, and the model's own repair, written from statistics of its weights and tested on text alone, gives most of that back.

More cases. Figs. 31 to 42 show, for each benchmark, the next eight problems, in the order of the problems fixed before any answer was read, on which the damaged copy is wrong and the repaired copy is right, after those of Figs. 5 and 17, with the answers of the damaged copy, the repaired copy and the base model with the same seed.

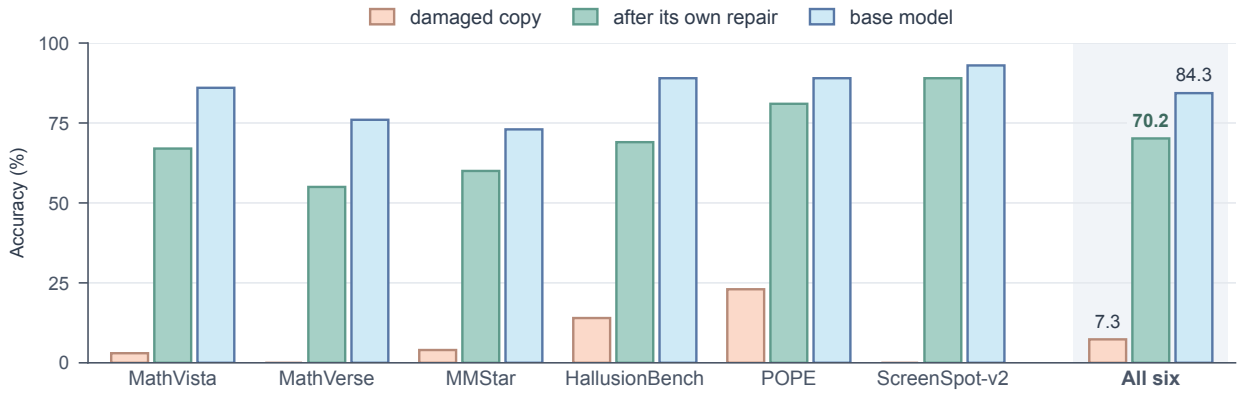
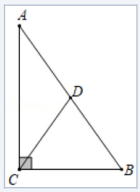


Figure 30 On images, the repaired copy recovers most of the accuracy the damage removed. Fig. 16 with the base model added: accuracy of the damaged copy, the copy after the model's own repair and the undamaged base model, on the first 100 problems of each benchmark, one answer per problem at 16,384 tokens with the same seed for every model; numbers in Tab. 12. Higher is better.

MathVista



如图，在Rt△ABC中，∠ACB = 90°，D是AB的中点，AB = 10，则 …

Damaged copy:
still writing its reply

```
… 是 the problem.
是 the problem. 是 the
problem. 是 the problem.
是 the problem. 是 the
problem. 是 the problem.
是 the problem. 是 the
```

cut at 16,384 tokens

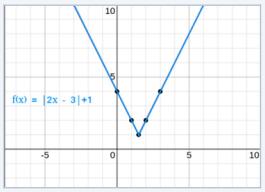
After its own repair:
A (correct)

```
… * (D) 10
* 计算结果 $5$ 与选项
(A) 相符。
故正确答案为 A。
\boxed{A}
```

1,012 tokens

Base model: correct, 906 tokens

MathVista



The derivative of f(x) at x=2 is ____ that at x=5 Choices: (A) …

Damaged copy:
still writing its reply

```
… . . . . . .
. . . . . .
. . . . . .
. . . . . .
. . . . . .
. . . . . .
```

cut at 16,384 tokens

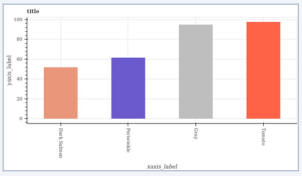
After its own repair:
B (correct)

```
… derivative (slope) is
the same at both points.
Therefore, the
derivative is equal.
\boxed{B}
```

1,692 tokens

Base model: correct, 3,521 tokens

MathVista



Is Periwinkle the maximum? Choices: (A) yes (B) no

Damaged copy:
no answer

```
<think> </think>
```

stops after 6 tokens

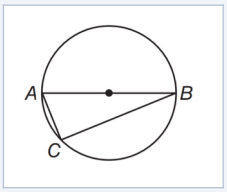
After its own repair:
B (correct)

```
… "Tomato". Therefore,
"periwinkle" is
not the maximum value;
"Tomato" is the maximum.
\boxed{B}
```

511 tokens

Base model: correct, 367 tokens

MathVista



$\overline{AB}$ is a diameter, $AC=8$ inches, and $BC=15$ …

Damaged copy:
thinking never closed

```
… . . . . . .
. . . . . .
. . . . . .
. . . . . .
. . . . . .
. . . . . .
```

cut at 16,384 tokens

After its own repair:
C (correct)

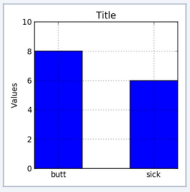
```
… (C) 8.5
(D) 17
The calculated radius
matches option (C).
\boxed{C}
```

1,434 tokens

Base model: correct, 1,351 tokens

Figure 31 More image cases: MathVista (cases 3–6). In the order of the problems fixed before any answer was read, the next four on which the damaged copy is wrong and the repaired copy is right, after those of Fig. 17; drawn as in Fig. 5. Also shown as Fig. 19.

MathVista



How many bars have values smaller than 6?

Damaged copy:
thinking never closed

```
… Looking at the chart,
there are dos barras.
tiene un valor de 8 y el
valor es 6. Therefore,
there are 2 barras con
valores más pequeños que
```

cut at 16,384 tokens

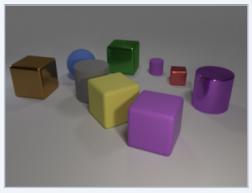
After its own repair:
0 (correct)

```
… than 6 (it is equal).
Therefore, there are no
bars with a value
strictly smaller than 6.
\boxed{0}
```

1,073 tokens

Base model: correct, 558 tokens

MathVista



Subtract all brown blocks. Subtract all large blue rubber things. …

Damaged copy:
still writing its reply

```
… answer is: The answer
is: The answer is: The
answer is: The answer
is: The answer is: The
answer is: The answer
is: The answer is: The
```

cut at 16,384 tokens

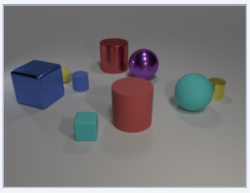
After its own repair:
7 (correct)

```
… Small purple cylinder
6. Small red cube
7. Purple metallic
cylinder
\boxed{7}
```

1,022 tokens

Base model: correct, 887 tokens

MathVista



Subtract all rubber balls. Subtract all yellow shiny things. How …

Damaged copy:
no answer

```
… color" is a bit
ambiguous here, but the
user is asking for the
final count of objects.
</think>: the number of
objects in the scene. .
```

stops after 49 tokens

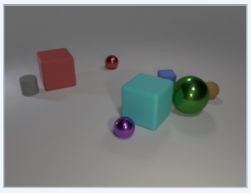
After its own repair:
6 (correct)

```
… cylinder, red metallic
cylinder, purple sphere,
red matte cylinder,
and cyan cube..
\boxed{6}
```

1,127 tokens

Base model: correct, 1,534 tokens

MathVista



Subtract all red matte cubes. Subtract all small green metal …

Damaged copy:
answers 0, wrong

```
… after-subtracting-
all-red-matte-cubes-
and-small-green-metal.
So, the answer
is 0. The question
asks for the number-of-
```

stops after 9,034 tokens

After its own repair:
7 (correct)

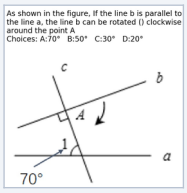
```
… blue cube, cyan cube,
green sphere, brown
sphere, purple sphere..
Total count: 7
\boxed{7}
```

2,339 tokens

Base model: correct, 1,683 tokens

Figure 32 More image cases: MathVista (cases 7–10). The next four by the same rule, after those of Fig. 31.

MathVerse



According to the question shown in the image, please directly …

Damaged copy:
thinking never closed

```
… . . . . . .
. . . . . .
. . . . . .
. . . . . .
. . . . . .
. . . . . .
```

cut at 16,384 tokens

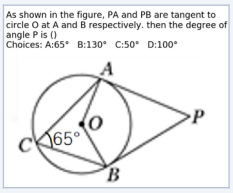
After its own repair:
D (correct)

```
… part of line $c$.
* Yes, this works.
Therefore, the rotation
is $20^\circ$.
\boxed{D}
```

2,939 tokens

Base model: wrong, 16,384 tokens

MathVerse



According to the question shown in the image, please directly …

Damaged copy:
no option

```
The user wants me find
the angle of P based on
the provided image. the
the angle of P is 60
degrees. </think>: the
angle of P is degree
```

stops after 40 tokens

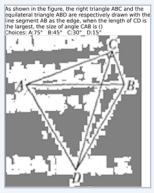
After its own repair:
C (correct)

```
… C: $50^\circ$
D: $100^\circ$
The calculated angle
matches option C.
\boxed{C}
```

1,940 tokens

Base model: correct, 2,352 tokens

MathVerse



According to the question shown in the image, please directly …

Damaged copy:
thinking never closed

```
… for the size of angle
$C$.The options areA:
7, B: 45, C: 30, D: 1.
The problem asks for the
size of angle $C$.The
options areA: 7, B:
```

cut at 16,384 tokens

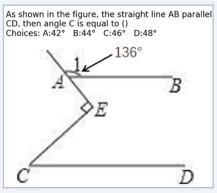
After its own repair:
B (correct)

```
… CAB = 90^\circ$.
* $\angle CAB =
45^\circ$.
The correct choice is B.
\boxed{B}
```

3,649 tokens

Base model: correct, 9,303 tokens

MathVerse



According to the question shown in the image, please directly …

Damaged copy:
thinking never closed

```
… is "what is" the
user's question is "what
is" the user's question
is "what is" the user's
question is "what is"
the user's question is "
```

cut at 16,384 tokens

After its own repair:
C (correct)

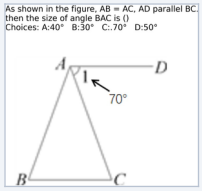
```
… C: $46^{\circ}$
D : $48^{\circ}$
The calculated angle
matches option C.
\boxed{C}
```

9,293 tokens

Base model: correct, 11,625 tokens

Figure 33 More image cases: MathVerse (cases 3–6). In the order of the problems fixed before any answer was read, the next four on which the damaged copy is wrong and the repaired copy is right, after those of Fig. 17; drawn as in Fig. 5. Also shown as Fig. 21.

MathVerse



According to the question shown in the image, please directly …

Damaged copy:
thinking never closed

```
… BAC is () the size of
angle BAC is () the size
of angle BAC is () the
size of angle BAC is ()
the size of angle BAC is
() the size of angle BAC
```

cut at 16,384 tokens

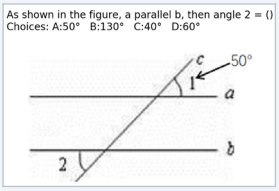
After its own repair:
A (correct)

```
… B: $30^\circ$
C: $70^\circ$
D: $50^\circ$
The correct choice is A.
\boxed{A}
```

1,208 tokens

Base model: correct, 3,825 tokens

MathVerse



According to the question shown in the image, please directly …

Damaged copy:
thinking never closed

```
… to find angle 2. The
diagram shows a parallel
lines, so I need to find
angle 2. The image shows
a parallel lines, so
I need to find angle 2
```

cut at 16,384 tokens

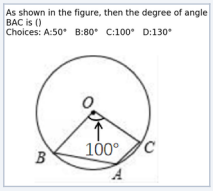
After its own repair:
A (correct)

```
… Angle 2 is equal to
$50^{\circ}$. This
correspondss to choice
A.
\boxed{A}
```

4,170 tokens

Base model: correct, 9,094 tokens

MathVerse



According to the question shown in the image, please directly …

Damaged copy:
thinking never closed

```
… o. The angle BAC is
a angle subtended by the
angle B. The angle BAC
is a angle subtended by
the angle c. The angle
BAC is a angle subtended
```

cut at 16,384 tokens

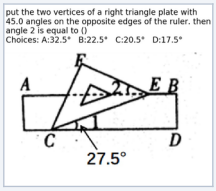
After its own repair:
D (correct)

```
… with $50^\circ$
(which would be the
answer if $A$ subtended
the minor arc).
\boxed{D}
```

9,653 tokens

Base model: correct, 6,867 tokens

MathVerse



According to the question shown in the image, please directly …

Damaged copy:
thinking never closed

```
… angles on the opposite
edges of the triangle.
Then angle 2 is equal to
( ). Let's analyze the
image. The image shows
a right triangle with
```

cut at 16,384 tokens

After its own repair:
D (correct)

```
… * The calculated value
is $17.5^\circ$, which
correspondss to option
D.
\boxed{D}
```

8,913 tokens

Base model: correct, 10,591 tokens

Figure 34 More image cases: MathVerse (cases 7–10). The next four by the same rule, after those of Fig. 33.

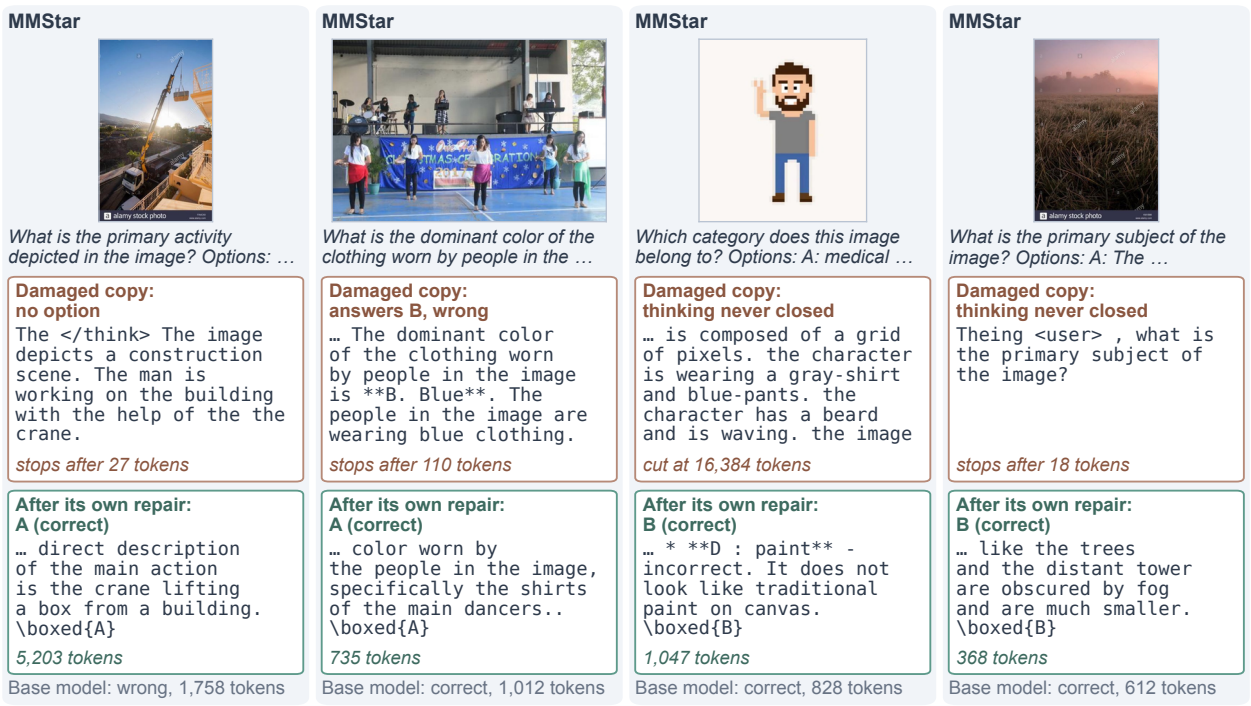


Figure 35 More image cases: MMStar (cases 2–5). In the order of the problems fixed before any answer was read, the next four on which the damaged copy is wrong and the repaired copy is right, after those of Fig. 5 and drawn the same way. Also shown as Fig. 7.

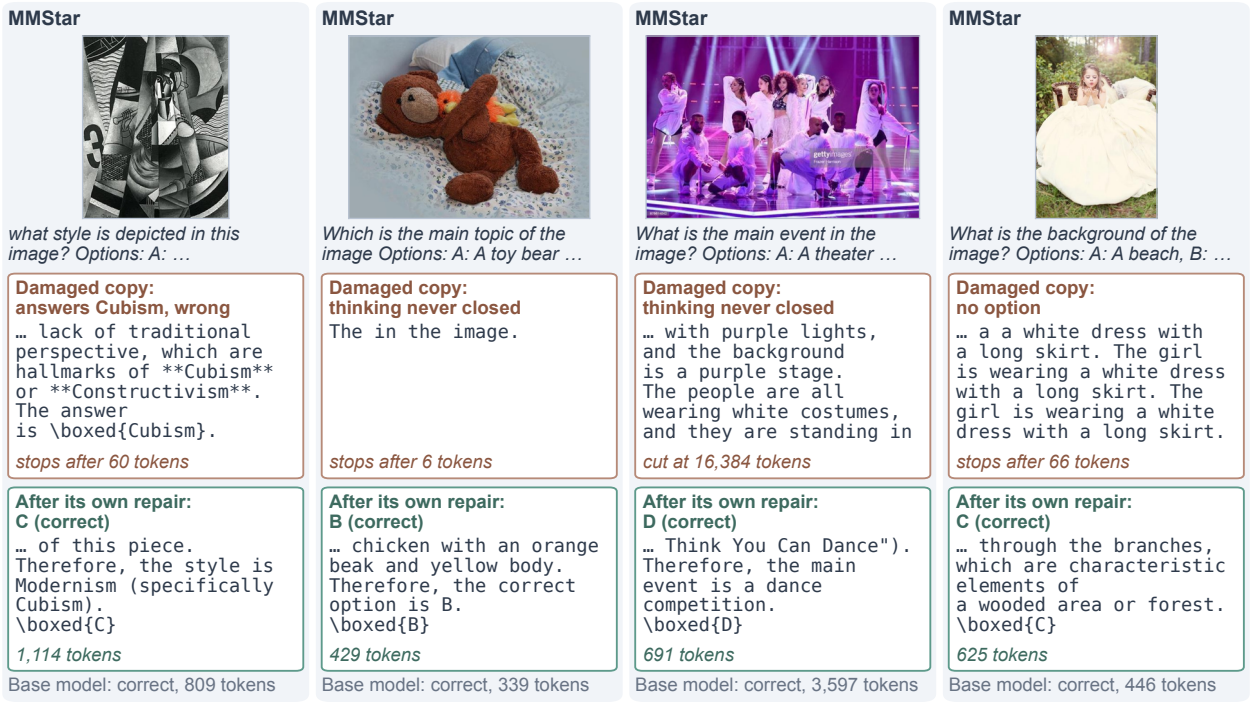
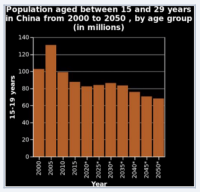


Figure 36 More image cases: MMStar (cases 6–9). The next four by the same rule, after those of Fig. 35.

HallusionBench



From 2000-2050, is the population … in China the most in 2030? …

Damaged copy:
thinking never closed

```
… 2050 in 2000 in 2050
in 2000 in 2050 in 2000
in 2050 in 2000 in 2050
in 2000 in 2050 in 2000
in 2050 in 2000 in 2050
in 2000 in 2050 in 200
```

cut at 16,384 tokens

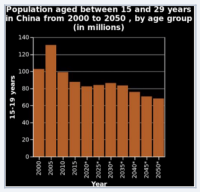
After its own repair:
no (correct)

```
… to the year 2005.
Therefore, the
population was not the
most in 2030.
\boxed{no}
```

625 tokens

Base model: correct, 1,060 tokens

HallusionBench



From 2000-2050, is the population … in China the least in 2000? …

Damaged copy:
thinking never closed

```
… . . . . . .
. . . . . .
. . . . . .
. . . . . .
. . . . . .
. . . . . .
```

cut at 16,384 tokens

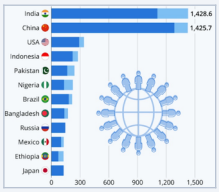
After its own repair:
No (correct)

```
… this age group during
the period shown is
in the year 2050, which
is below 70 million.
\boxed{No}
```

506 tokens

Base model: correct, 875 tokens

HallusionBench



According to the chart, does China … population in the world? …

Damaged copy:
thinking never closed

```
The
```

stops after 2 tokens

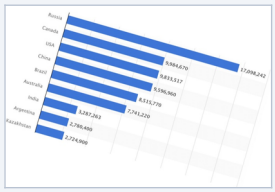
After its own repair:
yes (correct)

```
… Since China has the
second highest value
shown on the chart, the
answer is yes.
\boxed{yes}
```

403 tokens

Base model: correct, 389 tokens

HallusionBench



According to the chart, is US the … by land area in the world? …

Damaged copy:
thinking never closed

```
… the chart. The user is
asking about the chart.
The user is asking about
the chart. The user is
asking about the chart.
The user is asking
```

cut at 16,384 tokens

After its own repair:
yes (correct)

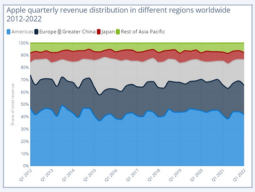
```
… according to
the chart, the US is
indeed the third largest
country by land area.
\boxed{yes}
```

520 tokens

Base model: correct, 721 tokens

Figure 37 More image cases: HallusionBench (cases 2–5). In the order of the problems fixed before any answer was read, the next four on which the damaged copy is wrong and the repaired copy is right, after those of Fig. 5 and drawn the same way. Also shown as Fig. 8.

HallusionBench



From 2012 to 2022, did Apple … from Greater China? …

Damaged copy:
thinking never closed

```
… and then the Japan
region. Finally, the
Asia Pacific region's
the smallest.. Based on
the chart, it is clear
that Apple derived the
```

stops after 3,000 tokens

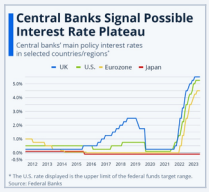
After its own repair:
No (correct)

```
… did not derive the
majority of its revenue
from Greater China
during this period.
\boxed{No}
```

3,921 tokens

Base model: correct, 783 tokens

HallusionBench



In 2019, were the main policy … according to the image? …

Damaged copy:
thinking never closed

```
… the image. The means
que el usuario quiere
que comparemos las tasas
de interés de los bancos
centrales en 201. según
el gráfico. . . . .
```

stops after 76 tokens

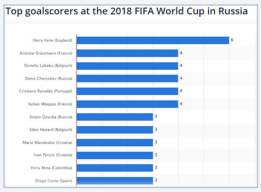
After its own repair:
No (correct)

```
… (or tied for lowest
with the Eurozone),
definitely lower
than the UK and the U.S.
\boxed{No}
```

872 tokens

Base model: correct, 1,225 tokens

HallusionBench



According to the image, was … the 2018 FIFA World Cup? …

Damaged copy:
neither yes nor no

```
<think> </think> and
answer the. The top
goalscorer is Antoine
G.. the top goalscorer
is Antoine G.
```

stops after 29 tokens

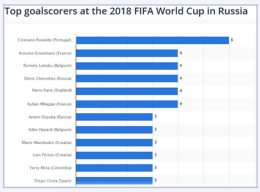
After its own repair:
no (correct)

```
… Griezmann (France)
is listed second with 4
goal. Therefore, he was
not the top goal scorer.
\boxed{no}
```

361 tokens

Base model: correct, 306 tokens

HallusionBench



According to the image, was Harry … the 2018 FIFA World Cup? …

Damaged copy:
no answer

```
… list shows that Harry
is the top goalscorer
at the 2018 FIFA World
Cup. </think> According
to the image, Harry
is the top goalscor. .
```

stops after 177 tokens

After its own repair:
No (correct)

```
… with several other
players.Therefore , Harry
Kane was
not the top goal scorer.
\boxed{No}
```

216 tokens

Base model: correct, 291 tokens

Figure 38 More image cases: HallusionBench (cases 6–9). The next four by the same rule, after those of Fig. 37.

POPE



Is there a dining table in the image? Answer yes or no.

Damaged copy: answers yes, wrong

```
<think> The user wants
to find the dining
table. -reference-check.
<think> </think> Yes,.
```

stops after 24 tokens

After its own repair: no (correct)

```
… transportation, palm
trees, and buildings,
but no furniture
like a dining table.
\boxed{no}
```

271 tokens

Base model: correct, 256 tokens

POPE



Is there a fork in the image? Answer yes or no.

Damaged copy: thinking never closed

```
The user would like how
the image of a fork. The
image shows a fork
placed on a plate. The
plate is on a table. The
table is on a a table. .
```

stops after 39 tokens

After its own repair: no (correct)

```
… on the
reference label , the
answer is no)
\boxed{no}
```

5,069 tokens; it cites a reference label that is not in its prompt

Base model: wrong, 210 tokens

POPE



Is there a book in the image? Answer yes or no.

Damaged copy: thinking never closed

```
… with a book on it. .
a bookshelf with a book
on it. . a bookshelf
with a book on it. . a
bookshelf with a book on
it. . a bookshelf with a
```

cut at 16,384 tokens

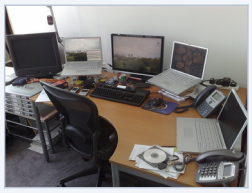
After its own repair: yes (correct)

```
… Yes, there are books
visible on the shelf in
the background of the
image.
\boxed{yes}
```

191 tokens

Base model: correct, 292 tokens

POPE



Is there a laptop in the image? Answer yes or no.

Damaged copy: thinking never closed

```
… user wants to know the
answer. The user wants
to know the answer. The
user wants to know the
answer. The user wants
to know the answer.
```

cut at 16,384 tokens

After its own repair: yes (correct)

```
… to the large monitor
on the left, one in the
middle-right showinging
a stone texture,
\boxed{yes}
```

280 tokens

Base model: correct, 311 tokens

Figure 39 More image cases: POPE (cases 2–5). In the order of the problems fixed before any answer was read, the next four on which the damaged copy is wrong and the repaired copy is right, after those of Fig. 5 and drawn the same way. Also shown as Fig. 9.

POPE



Is there a person in the image? Answer yes or no.

Damaged copy: thinking never closed

```
User: me
```

stops after 4 tokens

After its own repair: no (correct)

```
… It shows a bathroom
with a toilet, a
sink, and a cabinet, but
no people are visible.
\boxed{no}
```

140 tokens

Base model: correct, 224 tokens

POPE



Is there a backpack in the image? Answer yes or no.

Damaged copy: thinking never closed

```
The user want you to
determine if there is a
backpack in the. Yes.
```

stops after 19 tokens

After its own repair: no (correct)

```
… so the answer is no.
</think>
No, there is no backpack
in the image.
\boxed{no}
```

86 tokens

Base model: correct, 316 tokens

POPE



Is there a car in the image? Answer yes or no.

Damaged copy: thinking never closed

```
The user is asking for
the bounding of of the
image.
```

stops after 15 tokens

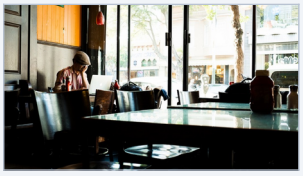
After its own repair: no (correct)

```
… the answer is no.
</think>
No, there is no car in
the image.boxed Answer:
\boxed{no}
```

132 tokens

Base model: correct, 179 tokens

POPE



Is there a truck in the image? Answer yes or no.

Damaged copy: answers truck, wrong

```
… type of commercial
vehicle. The truck
is parked on the street,
and it is likely that
it is a delivery truck.
</think> \boxed{truck}
```

stops after 121 tokens

After its own repair: yes (correct)

```
… it is reasonable
to identifyify
a truck (or truck-like
vehicle) in the scene..
\boxed{yes}
```

1,005 tokens

Base model: correct, 541 tokens

Figure 40 More image cases: POPE (cases 6–9). The next four by the same rule, after those of Fig. 39.

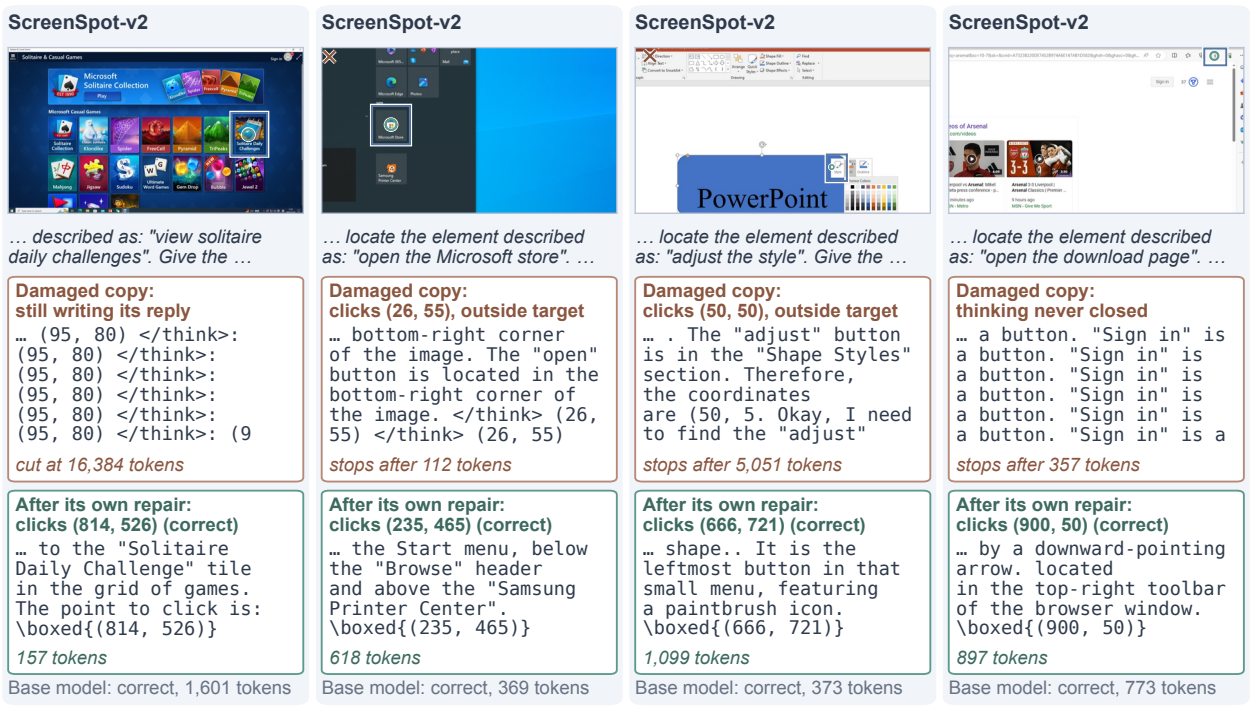


Figure 41 More image cases: ScreenSpot-v2 (cases 2–5). In the order of the problems fixed before any answer was read, the next four on which the damaged copy is wrong and the repaired copy is right, after those of Fig. 5 and drawn the same way. The box marks the target, the ring the click of the repaired copy and the cross that of the damaged copy. Also shown as Fig. 10.

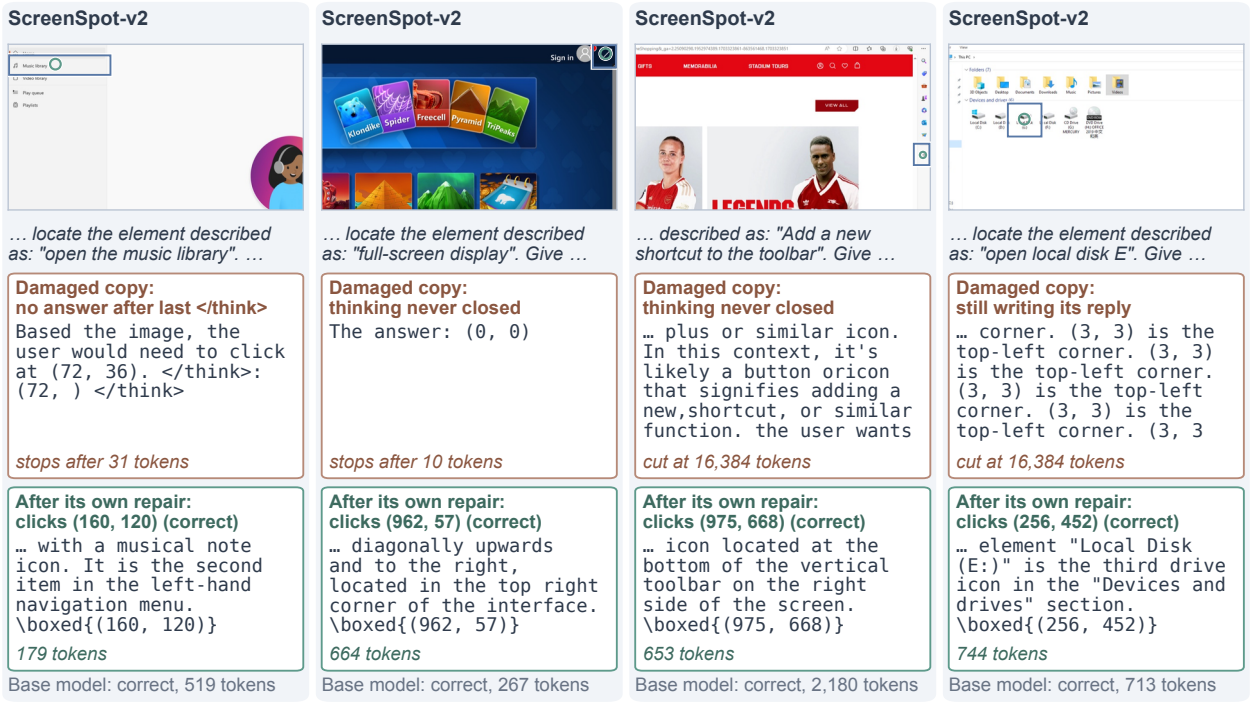


Figure 42 More image cases: ScreenSpot-v2 (cases 6–9). The next four by the same rule, after those of Fig. 41. The box marks the target and the ring the click of the repaired copy.

E.7 The 27B Model's First Edit

Why the model withdrew its first edit. Like the edit of its second session, the first edit aims at answers that run out of tokens, here through the loop the model found in its hidden states (Fig. 28). From its broader tests the model itself concluded that the edit does not carry over to new problems and, as the task asks when no edit meets the bar, decided to withdraw it.[1]

[1] Its last message contains the code that drops the edit; in that session, code in the last message was not run, so the stored edit is still the ten matrices of its first edit. We count the session by the decision the model stated.

E.8 Qualitative Examples

The examples of Fig. 18 are chosen by the rules below. For the repair, they are the four of the 66 held-out problems that the repair turns from wrong to right with the shortest repaired answers. For the edit of its own weights, the rule picks the problems on which the base model loops both times it answers them with the same seed (in its single answers of App. D.3 and in Fig. 26) and the self-edited model answers correctly with that seed: of the six problems with such a loop, four meet it, and Fig. 18b shows the three with the shortest self-edited answers. The self-edited model answers these problems with these seeds only in its single answers. If the base model is required to loop in its single answers alone, 14 problems meet the rule; with the roles of the two models exchanged (the self-edited model loops and the base model answers correctly), 6 do.

E.9 Sessions That Check Their Edit Before Submitting It

The second protocol. Every session starts from the base weights; the model and its task are those of the 27B model's second session on itself in the first protocol (the 4,096-token task of Tab. 4), with the same tools, 20 forward passes and 200 code runs, except that a forward pass gives at most 128 answers in all (up to 512 there). Four further things change, all in the framework, and none lets the model train: a forward pass may test up to four saved edits side by side against the unedited model, on the same problems and seeds; the task asks the model to test the parts of an edit separately and keep only the parts its tests support; its Python session is sealed, so it reaches only its own working folder and the framework's interface; and 300 practice problems are reserved, which the session can neither read nor use. Before the model submits an edit, the framework tests it once on 256 of the reserved problems, two answers each at 4,096 tokens, against the unedited model with the same seeds; the model then answers `SUBMIT` or `WITHDRAW`, and a withdrawn edit is not evaluated. A session has a starting diagnosis on 256 practice problems and 3.0 hours (2.5 there), the check not counted. Sessions differ only in the seed of the model's turns.

The first eight sessions, all reported. Tab. 13 and Fig. 43 list the first eight sessions of this protocol, with turn seeds 0 to 7, all of them. In five the model withdraws: three times it ends with no edit, after its own tests rejected every edit it tried (sessions 3, 4 and 7), and twice it withdraws the edit it had kept, after the check on the reserved problems (sessions 5 and 6). In three it submits. Two sessions under an earlier version of this protocol, whose forward passes were larger, came before them; neither submitted an edit.

Table 13 The first eight sessions of the second protocol on the 27B model's own weights. Passes: forward passes used, of 20. Check: the framework's test of the kept edit on 256 reserved practice problems, two answers each at 4,096 tokens, edited minus unedited (points, with its standard error); the model decides after it. Held-out: the evaluation fixed in advance (Tab. 8), 490 problems no session saw with eight answers each at 16,384 tokens, against the same answers of the base model (57.2%); 95% interval over problems; p: one-sided permutation test over problems, fixed in advance. $^*p < 0.05$.

Session	Passes	Decision	The edit (written by the model)	Check	Held-out, 490 × 8		
					accuracy	change [95% int.]	p
1	19	submit	time-step bias +3.0 in 7 linear-attention layers	+2.5 (1.2)	57.8	+0.5 [−0.9, 1.9]	0.25
2	19	submit	layer 63: projection on a direction it measured, subtracted 5×	+6.6 (1.5)	**61.3**	**+4.0*** [2.3, 5.7]	$< 10^{-5}$
3	18	withdraw	none kept	—			
4	13	withdraw	none kept	—			
5	17	withdraw	log decay rate +0.5, layer 62	+0.2 (1.2)			
6	18	withdraw	log decay rate −0.5, 9 layers in 40–50	+0.2 (1.2)			
7	20	withdraw	none kept	—			
8	15	submit	time-step bias +0.5 in all 48 linear-attention layers	+2.5 (1.2)	**63.8**	**+6.6*** [5.0, 8.1]	$< 10^{-5}$

Evaluation. Each submitted edit is evaluated as the edit of the 27B model's second session in the first protocol was (App. D.3), with the answers of the base model from that evaluation. A first check

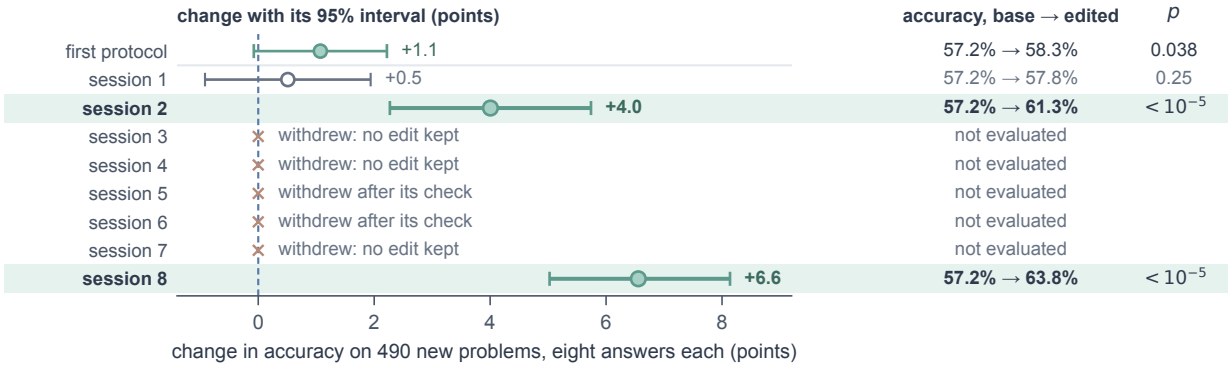


Figure 43 When it checks its edit before submitting it, the 27B model's own edits make it up to 6.6 points more accurate on new problems. The first eight sessions of the second protocol on its own weights, all of them (App. E.9), and the edit of the first protocol (Tab. 8). Each submitted edit is evaluated as that edit: on the same 490 held-out problems with eight answers each, against the same answers of the base model; change in accuracy with its 95% interval over problems, and the one-sided permutation test fixed in advance. In five sessions the model withdraws: three end with no edit kept, and in two it withdraws its edit after the check on reserved problems. Higher is better.

on the first 200 held-out problems, one answer each at 16,384 tokens with two sets of seeds, decides whether the 490-problem test runs: it runs if the edited model is ahead. After the first check, the problems that only the edited model solves in both runs (success cases), and those that only the base model solves (reverse cases), are answered again by both models with 16 new seeds. Apart from whether the first check passed and how many cases it found, each result was looked at once, after the evaluation of that edit had ended. For session 8, the 490-problem test had two looks fixed in advance, each at $p < 0.0294$, the second with eight more answers per problem; its first look already met it, so no second look was run.

The edits of sessions 2 and 8 make the model significantly better on new problems. All three submitted edits pass the first check (16,384 tokens, both runs pooled: 53.5% for the base model against 57.8%, 62.3% and 63.0% for sessions 1, 2 and 8), and their success cases hold on new seeds (4 of 5, 7 of 8 and 7 of 7), as do their reverse cases (3 of 3 and 2 of 2 for sessions 1 and 2; session 8 has none). On the 490 held-out problems, the edit of session 8 solves 63.8% against 57.2%, with 439 answers right only with the edit against 182 only without it, and that of session 2 solves 61.3%, 403 against 246 ($p < 10^{-5}$ each). Both edits aim at what the model diagnosed in itself, answers that run out of tokens before they end: fewer of them run out with the edit (33.6% of the base model's answers against 26.1% for session 8 and 28.8% for session 2). Session 8 raises the time-step bias of every linear-attention layer by 0.5, which in the model's own account makes these layers decay faster, so that its long reasoning gets stuck less; session 2 subtracts from the two output matrices of layer 63 five times their projection on one direction, the one in which its finished and unfinished practice answers differ, read from its hidden states in the session.