Redraft

Learning to Edit with a Small Local Model

LoRA, bounded strength search,
and what a paired test can establish

A technical companion to *Teaching a Small Model When to Leave My Writing Alone*. Seven experiments tested whether a local model could learn an editing contract without receiving the full writing skill at inference.
The latest candidate gained 47 acceptable edits on a paired 1,000-case test. Meaning errors remained at 48; 43 of the 47 net gains came from punctuation-focused variants.

Francesco Di Costanzo

10 November 2026 / Version 1.0

Original experimental evidence, with implementation inspection.
Synthetic evaluation; no independent human validation or installed daily editor.

1 Question, contract and evidence

Redraft tests whether a small model can learn a particular editing contract through examples. It receives supplied prose and optional audience or purpose context, then returns only a revision. The contract prioritises facts, attribution, uncertainty, conditions and commitments; preserves quotations, code and identifiers; removes unnecessary rhetoric; and leaves already-good prose alone. Avoiding em dashes outside protected literals is an additional owner preference. A question or instruction inside the draft is data to edit. Redraft is not a fact-checker, topic-to-document writer or personal memory system. [E1]

The implementation uses a frozen, four-bit Qwen3.5-4B base and a trained LoRA adapter. The detailed writing skill guides example authoring and output judgment, but is absent from training inputs and adapted inference. Requests do not update the weights. Synthetic demonstrations were used instead of personal correspondence. Codex authored, curated and judged the material in the same session. These records describe that process, not an independent assessment of the author's prose or a human evaluation panel.

The study took place on an Apple M3 MacBook Air with 16 GB unified memory. An external 1 TB SSD held weights and artifacts; it did not supply RAM. No Spark computation was used. The article and paper describe work completed through 10 November 2026. The new candidate remains experimental. The legacy command still selects experiment 003, and no daily application has been installed. [E2, 7]

1.1 What counts as evidence

Observed claims below come from frozen outputs, recorded judgments, aggregate audits or resource receipts. **Implemented** claims come from the inspected training, search and evaluation code, including the installed MLX implementation. General LoRA, Adam and CMA-ES references explain methods; they do not validate Redraft's results. **Proposed** next steps have not been run.

The experimental code and raw evaluation corpus are not supplied as public research artifacts with this paper. Appendix A provides a readable evidence locator, version information and artifact identities. Selected synthetic outputs are reproduced in Section 8. Readers can inspect the method and reported results, but cannot independently reproduce every score from this publication alone. The experimental reports and this paper share the same underlying evidence.

1.2 Decision rule

The owner adopted a practical target during development: at least 75% acceptable revisions and at most 5% meaning-error cases. A meaning error always makes an output unacceptable. Experiment 007 adds a paired promotion condition: no decline in total acceptance or increase in total meaning errors, with a strict improvement in at least one. It does not constrain every subgroup or error severity. Passing it supports retaining an experimental candidate, not unattended use. [E7]

2 Development history and valid comparisons

The initial 2B screen produced 24 acceptable edits with a minimal wrapper and 30 with the corrected full skill on 100 development cases. A short adapter-update probe worked, but the original quality prerequisite stopped substantive training. The owner then selected 4B and changed the question: a starting model need only be capable of learning. A concise specification produced 34 acceptable edits on the same development set, with no confirmed critical failures and one unresolved case. A 2,048-token training probe exceeded swap limits. [E2]

Run	Change	Final comparison within that run
003	300 pairs; 900 updates; rank-eight LoRA	50 cases: base 19 acceptable / 2 meaning errors; trained 37 / 4.
004	69 configurations, including 32 CMA-ES proposals	200 cases: base 57 / 19; full adapter 137 / 30; selected half strength 101 / 4, plus one unresolved.
005	Bounded search under the new 75% / 5% target	1,000 cases: selected strengths 527 / 49; half strength 479 / 19; full strength 314 / 332. Target missed.
006	Fold strengths into an adapter copy; train on 200 new + 400 replay pairs	1,000 new cases: new step 400 scored 800 / 43; searched control 703 / 33; half strength 642 / 16.
007	Restrained continuation on 200 new + 200 replay pairs	1,000 new cases: new step 100 scored 904 / 48; unchanged 006 control 857 / 48.

Table 1. Each row has its own final dataset. Counts are acceptable outputs / meaning-error cases. Meaning errors are included among failures; the two columns do not partition all outcomes. [E2–7]

Experiment 004 selected checkpoint 700 at uniform half strength. The evolutionary candidates did not beat it on selection. Experiment 005 evaluated 32 actual CMA-ES proposals; two planned branches were curtailed to preserve final-test time. Search improved the balance on that run's paired test but did not reach 75% acceptance. More combinations were neither guaranteed to work nor proven useless.

Experiment 006 completed 1,600 updates at learning rate 2×10^{-5}, yet selected the first checkpoint, at 400 updates. It achieved the first observed pass of 75% / 5%, while increasing meaning errors and incomplete outputs relative to the searched control. Experiment 007 reduced the learning rate and again selected an early checkpoint. These are observations from particular runs, not a universal claim that less training is better.

Do not plot 52.7%, 80.0% and 90.4% as a learning curve on one benchmark. The same experiment-006 adapter scored 80.0% on its own final set and 85.7% on experiment 007's final set. Paired comparisons within a row isolate the candidate difference for those drafts; cross-row subtraction does not.

3 The adapter and its mathematics

LoRA adds a low-rank update to selected frozen linear calculations [1]. Redraft uses MLX-LM's `LoRALinear`, with row-vector input convention [2]. For a batch of internal activations x, let

$$x \in \mathbb{R}^{n \times d_{\text{in}}}, \qquad W \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}},$$
$$A \in \mathbb{R}^{d_{\text{in}} \times r}, \qquad B \in \mathbb{R}^{r \times d_{\text{out}}}.$$

With zero dropout in these runs, the mathematical calculation is

$$y = xW^{\mathsf{T}} + b + s(xA)B. \tag{1}$$

W and any base bias b stay frozen. Only A and B learn. The implementation casts the scaled update to the input dtype before addition; the equation abstracts those finite-precision operations. The base calculation uses four-bit affine quantisation, group size 64. This is not a claim of NF4, double quantisation or a paged optimiser. [E3, 8]

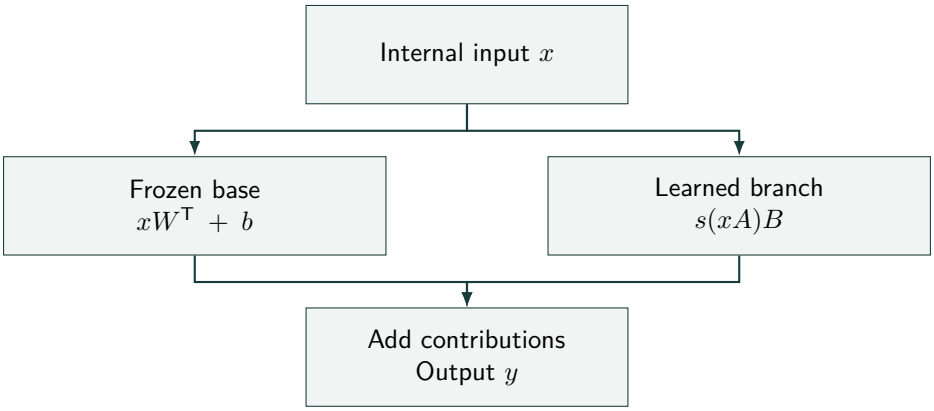


Figure 1. One adapted projection. Both branches run at inference; gradients update only the learned branch during training. This is not an ensemble of two models.

Equivalently, $W' = W + sB^{\mathsf{T}}A^{\mathsf{T}}$, and $\text{rank}(W' - W) \leq r$. Each projection adds $r(d_{\text{in}} + d_{\text{out}})$ parameters. Here $r = 8$. The final eight of 32 layers each adapt the MLP gate, up and down projections. Gate/up map 2,560 to 9,216 dimensions; down reverses them. Therefore

$$8 \text{ layers} \times 3 \text{ projections} \times 8(2560 + 9216) = 2{,}260{,}992.$$

There are 48 trainable A/B tensors. The installed implementation uses the configured `scale` directly: $s = 16$, not $16/r$. In conventional α/r notation, matching this multiplier would require $\alpha = 128$. A starts random and B starts at zero in a newly created adapter, initially adding nothing. Continuations instead start from learned tensors. No new base-model fusion or daily runtime was installed. [E8]

4 Teaching objective and memory constraints

Training compares predictions with a curated revision. At each position it supplies the known preceding reference tokens, a method called teacher forcing. Prompt tokens remain available as context, but only completion tokens, including the end-of-turn token, contribute to the objective. For mask m_t and adapter parameters θ,

$$\mathcal{L}(\theta) = -\frac{1}{\sum_t m_t} \sum_t m_t \log p_\theta(y_t \mid d, c, y_{<t}), \tag{2}$$

where d is the draft, c optional context and y the reference revision. A lower loss means the reference continuations are more likely under this conditioning. It is not a direct score for preserving meaning in a freely generated edit. Editorial judgments select checkpoints separately. No reward-model training, RLHF or DPO was used. [E3, 7, 9]

Experiment 007 starts from a copy of experiment 006 step 400, with fresh Adam state [4]. It uses learning rate 5×10^{-6}, microbatch one, global gradient-norm clipping at 1.0, seed 707 and gradient checkpointing. For an unclipped gradient g, clipping supplies

$$\widetilde{g} = g \min\{1, 1/(\|g\|_2 + 10^{-6})\}$$

to the optimiser. The 10^{-6} stabiliser matches the installed MLX function. This bounds the gradient norm, not the semantic size of a revision. All 600 updates completed. Checkpoints were saved at 100, 200, 400 and 600; two checks without at least 0.001 validation-loss improvement would have stopped training. That condition did not occur.

Checkpoint	Diagnostic loss	Acceptable / 200	Meaning errors
Unchanged control	0.076682	178	5
100 updates	0.061890	192	3
200 updates	0.060805	192	3
400 updates	0.055758	191	4
600 updates	0.054432	189	7

Table 2. Loss uses historical diagnostic validation; generated-edit selection uses a separate 200-case set. Step 100 won the stable-ID tie-break before any final outputs existed.

4.1 Reducing allocation without removing the prompt

The inspected objective computes the transformer over the full input, then applies the vocabulary projection only to positions with scored completion targets, in checkpointed chunks of 16. Checkpointing recomputes intermediate work during differentiation rather than retaining all of it. The vocabulary contains 248,320 entries, making the avoided unscored projections material. Prompt context and the end-of-turn target remain present.

Original and revised loss, and all 48 adapter-gradient tensors, matched exactly on the equivalence fixture. This is a measured case, not a proof for every floating-point input. An initial attempt passed integer targets as differentiable checkpoint arguments and failed; the correction captured them as constants. A 256-token probe still exceeded the allocation cap. The accepted ceiling was 192 tokens for the *entire* training sequence, with actual sequences of 60–188 tokens. It did not provide long-document training. [E7, 9]

5 Searching influence, then learning again

Training adjusts the A/B entries. The earlier strength search instead holds those learned entries fixed and chooses eight numbers $g_\ell \in [0, 1]$, one for each adapted layer, shared by its gate/up/down projections. The contribution in layer ℓ becomes

$$\Delta y_\ell = 16 g_\ell (x A_\ell) B_\ell. \tag{3}$$

A searched candidate is a checkpoint *and* its eight strengths. Loading the checkpoint alone at its original scale does not reproduce the searched candidate. There is no independently interpretable "humour dial" or "truth dial" in these eight values. [E4, 5, 10]

The search used pycma 4.4.0's CMA-ES. Its sampling can be described schematically as

$$u_k = m + \sigma C^{1/2} z_k, \quad z_k \sim \mathcal{N}(0, I), \qquad g_k = T_{[0,1]}(u_k).$$

m locates the current search distribution, C describes correlated directions, and σ its overall scale. Candidate rankings guide weighted recombination and updates of that distribution. T is pycma's `BoundTransform`, a smooth piecewise linear/quadratic boundary transform, not simple componentwise clipping. This expression explains sampling; it is not a reimplementation of all covariance and step-size updates [5, 6].

Experiment 005 used population eight, initial $\sigma = 0.15$, at most four generations per branch, and a fixed constrained editorial ordering. Candidates within the 5% meaning-error limit were ranked principally by acceptable count; over-limit candidates were ranked principally by meaning errors. The optimiser received ordinal ranks, not token loss. Selection and final evaluation were separate from the search feedback. The executed search made 32 proposals; time reserved for final comparison curtailed two planned branches. It was neither exhaustive nor a proof that sufficient search would reach the target. [E5, 10]

5.1 Encoding the searched seed

Experiment 006 prepared an adapter copy by multiplying each layer's B tensor by its chosen strength:

$$B'_\ell = g_\ell B_\ell, \qquad 16(x A_\ell) B'_\ell = 16 g_\ell (x A_\ell) B_\ell.$$

The runtime scale then returns to 16. This is mathematically the same update, subject to finite-precision rounding. Four greedy outputs matched before and after encoding, and on reload. The base remained frozen. This operation differs from fusing the adapter into the base weights: only B was rescaled, and subsequent training updated A and B. Experiment 007 continued from the resulting trained experiment-006 candidate. [E6, 10]

Search can rebalance changes already represented in a learned adapter. Further supervised training can learn different changes from new examples. Their different search spaces explain the engineering choice, but the experiments do not provide a controlled comparison of all possible training and search budgets.

6 Data construction, review and uncertainty

Experiment 007 used four distinct data roles. Related scenario variants stayed together within a split. The material was frozen before substantive training; candidate selection was frozen before final generation. [E7]

Role	Cases	Composition and interpretation
Training	400	200 new pairs from 50 families plus 200 replay pairs; 122 unchanged references.
Loss validation	80	Historical, already exposed; diagnostic early stopping only.
Selection	200	50 new families, separate from training and final; chooses among four checkpoints and the control.
Final	1,000	200 other families, each with clean, headed, email, editable-dash and optional-context variants.

The final set contains 224 unchanged references, 245 protected-literal cases, 200 editable-dash variants and 200 optional-context variants. One hundred drafts exceed 200 words, drawn from just 20 families; the maximum is 248 words. Actual teaching drafts are 7–52 words, or 8–63 tokens. Longer evaluation is extrapolation beyond the teaching lengths.

The overlap screen compared fresh cores and references against 3,781 historical source drafts and other splits. It rejected exact matches and pairs jointly meeting lexical Jaccard ≥ 0.5 and character SequenceMatcher ≥ 0.7. Passing that screen does not establish independence of editorial patterns. Many short drafts are already clear; an unchanged output can pass even when a reference offers optional concision.

6.1 Judgment and paired estimands

An exact curated-target match inherits the curation judgment, ignoring only trailing line feeds. Other unique combinations of source, context, output and finish status receive editorial review. Across selection and final, 3,000 outputs resolve to 1,438 content-bound judgments: 435 exact targets and 1,003 editorial reviews, with none unresolved. Review views hide names but do not guarantee blinding. Codex was both author and judge.

Let A_i indicate acceptance and S_i a meaning error. $A_i S_i = 0$. Failure categories can overlap. The reported rates and paired difference are

$$\widehat{a} = \frac{1}{n}\sum_i A_i, \qquad \widehat{s} = \frac{1}{n}\sum_i S_i, \qquad \widehat{\Delta} = \frac{1}{n}\sum_i (A_i^{\text{new}} - A_i^{\text{control}}).$$

Promotion requires $\widehat{a}_{\text{new}} \geq .75$, $\widehat{s}_{\text{new}} \leq .05$, no unresolved cases, $\widehat{a}_{\text{new}} \geq \widehat{a}_{\text{control}}$ and $\widehat{s}_{\text{new}} \leq \widehat{s}_{\text{control}}$, with at least one strict paired inequality.

The bootstrap resamples whole families with replacement, using 2,000 draws: seed 50595 for rates and 60695 for paired acceptance. For equal five-case families, the paired statistic is the average within-family difference. Reported bounds use ordered draws 50 and 1,950. This respects within-family clustering but not every shared template or judge bias. Wilson intervals in the underlying aggregate assume independent cases and are not the preferred summary here. Greedy temperature-zero decoding was fixed; these intervals concern sampled cases, not random decoding. [E7, 11]

7 Final paired results and the promotion limit

Both candidates used identical prompt hashes, token counts, inference flags and output allowances on all 1,000 final cases. There was no full-skill injection, repair pass or hidden retry to improve an edit. Four selected outputs matched after a fresh-process reload. [E7]

Measure	006 control	007 step 100
Acceptable	857	904
Meaning-error cases	48	48
Material / subtle	9 / 39	12 / 36
Incomplete	25	10
Protected-literal failures	5	13
Actual editable-dash misses	51	2
Exactly unchanged clean references / 224	203	200

Table 3. Failure categories overlap. The new candidate gains 78 acceptable cases and loses 31, net +47 or +4.7 percentage points. It removes 18 meaning errors and introduces 18 different ones. Equal totals do not imply unchanged failures.

The family-bootstrap 95% interval for the paired acceptance gain is +2.6 to +6.9 points. For the new candidate's individual rates, intervals are 87.6–93.1% acceptance and 2.9–7.0% meaning errors. The observed 4.8% error rate passes 5%; its interval does not establish that the underlying future rate is below 5%.

Overlapping slice	Cases	Control acceptable	New acceptable
Editable-dash variants	200	142	185
All other variants	800	715	719
Puffery-tagged families	45	25	25
Headed documents	200	178	172
Drafts over 200 words	100	71	80

Table 4. These slices overlap and are not additional independent tests. Forty-three of the 47 net gains come from editable-dash variants. The other 800 cases gain four, with a paired interval of −1.625 to +2.75 points. That result does not establish improvement outside the punctuation-focused group.

Headed-document meaning errors rise from ten to 16. Both candidates make 17 meaning errors in the longer subset; eight of the new candidate's incomplete longer outputs repeat a title and omit the body. Protected-literal failures and material errors also rise. The frozen promotion rule passes because it constrains total counts, not these regressions. The retained candidate is therefore not uniformly safer.

The frozen dash checker reports 52 control flags and three new flags. Each includes one corrupted protected quotation whose punctuation no longer matches its excluded literal. Actual editable-dash misses are 51 and two. The supplementary audit preserves this distinction without changing primary judgments: those corrupted quotations remain literal failures either way. [E7]

8 Recorded synthetic examples

These are exact source/output pairs from the final selected adapter, aligned by case identity and content-bound judgment. They are not proposed reference answers. Names and situations in the dataset are synthetic; no personal correspondence is reproduced. Line breaks below preserve the meaningful paragraph structure. [E12]

8.1 Accepted: uncertainty retained

Source, no optional context

> A borrowed metronome — The metronome may be available after the lesson, but I have not asked its owner or confirmed availability.

Actual output

> A borrowed metronome. The metronome may be available after the lesson, but I have not asked its owner or confirmed availability.

The punctuation changes while possible availability and the absence of owner confirmation remain. This demonstrates one accepted case, not general preservation reliability.

8.2 Meaning failure: a useful heading disappears

Source, no optional context

> The faded arrow
>
> The arrow points towards the courtyard. It does not say which door to use. I have not checked whether both doors open.

Actual output

> The arrow points towards the courtyard. It does not say which door to use. I have not checked whether both doors open.

Removing the heading loses the supplied fact that the arrow is faded. The recorded judgment is a subtle meaning failure despite otherwise faithful body text.

8.3 Literal failure: the exception to the punctuation rule

Source

> The card beneath the vase reads “Borrowed—not for sale”. The vase itself has no writing on it.

Actual output

> The card beneath the vase reads “Borrowed. Not for sale”. The vase itself has no writing on it.

Optional purpose was “Make this supplied text clear while keeping its meaning.” Audience was “The person already involved in this situation.” The output changes punctuation and capitalisation inside a protected quotation. It fails the literal boundary even though the wording remains intelligible. The article’s policy against editable dashes must not be applied to the quoted source when presenting this evidence.

9 Resources, operational status and next test

Measurements refer to experiment 007, not the entire seven-run development effort. They include failed model probes where specified. Authoring and editorial-review wall time were not separately instrumented, so these are not total labour-cost measurements. [E7]

Measurement	Observed
All model-process time, including failed probes	8,045.89 s (2 h 14 m 6 s)
Training loop	1,499.35 s (about 25 min)
Final paired comparison	4,995.10 s
Training MLX peak allocation	7.018 GiB
Final-process MLX high-water allocation	3.162 GiB
Maximum sampled process-tree RSS across stages	3.073 GiB
Added swap: training / selection / final	Zero / zero / zero
Rejected 256-token probe	8.537 GiB MLX; +982.06 MB swap
New-adapter generation median / p95	1.477 s / 11.020 s

Table 5. MLX and RSS counters overlap and must not be added. The final-process high-water allocation is shared across that process, not a separate observation for each adapter. Generation timings describe the tested mixture, not cold-start latency in a finished app.

The accepted run used six disposable gradient attempts, 600 substantive updates and 3,032 generation attempts: 32 diagnostics/reload checks, 1,000 selection outputs and 2,000 final outputs. Budgets were 12 hours of model-process execution and 3,200 generations, with 2,000 final generations reserved. The rejected probe exceeded the separate 8 GiB allocation cap even though its added swap stayed just below the 1,024 MB watchdog threshold. No new hardware, cloud GPU or dependencies were added for experiment 007.

The selected adapter, prior winner, base and frozen input identities were checked in the final audit. Base weights and the previous winner remained unchanged. Reload checks support reproduction of four chosen outputs; they are not a complete cross-hardware determinism guarantee. The selected adapter has not been fused into a new base or installed as a daily editor.

9.1 What remains unestablished

There is no independent human quality estimate, general long-document guarantee, or paired comparison against a hosted model supplied with the full skill. The study does not establish financial savings or reduced owner review time. The short synthetic training distribution, related final variants and same-session author/judge restrict generalisation. A model may improve a visible stylistic preference while becoming less faithful in a subgroup.

A next test should freeze fresh scenario families, add independent review, and measure correction effort alongside acceptance and meaning errors. Heading/body preservation and exact quotation handling deserve explicit gates. The observed 75% / 5% target can remain an experimental criterion, but its promotion rule should be reconsidered if severity or particular document types matter more than total counts. These are proposals; no further model job was run to prepare this publication.

A Evidence locator and artifact identity

The paper was checked against the author's Redraft development snapshot `aa962eec0cad261a7a62d8f58d263bedf4afd1ea`. The repository was not anonymously accessible during publication checks. The locators below describe the inspected records without requiring private access or exposing local storage paths. Aggregate evidence is reported in this paper; raw outputs, the full datasets and weights are not distributed with it.

ID	Record or implementation	Supports
E1	Editing contract v1, 8 November	Preservation, literal boundaries, supplied-prose scope.
E2	Development history; reports 001–002	Initial gate, 2B/4B move, hardware and failed long-sequence probe.
E3	Experiment 003 report	First supervised training, LoRA scope, 50-case comparison.
E4	Experiment 004 report	69 configurations, selected half strength, 200-case results.
E5	Experiment 005 report and aggregates	Executed search, constrained ranking, failed 75% target.
E6	Experiment 006 report and execution	Seed folding, continuation, fresh paired results.
E7	Experiment 007 plan, execution, report, results, final and supplementary audits	Frozen procedure, selected identity, all final counts and subgroup/resource caveats.
E8	Installed MLX-LM `LoRALinear`; converted model configuration	Shapes, initialisation, direct scale, quantisation and parameter count.
E9	`exp007_loss.py`, `exp007_probe.py`, `exp007_train.py`, `pilot_common.py`	Completion objective, equivalence test, memory optimisation and training.
E10	Search/common/load code, `exp005_evolve.py`, `exp006_seed.py`; installed pycma defaults	Eight strengths, bounded CMA-ES and B-tensor folding.
E11	Experiment 005–007 report code	Review aggregation, family bootstrap and paired promotion inequalities.
E12	Frozen final inputs, new-step-0100 outputs and content-bound judgments	The three exact synthetic examples in Section 8.

Installed packages inspected: MLX 0.32.4, MLX-LM 0.32.2; strength search used pycma 4.4.0. The base derives from the official Qwen3.5-4B artifact [3]. Those installed versions govern the implementation claims; upstream pages linked below may change.

B Checksums and public references

Selected adapter SHA-256

`67a9a86908988510586da6b0def2b3033eeb1b4b93d0be436e188afb7b63e0d2`

Frozen dataset identity SHA-256

`06c099818dd80387b0c9f779c1a577c97471d87e9fa2b0ad2a960de9e23f5382`

Frozen experiment identity SHA-256

`d40d40245c88dd9e1a79f2c8ecc375619f1082fa560121e5d04a649f2234a1c6`

These identify retained artifacts; hashes alone do not make unavailable data reproducible. The three quoted examples correspond to final cases 021-v4, 039-v2 and 031-v5 respectively. Their content-bound judgment SHA-256 values are:

Accepted metronome: `ec15f0b6f48416ea5b3b1225708b479f4e8f63c92404dc84a09e918ad7f3746b`

Faded-arrow failure: `bc968ea25fbf011a498ef65fcc104a290b2b5a3b04c8d0e6203571dc06bc6a59`

Quotation failure: `1c612d9978c6d3417bd82ce0d89a3e3095677428f7d1a8f96d4f858bbce5ab18`

References

[1] E. J. Hu et al. *LoRA: Low-Rank Adaptation of Large Language Models*. 2021. https://arxiv.org/abs/2106.09685

[2] MLX-LM contributors. *LoRALinear implementation*. Official upstream source; installed version inspected separately. https://github.com/ml-explore/mlx-lm/blob/main/mlx_lm/tuner/lora.py

[3] Qwen. *Qwen3.5-4B*. Official model card. https://huggingface.co/Qwen/Qwen3.5-4B

[4] D. P. Kingma and J. Ba. *Adam: A Method for Stochastic Optimization*. 2014, revised 2017. https://arxiv.org/abs/1412.6980

[5] N. Hansen. *The CMA Evolution Strategy: A Tutorial*. 2016, revised 2023. https://arxiv.org/abs/1604.00772

[6] pycma contributors. *BoundTransform*. Official API documentation. https://cma-es.github.io/apidocs-pycma/cma.boundary_handler.BoundTransform.html

[7] R. Takehi, J. Teoh, A. Zhang and O. Khattab. *Sparks of Weight-Space Programming: LLMs can natively read & write neural network weights*. 2 November 2026. Inspiration for the project, not the method implemented or validation of its results. https://rikiyat.github.io/blog/weight-space-programming/

Companion article and inline figures: https://fdicostanzo.com/posts/2026/20261110_teaching_a_small_model_to_edit. This paper reports the author's evidence and separates it from the public method references above.