Institute of Foundation Models
Mohamed bin Zayed University of Artificial Intelligence

Latent Feedback Is a Compute and Data Amplifier

Scaling Law for Latent Feedback Transformers on Repeated Data

Benhao Huang[‡*], **Chaitanya Dwivedi**[‡], **Xuezhe Ma**[‡†], **Alexander Moreno**[‡†], **Hector Liu**[‡]
[‡]Institute of Foundation Models, [†]USC, [*]CMU

Code: https://github.com/ifm-ai/LoRM

Date: November 11, 2026

arXiv:2611.20404v1 [cs.LG] 16 Nov 2026

Overview

When high-quality data runs out, a common practice is to train on it for several epochs, but each repeat is worth less than the last. Latent feedback, which feeds each token's top-layer state into the next token's input, is another way to spend that compute, yet it has not been tested against a vanilla Transformer that spends the same compute on more epochs. We run this test across a ladder of model sizes, unique data and epochs, tuning both models at every size and amount of unique data. Latent feedback makes every FLOP and every token count for more except when unique data is scarce relative to model size, a regime that today's models are far from. We find that latent feedback serves as an amplifier: with unique data fixed, it reaches a Transformer's loss with less compute; with epochs fixed, it reaches that loss with less unique data. Furthermore, both savings grow with unique data per parameter and with repetition, and downstream tasks agree. To go beyond our grid, we fit one scaling law to both models, in which latent feedback enters as a Transformer with more parameters; it predicts that both savings keep growing with unique data and repetition. Fitted on models from 60M to 271M parameters, the law predicts the losses at a held-out 778M size to within 0.3% and the gap between the two models to within 0.05%. At most, latent feedback saves 48% of the compute and 47% of the unique data. When data, not compute, is the limit, latent feedback is the better use of a FLOP for a model of a given size.

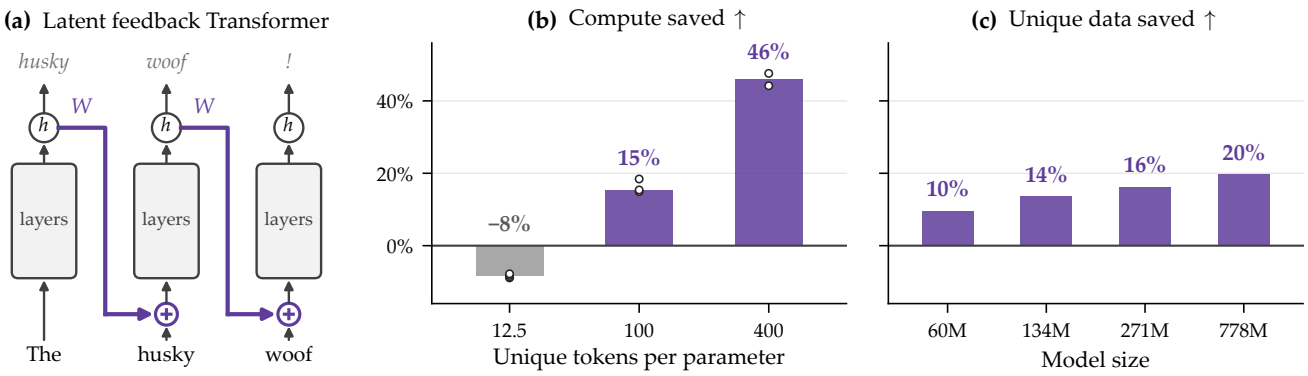

Figure 1: Latent feedback saves compute and unique data. (a) Latent feedback: each token's top-layer state h enters the next token's input through a map W. (b) The share of compute that latent feedback saves at equal unique data, at a budget of two Transformer epochs; each bar is the median over model sizes, and each dot is one model size. (c) The share of unique data it saves when both models train one epoch on 50 unique tokens per parameter.

1 Introduction

High-quality text, not compute, now limits pretraining of large language models [1, 2]. Even web text, mostly **not** high quality, grows far more slowly than pretraining compute [3, 4]. So a common practice is to train on the best sources for several epochs [5–10]. However, each repeat is worth less than the last, so additional epochs bring diminishing returns [7, 10–12].

Latent feedback offers another use for that compute. A latent feedback Transformer feeds each token's top-layer state into the next token's input. Prior work reports improved performance for extra training FLOPs [13, 14]. Yet these comparisons fix the training tokens, and none repeats data. With repeated data, the fair baseline is a Transformer that spends the same FLOPs on more epochs. Whether latent feedback's gain survives at scale, against a vanilla Transformer given the same FLOPs and unique data, is still largely unknown. We therefore ask:

When data runs out, can latent feedback make every FLOP and every token count for more?

We answer it on a ladder of models from 60M to 778M parameters, each trained on 12.5 to 400 unique tokens per non-embedding parameter (UTPP) for one to four epochs, that is up to 1,600 training tokens per parameter with repetitions (TPP). Both models are tuned at every size and amount of unique data, and the pools of unique tokens are nested and deduplicated. The answer is yes, once each parameter has enough unique data, and today's models have far more than enough. We measure the answer at a fixed model size in two ways, and then fit a law to extrapolate it.

- **Compute.** With unique data fixed, the Transformer must train more epochs to match latent feedback; it needs κ times the FLOPs (Figure 2 (a)). The saving grows with UTPP and with epochs: at two epochs, latent feedback costs 8% more at 12.5 UTPP, but saves 15% to 18% at 100 and 44% to 48% at 400. Above about 75 UTPP it saves compute at every budget we train, and at two epochs already above 33. Practice is far past this point: Chinchilla used 20 TPP [6], whereas today the median 1B–10B dense model trains on 1,040, at least 260 UTPP even over four epochs [15–17].[1]
- **Data.** With epochs fixed, the Transformer needs μ times the unique data to match latent feedback with the same parameters and key-value (KV) cache (Figure 3 (a)). At one epoch on 50 UTPP, latent feedback saves 16% of the unique data at 271M and 20% at 778M; with three epochs on 100 UTPP, it saves 41% at 271M.
- **A fitted law for extrapolation.** To go beyond our grid, we fit one scaling law to both models, with latent feedback entering as a Transformer with v times the parameters (Section 5.3). Fitted on 60M to 271M, it predicts the held-out 778M losses to within 0.3% and the gap between the models to within 0.05%. It then draws the boundary between the two regimes, a FLOP spent on another epoch or on latent feedback, and predicts that both savings grow with unique data and repetition beyond the grid.

2 Background

Latent Feedback Transformers feed each position's top-layer state into the input of the next position [13, 14]. At decoding time, position t enters the blocks f with input u_t and leaves as h_t:

$$u_t = e_t + W\ \mathrm{rms}(h_{t-1}). \tag{1}$$

[1] Median over the 1B–10B dense models among 257 base models whose reports state both their parameter and training token counts; Appendix B gives the full statistics.

However, teacher-forced training cannot run the positions one by one to get top-layer states, so it runs passes over the whole sequence as fixed-point iterations:

$$\left[\, u_t^{(p+1)} = e_t + W\ \mathrm{rms}\left(h_{t-1}^{(p)}\right), \quad h_{t-1}^{(p)} = f\left(u_{t-1}^{(p)}\right) \,\right] \xrightarrow{\;p\to\infty\;} u_t^{(*)} = e_t + W\ \mathrm{rms}\left(h_{t-1}^{(*)}\right), \tag{2}$$

The first pass reads no feedback, $u_t^{(1)} = e_t$, and pass $p = 2, 3, \ldots$ reads the previous pass's top states, ideally till fixed points $h_{t-1}^{(*)} = h_{t-1}$. In practice, two to three passes come close to the fixed point [14].

At inference, autoregressive decoding already has the previous token's top state when it predicts the current token, so each generated token costs what a Transformer's does, though prefilling a prompt needs additional passes of Equation 2. The KV cache is a Transformer's plus one vector.

Training. Each update runs one pass with probability p and two passes otherwise (Algorithm 1). The loss and gradient are computed only on the last pass. The map W is drawn from a normal distribution with standard deviation 0.02. Because of the additional forward passes, a token with feedback costs $c > 1$ times a Transformer's training FLOPs. So p trades that cost against the feedback each update sees.

Algorithm 1 One training update

```
x = embed(tokens)
if rand() > p:
    with no_grad():
        s = blocks(x)
    # token t reads s[t-1]
    x = x + W(rms_norm(shift(s)))
loss = ce(head(blocks(x)), y)
loss.backward()
```

3 Scaling law

Budget. A run has N non-embedding parameters [18–20], U unique tokens and a compute budget F; models are named by their total parameters. We count F in Transformer epochs, its FLOPs over those of one Transformer epoch on U, so a Transformer trains $R = F$ epochs. Latent feedback with a training schedule s that costs $c_s > 1$ Transformer FLOPs per token trains only $R = F/c_s$ epochs. So at equal compute, latent feedback sees fewer repeats than the Transformer. Appendix A lists the notation.

Law. We extend the data-constrained law of Muennighoff et al. [12]:

$$L_s(N, U, F) = E + \left[\left(\frac{A}{(v_s N)^\alpha}\right)^q + \left(\frac{B}{D_s^\beta}\right)^q\right]^{1/q}, \qquad D_s = U\left[1 + R^*\left(1 - e^{-(R-1)/R^*}\right)\right], \qquad R = \frac{F}{c_s}. \tag{3}$$

where E is the irreducible loss, $A/(v_s N)^\alpha$ is the capacity term, and B/D_s^β is the data term. The factor v_s counts the effective parameters of training schedule s, with $v = 1$ for the Transformer. D_s is the effective number of training tokens, where R^* sets how fast repeats lose value, and q couples the two terms.[2]

Repeats. Epoch k is worth about $e^{-(k-1)/R^*}$ of a fresh one, so the effective data D_s saturates at $U(1 + R^*)$. Below one epoch we set $D_s = RU$. The two pieces join smoothly: near $R = 1$, $1 - e^{-(R-1)/R^*} \approx (R-1)/R^*$, so both give $D_s \approx RU$ with the same slope. Following Qin et al. [22], we let $R^* \propto (U/N)^\rho$.

Capacity. A schedule reruns the same blocks on a token, as a recurrent-depth model reruns its layers [23, 24]. Scaling laws for such models count the extra passes as effective parameters [25, 26]. We follow their

[2] Table 9 in Appendix E compares three ways to combine the two terms. Simply adding them ($q = 1$) leaves an error about three times larger. Our fit gives $q \approx 3.7$, so the larger term dominates the loss: whichever of model size and data is scarcer limits the run. This form is the law of Videau et al. [21] with $k = 1/q$, written so that α and β stay each term's own exponent. Tying q to $1/\beta$, as Kaplan et al. [18] do, leaves an error about six times larger.

practice and introduce v_s that multiplies the parameter count N,

$$v_s = v_0 \Big(\frac{D_s}{100\,N}\Big)^{\omega} \Big(\frac{N}{10^8}\Big)^{\gamma}, \tag{4}$$

where D_s is the effective data of Equation 3 and v_0, ω and γ are fitted.[3] The constants 100 and 10^8 serve as a reference point: v_0 is the factor of a model with 10^8 parameters trained on 100 effective tokens per parameter, and ω and γ give how it grows with data per parameter and with size from there.

The law predicts a crossover. To first order in $\ln v_s$ and $\ln c_s$, a schedule's gap to the Transformer at equal F is (Derivations in Appendix F)

$$L_s - L_T \approx S\Big[-\,w\,\alpha \ln v_s \;+\; (1-w)\,\bar{\varepsilon}\,\beta \ln c_s\Big], \qquad w = \frac{X^q}{X^q + Y^q}, \qquad \bar{\varepsilon} = \frac{\ln D_T - \ln D_s}{\ln c_s} \in (0,1], \tag{5}$$

where L_T is the Transformer's loss (Equation 3 with $v = c = 1$); $X = A/N^{\alpha}$ and $Y = B/D_T^{\beta}$ are its two terms, $S = (X^q + Y^q)^{1/q}$, and D_T and D_s are the effective data of the Transformer and of the schedule.

The first term is the extra capacity, weighted by w, which nears 1 when the capacity term dominates. The second term, weighted by $1 - w$, is the data that latent feedback gives up. At equal F it trains c_s times fewer epochs, so its effective data D_s is smaller than the Transformer's D_T. How much smaller depends on what the dropped epochs were worth, which $\bar{\varepsilon}$ measures: $\bar{\varepsilon} = 1$ when each dropped epoch was worth a fresh one, so $D_s = D_T/c_s$, and $\bar{\varepsilon} \to 0$ when repeats are worth nothing, so dropping them costs no effective data. With U fixed, D grows with F (Equation 3), so w rises toward 1 and repeats saturate ($\bar{\varepsilon} \to 0$). As the budget grows, the capacity term thus counts for more and the data term for less. So feedback that costs more compute per token, may lose at small budgets but win past a crossover budget F^*.

Two multipliers express latent feedback's advantage as the compute or the unique data it saves. Both compare the two models at the same size N and the same validation loss. The *compute multiplier* κ is the Transformer's budget over latent feedback's, at equal unique data U, so latent feedback saves a share $1 - 1/\kappa$ of the compute. The *data multiplier* μ is the Transformer's unique data over latent feedback's, at equal epochs R, so latent feedback saves a share $1 - 1/\mu$ of the data:

$$\kappa = \frac{F_T}{F} \text{ with } L_T(N, U, F_T) = L_s(N, U, F), \qquad \mu = \frac{U_T}{U} \text{ with } L_T(N, U_T, R) = L_s(N, U, c_s R). \tag{6}$$

Fitting. We fit all cells of both models jointly by least squares. Latent feedback differs from the Transformer only through its capacity factor v_s, with three fitted coefficients (v_0, ω and γ), and its measured cost c_s; all other coefficients are shared.[4] To separate the effect of the size and data terms, our cells vary training tokens per parameter from 12.5 to 1,600 [28].

4 Setup

Models. Both models, the Transformer and the one with latent feedback, share one backbone at 60M, 134M, 271M and 778M parameters. The backbone is a modern decoder: pre-norm RMSNorm [29], rotary positions [30], QK-norm [31] and SwiGLU [32], with untied embeddings and OLMo 2's initializa-

[3] Appendix E compares this form with a constant v and with a v that grows only with data per parameter; this form fits best and predicts the held-out 778M best (Table 9).

[4] A two-phase fit, which first fits E, A, α, B, β on one-epoch Transformer runs [12, 27], fits worse than our joint fit in both the additive and the coupled form (Table 9 in Appendix E).

tion [33]. Its matrices train with NorMuon [34], a normalized variant of Muon [35], and its embeddings and output head with AdamW [36]; see Table 10 in Appendix G for details. Latent feedback uses $p = 0.15$ throughout[5], so its cost c is about 1.20 to 1.27 across the four sizes.

Data. We take the Dolma 3 sample [9] and process it ourselves: we deduplicate it and remove the documents that share a paragraph with Paloma [37]. This keeps Paloma a clean held-out benchmark for a fair evaluation [38, 39]. We order its documents so that every prefix has the whole sample's mixture of sources and topics with same ratios, and each run's pool is the first U tokens of this order (Appendix I). Pools are thus nested and differ only in size, from $U = 12.5N$ to $400N$.

Tuning. We call one setting of model size N, unique data U, and budget F a cell. For each model, size and amount of unique data, we tune a main run over the learning rate η and the timescale $\tau = 1/(\eta\lambda T)$ [40, 41], where λ is the weight-decay and T the number of steps; other F branch off the main run as decay runs and keep its hyperparameters. Each grid is centered on a predicted start point from earlier runs and extended wherever the best point sits on its edge; a cell's value is the minimum of a quadratic fitted to its points. Where unique data is small, we also shrink the batch to stay near its optimum [41]. Each main run uses a warmup-stable-decay schedule [42] (Appendix G).

Evaluation. We measure on a held-out validation set of 47,274 documents (100M tokens), taken from the front of the same stratified order of the deduplicated Dolma 3 sample. It keeps the mixture ratio and excludes training documents. For downstream tasks, we score the Olmo 3 Base Easy suite [9], 109 tasks in 37 suites, as bits per byte (BPB) of the correct continuation on OLMES requests [43]. Each suite is weighted by the inverse of its seed variance, so noisy suites do not dominate (Appendix H).

Decoding latent feedback token by token (Equation 1) is slow. We score it instead with 4 parallel passes of Equation 2, whose states $h^{(4)}$ nearly match the decoding states, so more passes leave the loss almost unchanged (Table 5 in Appendix C).

5 Experiment and result

This section reads the grid of runs from Section 4 on validation loss in three steps, and then checks it on downstream tasks. First, with unique data fixed, we measure the compute that latent feedback saves (κ, Section 5.1). Second, with the epochs fixed, we measure the unique data that it saves (μ, Section 5.2). Third, we fit the law of Section 3 (Equation 3) jointly to both models, test it on the held-out 778M, and use it to explain the two savings and to predict them beyond the grid (Section 5.3). Finally, we check that the comparison is robust: repeats do not overfit, gaps are stable across seeds, and downstream tasks agree with the validation loss (Section 5.4).

5.1 Equal unique data: the compute multiplier

At equal size and unique data, latent feedback reaches a Transformer's loss with less compute (Figure 2 (a)). Both models train the same model size N on the same U unique tokens, and only the budget F varies: the Transformer trains F epochs and latent feedback, about 1.2 times as costly per token, trains F/c_s. Latent feedback saves a share $1 - 1/\kappa$ of the compute (Equation 6).

The saving grows with UTPP and with epochs (Figure 2). Along UTPP, latent feedback costs more compute than the Transformer when unique data is scarce and saves compute when it is plentiful. The unique data at which both models need the same compute ($\kappa = 1$) falls as the budget grows: about 75 UTPP at $F = 1$ and 33 at $F = 2$. Above ≈75 UTPP, latent feedback saves compute at every budget we

[5] We find that its loss at equal compute barely moves with p when around 0.15 (Appendix G).

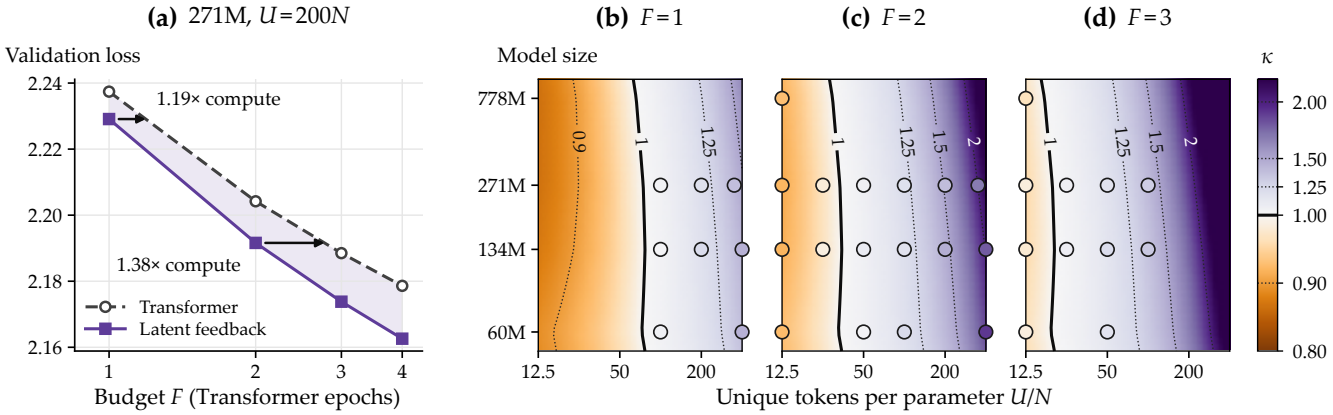

Figure 2: The compute multiplier κ across unique data and model size. (a) 271M with $U = 200N$: each arrow runs from a latent feedback cell to the budget at which the Transformer reaches its loss, labelled with κ. (b)–(d) Dots are measured cells, filled with the κ read off the Transformer's losses at $F = 1$ to 4 interpolated in $\log F$; everywhere else the map shows the law of Section 5.3. κ grows with unique tokens per parameter and with the budget F, and its contours run nearly vertically: it depends on U/N far more than on size.

train. Along epochs, the saving grows with the budget F. At 400 UTPP and $F = 2$, latent feedback saves nearly half the compute ($\kappa = 1.91$ at 60M); Table 7 in Appendix D lists every measured κ.

5.2 Equal epochs: the data multiplier

At equal size and epochs, latent feedback reaches a Transformer's loss with up to half the unique data (Figure 3 (a)). With size and epochs equal, both serve with the same parameters and KV cache; latent feedback pays only about 1.2 times the training FLOPs per token, and saves $1 - 1/\mu$ of the unique data. At 271M and two epochs, for example, $\mu = 1.34$ at $U = 50N$, so latent feedback needs 25% less unique data than the vanilla Transformer (Figure 3 (a)).

The saving grows with U/N and with epochs (Figure 3). At the smallest U the two models need about the same data. The largest measured saving at 271M is 41% of the unique data, with $U = 100N$ and three epochs ($\mu = 1.69$). At one epoch it also grows with size. Table 7 in Appendix D lists every measured μ.

5.3 One law for both models

The two multipliers above are measured only inside the grid, up to 778M, 400 UTPP and four epochs. To explain them and to go beyond the grid, we fit Equation 3 jointly to both models' losses at 60M, 134M and 271M, which gives

$$L = 0.5 + \left[\left(\frac{1.75}{(v\,n)^{0.112}}\right)^{3.75} + \left(\frac{1.40}{d^{0.126}}\right)^{3.75}\right]^{1/3.75}, \qquad R^* = 14.6\left(\frac{u}{n}\right)^{0.45}, \tag{7}$$

$$v = 1 \ \text{ for the Transformer,} \qquad v = 1.14\left(\frac{d}{n}\right)^{0.010} n^{0.007} \ \text{ for latent feedback,}$$

with $n = N/10^8$, $u = U/10^{10}$ and $d = D/10^{10}$, where the effective data D follows Equation 3.[6]

The law predicts a held-out size (Table 1). The 778M model is 3.3 times the size of the largest fitted one, yet the law predicts its losses with a mean absolute error of 0.13% and at most 0.29%. This error is below those that recent scaling laws report on their own held-out runs, 0.50% for Farseer and 0.60% for Skaling [21, 44]. The law also picks the winner: it is off by about the same amount for both models, so it

[6] Table 4 in Appendix A lists every coefficient and its meaning and fitted value.

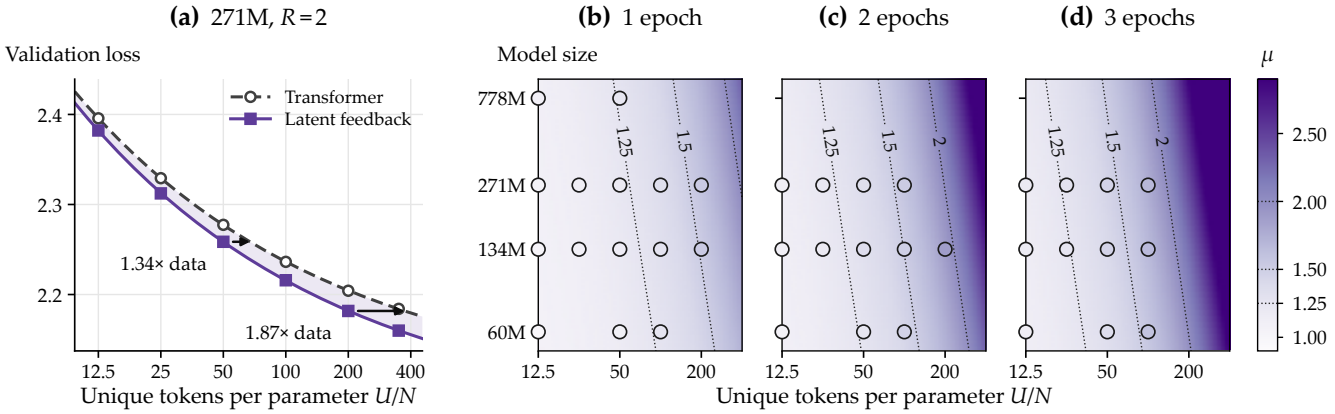

Figure 3: The data multiplier μ across unique data and model size. (a) 271M at two epochs: each arrow is the unique data latent feedback saves; the arrow at $U = 200N$ is extrapolated. (b)–(d) Dots are measured cells, read off the Transformer's losses across U interpolated in $\log U$ (Appendix D); everywhere else, the map shows the law of Section 5.3. μ grows with unique data and with epochs, and depends on U/N far more than on size.

predicts the gap to within 0.05% in every held-out row, and gives the measured sign in 6 of 7, with the seventh a tie, and in all but two of the fitted cells, both at $U = 12.5N$ and $F = 4$, where the predicted gap is within 0.001 of zero.

With the law checked at 778M, we use it to explain the two savings, to draw the boundary between the two regimes, and to predict the savings beyond the grid.

First, the law explains the two savings: latent feedback trades epochs for capacity. Equation 7 has a capacity term and a data term, and latent feedback moves both. It lowers the capacity term, because it fits as a Transformer with 1.14 times the parameters, a factor that grows slowly with data per parameter and with size. However, it raises the data term, because its cost leaves it about 1.2 times fewer epochs, and repeated epochs still count: in the coupled fits, four epochs at 134M are worth 3.0 to 3.9 fresh ones. Latent feedback therefore wins where capacity limits the loss, and loses where data limits it. This is why both savings grow with unique data and with more epochs: each moves the Transformer toward its capacity floor, which latent feedback lowers.

We want to point out the savings hold at a fixed model size. Latent feedback spends c times a Transformer's training FLOPs per token, as much as a Transformer c times larger, yet the law counts it as a Transformer only v times larger, and v stays below c (1.10 to 1.19 against 1.20 to 1.27). So when model size is free, simply spending the same FLOPs on a Transformer c times larger, with the same unique data, gives a loss 0.006 to 0.012 lower in every fitted cell. That Transformer, however, costs c times more per decoding step and holds a larger KV cache, while latent feedback serves at the cost of the smaller model.

Second, the law draws the boundary between the two regimes (the line $\kappa = 1$ in Figure 2).[7] On one side of this line, a FLOP is worth more as another Transformer epoch; on the other, it is worth more as latent feedback. The boundary moves toward less unique data as the budget grows: latent feedback wins from about 75 UTPP at $F = 1$, 33 at $F = 2$, 20 at $F = 3$ and 12 at $F = 4$.

Third, the law predicts that both savings keep growing beyond the grid (Figure 4; Table 8 in Appendix E). Each panel extends one variable beyond the measured range. Along size, the savings rise because $\alpha < \beta$ makes larger models more capacity-limited, and further through γ: at 100B, 100 UTPP

[7] Figure 7 in Appendix E maps the loss gap behind it.

Table 1: Held-out test of the law at 778M. Fitted on 60M to 271M only. Error: (predicted − measured) / measured.

		Transformer			Latent feedback		
U	F	Measured	Predicted	Error (%)	Measured	Predicted	Error (%)
$12.5N$	1	2.2054	2.2112	+0.26	2.2211	2.2276	+0.29
$12.5N$	2	2.1434	2.1460	+0.12	2.1501	2.1532	+0.15
$12.5N$	3	2.1136	2.1172	+0.17	2.1162	2.1195	+0.16
$12.5N$	4	2.0950	2.1005	+0.26	2.0950	2.0996	+0.22
$50N$	1	2.0878	2.0873	−0.03	2.0893	2.0897	+0.02
$50N$	2	2.0452	2.0454	+0.01	2.0421	2.0413	−0.04
$100N$	1	2.0448	2.0441	−0.03	2.0413	2.0406	−0.03

and $F = 2$, the law predicts $\kappa = 1.83$, and 1.41 with $\gamma = 0$. Within 60M to 271M, γ is not significantly above zero ($p = 0.054$), but the held-out 778M favors it: without γ, the law underrates latent feedback in all seven held-out rows. The dashed curves in Figure 4 show the savings if v does not grow with size. More unique data per parameter raises them fast, until the Transformer cannot catch up at any budget: it nears its capacity floor, while latent feedback's floor is that of a Transformer about 1.2 times larger. More repeats push the compute saving the same way, since repeats saturate for both models and only the capacity gap remains; the data saving instead levels off at a height set by the unique data. The curves only shift with the other variables, so latent feedback passes the crossover at any unique data once a run repeats enough, as long as repeats do not overfit (true up to four epochs; Figure 6 in Appendix D).

5.4 Robustness: repeats, seeds and downstream tasks

Repeats up to four epochs do not overfit either model (Figure 6 in Appendix D). Even at the scarcest unique data, validation loss keeps falling through four epochs at every size. The gaps are stable across seeds: at 134M, a gap's standard deviation is below 0.002 and none changes sign (Appendix D).

Downstream tasks agree well with the validation loss (Table 2). In Table 2, the Transformer has the lower BPB only at the smallest unique data ($U \leq 25N$) and budgets of $F \leq 2$; latent feedback has it everywhere else, so the winner switches at almost the same place as in Figure 2. At $F = 4$, latent feedback has the lower BPB in all 17 finished cells. On downstream tasks, latent feedback starts to win slightly earlier than on validation loss, and the two metrics pick the same winner in 59 of 71 paired cells (Table 6 in Appendix D gives the validation losses).

6 Related work

Spending compute on scarce data. With unique data scarce, prior work fits a scaling law and, given unique data, compute, and a fixed architecture, chooses the model size and number of epochs [12, 22, 27]. When taking architecture into account, sparsity in models delays the point where repeats saturate [45]; masked diffusion [46] overtakes autoregressive training once compute is abundant, across unique data and epochs. However, its gain needs excessive epochs: on 100M unique tokens its best loss comes at 500 epochs, against 50 for autoregressive training, whose hyperparameters follow a prior recipe instead of being tuned per run. The looped model Parcae [47] spends extra FLOPs on loops at a fixed parameter count, yet weighs them against fresh tokens rather than epochs. Latent feedback pays off within four epochs once each parameter has enough unique data, and we price it against epochs on fixed unique data, tuning both models at every size and amount of unique data.

Recurrence in depth. Looped models buy more compute per token by running their layers several times within each token [23, 24, 48], and their scaling laws count the extra loops as effective parameters [25, 26,

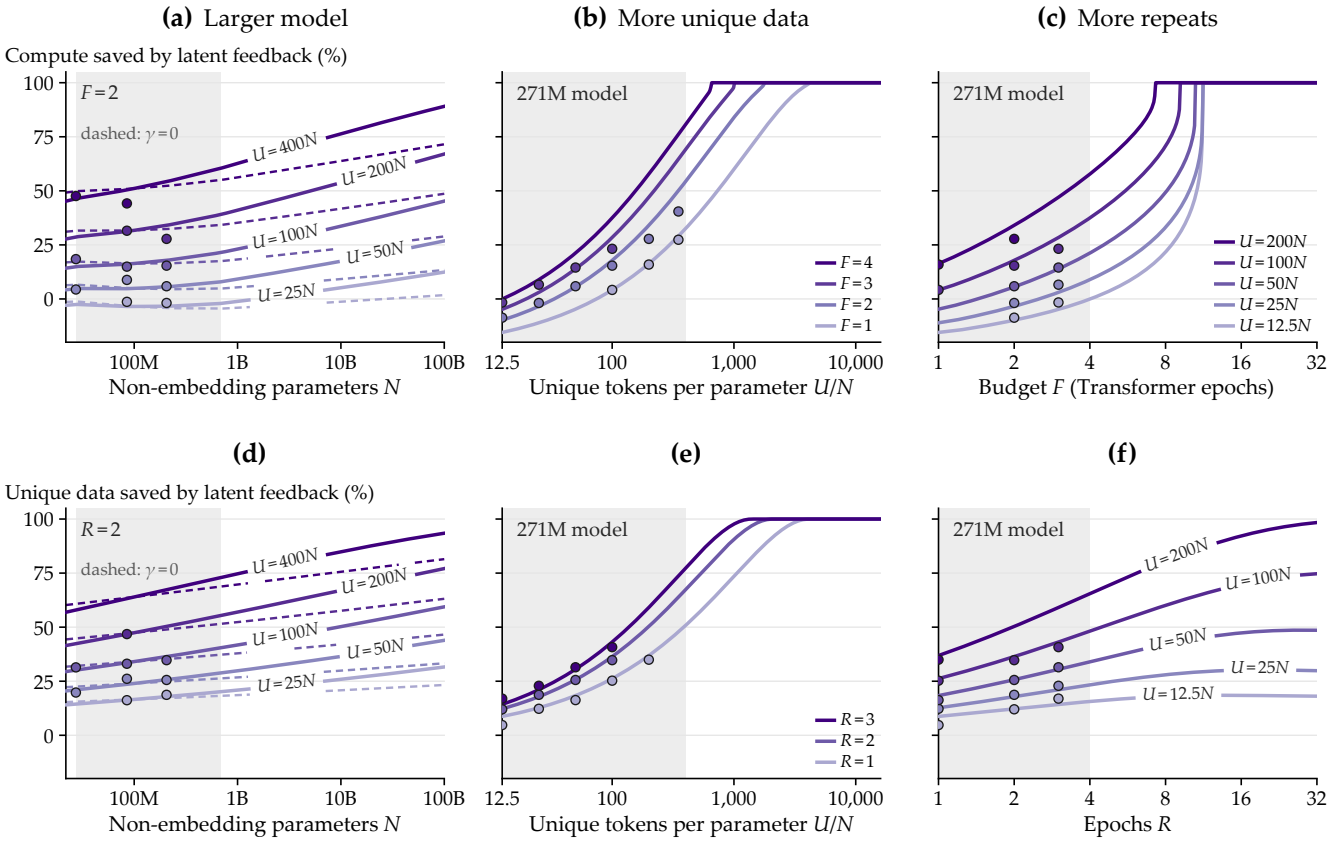

Figure 4: The savings beyond the grid, as the law of Section 5.3 predicts them. Every curve comes from one fit of the law to all cells; the dots are the few of those cells that lie on each curve. Each panel extends one variable beyond its measured range (shaded), with one curve per value of a second variable; in (a) and (d), dashed curves set $\gamma = 0$, so v does not grow with size. Top: compute saved at equal unique data. Bottom: unique data saved at equal epochs. A saving of 100% means the Transformer cannot reach latent feedback's loss with any number of epochs or any amount of unique data.

47]. Our law takes the same view: latent feedback enters it as a Transformer with v times the parameters. However, a looped model still pays its extra compute again at every decoding step, even if it allows KV cache sharing [49]. Latent feedback recurs across tokens instead: it pays for its extra passes in training and in prefill, and generates each token like a Transformer with one more cached vector.

Feedback across tokens. Feedback memory [13] first fed deep states to later tokens but trained token by token; recent models feed back a deeper layer [14, 50–52] or a frozen teacher's state [53]. At equal compute or training time, some beat a Transformer [14, 50] and others trail it [51, 53]. Yet none varies data repetition: even Wang et al. [14], who motivate feedback by scarce unique data, give the Transformer more tokens rather than more epochs. We compare against more epochs, across model size, unique data and repetition, and fit all three in one scaling law.

7 Conclusion

We show that latent feedback makes every FLOP and every token count for more: once each parameter has enough unique data, it reaches a given loss with less compute than more epochs, and with less unique data at equal epochs. Both savings grow with unique data and repetition, and one scaling law, fitted to latent feedback and the Transformer together, predicts them at a held-out size. As a takeaway, when

Table 2: Downstream tasks at equal model size, unique data and compute. Bits per byte (BPB ↓) of both models at budgets of $F = 1$ to 4 Transformer epochs. The lower of each pair is shaded. BPB averages 37 task suites, each weighted by its inverse seed variance (Appendix H).

		$F = 1$		$F = 2$		$F = 3$		$F = 4$	
Size	U	Transformer	Feedback	Transformer	Feedback	Transformer	Feedback	Transformer	Feedback
60M	12.5N	1.2324	1.2453	1.1623	1.1680	1.1399	1.1388	1.1279	1.1186
	50N	1.1051	1.1014	1.0664	1.0539	1.0511	1.0371	1.0413	1.0278
	100N	1.0673	1.0566	1.0373	1.0205	1.0220	1.0072	1.0163	0.9999
	400N	1.0116	0.9917	0.9970	0.9708	0.9867	0.9634	0.9822	0.9615
134M	12.5N	1.0374	1.0514	0.9907	0.9889	0.9677	0.9624	0.9559	0.9466
	25N	0.9865	0.9835	0.9428	0.9423	0.9263	0.9225	0.9151	0.9088
	50N	0.9444	0.9380	0.9160	0.9062	0.9013	0.8880	0.8930	0.8766
	100N	0.9071	0.9020	0.8849	0.8767	0.8739	0.8602	0.8648	0.8544
	200N	0.8838	0.8725	0.8637	0.8484	0.8527	0.8389	0.8501	0.8309
	400N	0.8634	0.8541	0.8506	0.8348	0.8412	0.8282	0.8370	0.8213
271M	12.5N	0.9096	0.9173	0.8665	0.8675	0.8508	0.8470	0.8380	0.8337
	25N	0.8678	0.8690	0.8325	0.8282	0.8179	0.8119	0.8073	0.8012
	50N	0.8336	0.8309	0.8045	0.8009	0.7902	0.7854	0.7844	0.7729
	100N	0.8021	0.7980	0.7792	0.7723	0.7675	0.7582	0.7619	0.7502
	200N	0.7783	0.7712	0.7603	0.7499	0.7499	0.7392	0.7430	0.7308
	350N	0.7652	0.7525	0.7469	0.7365	0.7394	0.7265	0.7375	0.7217
778M	12.5N	0.7670	0.7715	0.7301	0.7322	0.7134	0.7106	0.7009	0.6985
	50N	0.6959	0.6949	0.6673	0.6662	—	—	—	—
	100N	0.6683	0.6613	—	—	—	—	—	—

high-quality data, not compute, limits a run of a given model size, feeding back the top-layer state uses compute better than training more epochs.

Limitation. So far, our ladder stops at a relatively small scale and uses a single backbone; larger models and data budgets, as well as other architectures may shift the crossover. We also train latent feedback with a single fixed schedule, and further optimizations are possible. Lastly, a comparison with the recurrent depth models will be of great value. We leave these to future work.

Acknowledgment. We sincerely thank Zitian Gao and Xi Wang for many valuable discussions.

References

[1] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. PaLM: Scaling language modeling with pathways. arXiv:2204.02311, 2022.

[2] Kimi Team, Yifan Bai, Yiping Bao, Y. Charles, Cheng Chen, Guanduo Chen, Haiting Chen, Huarong Chen, Jiahao Chen, Ningxin Chen, et al. Kimi K2: Open agentic intelligence. arXiv:2507.20534, 2025.

[3] Konwoo Kim, Suhas Kotha, Percy Liang, and Tatsunori Hashimoto. Pre-training under infinite compute. arXiv:2509.14786, 2025.

[4] Pablo Villalobos, Anson Ho, Jaime Sevilla, Tamay Besiroglu, Lennart Heim, and Marius Hobbhahn. Will we run out of data? limits of LLM scaling based on human-generated data. arXiv:2211.04325, 2024. Version 2.

[5] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. arXiv:2005.14165, 2020.

[6] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford,

Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.

[7] Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, et al. StarCoder 2 and the stack v2: The next generation. arXiv:2402.19173, 2024.

[8] Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, et al. Phi-4 technical report. arXiv:2412.08905, 2024.

[9] Team Olmo, Allyson Ettinger, Amanda Bertsch, Bailey Kuehl, David Graham, David Heineman, Dirk Groeneveld, Faeze Brahman, Finbarr Timbers, Hamish Ivison, et al. Olmo 3. arXiv:2512.13961, 2025.

[10] The Microsoft AI Team. MAI-Thinking-1: Building a hill-climbing machine. https://microsoft.ai/pdf/mai-thinking-1.pdf, 2026. Technical report.

[11] Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Gabriel Martín Blázquez, Guilherme Penedo, Lewis Tunstall, Andrés Marafioti, Hynek Kydlíček, Agustín Piqueres Lajarín, Vaibhav Srivastav, et al. SmolLM2: When smol goes big – data-centric training of a small language model. arXiv:2502.02737, 2025.

[12] Niklas Muennighoff, Alexander M. Rush, Boaz Barak, Teven Le Scao, Aleksandra Piktus, Nouamane Tazi, Sampo Pyysalo, Thomas Wolf, and Colin Raffel. Scaling data-constrained language models. arXiv:2305.16264, 2023.

[13] Angela Fan, Thibaut Lavril, Edouard Grave, Armand Joulin, and Sainbayar Sukhbaatar. Addressing some limitations of transformers with feedback memory. arXiv:2002.09402, 2020.

[14] Xi Wang, Ziyang Cai, Zheng Zhan, Harry Dong, Ying Fan, Gustavo de Rosa, Tim Pearce, and John Langford. Full-bandwidth transformer. arXiv:2608.08888, 2026.

[15] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, et al. The Llama 3 herd of models. arXiv:2407.21783, 2024.

[16] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, et al. Qwen3 technical report. arXiv:2505.09388, 2025.

[17] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, et al. Gemma 3 technical report. arXiv:2503.19786, 2025.

[18] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. arXiv:2001.08361, 2020.

[19] Houyi Li et al. Predictable scale: Part I, Step Law – optimal hyperparameter scaling law in large language model pretraining. arXiv:2503.04715, 2025.

[20] Kimi Team. Kimi Linear: An expressive, efficient attention architecture. arXiv:2510.26692, 2025.

[21] Mathurin Videau, Badr Youbi-Idrissi, David Lopez-Paz, and Kartik Ahuja. Skaling: Chinchilla's exponents meet Kaplan's coupling. arXiv:2608.07222, 2026.

[22] Tian Qin, Kimia Hamidieh, and David Alvarez-Melis. Bridging compute- and data-optimal pretraining. arXiv:2607.25271, 2026.

[23] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers. In *International Conference on Learning Representations*, 2019.

[24] Jonas Geiping et al. Scaling up test-time compute with latent reasoning: A recurrent depth approach. arXiv:2502.05171, 2025.

[25] Kristian Schwethelm, Daniel Rückert, et al. How much is one recurrence worth? iso-depth scaling laws for looped language models. arXiv:2604.21106, 2026.

[26] Xuyuan Xiong, Siyu Chen, Zihan Shao, Yue Su, Weiyang Jin, Ping Luo, Zhuoran Yang, and Xihui Liu. Truncate or not? a scaling study of looped transformers with truncated backpropagation. arXiv:2610.26030, 2026.

[27] Justin Lovelace, Christian Belardi, Srivatsa Kundurthy, Shriya Sudhakar, and Kilian Q. Weinberger. Prescriptive scaling laws for data constrained training. arXiv:2605.01640, 2026.

[28] Joshua Shay Kricheli, Alexander Lawrence Reid, Soumajyoti Sarkar, Venkata Gandikota, and Paulo Shakarian. Tokens-per-parameter coverage is critical for robust LLM scaling law extrapolation. arXiv:2605.08541, 2026.

[29] Biao Zhang and Rico Sennrich. Root mean square layer normalization. In *Advances in Neural Information*

Processing Systems, 2019.

[30] Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. RoFormer: Enhanced transformer with rotary position embedding. arXiv:2104.09864, 2021.

[31] Alex Henry, Prudhvi Raj Dachapally, Shubham Shantaram Pawar, and Yuxuan Chen. Query-key normalization for transformers. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2020.

[32] Noam Shazeer. GLU variants improve transformer. arXiv:2002.05202, 2020.

[33] Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, et al. 2 OLMo 2 furious. *arXiv preprint arXiv:2501.00656*, 2025.

[34] Zichong Li, Liming Liu, Chen Liang, Weizhu Chen, and Tuo Zhao. NorMuon: Making Muon more efficient and scalable. arXiv:2510.05491, 2025.

[35] Keller Jordan, Yuchen Jin, Vlado Boza, Jiacheng You, Franz Cesista, Laker Newhouse, and Jeremy Bernstein. Muon: An optimizer for hidden layers in neural networks. https://kellerjordan.github.io/posts/muon/, 2024.

[36] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

[37] Ian Magnusson, Akshita Bhagia, Valentin Hofmann, Luca Soldaini, Ananya Harsh Jha, Oyvind Tafjord, Dustin Schwenk, Evan Pete Walsh, Yanai Elazar, Kyle Lo, Dirk Groeneveld, Iz Beltagy, Hannaneh Hajishirzi, Noah A. Smith, Kyle Richardson, and Jesse Dodge. Paloma: A benchmark for evaluating language model fit. arXiv:2312.10523, 2023.

[38] Luca Soldaini, Rodney Kinney, Akshita Bhagia, et al. Dolma: an open corpus of three trillion tokens for language model pretraining research. arXiv:2402.00159, 2024.

[39] Dirk Groeneveld, Iz Beltagy, Pete Walsh, et al. OLMo: Accelerating the science of language models. arXiv:2402.00838, 2024.

[40] Xi Wang and Laurence Aitchison. How to set AdamW's weight decay as you scale model and dataset size. arXiv:2405.13698, 2024.

[41] Shane Bergsma, Nolan Dey, Gurpreet Gosal, Gavia Gray, Daria Soboleva, and Joel Hestness. Power lines: Scaling laws for weight decay and batch size in LLM pre-training. arXiv:2505.13738, 2025.

[42] Alexander Hägele, Elie Bakouch, Atli Kosson, Loubna Ben Allal, Leandro Von Werra, and Martin Jaggi. Scaling laws and compute-optimal training beyond fixed training durations. arXiv:2405.18392, 2024.

[43] Yuling Gu, Oyvind Tafjord, Bailey Kuehl, Dany Haddad, Jesse Dodge, and Hannaneh Hajishirzi. OLMES: A standard for language model evaluations. arXiv:2406.08446, 2024.

[44] Houyi Li et al. Predictable scale: Part II, Farseer: A refined scaling law in large language models. arXiv:2506.10972, 2025.

[45] Boqian Wu, Qiao Xiao, Patrik Okanovic, Tomasz Sternal, Maurice van Keulen, Mykola Pechenizkiy, Elena Mocanu, Torsten Hoefler, and Decebal Constantin Mocanu. When data is scarce: Scaling sparse language models with repeated training. In *International Conference on Machine Learning*, 2026. arXiv:2606.01155.

[46] Mihir Prabhudesai, Mengning Wu, Amir Zadeh, Katerina Fragkiadaki, and Deepak Pathak. Diffusion beats autoregressive in data-constrained settings. arXiv:2507.15857, 2025.

[47] Hayden Prairie, Zachary Novack, Taylor Berg-Kirkpatrick, and Daniel Y. Fu. Parcae: Scaling laws for stable looped language models. arXiv:2604.12946, 2026.

[48] Benhao Huang, Chufan Shi, Junlin Chen, Shicheng Wen, Zhengzhong Liu, Eric Xing, and Xuezhe Ma. Towards looped models done right, part i: Topology, input injection, recurrent-state design. *arXiv preprint arXiv:2610.23036*, 2026.

[49] Benhao Huang, Chufan Shi, Junlin Chen, Shicheng Wen, Zhengzhong Liu, Eric Xing, and Xuezhe Ma. Towards looped models done right, part ii: Rethinking at fixed points. *arXiv preprint arXiv:2610.06833*, 2026.

[50] Zeyi Huang, Xuehai He, LiLiang Ren, Yiping Wang, Baolin Peng, Hao Cheng, Shuohang Wang, Pengcheng He, Jianfeng Gao, Yong Jae Lee, and Yelong Shen. Latent recurrent transformer: Architecture exploration, training strategies, and scaling behavior. arXiv:2605.26797, 2026.

[51] Ziyang Cai, Xingyu Zhu, Yihe Dong, Yinghui He, and Sanjeev Arora. T^2MLR: Transformer with temporal middle-layer recurrence. arXiv:2607.15178, 2026.

[52] Michael C Mozer, Shoaib Ahmed Siddiqui, Danny Sawyer, Sunny Sanyal, and Rosanne Liu. Recirculation. *arXiv preprint arXiv:2608.17981*, 2026.

[53] Dor Tirosh, Ido Amos, and Mor Geva. Pretraining latent information feedback transformers with teacher supervision. arXiv:2609.38149, 2026.
[54] Qwen Team, An Yang, Baosong Yang, Beichen Zhang, et al. Qwen2.5 technical report. arXiv:2412.15115, 2024.
[55] Meta. Llama 3.2 model card. https://github.com/meta-llama/llama-models/blob/main/models/llama3_2/MODEL_CARD.md, 2024.
[56] Aili Chen, Aonian Li, Baichuan Zhou, Bangwei Gong, Binyang Jiang, et al. The MiniMax-M2 series: Mini activations unleashing max real-world intelligence. arXiv:2605.26494, 2026.
[57] Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. TinyLlama: An open-source small language model. arXiv:2401.02385, 2024.
[58] Meta. Llama 3 model card. https://github.com/meta-llama/llama-models/blob/main/models/llama3/MODEL_CARD.md, 2024.
[59] NVIDIA, Aaron Blakeman, Aarti Basant, Abhinav Khattar, et al. Nemotron-H: A family of accurate and efficient hybrid Mamba-Transformer models. arXiv:2504.03624, 2025.
[60] MiniCPM Team, Chaojun Xiao, Yuxuan Li, Xu Han, Yuzhuo Bai, et al. MiniCPM4: Ultra-efficient LLMs on end devices. arXiv:2506.07900, 2025.
[61] Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, et al. Phi-3 technical report: A highly capable language model locally on your phone. arXiv:2404.14219, 2024.
[62] Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, et al. Gemma 2: Improving open language models at a practical size. arXiv:2408.00118, 2024.
[63] GLM-4.5 Team, Aohan Zeng, Xin Lv, Qinkai Zheng, et al. GLM-4.5: Agentic, reasoning, and coding (ARC) foundation models. arXiv:2508.06471, 2025.
[64] DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, et al. DeepSeek-V3 technical report. arXiv:2412.19437, 2024.
[65] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, et al. Llama 2: Open foundation and fine-tuned chat models. arXiv:2307.09288, 2023.
[66] Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O'Brien, et al. Pythia: A suite for analyzing large language models across training and scaling. arXiv:2304.01373, 2023.
[67] DeepSeek-AI, Xiao Bi, Deli Chen, Guanting Chen, et al. DeepSeek LLM: Scaling open-source language models with longtermism. arXiv:2401.02954, 2024.
[68] Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, et al. The Falcon series of open language models. arXiv:2311.16867, 2023.
[69] Martin Marek, Sanae Lotfi, Aditya Somasundaram, Andrew Gordon Wilson, and Micah Goldblum. Small batch size training for language models: When vanilla SGD works, and why gradient accumulation is wasteful. arXiv:2507.07101, 2025.
[70] Tim Pearce and Jinyeop Song. Reconciling Kaplan and Chinchilla scaling laws. arXiv:2406.12907, 2024.

A Notation

Table 3 lists the quantities of the main text, and Table 4 the symbols of the fitted law.

Table 3: Quantities of the main text, in the order they are defined.

Quantity	Meaning	Defined in
TPP	Training tokens per parameter, repeats included	Section 1
UTPP	Unique tokens per non-embedding parameter, U/N	Section 1
Latent feedback Transformer	Algorithm 1: one or two passes, no gradient in the first	Section 2
p	Probability that an update runs one pass	Section 2
F	Compute budget, in Transformer epochs	Section 3
c_s	Training FLOPs per token of schedule s over a Transformer's	Section 3
R	Epochs trained within the budget, F/c_s	Section 3
F^*	Crossover budget, above which latent feedback wins	Section 3
κ	Compute multiplier, at equal size and unique data	Section 3
μ	Data multiplier, at equal size and epochs	Section 3

Table 4: Symbols of Equation 8, with the joint fit on the 128 cells. Appendix E gives the equation.

Symbol	Meaning	Value or source
Inputs		
s	Training schedule	Transformer, or latent feedback at $p = 0.15$
N	Non-embedding parameters	27.3M, 85.0M, 205.6M (60M, 134M, 271M)
U	Unique tokens	$12.5N$ to $400N$
F	Compute budget, in Transformer epochs	1 to 4
Fitted coefficients		
E	Irreducible loss	0.5 (its lower bound)
A	Scale of the parameter term	1.755
α	Exponent of the parameter term	0.112
B	Scale of the data term	1.401
β	Exponent of the data term	0.126
q	Coupling of the two terms; $q = 1$ is additive	3.748
R_T^*	R^* at $U = 100N$	14.60
ρ	Growth of R^* with U/N	$+0.450$
v_0	Capacity factor v_s at the reference run	1.139
ω	Growth of v_s with effective tokens per parameter	0.010
γ	Growth of v_s with model size	0.007
Fixed constants		
c_s	Training FLOPs per token over a Transformer's	1.197, 1.232, 1.251 (60M, 134M, 271M)
N_0, D_0	Reference run: parameters and effective tokens	10^8 and 10^{10}
Derived		
R	Epochs within the budget F	F/c_s
n, u, d_s	N, U and D_s in reference units	N/N_0, U/D_0, D_s/D_0
D_s	Effective data, in fresh tokens	Equation 8
L_s	Predicted validation loss at the tuned optimum	Equation 8

B Repeated data in practice

Models train far past the compute-optimal ratio (Figure 5). Chinchilla trained on 20 tokens per parameter [6]. We checked 257 base models whose reports state both counts, each against its source. The median dense model trains on about 6,800 tokens per parameter below 1B. It uses 1,040 from 1B to 10B, and 210 from 10B to 100B. The median mixture-of-experts (MoE) model uses 1,640 per active parameter.

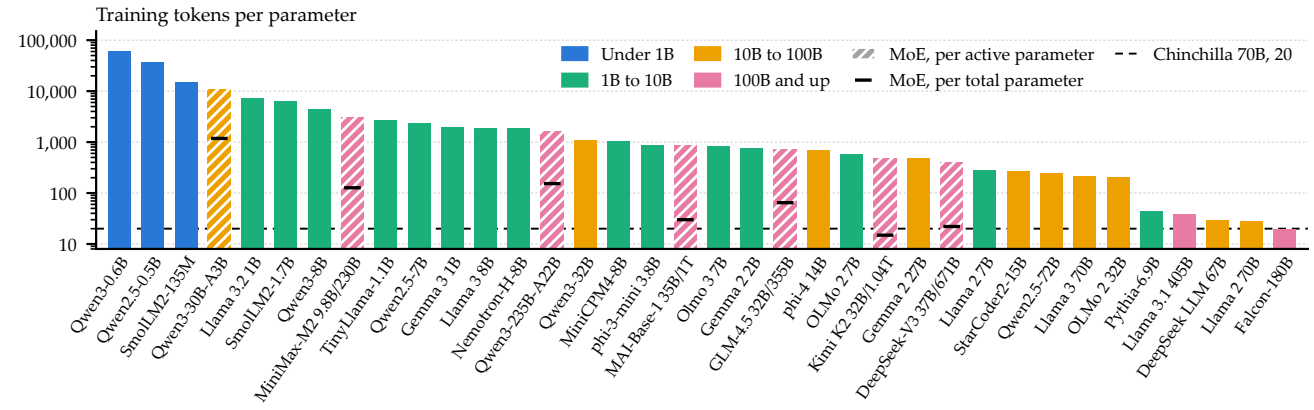

Figure 5: Training tokens per parameter of base models. Sources are each developer's report or model card [2, 7–11, 15–17, 33, 54–68]. Colour gives the total size, and hatched bars are MoE models. MoE bars count active parameters; the tick marks total parameters.

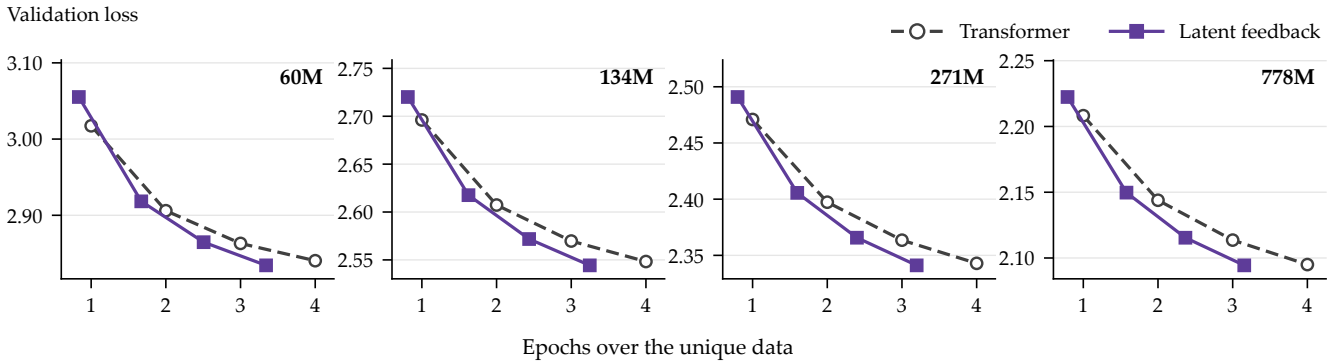

Figure 6: Repeats up to four epochs do not overfit. Validation loss after one to four epochs at the scarcest unique data, $U = 12.5N$, one panel per model size; both curves keep falling. Latent feedback costs $c > 1$ Transformer FLOPs per token, so at equal compute it trains fewer epochs and its points sit to the left.

C Parallel passes

Each pass feeds position t only the states of earlier positions, so after p passes the first p positions equal decoding, and the fixed point of Equation 2 is unique and equals decoding. More parallel passes than the 4 we evaluate with barely move the loss, far less than the gaps between the two models (Table 5).

Table 5: Loss of latent feedback scored with P parallel passes of Equation 2, on ten runs of an earlier ladder (one 271M, nine 778M) and a 10M-token Dolma 3 validation slice. Differences are taken per run against 16 passes.

Passes P	Mean loss	Mean difference ↓	Largest difference ↓
4	2.02863	0.00007	0.00013
8	2.02869	0.00001	0.00004
16	2.02870	–	–

D Every cell

Table 6 lists the validation loss of every cell at $F = 1$ to 4, and Table 2 its downstream BPB.

Seeds. At 134M we ran three seeds at $U = 12.5N$ and $100N$, for $F = 1$ to 4. The standard deviation of a

cell's loss is 0.0003 to 0.0014. That of a gap is 0.0004 to 0.0017, and every gap keeps its sign.

Multipliers. Table 7 (top) lists the measured κ. Each latent feedback cell's loss is read off the Transformer's losses at the same size and U, interpolated in $\log F$ between $F = 1$ and 4. Table 7 (bottom) lists the measured μ. Latent feedback's loss at $R = 1, 2, 3$ is interpolated in $\log R$ between its branches at $R = F/c$. We then read μ off the Transformer's losses across U, interpolated in $\log U$ between levels.

Table 6: Validation loss ↓ of every cell at $F = 1$ to 4. Both models, at equal size, unique data and budget F; Feedback is the latent feedback Transformer. The lower of each pair is shaded: purple where latent feedback leads, grey where the Transformer does. 778M is held out of the law; —: its grid has no such cell.

		$F = 1$		$F = 2$		$F = 3$		$F = 4$	
Size	U	Transformer	Feedback	Transformer	Feedback	Transformer	Feedback	Transformer	Feedback
60M	$12.5N$	3.0141	3.0508	2.9060	2.9183	2.8631	2.8647	2.8404	2.8343
	$50N$	2.8110	2.8199	2.7456	2.7424	2.7166	2.7074	2.6996	2.6866
	$100N$	2.7428	2.7387	2.6918	2.6803	2.6689	2.6536	2.6553	2.6376
	$400N$	2.6507	2.6343	2.6199	2.5987	2.6059	2.5826	2.5973	2.5737
134M	$12.5N$	2.6927	2.7202	2.6073	2.6178	2.5697	2.5718	2.5483	2.5440
	$25N$	2.5999	2.6126	2.5320	2.5333	2.5014	2.4964	2.4826	2.4743
	$50N$	2.5281	2.5341	2.4716	2.4656	2.4452	2.4343	2.4292	2.4159
	$100N$	2.4685	2.4670	2.4238	2.4155	2.4030	2.3905	2.3904	2.3750
	$200N$	2.4234	2.4132	2.3880	2.3722	2.3711	2.3534	2.3610	2.3421
	$400N$	2.3890	2.3741	2.3610	2.3428	2.3478	2.3281	2.3397	2.3190
271M	$12.5N$	2.4691	2.4907	2.3957	2.4045	2.3629	2.3642	2.3428	2.3394
	$25N$	2.3893	2.3993	2.3292	2.3308	2.3015	2.2974	2.2842	2.2777
	$50N$	2.3283	2.3323	2.2774	2.2739	2.2540	2.2462	2.2397	2.2288
	$100N$	2.2769	2.2744	2.2363	2.2283	2.2169	2.2058	2.2048	2.1921
	$200N$	2.2374	2.2291	2.2042	2.1916	2.1885	2.1738	2.1786	2.1626
	$350N$	2.2114	2.1988	2.1841	2.1681	2.1712	2.1533	2.1633	2.1444
778M	$12.5N$	2.2054	2.2211	2.1434	2.1501	2.1136	2.1162	2.0950	2.0950
	$50N$	2.0878	2.0893	2.0452	2.0421	—	—	—	—
	$100N$	2.0448	2.0413	—	—	—	—	—	—

Table 7: Measured multipliers. Top: compute multiplier $\kappa\uparrow$ at equal size and unique data. Bottom: data multiplier $\mu\uparrow$ at equal size and epochs. —: no such cell, or the read leaves the measured range (one to four Transformer epochs for κ, the largest U for μ).

	$F=1$				$F=2$				$F=3$			
UTPP	60M	134M	271M	778M	60M	134M	271M	778M	60M	134M	271M	778M
12.5	—	—	—	—	0.92	0.92	0.92	0.93	0.98	0.98	0.98	0.97
25	—	—	—	—	—	0.99	0.98	—	—	1.08	1.07	—
50	—	—	—	—	1.05	1.10	1.06	—	1.17	1.22	1.17	—
100	1.06	1.02	1.04	—	1.23	1.18	1.18	—	—	1.33	1.30	—
200	—	1.22	1.19	—	—	1.46	1.38	—	—	—	—	—
350	—	—	1.38	—	—	—	1.68	—	—	—	—	—
400	1.45	1.45	—	—	1.91	1.79	—	—	—	—	—	—

	1 epoch				2 epochs			3 epochs		
U	60M	134M	271M	778M	60M	134M	271M	60M	134M	271M
$12.5N$	0.99	1.02	1.05	1.09	1.08	1.12	1.14	1.14	1.18	1.20
$25N$	—	1.10	1.14	—	—	1.19	1.23	—	1.26	1.30
$50N$	1.11	1.16	1.19	1.25	1.25	1.35	1.34	1.35	1.47	1.46
$100N$	1.27	1.29	1.34	—	1.46	1.49	1.53	1.60	1.67	1.69
$200N$	—	1.54	1.54	—	—	1.88	—	—	—	—

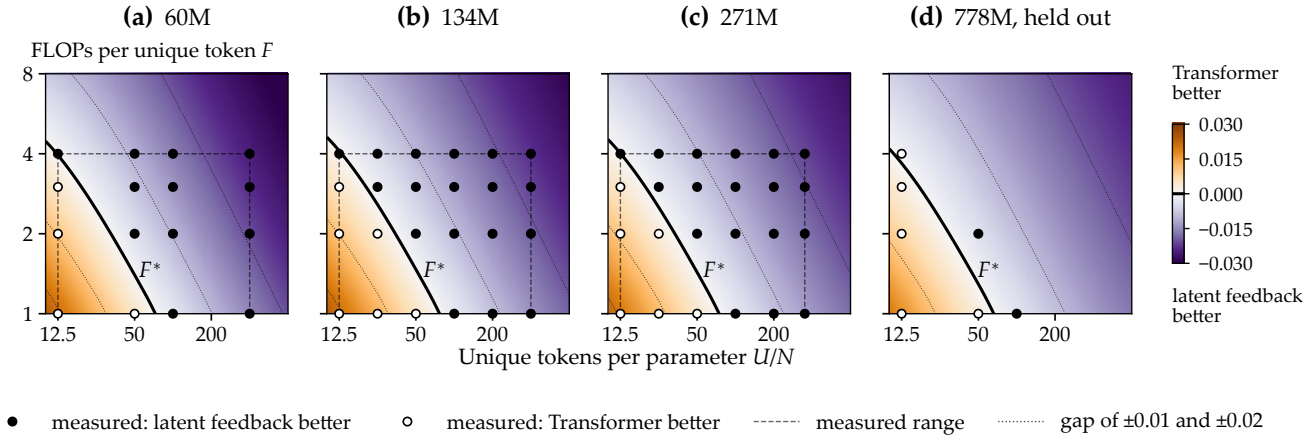

Figure 7: The loss gap at equal compute across unique data and budget. Markers are measured cells; everywhere else, the map is extrapolated with the law of Section 5.3. Latent feedback wins above the solid line, the crossover F^*, which falls as unique data grows and hardly moves with size.

E The fitted law

Equation 8 is Equation 3 in its best form (Table 9). It is the coupled form, with latent feedback's capacity growing with TPP and with size. It is fitted jointly to both models' 128 cells at 60M, 134M and 271M. Each symbol has one of four kinds (Table 4). **Inputs** are chosen by whoever uses the fitted law. **Fitted coefficients** are estimated once, jointly, from the 128 cells. **Fixed constants** are measured or set by design and never fitted. Black symbols are **derived**: computed from the other three kinds.

$$L_s(N,U,F) = E + \left[\left(\frac{A}{(v_s\, n)^{\alpha}}\right)^{q} + \left(\frac{B}{d_s^{\beta}}\right)^{q}\right]^{1/q}, \qquad v_s = v_0\left(\frac{d_s}{n}\right)^{\omega} n^{\gamma}, \tag{8}$$

$$D_s = U\left[1 + R^*\left(1 - e^{-(R-1)/R^*}\right)\right], \qquad R = \frac{F}{c_s}, \qquad R^* = R_T^*\left(\frac{u}{n}\right)^{\rho},$$

$$n = \frac{N}{N_0}, \qquad u = \frac{U}{D_0}, \qquad d_s = \frac{D_s}{D_0}, \qquad N_0 = 10^8, \qquad D_0 = 100\, N_0 = 10^{10}.$$

Below one epoch the bracket is replaced by R. The Transformer T has $v_T = 1$. Every size is measured against one reference run: $N_0 = 10^8$ non-embedding parameters trained on $D_0 = 10^{10}$ effective tokens, that is, 100 tokens per parameter. So d_s/n counts effective tokens per parameter in units of 100, and A, B, v_0 and R_T^* are the values of their terms at the reference run. The reference is a choice of units: moving it rescales these four coefficients and changes no prediction.

- **One fit per model family.** E, A, α, B, β and q belong to one shape ladder, dataset, tokenizer, optimizer and tuning recipe. Every schedule shares them, and a new family refits them. The batch is not fixed: at small U it follows the optimum of Bergsma et al. [41]. Its β_2 then moves to keep its half-life in tokens [69] (Appendix G).
- **The count of N.** We count non-embedding parameters in both U/N and the size term [18–20, 44]. Pearce and Song [70] trace part of the disagreement between Kaplan and Chinchilla to this count. The count matters at small scale: at 60M, embeddings are 55% of the parameters. Counting all parameters, or adding the output head, fits the Transformer cells about as well. The joint coupled RMSE is then 0.0014 against 0.0018, but α moves to 0.24 or 0.18.

Table 8: The fitted law beyond the grid: predicted losses along one curve of Figure 4 (a) to (c), with κ. Each block extends one axis from a grid cell (its first row); sizes above 778M count non-embedding parameters. ∞: no number of epochs lets the Transformer match.

Axis extended	Size	U	F	Transformer ↓	Feedback ↓	κ ↑
Larger model	271M	$100N$	2	2.2356	2.2266	1.22
	778M	$100N$	2	2.0116	2.0025	1.27
	3B	$100N$	2	1.7735	1.7639	1.39
	10B	$100N$	2	1.6085	1.5988	1.51
	100B	$100N$	2	1.3505	1.3410	1.83
More unique data	271M	$350N$	1	2.2095	2.1967	1.41
	271M	$1{,}000N$	1	2.1754	2.1561	2.50
	271M	$4{,}000N$	1	2.1485	2.1227	>10
	271M	$16{,}000N$	1	2.1341	2.1035	∞
More repeats	271M	$50N$	4	2.2399	2.2302	1.27
	271M	$50N$	8	2.2150	2.1999	1.97
	271M	$50N$	16	2.1997	2.1803	∞
	271M	$50N$	32	2.1926	2.1700	∞

Table 9: Forms of Equation 3 on the cells at 60M, 134M and 271M, with $E \geq 0.5$. RMSE over the cells and over the 64 equal-FLOPs gaps, in sample and leave-one-out (LOO). v(TPP): v grows as a power of TPP; v(TPP, N): also of N. Skaling equals the coupled form ($k = 1/q$); Kaplan ties q to $1/\beta$.

Form	Fit	Cells	Coefficients	RMSE ↓	LOO ↓	Gap RMS ↓	Gap LOO ↓
Transformer only							
Additive	two-phase	64	7	0.0086	0.0103	—	—
Additive	joint	64	7	0.0070	0.0080	—	—
Coupled	two-phase	64	8	0.0034	0.0037	—	—
Coupled	joint	64	8	**0.0019**	**0.0024**	—	—
Both arms							
Additive	joint	128	8	0.0079	0.0087	0.0068	0.0082
Additive, v(TPP)	joint	128	9	0.0076	0.0084	0.0050	0.0060
Coupled	joint	128	9	0.0027	0.0032	0.0035	0.0044
Coupled, v(TPP)	joint	128	10	**0.0025**	**0.0030**	0.0031	0.0039
Coupled, v(TPP, N)	joint	128	11	**0.0025**	**0.0030**	**0.0029**	**0.0037**
Skaling, v(TPP, N)	joint	128	11	**0.0025**	**0.0030**	**0.0029**	**0.0037**
Kaplan, v(TPP, N)	joint	128	10	0.0149	0.0161	0.0139	0.0148

- **The size term.** Fitting γ barely changes the in-sample error, and an F-test of γ against $\gamma = 0$ gives $p = 0.054$. Without γ, however, the law underrates latent feedback in all seven held-out 778M rows, and its gap error grows 2.5 times. These rows come from three main runs, so they give only three independent signs, not seven. We are training larger models to test the size term further.
- **The loss floor.** Three sizes do not pin it: E sits at its lower bound of 0.5 in every form. Freed from its bound, E falls to 0 to 0.3 at equal error. The exponent α then moves from 0.11 to 0.07 to 0.10.

F Derivation of the crossover

This appendix derives Equation 5, gives the exact condition for the crossover budget F^*, and shows that the law has at most one crossover.

Setup. Fix N, U and F. The Transformer has $v = 1$ and trains $R = F$ epochs; its terms are $X = A/N^\alpha$ and

$Y = B/D(F)^\beta$, and its reducible loss is $S = (X^q + Y^q)^{1/q}$. A schedule s multiplies each term by a factor:

$$X_s = X\, v_s^{-\alpha}, \qquad Y_s = Y\left(\frac{D(F)}{D(F/c_s)}\right)^\beta = Y\, c_s^{\bar\varepsilon\beta}, \qquad \bar\varepsilon = \frac{1}{\ln c_s}\int_{F/c_s}^{F} \varepsilon(R)\,\mathrm{d}\ln R. \tag{9}$$

Here $\varepsilon(R) = \mathrm{d}\ln D/\mathrm{d}\ln R$ is the elasticity of effective data to epochs. So the $\bar\varepsilon$ of Equation 5, with $D_T = D(F)$ and $D_s = D(F/c_s)$, is the mean of this elasticity over the epochs that the schedule gives up.

First order. Because $\ln S = \frac{1}{q}\ln(X^q + Y^q)$, its derivatives are $\partial\ln S/\partial\ln X = w$ and $\partial\ln S/\partial\ln Y = 1 - w$, with $w = X^q/(X^q + Y^q)$. So, to first order in $\ln v_s$ and $\ln c_s$,

$$\ln\frac{S_s}{S} \approx -\,w\,\alpha\ln v_s \;+\; (1-w)\,\bar\varepsilon\,\beta\ln c_s. \tag{10}$$

Equation 5 follows, since $L_s - L_T = S_s - S \approx S\ln(S_s/S)$.

Elasticity. From Equation 3, with $x = (R-1)/R^*$,

$$\varepsilon(R) = \frac{R\,e^{-x}}{1 + R^*(1 - e^{-x})} \;\text{ for } R \ge 1, \qquad \varepsilon(R) = 1 \;\text{ for } R \le 1. \tag{11}$$

It is 1 at one epoch and falls to 0 as repeats saturate. It never rises on the way: $\mathrm{d}\ln\varepsilon/\mathrm{d}\ln R = 1 - R/R^* - \varepsilon$, and $\varepsilon \ge e^{-x} \ge 1 - x > 1 - R/R^*$, because the denominator is at most $1 + R^*x = R$.

Exact crossover. The two losses are equal when $X_s^q + Y_s^q = X^q + Y^q$. With Equation 9, this is

$$\left(\frac{X}{Y}\right)^q = \frac{c_s^{q\bar\varepsilon\beta} - 1}{1 - v_s^{-q\alpha}}. \tag{12}$$

No approximation enters. To first order, its right side is $\bar\varepsilon\,\beta\ln c_s/(\alpha\ln v_s)$: data cost over capacity gain.

One crossover. Hold N and U and raise F. The left side of Equation 12 rises, because $D(F)$ grows, so Y falls while X stays. The right side falls: $\bar\varepsilon$ falls because ε does, and v_s does not shrink, since $\omega > 0$. Therefore the two sides meet at most once, and the schedule loses below F^* and wins above it. Moreover, the right side tends to 0 as repeats saturate while the left side stays positive. So under the law, a costlier schedule with $v_s > 1$ wins at a large enough budget.

Data per parameter. More unique data lowers Y at every F, which raises the left side and moves F^* earlier. However, it also raises R^*, which slows saturation and moves F^* later. In the fitted law the first effect wins: at 134M, F^* falls from 4.1 at $U = 12.5N$ to 1.5 at $50N$ (Figure 7).

Accuracy. Equation 12 reproduces every crossover of Figure 7. Equation 5 stays within 0.002 of the fitted law's gap for $U = 12.5N$ to $100N$ and $F = 1$ to 4, always in favour of latent feedback. So its crossover comes 7% to 12% early: 3.8 against 4.1 at 134M and $U = 12.5N$. The bias arises because the first-order ratio lies below the exact right side of Equation 12.

G Model and training configuration

All sizes share one recipe (Table 10). They differ only in shape, learning rate, weight decay, schedule lengths and, at small U, the batch.

Learning rate. The learning rate η is that of the matrix optimizer. The embedding and the output head train at $0.905\,\eta$ and $0.08\,\eta$.

Weight decay. It follows a timescale rule $\tau = r\,a\,\mathrm{TPP}^b$ in training tokens per parameter (TPP). The rule has $(a, b) = (0.403, -0.483)$ for the Transformer and $(0.751, -0.577)$ for latent feedback. We re-centre r per arm on the 134M cells.

Tuning grid. Each model gets three learning rates and two or three τ levels per cell (Table 10). We add the next point wherever the best one sits on an edge of the grid. A cell's value is the minimum of a quadratic fitted to its points if that minimum lies inside the grid, and otherwise the best point.

Decay branches. Each branch keeps its main run's η and λ, as in other warmup-stable-decay scaling studies [42]. On an earlier pool we checked this against a separate ladder, whose runs were tuned on their own grids for about two epochs. In seven cells at two sizes and 1 to 2.3 epochs, five of the Transformer and two of latent feedback, the branches match those runs within 0.0012, with a mean difference of 0.0003.

Batch. The batch keeps each cell near the optimum of Bergsma et al. [41]. It stays between 0.46 and 1.72 times that optimum at every size. At 134M and 271M it is 0.25M tokens up to $U = 25N$ and 0.5M from $50N$; 778M uses 0.5M (Table 10). Both models share each cell's batch; one 60M cell run at both batches moves the gap by 0.002. We set β_2 by its half-life in tokens [69].

Table 10: Configuration of the runs; the latent feedback Transformer adds the feedback map W (Equation 2). Learning rate: centre of the grid, Transformer / latent feedback, one step of $\sqrt{2}$ higher at $U = 12.5N$ except at 134M. Warmup doubles with each halving of the batch.

	60M	134M	271M	778M
Parameters	60,041,728	134,108,928	271,094,784	777,862,656
Non-embedding parameters	27,273,728	84,956,928	205,558,784	679,558,656
Layers	8	12	16	24
Width	512	768	1024	1536
Attention heads	4	6	8	12
MLP hidden width	1536	2048	2816	4096
U/N	12.5, 50, 100, 400	12.5 to 400	12.5 to 350	12.5, 50, 100
Main run length F, by U/N	4	4	4	4, 2, 1
Tokens per update[a]	0.125M / 0.25M	0.25M / 0.5M	0.25M / 0.5M	0.5M
β_2, AdamW and NorMuon[a]	0.987 / 0.975	0.975 / 0.95	0.975 / 0.95	0.95
Warmup updates at 0.5M	115	256	517	1,484
Latent feedback: one-pass probability p			0.15	
Latent feedback: cost c	1.1970	1.2317	1.2514	1.2677
Learning rate η	0.0252	0.0163	0.0111 / 0.0157	0.0043 / 0.0061
Tuning grid in η[b]			$\eta \times \sqrt{2}^{\pm 1}$	
Tuning grid in τ[b]	$\tau \times 2^{\pm 1}$	$\tau \times 2^{\pm 1}$	$\tau \times 2^{-1}, 2^{-2}$	—
Data	Stratified Dolma 3 (Appendix I)			
Tokenizer	Llama 2 (32,000 tokens)			
Sequence length	2048 tokens			
Epoch order	One fixed order, the same in every epoch			
Block	Pre-norm RMSNorm, rotary positions, QK-norm, SwiGLU; head width 128			
Embeddings	Untied input and output, no weight decay			
Initialization	OLMo 2 [33]			
Matrix optimizer	NorMuon, momentum 0.95, from 0.85 over the first 300 updates at 0.5M			
Embedding and head	AdamW, $\beta_1 = 0.9$			
Schedule	Warmup-stable-decay; linear decay to zero over the last 20% of each budget			
Other	Gradient clipping at 1.0; bfloat16			

[a] At low U / above: low is $U = 12.5N$ at 60M, and $12.5N$ and $25N$ at 134M and 271M. [b] Plus the next point out wherever a cell's best point sits on an edge; τ points at the centre η, and at 271M only at $U = 12.5N$ and $50N$.

Choice of p. On earlier pools, p from 0.01 to 0.31 changed the loss at equal compute by under 0.003. With

$p = 0$, the passes fell into a two-cycle, so decoding failed.

H Downstream suites

Weighting the downstream suites. Table 2 lists the downstream BPB of every cell. A plain mean over the 37 suites is dominated by yes/no and arithmetic suites, whose seed noise reaches 0.24 BPB. So we weight each suite by the inverse of its seed variance. Let $b_{r,j}$ be run r's BPB on suite j, the mean over the suite's tasks. The seed variance σ_j^2 is the variance of $b_{r,j}$ across seeds 0 to 2, averaged over 16 groups of runs at 134M (both models, $U = 12.5N$ and $100N$, $F = 1$ to 4). A run's downstream score is then

$$b_r = \frac{\sum_j w_j \, b_{r,j}}{\sum_j w_j}, \qquad w_j = \frac{1}{\sigma_j^2}. \tag{13}$$

I Nested pools

Construction. Each raw file of the Dolma 3 sample holds one of 224 source-topic buckets. We keep the first copy of each exact duplicate, a document with the same token sequence. We drop documents that share a paragraph of 13+ words with Paloma's validation or test sets [37]. These are 0.16% of the documents. We shuffle documents within each bucket. We then merge the buckets so that each advances along the stream at the rate of its token share. Every prefix is then a stratified sample of the stream, so pools differ only in size. From 1B tokens on, a prefix's bucket mixture lies within 0.003 of the whole stream's in total variation.