arXiv:2611.22351v1 [cs.LG] 17 Nov 2026

Scaling Laws for Vision-Language Fusion

Seng Pei Liew Teppei Suzuki Seiya Tokui

Noetra Corp.

Abstract

Vision-language models (VLMs) have emerged as central foundation models combining multimodal perception with text comprehension. While VLMs are traditionally built via late fusion by continued pre-training an off-the-shelf language model on multimodal (image-text) data, mixing modalities from the start of pre-training (early fusion) has recently gained traction. However, the fundamental trade-offs between these paradigms remain poorly understood, resulting in expensive and heuristic-driven training recipes. In this paper, we systematically analyze the learning dynamics of these multimodal pre-training processes, showing that validation losses during vision-language fusion can be modeled by separating the base learning trajectory, where different data types contribute at distinct rates, from a distribution shift driven by the competition between modalities. We derive a scaling law that unifiedly describes these dynamics, enabling accurate loss predictions across mixture ratios, training steps, and model sizes. Crucially, our law characterizes the trade-off between early and late fusion: under a fixed data budget, early fusion preserves text performance while late fusion favors vision performance, yielding a trade-off that increasingly tilts toward early fusion as model size scales. Extensive experiments confirm that our scaling law reliably captures multimodal learning dynamics and governs fusion decisions.

1 Introduction

Vision-language models (VLMs), combining text comprehension with visual capabilities, have seen rapid growth and adoption in recent years [3, 25, 31, 42]. Traditionally, VLMs are constructed by augmenting a pre-trained LLM with a pre-trained visual encoder (e.g., CLIP [37]), and further pre-training it on multimodal data, specifically image-text data (*late fusion*) [21, 30, 33]. More recently, the community has begun to explore training VLMs from scratch or in the early stages, by mixing all modalities in the beginning (*early fusion* or *native multimodal pre-training*) [34, 44, 49]. These two paradigms represent contrasting approaches to *fusion timing*, differing in when multimodal data are injected during pre-training.

Given the prohibitive computational cost of large-scale multimodal pre-training, accurately forecasting final performance and resource efficiency prior to execution has become essential. To this end, scaling laws [4, 12, 13, 38], originally studied at large scales in language modeling [14, 19], have recently been extended to multimodal domains, where a growing body of work has sought to characterize scaling behaviors under this paradigm, providing valuable guidelines for allocating compute, balancing model parameters, and tuning pre-training mixtures [1, 20, 49].

However, a unified scaling law capable of characterizing both early and late fusion within a single formulation remains absent, leaving practitioners unable to systematically determine which paradigm to adopt under a given compute or data budget. Existing formulations treat early and late fusion in isolation and fail to capture how cross-modal interactions and distribution shifts evolve as fusion timing varies. Consequently, deciding whether to invest expensive pre-training compute into native multimodal or continued late-fusion pre-training remains largely heuristic.

Contributions. In this paper, we bridge this gap by establishing a unified scaling law for multimodal pre-training that directly models the trade-offs across different fusion timings/strategies. At the core of our scaling law is the decomposition of validation loss into two terms: a base learning trajectory where different data types advance progress at distinct exchange rates, and an additive distribution shift governed by competition between modalities. This formulation spans a continuous spectrum of distribution shifts: in the

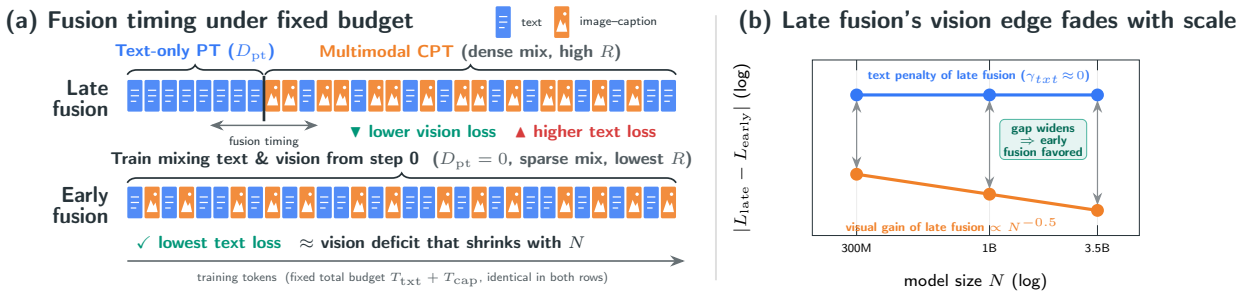


Figure 1: **Schematic of early versus late fusion in our fusion timing experiment and their scale-dependent performance trade-off. (a)** Late fusion begins with text-only pre-training before switching to dense multimodal continued pre-training (high mixture ratio R), inducing a large distribution shift that favors vision loss at the cost of text degradation. In contrast, early fusion distributes multimodal data across all steps (lowest R), avoiding distribution shift to protect text capability at the expense of visual performance. **(b)** Across model size, the text-loss penalty paid by late fusion is essentially independent of model size ($\gamma_{\text{txt}} \approx 0$), while its vision-loss advantage shrinks: $\propto N^{-0.5}$, so the trade-off tilts increasingly toward early fusion as models scale (see also Tab. 3). Combined with the observation that vision loss reduction does not translate to vision benchmark gains (see Sec. 6.2), early fusion becomes the preferred recipe at scale.

small-shift regime (*image-text interleaved documents*), cross-modal dynamics resemble standard LLM continued pre-training, whereas in the large-shift regime (*image-caption data*), cross-modal transfer diminishes sharply (Sec. 4).

Critically, our framework reveals that studying fusion timings under a fixed data budget reduces to modeling distribution shift governed by the post-fusion mixture ratio. At any given model size, this reveals an inherent trade-off: early fusion preserves text performance, whereas late fusion favors vision performance. Importantly, however, our law reveals that the amplitude of this distribution shift decays with model size (Sec. 5.1), meaning that the vision advantage of late fusion diminishes at scale, making early fusion increasingly attractive for larger models (Sec. 6.2). Figure 1 illustrates this setup and outcome.

Furthermore, we find that current vision benchmarks mostly target instruction-following models, confounding pre-training studies like ours with factors like instruction-following capabilities. We resolve this by providing new benchmarks to evaluate pre-trained models without requiring instruction-following capabilities and complication due to supervised fine-tuning, enabling a clean evaluation of the scaling law implications beyond validation losses (Sec. 6.1).

2 Related Work

Fusion and Scaling Laws. Prior fusion timing comparison studies focus on attention structures [10] and high-level empirical performance [44]; they do not provide a quantitative study of the trade-offs between early and late fusion, which is the focus of our work. Multimodal scaling laws that compare different training paradigms similar to ours are [6, 40], but they mainly compare pre-trained visual encoder with encoder-free architecture mixing data modalities altogether from the start. We instead use a pre-trained encoder, an approach that has seen adoption by popular open-weight models [22, 34, 44, 45], to study fusion timings.

Continued Pre-training. Late fusion of VLMs may be seen as a form of transfer learning or more generally two-stage learning, of which the scaling behavior in the context of LLMs' final performance has been relatively well-studied [2, 11, 28, 29, 55]. Particularly, continued pre-training, of which late fusion of VLMs is most similar to, has become standard to enhance downstream performance in math, coding, and other domains [9, 18, 41, 50, 53]. Our scaling-law approach bears similarity to this approach, but with a particular focus on multimodal data, a different regime of distribution shift than the typical domain transfer setting, and the fusion timing effects.

3 Preliminaries

Basic Notation. Throughout this work, we refer to the text-only LLM training as pre-training (PT) and multimodal continued pre-training as continued pre-training (CPT, or late fusion). Early fusion (or native multimodal pre-training) corresponds to the special case where multimodal training begins from the first step.

Control Variables and Observables. We denote the PT and CPT tokens by D_{pt} and D_{cpt}, and their corresponding step counts by S_{pt} and S_{cpt}. R represents the multimodal token fraction in the CPT data mixed with text-only data. The total model size is denoted by N. These are the control variables of our study.

The main observable is the cross-entropy validation loss L_m across modalities $m \in \{\text{txt}, \text{int}, \text{cap}\}$, representing the text, image-text interleaved data (interleaved in short), and image-caption data (caption in short) respectively. We sometimes omit the modality subscript m for notational convenience but restore it (e.g. L_m) when comparing or jointly analyzing modalities explicitly.

Fitted Parameters. When formulating our scaling laws, parameters are determined via empirical fitting unless stated otherwise. By convention, Greek letters $(\alpha, \beta, \gamma, \eta)$ denote scaling exponents, while Latin letters (L_0, A, B) denote baseline offsets and scaling coefficients. All parameters are fitted independently per evaluation modality.

Experimental Setup. We train a Gemma-3-style LLM [43] (189M to 3.5B) as the backbone of the VLM, and augment it with a SigLIP2 encoder [47] to process images. We use mainly a 300M VLM for the non-model-size scaling studies. FineWeb-Edu [36] is the text-only data source for PT; OBELICS and OmniCorpus are used as the interleaved source [24, 26]; the caption data in use are from the ReLAION dataset [39]. Full hyperparameters and architectural details are in Appendix A.

4 Scaling Vision-Language Fusion

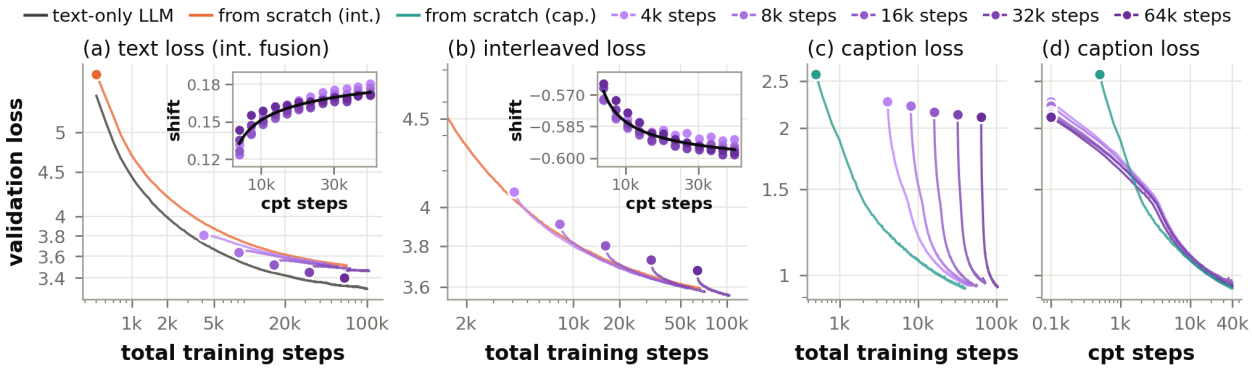


Figure 2: **Base decomposition and cross-modal exchange rates. (a,b)** Text and interleaved validation loss curves for interleaved fusion, respectively, across various fusion points D_{pt}, showing the transition towards the from-scratch trajectory. The insets illustrate the shift term ΔL growing and saturating over CPT steps independent of D_{pt}. **(c)** Caption validation loss curves with respect to total training steps across various fusion points D_{pt}. **(d)** When plotted against CPT steps, caption loss curves collapse onto each other, indicating that text tokens play dimunitive role in caption loss.

In this Section, we study the learning dynamics of multimodal pre-training, considering the fusion with interleaved and caption data *separately*, which represent different regimes of distribution shift, and make comparisons between them.

Interleaved data closely resemble the text-only web corpus used in pre-training, leading to a milder distribution shift. In practice, however, caption data are the primary source for infusing visual capabilities

into VLMs [7, 30, 33]. As captions consist of short alt-texts rather than curated documents, they introduce a pronounced distribution shift during CPT.

4.1 Base Decomposition and Exchange Rates

To establish basic understanding of how validation loss behaves before attempting to formulate the scaling laws, we perform PT on the LLM for various D_{pt}, and perform CPT using interleaved or caption data ($R = 1$) on these checkpoints. To gain understanding of the behavior, we also pre-train the VLM with interleaved/caption data from scratch and plot its curve, as well as the text-only LLM loss curve. These serve as control curves for comparison with the true CPT curves.

Observation. When multimodal data is introduced during CPT, validation loss deviates from the text pre-training trajectory before settling into a power-law descent. On interleaved data, CPT curves transition smoothly from the pre-training baseline toward the multimodal descent (Figure 2a). The *distribution shift* ΔL (loss curve minus the base power law) saturates over CPT steps with negligible dependence on D_{pt} (Figure 2a, inset), motivating the additive decomposition into a base learning trajectory and a transient shift term.

$$L = L_0 + A\,\mathcal{D}_{\text{base}}^{-\alpha} + \Delta L(D_{\text{cpt}}) \tag{1}$$

with $\mathcal{D}_{\text{base}} = D_{\text{pt}} + D_{\text{cpt}}$ the *base term*, L_0 the irreducible loss, A the scaling coefficient, and α the scaling exponent, all treated as empirical fitting parameters. These loss behaviors are consistent with those described in the LLM CPT literature (e.g., [17]).

On the other hand, caption data show different behavior. As shown in Figure 2b, increasing text pre-training D_{pt} provides negligible improvement for caption loss. When plotted against CPT steps S_{cpt} (Figure 2c), runs across different D_{pt} collapse onto the from-scratch trajectory, indicating that text tokens do not buy much caption loss improvement.

Formulation. To capture this, we modify $\mathcal{D}_{\text{base},m}$ via a *cross-modal exchange rate* $w_m \in [0, \infty)$, where learning progresses at full strength on native tokens but is modified by w_m on out-of-domain tokens. Furthermore, as in practice, multimodal CPT blends multimodal data with text at a multimodal token fraction $R \in [0, 1]$, a total CPT budget of D_{cpt} tokens is partitioned into RD_{cpt} multimodal tokens and $(1 - R)D_{\text{cpt}}$ text tokens. Summing these contributions leads us to propose the following base term for caption/interleaved and text losses:

$$\begin{aligned} \mathcal{D}_{\text{base,cap/int}} &= w_{\text{cap/int}} D_{\text{pt}} + \big(R + w_{\text{cap/int}}(1 - R)\big) D_{\text{cpt}}, \\ \mathcal{D}_{\text{base,txt}} &= D_{\text{pt}} + \big(w_{\text{txt}} R + (1 - R)\big) D_{\text{cpt}}. \end{aligned} \tag{2}$$

	Interleaved fusion		Caption fusion	
	L_{int}	L_{txt}	L_{cap}	L_{txt}
L_0	3.818	3.244	1.392	3.243
A	90.8	89.6	694	84.5
α	0.627	0.630	0.894	0.622
w	1.34	0.765	0.013	0.438
B_0	0.416	0.182	0.501	0.741
k_m	36	0.185	1.95	0.105
τ (steps)	339	7.4	1	32.8
β	0.170	0.299	0.353	0.463
fit RMSE	0.0022	0.0018	0.0099	0.0017
held-out RMSE	0.0031	0.0015	0.0084	0.0036

Table 1: Parameter fits of Equation (5). D_{pt}, D_{cpt} and τ are in steps of 2^{20} tokens.

As shown later in the full fit (Table 1), $w_{\text{int}} \approx 1.34$, confirming that text and interleaved tokens are near-substitutable in the low-shift regime. In contrast, $w_{\text{cap}} \approx 0.013$, explaining why caption learning progresses mainly on caption tokens (RD_{cpt}) rather than total steps.

4.2 Mixture Shift and Fractional Competition

We next consider more closely the distribution shift term ΔL and the multimodal fraction R. We first perform interleaved fusion with various R at a fusion point. As expected, a smaller R causes a smaller shift for text loss, and the opposite occurs for interleaved loss (Fig. 4a,b).

Formulation. Furthermore, we see that R curves are parallel to each other throughout D_{cpt}. This indicates that there is no direct R-dependency on the base term (e.g., multiplicative R factor would lead to curves with different R's converging with D_{cpt}).

The parallelism of the curves also suggests that the R-dependency can be factored out of the shift term: $\Delta L = \sigma \cdot B_0 \cdot \rho(R) \cdot \Phi(D_{\text{cpt}})$, where $\sigma \in \{+1, -1\}$ depending on the modality in consideration, B_0 a parameter to be fitted; $\Phi(D_{\text{cpt}})$ represents dependency on CPT steps, and $\rho(R)$ represents the R-dependent shift amplitude. We propose to parameterize $\rho_m(R)$ as a *fractional competition between modalities*:

$$\rho_m(R) = \frac{k_m R}{k_m R + (1 - R)}, \tag{3}$$

where k_m represents the competitive pull of multimodal tokens relative to text. Since $\rho_m(0) = 0$ and $\rho_m(1) = 1$, $B_{0,m}$ directly represents the saturated shift amplitude. Furthermore, it is concave for $k > 1$ and convex for $k < 1$, one expression covers both curvatures (text loss is convex; multimodal loss is concave).

In Fig. 4c,d, we show that this form fits well to the R-dependent shift amplitude for both caption and text losses, except at the high R regime for text loss where the text capabilites collapse. [1] As CPT without text is undesirable in practice (destroying core text capabilities), we restrict the fitting domain of the caption-fusion L_{txt} scaling law to $R \leq 0.82$, treating the collapse at $R = 1$ as the extreme boundary where the scaling law breaks down.

The Φ Term. We desire a functional form for $\Phi(D_{\text{cpt}})$ that grows from zero and saturates to one over CPT steps, capturing the buildup of the distribution shift (Fig. 2). A simple form with this behavior would be $\Phi(D_{cpt}) = 1 - (1 + D_{cpt}/\tau)^{-\beta}$, with β,τ fitting parameters, τ being the saturation time. However, to elegantly relate it to the fraction term and R, we find a mathematical analogue in chemical kinetics [23].

In the (pseudo) reversible first-order reaction $A \underset{k_b}{\overset{k_f}{\rightleftharpoons}} B$, the equilibrium fraction $k_f/(k_f + k_b)$ corresponds to our shift amplitude $B_0\rho(R)$, and the reaction rate provides a principled way to relate $\Phi(D_{\text{cpt}})$ to R. Borrowing this structure from chemistry, we propose $\Phi(D_{\text{cpt}})$ to be: [2]

$$\Phi(D_{\text{cpt}}) = 1 - \left(1 + (kR + 1 - R)\frac{D_{\text{cpt}}}{\tau}\right)^{-\beta} \tag{4}$$

4.3 The Unified Scaling Law

Assembling all the components, we have the final form of the unified scaling law:

$$L = L_0 + A\,\mathcal{D}_{\text{base}}^{-\alpha} + \sigma\, B_0\, \rho(R)\, \Phi(D_{\text{cpt}}) \tag{5}$$

where $\sigma_{\text{cap}} = \sigma_{\text{int}} = -1$, $\sigma_{\text{txt}} = +1$, and $\mathcal{D}_{\text{base}}$, ρ and Φ are defined in Equations 2, 3 and 4 respectively. Tab. 1 shows the fitted parameters.

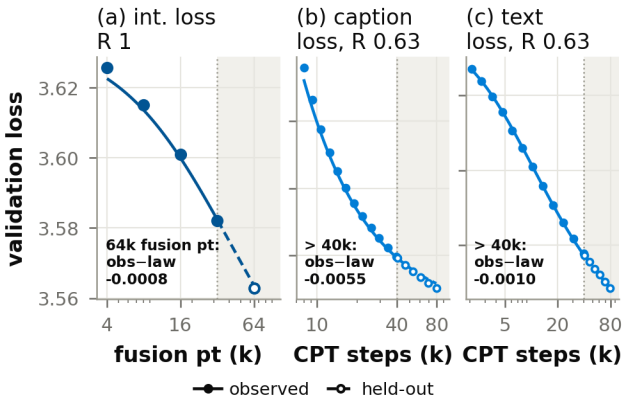


Figure 3: **Validation of the scaling laws on held-out configurations beyond the fitting domain.** The law extrapolates to later fusion points and longer CPT horizons. More results in Appendix B.

Comparison with LLM CPT. Multimodal fusion exhibits a few empirical differences from unimodal text training:

- The parameter w_m does not appear in standard LLM CPT laws [48]. While our interleaved scaling law resembles LLM CPT when $w_{\text{int}} \approx 1$, caption scaling exhibits a negligible $w_{\text{cap}} \approx 0$, capturing a breakdown in cross-modal transfer under large distribution shifts.
- Training trajectories are inherently noisier than LLM CPT, and exhibit clearly an initial transient phase during early iterations (Fig. 2d). We omit this transient period and focus on modeling the subsequent power-law regime, consistent with standard scaling law formulations [19].

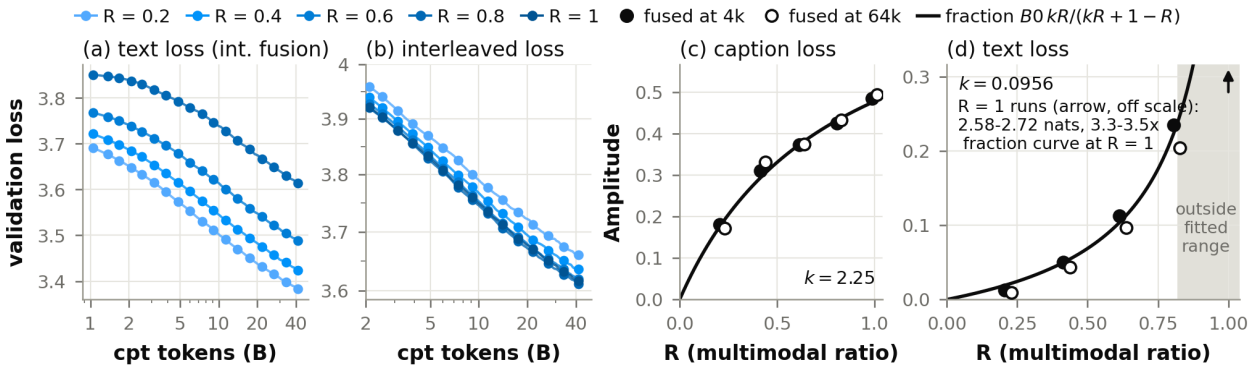


Figure 4: **Mixture Ratio Dependency. (a,b)** Text and interleaved validation loss curves for interleaved fusion varying the mixture ratio R. **(c,d)** Fitted shift amplitude as a function of R using the proposed fraction form for caption fusion, showing good fits at step 4,000 and robust validations at step 64,000.

Validation beyond the fit domain and Early Fusion. We further show that the scaling laws fit well on held-out validation data. Validation performed includes fusion extrapolated to larger D_{pt} and extrapolation of CPT to larger D_{cpt} (Fig. 3). The overall RMSE is also shown in Tab. 1. Additionally, our scaling laws apply to early fusion, i.e., $D_{\text{pt}} \to 0$. See Appendix B for more validation, extrapolation/interpolation and early fusion results.

Finding 1

Both large and small distribution shift in multimodal fusion can be described with a unified scaling law. The base term unifies distinct modalities via cross-modal exchange rates, while the shift term is governed by a cross-modal competition in the form of a fraction.

5 Law Extensions

5.1 Extension to Model Size

We next consider extending the law to include (total VLM) model size, N, where we previously held fixed. Figure 5a shows that the loss decreases monotonically with size as expected.

To describe this in the scaling law, the standard extension is incorporating a Chinchilla N term [14] to the law: $C\,N^{-\eta}$. Under this extension, the loss gap between different R's would be invariant to model size ($L(R_1, N) - L(R_2, N) = \text{constant w.r.t. } N$).

Empirical results directly refute this hypothesis (Figure 5b). On the caption axis, we see that the gap between $R = 0.4$, $R = 1$ and $R = 0.6$, $R = 1$ decreases with model size. Conversely, on the text axis, the gap remains essentially unchanged.

We therefore parametrize the shift amplitude as a power law of model size, $B_0(N) = B_0\,N^{-\gamma}$, on top of the Chinchilla extension:

$$L = L_0 + A\,\mathcal{D}_{\text{base}}^{-\alpha} + \sigma B_0\,N^{-\gamma}\,\rho(R)\,\Phi(D_{\text{cpt}}) + C\,N^{-\eta}. \tag{6}$$

Fitting this scaling law across model sizes yields $\gamma_{\text{cap}} = 0.5$, and we show that it is extrapolatable leaving out and validating the 3.5B endpoint (Figure 5c). For text loss, we find $\gamma_{\text{txt}} = 0.02$, practically independent of model size. We thus fix $\gamma_{\text{txt}} = 0$.

[1] In the same Fig. we further validate the fits at a different, larger value of the fusion point. Furthermore, in Appendix B, we compare our proposal with exponential forms [51] showing that the fits are equally well and ours fits the high R regime better.

[2] See Appendix D for the chemical kinetics analogy details.

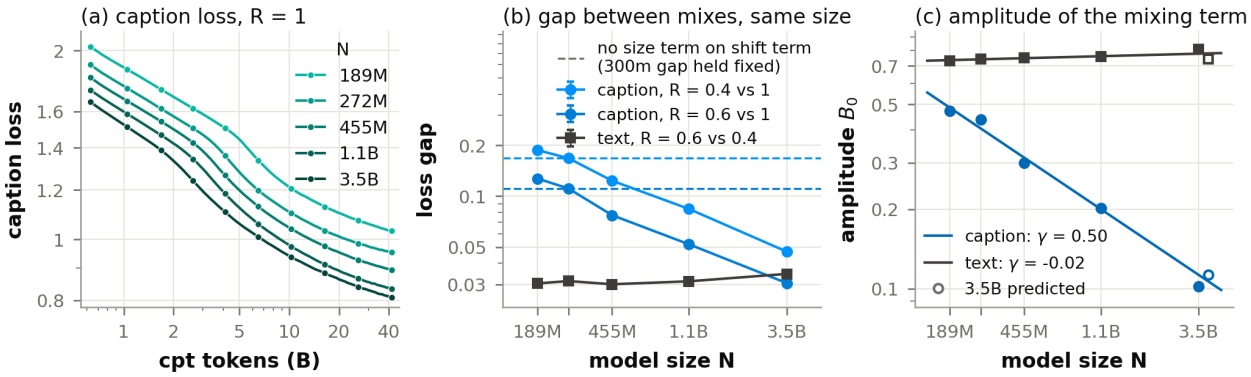


Figure 5: **Scaling laws extended to varying model sizes (N). (a)** While caption loss decreases monotonically with scale , **(b)** the gap between different data mixtures ($L(R_1) - L(R_2)$) shrinks for captions but remains constant for text. **(c)** The amplitude B_0 of the shift term follows a power-law decay with model size for captions but is size-independent for text.

Finding 2

The loss scales with an additive power-law term with respect to the model size. Furthermore, the shift term decays with model size ($N^{-0.5}$).

5.2 Extension to Fusion with All Modalities

f	Shares ($s_t/s_c/s_i$)	Caption Res. int = txt	Caption Res. w_{int}=0.18	Text Res. int = cap	Text Res. int = txt
Caption-free ($s_{cap} = 0$)					
1.00	.85 / .00 / .15	–	–	−0.012	+0.007
1.00	.40 / .00 / .60	–	–	−0.077	+0.048
Low interleaved ($f \leq 0.25$): baseline law holds					
0.22	.38 / .48 / .14	−0.016	−0.011	−0.025	+0.023
High interleaved ($f > 0.25$): exchange rate w_{int} needed					
0.33	.40 / .40 / .20	−0.029	−0.020	−0.034	+0.027
0.75	.20 / .20 / .60	−0.057	+0.014	−0.140	+0.082

Table 2: **Evaluation residuals (nats) for three-modality mixtures across interleaved share f (abbreviated).** Caption loss evaluating interleaved tokens as text versus exchange rate w_{int}=0.18. Text loss evaluating them as caption and text bounds the true loss. Full results in Appendix B.

We extend the study to training on text, caption and interleaved data at once. We consider a fixed caption/interleaved blend, with the interleaved fraction of the multimodal side denoted by f (so the ratios are $s_{int} = f\,R$, $s_{cap} = (1-f)R$ and $s_{txt} = 1 - R$.), treating it only as a small perturbation to the caption fusion scaling law, for practical reasons discussed in the beginning of Sec. 4 (interleaved data do not infuse strong visual capabilities compared to caption data).

The simplest extension is treating interleaved tokens as part of text tokens; as we have shown in Sec. 4, the distribution shift is mild compared to caption data. Simply counting the interleaved token as text, we find that the scaling law for caption loss holds for $f \lesssim 0.25$ (Tab. 6). For text loss, treating the interleaved token as caption and text respectively upper and lower bound the observed values.

The prediction of caption law can be slightly improved with the following method. Let w_{int} be the exchange rate of an interleaved token with respect to caption. We can reuse our two-source scaling law by assigning interleaved tokens an exchange rate w_{int} such that $\mathcal{D}_{base,cap}$ is modified to:

$$w_{cap}\, D_{pt} + \left[(1-f) + f\, w_{int}\right] R\, D_{cpt} + w_{cap}\,(1-R)\, D_{cpt},$$

which adds a single parameter w_{int} to the caption law. Fitting yields $w_{int} \approx 0.18$, meaning that an interleaved token is therefore worth about a fifth of a caption token for the caption loss. This yields better fitting at low and high f's (Tab. 2).

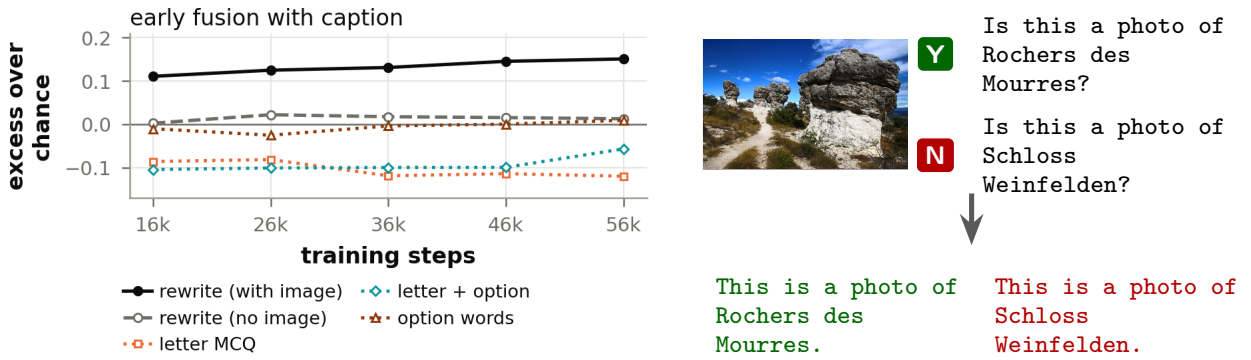


Figure 6: **(left) MME benchmark accuracies over training for from-scratch multimodal (caption) pre-training of 1b VLM.** We compare our rewriting strategy (with and without image prefix), with naive scoring with continuation formats (letter, letter + option, option words). Our strategy maintains a positive visual grounding gap (scores with image prefix higher than without) and steady learning progression, while naive formats collapse to chance or produce erratic scores due to prior biases. More results in Appendix E. **(right) Illustration of the rewriting strategy used for downstream evaluation of base VLMs.** Multiple-choice questions are converted into statements conditioned on the image prefix to enable log-likelihood ranking and avoid answer-prior collapse.

Finding 3

The caption scaling law extends to three-modality mixtures by treating interleaved data as text at low shares ($f \leq 0.25$), or by assigning it an effective exchange rate ($w_{\text{int}} = 0.18$) as the interleaved fraction scales higher ($f > 0.25$).

6 Early/Late Fusion Trade-offs

6.1 Base VLM Evaluation

In addition to validation loss, we seek to evaluate the multimodal pre-trained models via downstream vision benchmarks. In text-only language modeling, pre-trained only base checkpoints are evaluated without requiring instruction-following (IF) abilities by measuring the cloze-style (continuation) reference answer's perplexity or ranking candidate sentences (LAMBADA [35], HellaSwag [54]).

In multimodal learning, however, standard visual question answering (VQA) and reasoning benchmarks (MME [8], GQA [15], MMMU [52]) are predominantly designed for IF models. Base VLMs cannot be evaluated directly on them, and supervised fine-tuning (SFT) is not an appropriate remedy either, as it introduces substantial confounders (SFT data mixture, format bias, new vision-text alignment, etc.) that obscure the pre-training fusion dynamics under study.

To evaluate base models without IF, an intuitive alternative is naive text continuation perplexity (evaluating the reference answer's perplexity given an image and question prompt). This works partially for DocVQA [32] but fails on multiple-choice and binary tasks (MME, GQA); the model's output is dominated by drifting answer priors (e.g., alternating between predicting exclusively "Yes" and "No", or collapsing onto a single option letter) rather than visual understanding, yielding ungradable outputs or chance-level accuracy (Fig. 6).

Our Rewriting Strategy. To address this, we map each question-candidate pair into a declarative assertion conditioned on the image prefix (Fig. 6), and score them via byte-normalized log-likelihood ranking. This removes format-following demands while providing sufficient context to eliminate answer-prior collapse. We find that under this protocol, models evaluated on MME and GQA demonstrate visual learning above-chance scoring along the training trajectory while maintaining positive visual grounding gap (Fig. 6). See Appendix E for details of the rewriting strategy and other results.

Finding 4

Converting multiple-choice benchmarks to log-likelihood ranking benchmarks helps avoid answer-prior collapse and extract the visual capabilities of a base VLM without requiring instruction-following tuning.

6.2 Optimizing Fusion Timing

Armed with scaling laws and tools developed thus far, we here analyze the trade-offs across different fusiong timings. Specifically, we are interested in how the fusion timing D_{pt} affects the performance of the model under a fixed budget of text and caption tokens (T_{txt}, T_{cap} respectively). Formally, given the budget, we want to find the optimal fusion timing parameterized by the text-only pre-training duration $D_{\text{pt}} \in [0, T_{\text{txt}}]$ that minimizes the caption loss L_{cap} and/or text loss L_{txt}:

$$\arg\min_{D_{\text{pt}} \in [0, T_{\text{txt}}]} L_{\text{cap}}(D_{\text{pt}}) \quad \text{and/or} \quad \arg\min_{D_{\text{pt}} \in [0, T_{\text{txt}}]} L_{\text{txt}}(D_{\text{pt}})$$

Under fixed token budgets $(T_{\text{txt}}, T_{\text{cap}})$, the base term simplifies to $\mathcal{D}_{\text{base,cap}} = w_{\text{cap}} T_{\text{txt}} + T_{\text{cap}}$ (and analogously for text), remaining constant regardless of fusion timing D_{pt}. Therefore, the trade-off is determined solely by the shift term.

As proven in Appendix F, the caption shift magnitude $|\Delta L_{\text{cap}}|$ (and analogously for text) is strictly monotonically increasing with D_{pt}, pushing optimal solutions to the edge:

$$\arg\min_{D_{\text{pt}}} L_{\text{cap}} = T_{\text{txt}} \implies D^*_{\text{cpt}} = T_{\text{cap}},\ R^* = 1 \tag{7}$$

$$\arg\min_{D_{\text{pt}}} L_{\text{txt}} = 0 \implies D^*_{\text{cpt}} = T_{\text{txt}} + T_{\text{cap}},\ R^* = \frac{T_{\text{cap}}}{T_{\text{txt}} + T_{\text{cap}}} \tag{8}$$

In other words, late fusion (D_{cpt} shortest and R densest) favors lower caption loss while early fusion (D_{cpt} longest and R sparsest) favors lower text loss.

We further note that $B_{0,\text{cap}}(N)$ decays with model size (Sec. 5.1). This means that the late fusion advantage, captured by the shift term, diminishes with larger model size, and early fusion becomes more favorable for larger models.

				Val Loss (nats)		Text Benchmarks		Vision Benchmarks		
Scale	**Schedule**	D_{pt}	R	**Caption ↓**	**Text ↓**	**HSwag ↑**	**LAMB ↓**	**MME ↑**	**GQA ↑**	**DocVQA ↓**
300M	S0 (Early)	0	0.45	1.1488	3.4095	0.3691	69.8	0.6701	0.4103	3.88
	S1	4k	0.48	1.1418	3.4194	0.3668	73.0	0.6796	0.3694	3.73
	S3	16k	0.63	1.0983	3.4574	0.3613	89.3	0.6779	0.3891	4.13
	S4	24k	0.78	1.0486	3.5341	0.3591	80.0	0.6883	0.3969	3.78
	S5 (Late)	32k	1.00	0.9902	6.3135	0.2938	2466.4	0.6770	0.3869	3.95
	Spearman r_s			–	–	−0.94	+0.94	−0.20	+0.37	−0.31
1B	S0 (Early)	0	0.45	0.9659	3.0258	0.4862	27.0	0.7064	0.4134	3.10
	S1	4k	0.48	0.9600	3.0360	0.4795	27.3	0.7055	0.4236	2.91
	S3	16k	0.63	0.9304	3.0791	0.4702	24.8	0.7116	0.4158	2.94
	S4	24k	0.78	0.9029	3.1589	0.4569	32.3	0.7090	0.4159	2.67
	S5 (Late)	32k	1.00	0.8705	5.8537	0.3144	557.6	0.6839	0.4209	2.80
	Spearman r_s			–	–	−1.00	+0.70	+0.20	−0.40	+0.80

Table 3: **Validation losses and downstream benchmark performance across fusion timings (abbreviated).** Caption loss decreases with delayed fusion, while text loss increases. Text benchmarks closely track validation text loss (collapsing at S5, $R = 1.0$), whereas downstream vision benchmarks fail to track validation caption loss. Full results in Appendix E.

Empirical Evidence. We run fusion timing experiments to test if the above conclusion holds in practice. We fix the total budget to 56B tokens, and sweep the fusion timing, including early fusion (labeled S0 to S5),

across 300M and 1B models. Besides measuring the validation losses, we evaluate the models' downstream performance using methods proposed in Sec. 6.1.

Tab. 3 shows the main results. Validation loss behaviors conform to our predictions, as caption (text) loss decreases (increases) with delayed fusion timing. Moreover, we notice that the caption loss gap between early and late fusion (S0-S5) narrows from 0.159 to 0.095 (40% reduction), while the text loss gap changes mildly, from 2.904 to 2.828 (2.6% relative change), which are in accordance to our law prediction. Notably, the vanishing absolute loss gap also translates to a vanishing *token-equivalent* gap: since the base learning trajectory (main driver of loss decrease via extra tokens) is independent of model size (Eq. (5)), the extra caption token budget required by early fusion to catch up with late fusion's vision performance drops with model size from $\sim$ 39B tokens (300M model) to $\sim$ 23B (1B model) in our example. On the other hand, as the text shift amplitude is independent of model size, the number of tokens required by late fusion to match early fusion's text loss remains constant at any scale.

For downstream evaluation, results show a notable divergence between text and vision benchmarks. On text benchmarks, downstream capabilities closely reflect validation text loss. Across both 300M and 1B scales, accuracy benchmarks correlate strongly and negatively with text loss (see the Spearman's rank correlation coefficients in Tab. 3). The LAMBADA perplexity correlates positively with text loss, and notably explodes by orders of magnitude at S5 ($R = 1$), in line with our characterization of the $R = 1$ collapse regime in Section 4.

On the other hand, downstream performance of vision benchmarks fails to track validation caption loss. MME has little correlation with caption loss, and all three benchmarks' correlations are inconsistent across scale. We do observe that benchmark scores improve when both text and caption losses decrease simultaneously across scale, indicating the loss is still a valid performance proxy only when both modalities are jointly improved. One possible explanation for this is that VQA requires both language and vision capabilities, and optimizing one of them while degrading the other may not be optimal for downstream performance. Taken together, these findings strongly reinforce early fusion as the preferred recipe as late fusion penalizes language capabilities without delivering tangible downstream vision gains.

Finding 5

Studying fusion timing under a fixed budget reduces to analyzing the shift term, where in theory and in practice, early fusion minimizes text loss while late fusion minimizes vision loss, with the vision advantage of late fusion fading with model size. Since text benchmarks mirror the text loss behavior while vision benchmark scores do not reflect caption loss improvement, early fusion is overall the preferred pre-training recipe.

7 Conclusion

In this paper, we have presented scaling laws for vision-language fusion, by decomposing the learning dynamics into a base modality progress and an additive competitive shift, unifying both low-shift (interleaved) and high-shift (caption) regimes. The scaling laws also reveal that early fusion minimizes text loss while late fusion minimizes vision loss. However, as the late-fusion vision gains decay with model size and fails to translate to downstream gains, early fusion proves superior in practice.

Limitations and Future Work. As in standard scaling law literature, the exact fitted parameters depend on design choices (data, tokenization, etc.) that may not transfer; thus, our formulation may better be seen as a framework for practitioners to characterize and fit on their own training rather than rigid constants. Additionally, our formulation omits regimes not viable in practice, i.e., extreme caption ratios ($R \to 1$) that cause language collapse (Sec. 4.2). Finally, scaling to larger models and broader modalities remains exciting future work.

Acknowledgement

This research was conducted by Noetra Corp. under FRONTia, a Japanese national program led by the Ministry of Economy, Trade and Industry of Japan (METI) and the New Energy and Industrial Technology

Development Organization (NEDO), aimed at developing a domestic multimodal foundation model for AI robots and physical AI.

References

[1] Armen Aghajanyan, Lili Yu, Alexis Conneau, Wei-Ning Hsu, Karen Hambardzumyan, Susan Zhang, Stephen Roller, Naman Goyal, Omer Levy, and Luke Zettlemoyer. Scaling laws for generative mixed-modal language models. In *International Conference on Machine Learning*, pages 265–279. PMLR, 2023.

[2] Kosuke Akimoto, Taiki Miyagawa, and Masafumi Oyamada. $m3$ scaling law: Optimizing multi-epoch, multi-lingual, and multi-stage training for low-resource language models. *arXiv preprint arXiv:2410.12325*, 2024.

[3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022.

[4] Shun-ichi Amari, Naotake Fujita, and Shigeru Shinomoto. Four types of learning curves. *Neural Computation*, 4(4):605–618, 1992. doi: 10.1162/neco.1992.4.4.605.

[5] Tamay Besiroglu, Ege Erdil, Matthew Barnett, and Josh You. Chinchilla scaling: A replication attempt. *arXiv preprint arXiv:2404.10102*, 2024.

[6] Lin Chen, Bolin Ni, Qi Yang, Lan Jiang, Kun Ding, Xiaoran Fan, Hower Yang, Ying Wang, and Shiming Xiang. How far are we from removing the visual encoder? scaling laws for encoder-free multimodal pretraining, 2026. URL https://arxiv.org/abs/2609.35457.

[7] Matteo Farina, Vishaal Udandarao, Thao Nguyen, Selim Kuzucu, Maximilian Böther, Andreas Hochlehnert, Adhiraj Ghosh, Marianna Nezhurina, Karsten Roth, Joschka Struber, Yuhui Zhang, Sebastian Dziadzio, Elaine Sui, Soumya Jahagirdar, Dhruba Ghosh, Hasan Hammoud, Thomas De Min, Simone Caldarella, Jehanzeb Mirza, Sedrick Keh, Mehdi Cherti, Hilde Kuehne, Bernt Schiele, Serena Yeung-Levy, Muhammad Ferjad Naeem, Federico Tombari, Ana Klimovic, Elisa Ricci, Matthias Bethge, Sewoong Oh, Ameya Prabhu, Alessio Tonioni, Jenia Jitsev, Massimiliano Mancini, Ludwig Schmidt, and Nikhil Parthasarathy. Datacomp-vlm: Improved open datasets for vision-language models, 2026. URL https://arxiv.org/abs/2606.28551.

[8] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *Advances in Neural Information Processing Systems*, 38, 2026.

[9] Zheng Gong, Kun Zhou, Xin Zhao, Jing Sha, Shijin Wang, and Ji-Rong Wen. Continual pre-training of language models for math problem understanding with syntax-aware memory network, 2022. URL https://aclanthology.org/2022.acl-long.408/.

[10] Junlin Han, Shengbang Tong, David Fan, Minghao Chen, Philip Torr, Filippos Kokkinos, and Mike Lewis. Towards physics of multimodal pretraining: Knowledge flow, modality synergy, early unification, and recipes, 2026. URL https://arxiv.org/abs/2608.05000.

[11] Danny Hernandez, Jared Kaplan, Tom Henighan, and Sam McCandlish. Scaling laws for transfer. *arXiv preprint arXiv:2102.01293*, 2021.

[12] Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409*, 2017.

[13] Joel Hestness, Newsha Ardalani, and Gregory Diamos. Beyond human-level accuracy: Computational challenges in deep learning. pages 1–14, 2019.

[14] 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. *Proceedings of the 36th International Conference on Neural Information Processing Systems*, pages 30016–30030, 2022.

[15] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6693–6702. IEEE, 2019.

[16] Marcus Hutter. Learning curve theory. *arXiv preprint arXiv:2102.04074*, 2021.

[17] Adam Ibrahim, Benjamin Thérien, Kshitij Gupta, Mats Leon Richter, Quentin Gregory Anthony, Eugene Belilovsky, Timothée Lesort, and Irina Rish. Simple and scalable strategies to continually pre-train large language models. *Transactions on Machine Learning Research*.

[18] Joel Jang, Seonghyeon Ye, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun Kim, Stanley Jungkyu Choi, and Minjoon Seo. Towards continual knowledge learning of language models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL https://openreview.net/forum?id=vfsRB5MImo9.

[19] 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 preprint arXiv:2001.08361*, 2020.

[20] Aditi Khandelwal, Ayush K Tarun, Yixuan Xu, Imanol Schlag, Golnoosh Farnadi, Siva Reddy, Luke Zettlemoyer, and Antoine Bosselut. Scaling laws for multimodal data mixtures. *arXiv preprint arXiv:2610.14423*, 2026.

[21] Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. Grounding language models to images for multimodal generation. *arXiv preprint arXiv:2301.13823*, 2023.

[22] Thinking Machines Lab. Introducing inkling 2 small, 2026. URL https://thinkingmachines.ai/news/inkling-2-small/.

[23] Keith J. Laidler. *Chemical Kinetics*. Harper & Row, New York, 3rd edition, 1987.

[24] Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander M. Rush, Douwe Kiela, Matthieu Cord, and Victor Sanh. OBELICS: An open web-scale filtered dataset of interleaved image-text documents. In *Advances in Neural Information Processing Systems*, volume 36, 2024.

[25] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.

[26] Qingyun Li, Zhe Chen, Weiyun Wang, et al. Omnicorpus: A unified multimodal corpus of 10 billion-level images interleaved with text. 2024.

[27] Wanchao Liang, Tianyu Liu, Less Wright, Will Constable, Andrew Gu, Chien-Chin Huang, Iris Zhang, Wei Feng, Howard Huang, Junjie Wang, et al. Torchtitan: One-stop pytorch native solution for production ready llm pretraining. In *International Conference on Learning Representations*, volume 2025, pages 92161–92181, 2025.

[28] Seng Pei Liew and Takuya Kato. Reusing overtrained language models saturates scaling. *arXiv preprint arXiv:2510.06548*, 2025.

[29] Seng Pei Liew, Takuya Kato, and Sho Takase. Scaling laws for upcycling mixture-of-experts language models. *arXiv preprint arXiv:2502.03009*, 2025.

[30] Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. Vila: On pretraining for visual language models. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26679–26689. IEEE, 2024.

[31] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. 2023.

[32] Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Docvqa: A dataset for vqa on document images. In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2199–2208. IEEE, 2021.

[33] Brandon McKinzie, Zhe Gan, Jean-Philippe Fauconnier, Sam Dodge, Bowen Zhang, Philipp Dufter, Dhruti Shah, Xianzhi Du, Futang Peng, Floris Weers, Anton Belyi, Haotian Zhang, Karanjeet Singh, Doug Kang, Ankur Jain, Hongyu Hè, Max Schwarzer, Tom Gunter, Xiang Kong, Aonan Zhang, Jianyu Wang, Chong Wang, Nan Du, Tao Lei, Sam Wiseman, Guoli Yin, Mark Lee, Zirui Wang, Ruoming Pang, Peter Grasch, Alexander Toshev, and Yinfei Yang. Mm1: Methods, analysis & insights from multimodal llm pre-training, 2024. URL https://arxiv.org/abs/2403.09611.

[34] Mistral AI, Abdelaziz Bounhar, Abhav Mehrotra, et al. Mistral large 4, 2026. URL https://www.alphaxiv.org/abs/2611.mistral-large-4.

[35] Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. The LAMBADA dataset: Word prediction requiring a broad discourse context. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1525–1534, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1144. URL https://aclanthology.org/P16-1144.

[36] Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. *arXiv preprint arXiv:2306.01116*, 2023.

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

[38] Jonathan S Rosenfeld, Amir Rosenfeld, Yonatan Belinkov, and Nir Shavit. A constructive prediction of the generalization error across scales.

[39] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. Laion-5b: An open large-scale dataset for training next generation image-text models. 2022.

[40] Mustafa Shukor, Enrico Fini, Victor Guilherme Turrisi da Costa, Matthieu Cord, Joshua Susskind, and Alaaeldin El-Nouby. Scaling laws for native multimodal models. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12–23. IEEE, 2025.

[41] Yu Sun, Shuohuan Wang, Yu-Kun Li, Shikun Feng, Hao Tian, Hua Wu, and Haifeng Wang. ERNIE 2.0: A continual pre-training framework for language understanding. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8968–8975. AAAI Press, 2020. URL https://doi.org/10.1609/aaai.v34i05.6428.

[42] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

[43] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.

[44] Kimi Team, Tongtong Bai, Yifan Bai, Yiping Bao, SH Cai, Yuan Cao, Ziwei Chai, Y Charles, HS Che, Cheng Chen, et al. Kimi k2. 5: Visual agentic intelligence. *arXiv preprint arXiv:2602.02276*, 2026.

[45] Qwen Team. Qwen3. 8-omni: Towards native omni-modal agents. *arXiv preprint arXiv:2609.25611*, 2026.

[46] Howe Tissue, Venus Wang, and Lu Wang. Scaling law with learning rate annealing, 2024. *URL https://arxiv. org/abs/2408.11029*, 3.

[47] Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, et al. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*, 2025.

[48] Xingjin Wang, Howe Tissue, Lu Wang, Linjing Li, and Daniel Dajun Zeng. Learning dynamics in continual pre-training for large language models, 2025. URL https://arxiv.org/abs/2505.07796.

[49] Haoyuan Wu, Aoqi Wu, Hai Wang, Jiajia Wu, Jinxiang Ou, and Bei Yu. Scaling native multimodal pre-training from scratch. *arXiv preprint arXiv:2607.22043*, 2026.

[50] Prateek Yadav, Qing Sun, Hantian Ding, Xiaopeng Li, Dejiao Zhang, Ming Tan, Parminder Bhatia, Xiaofei Ma, Ramesh Nallapati, Murali Krishna Ramanathan, Mohit Bansal, and Bing Xiang. Exploring continual learning for code generation models. In Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 782–792. Association for Computational Linguistics, 2023. URL https://doi.org/10.18653/v1/2023.acl-short.68.

[51] Jiasheng Ye, Peiju Liu, Tianxiang Sun, Jun Zhan, Yunhua Zhou, and Xipeng Qiu. Data mixing laws: Optimizing data mixtures by predicting language modeling performance, 2025. URL https://arxiv.org/abs/2403.16952.

[52] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9556–9567, 2024.

[53] Daoguang Zan, Bei Chen, Dejian Yang, Zeqi Lin, Minsu Kim, Bei Guan, Yongji Wang, Weizhu Chen, and Jian-Guang Lou. CERT: continual pre-training on sketches for library-oriented code generation. In Luc De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 2369–2375. ijcai.org, 2022. URL https://doi.org/10.24963/ijcai.2022/329.

[54] Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence? *CoRR*, abs/1905.07830, 2019. URL http://arxiv.org/abs/1905.07830.

[55] Biao Zhang, Zhongtao Liu, Colin Cherry, and Orhan Firat. When scaling meets llm finetuning: The effect of data, model and finetuning method.

A Setup Details

Label	Dim	FFN	LLM	Non-Emb	VLM
200M	400	1280	92.8M	49.8M	189.3M
300M	640	2048	175.0M	106.2M	272.2M
500M	1040	3328	356.3M	244.5M	454.7M
1B	2000	6400	1017.0M	801.9M	1118.4M
3.5B	4000	12800	3416.4M	2986.3M	3523.9M

Table 4: **Model architecture across scales used in the paper.** Columns give LLM parameters, non-embedding parameters, and total VLM parameters (including the vision encoder and projector).

We describe in more detail the training setup of our study.

Model. We adopt Gemma-3 as the base architecture for our study. The number of layers is fixed to 18; head dimension is 256 with 4 query and 2 KV heads. We scale the model size via its feedforward neural network (FFN) layer width (Figure 4). The embedding is tied, and the vocabulary size is 107,520. The SigLIP2 ViT-B encoder in use converts each 16×16 patch to an embedding.

Data. For text-only pre-training, we use the original FineWeb-Edu corpus without modification. For interleaved data (OBELICS and OmniCorpus), we apply text-quality filtering following the FineWeb-Edu classification heuristics to match FineWeb-Edu corpus quality. For image-caption data (ReLAION), we rewrite the alt-text annotations to produce higher-quality captions.

Hyperparameters and Training Details. We summarize the hyperparameters used for training and other training details in Table 5.

We train with mixed precision and FSDP using a codebase based on `torchtitan` [27]. Additionally, sequence packing is performed to retain token processing efficiency.

During CPT, all components' neural network parameters are unfrozen. We perform warmup only in the beginning of PT and use a constant learning rate throughout the remaining PT and CPT phases. See Appendix C for discussion of learning rate annealing.

Fitting. We mainly follow the procedure described in [5] to perform fitting. Particularly, the logarithm of loss is fitted with Huber loss with $\delta = 10^{-3}$. L-BFGS-B is used to optimize the objective. Instead of grid search, we initialize the coefficients randomly a few hundred times and run L-BFGS-B to optimize the objective.

B More Results

B.1 More on Scaling Law Formulation

More on the Cross-Modal Exchange Rate.

Hyperparameter	Value
Peak Learning Rate	4×10^{-4}
Learning Rate Schedule (PT)	500-step linear warmup, then flat
Learning Rate Schedule (CPT)	Flat from step 1 (no warmup)
Batch Size	$\approx$ 1.05M tokens (256 seq $\times$ 4096)
Context Length	4,096 tokens
Optimizer	AdamW
Adam β_1, β_2	0.9, 0.95
Adam ϵ	10^{-8}
Weight Decay	0.1
Gradient Clipping Norm	1.0

Table 5: **Training hyperparameters for the language (PT) and vision-language continued pre-training (CPT) pipelines.** These values remain fixed across all evaluated model sizes.

We give more justification for the use of RD_{cpt} as the effective caption token count in the scaling law formulation. Note that the loss curve in Fig. 7b (same as Fig. 2c, repeated for convenience) has a bend transitioning to the power-law regime, where the the law starts describing the learning. If text tokens after fusion advanced learning as much as caption tokens do, runs with different R would reach this bend after a similar number of CPT steps.

However, as shown in Fig. 7c, the bend spreads from around 4,000 steps to 30,000 steps. When we count the steps in caption tokens, $R\,D_{\text{cpt}}$, as in the fourth panel in Fig. 7d, the spread is much smaller (4,000 to 6,000 caption-only steps). Furthermore,

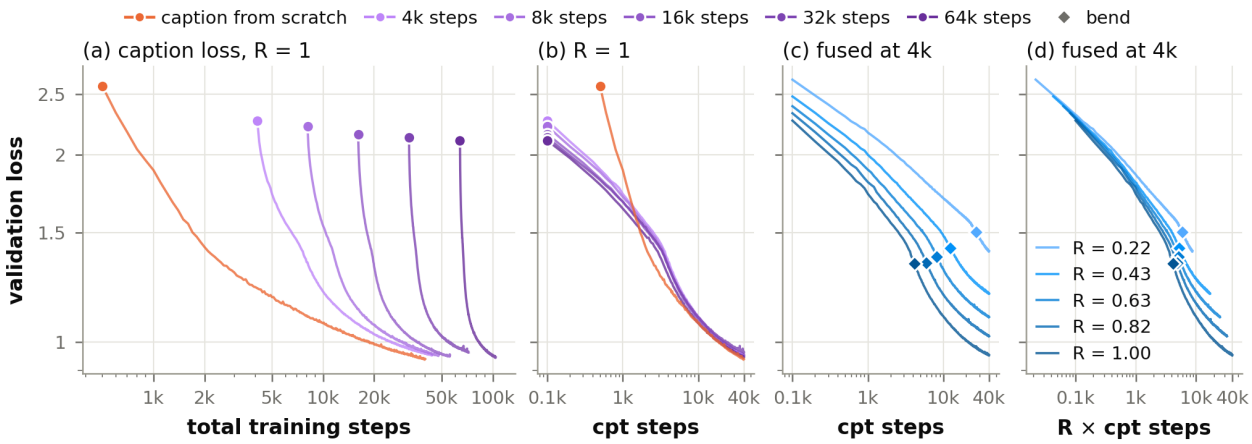


Figure 7: **Caption validation loss dynamics across training stages.** The two panels on the left show that text pre-training provides negligible benefit to caption loss. The two panels on the right show that the loss gap across varying mixture ratios (R) remains almost constant when measured by effective caption tokens ($R\times$ CPT steps) rather than raw CPT steps.

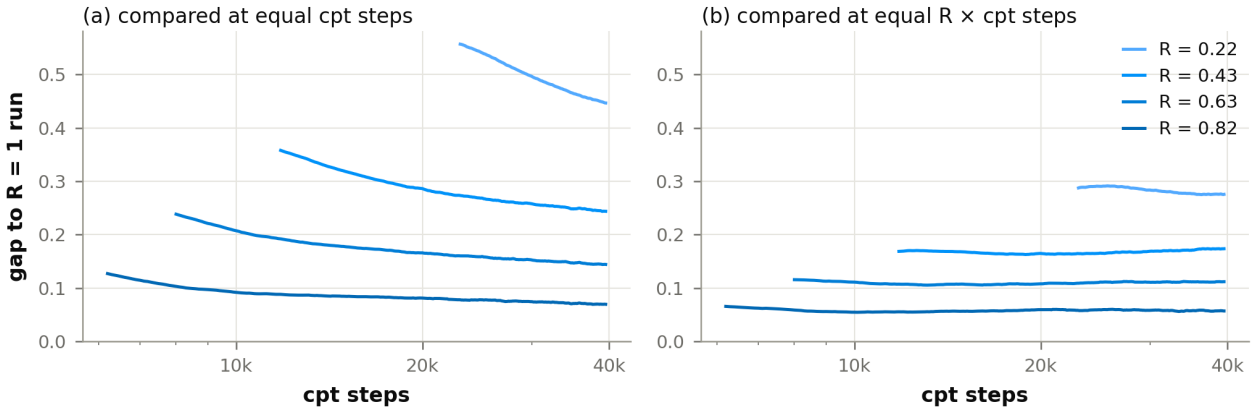


Figure 8: **The validation loss gap between mixed runs ($R < 1$) and the multimodal-only run ($R = 1$).** The gap shrinks when compared at equal CPT steps (left) but remains nearly constant when aligned by effective caption tokens ($R\times$ CPT steps) (right).

the gap of the loss curves among different R's stay almost constant at equal $R\,D_{\text{cpt}}$, whereas the gap shrinks at equal CPT steps (see Fig. 8 for comparison). The former observation is consistent with what we observe with the interleaved data in Section 4.

More on the Shift Term. The R-dependent $\rho(R)$ can be parameterized with exponential functions, following previous studies on data mixtures [20, 51]:

$$\rho_{\text{txt}}(R) = e^{a_{\text{txt}}R} - 1; \quad \rho_{\text{int}}(R) = 1 - e^{-a_{\text{int}}R} \tag{9}$$

with a_{txt} and a_{int} as fitting parameters. Note that the functional form is designed such that the law falls back to the base term without shift when $R = 0$.

Besides being more compact, our proposed fraction form also fits at least as well and generalizes better across data mixture regimes. The difference is most prominent at the edge of data mixture. As in Figure 9, exponential parameterization cannot fit simultaneously the mixed ($R < 1$) and unmixed runs ($R = 1$). In contrast, the fraction form tracks both regimes with better fidelity.

Nevertheless, incorporating the $R = 1$ collapse regime noticeably inflates residual error across the interior mixed runs. As CPT without text is undesirable in practice (destroying core text capabilities), we restrict

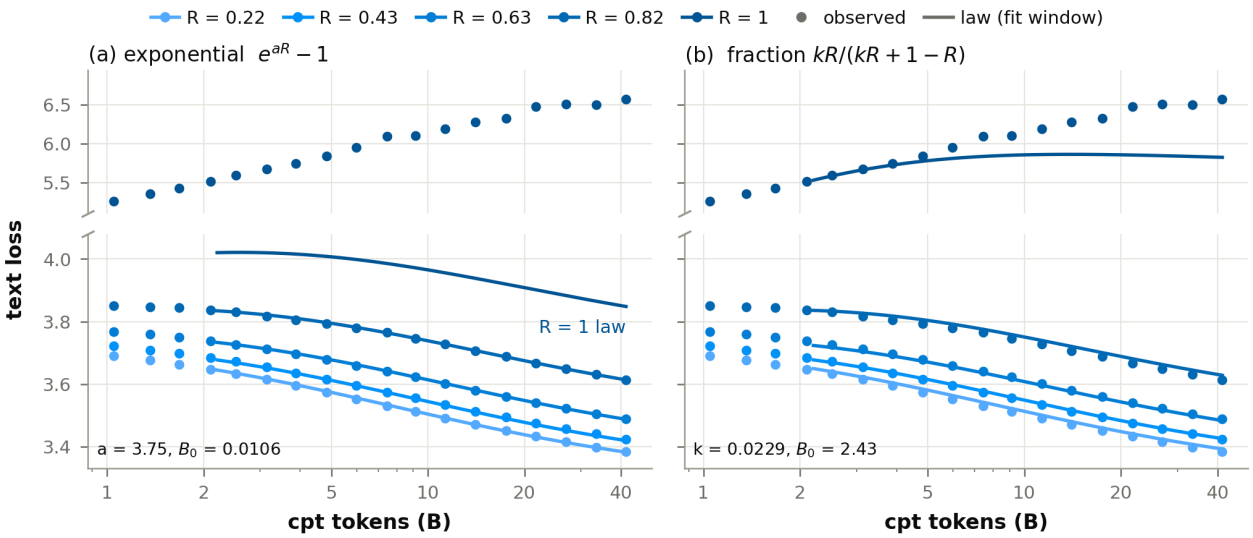


Figure 9: **Comparison of the exponential and fraction formulations for modeling the mixture ratio (R) dependency of the text loss shift term.** The fraction form $kR/(kR + 1 - R)$ better tracks both the mixed ($R < 1$) and unmixed ($R = 1$) training regimes.

the fitting domain of the text scaling law to $R \leq 0.82$, treating the collapse at $R = 1$ as the extreme boundary where the scaling law breaks down.

Within this range, the fraction form matches or outperforms the exponential form at describing the R dependency, while unifying all three evaluation axes under a single parameter $k = 22.7$, 2.25 and 0.0956 for interleaved, caption and text loss respectively (Fig. 10).

B.2 Validation Results, Early Fusion, and Three-Modality Mixtures

We run held-out validation to test our scaling laws rigorously, varying $D_{\text{pt}}, R, D_{\text{cpt}}$. The runs are:

- Interleaved held out evaluation `R0.40/4k`: validating interpolation at $R = 0.4$
- Interleaved held out evaluation `R1.00/64k`: validating extrapolation at $D_{\text{pt}} = 64{,}000$ steps.
- Text-caption held out evaluation `H1`: joint fusion and ratio interpolation at $D_{\text{pt}} = 12{,}000$ steps and $R = 0.722$
- Text-caption held out evaluation `H2`: late-fusion interpolation at $D_{\text{pt}} = 48{,}000$ steps with $R = 0.527$
- Text-caption held out evaluation `H3`: CPT extrapolation at $D_{\text{pt}} = 4{,}000$ steps with $R = 0.63$ and $D_{\text{cpt}} = 80{,}000$ steps.

The results are shown in Figs. 11 and 12. The held-out RMSE's reported in Table 1 are the average over these validation runs.

Early Fusion Results. Early fusion validation results are summarized in Figure 13. After the transient period (around 10,000 CPT steps), the early fusion runs ($D_{\text{pt}} = 0$) are well predicted by the scaling law.

More Results on the Three-Modality Mixtures. We give the full results of the three-modality mixtures in Table 6.

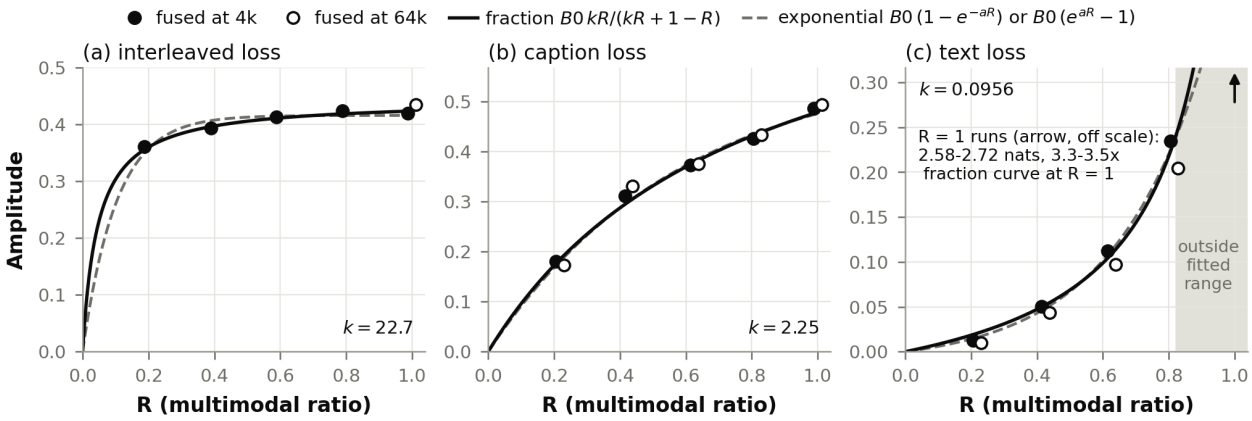


Figure 10: **Fitted shift amplitude as a function of the multimodal mixture ratio (R) using exponential and fraction forms.** The single fraction form unifiedly captures the concave curvature for interleaved/caption losses and the convex curvature for text loss using a single parameter k_m.

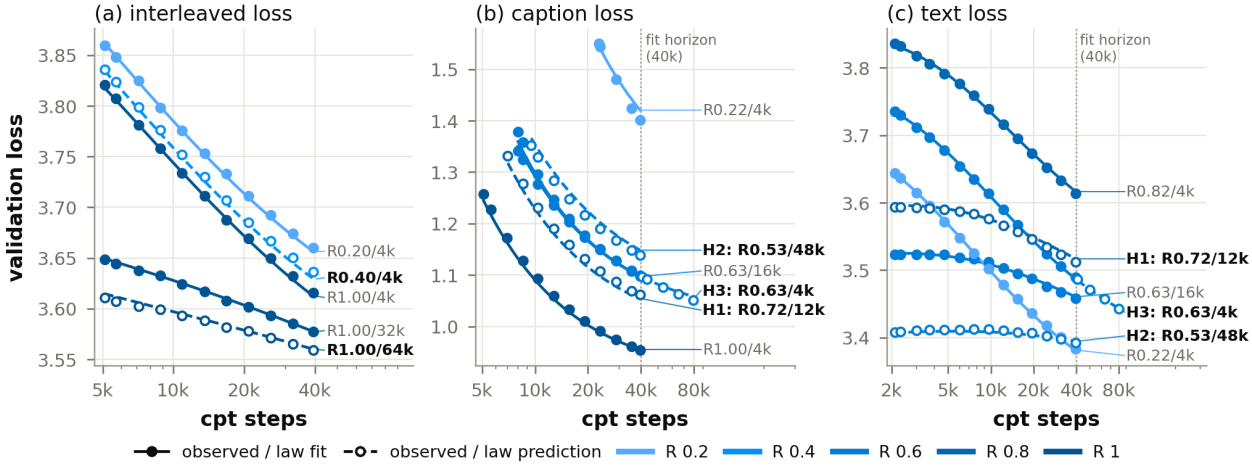


Figure 11: **Validation loss trajectories for both fitted runs and held-out configurations (extrapolated fusion points and CPT horizons) across interleaved, caption, and text modalities.**

C Extension to Learning Rate Annealing

The law was thus far fitted at a flat learning rate. Here, we extend it to describe learning rate decay/annealing dynamics. The result is an additive term to the scaling law using the annealing area S_2^{cpt} [46, 48]:

$$\begin{aligned} L &= L^{\text{flat}} - HN^{-\zeta} g(R)\, S_2^{\text{cpt}}, \\ S_2^{\text{cpt}} &= \sum_i m_i, \quad m_i = \lambda\, m_{i-1} + (\theta_{i-1} - \theta_i), \quad \lambda = 0.996, \end{aligned} \tag{10}$$

with θ the learning rate in units of its peak. $g(R) = 1 + b\,\rho(R)$ with $\rho(R) = kR/(kR+1-R)$ and k from the flat law (Table 1). On the text axis, $\zeta < 0$, meaning that while text-scarce regimes ($R \to 1$) increase steeply the text loss during flat pre-training, annealing recovers a substantial part of this deficit. Conversely, on the caption axis, annealing gains are size-invariant ($\zeta \approx 0$) but diminish mildly as $R \to 1$ ($b < 0$). We further find that warmup and annealing effects in PT are negligible. Note that annealing does not affect the fusion trade-off conclusion, though it makes the loss difference between early/late fusion smaller.

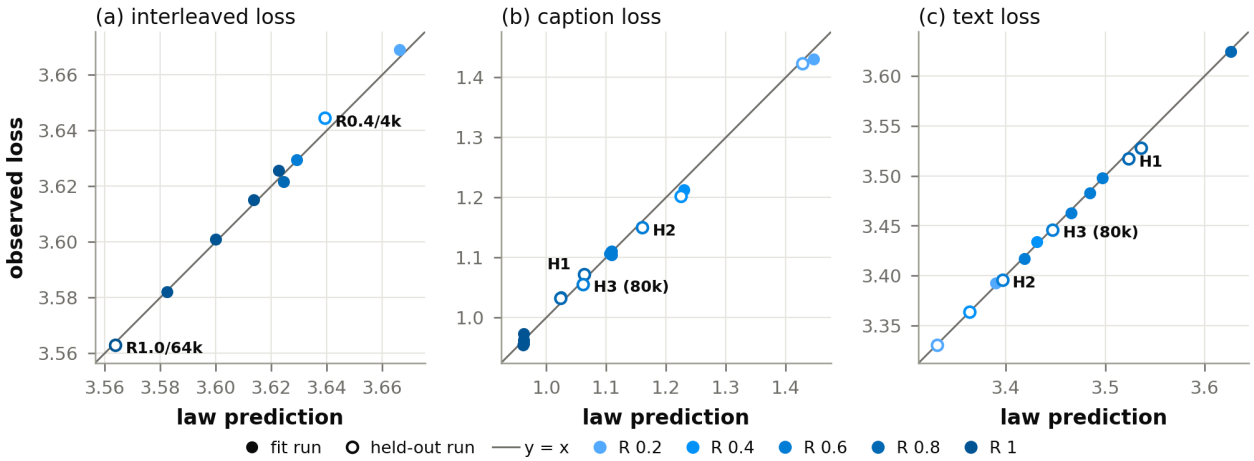


Figure 12: **Observed versus predicted validation losses for fitted and held-out validation runs.** The law predictions tightly align with observed losses across varying configurations.

D The Chemical Kinetics Analogy

We model the shift term in Equation 5 via the (pseudo) reversible first-order reaction kinetics. Particularly, we consider heterogeneous surface catalysis, where a continuous external gas stream (training data) reacts with a fixed solid catalyst surface (model parameters).

The model's fixed representational capacity acts as surface sites in either a text-aligned state (A) or a multimodal-aligned state (B), such that $[A]+[B]=1$ by conservation laws. Data are pumped in as a stream containing a molar fraction R of multimodal tokens (M) and $1-R$ of text tokens (T).

Transitions occur when the gas hits the surface:

$$A + M \xrightarrow{\kappa_M} B$$
$$B + T \xrightarrow{\kappa_T} A$$

where κ_M and κ_T are the intrinsic reactivity. Since the stream fractions $[M]=R$ and $[T]=1-R$ are externally fixed, mass conservation yields a pseudo-first-order differential equation for state B:

$$\begin{aligned}\frac{d[B]}{dt} &= \kappa_M R[A] - \kappa_T(1-R)[B] \\ &= \kappa_M R - \Big[\kappa_M R + \kappa_T(1-R)\Big][B]\end{aligned} \tag{11}$$

At dynamic equilibrium ($\frac{d[B]}{dt}=0$), the steady-state fraction $[B]_\infty$ is:

$$[B]_\infty = \frac{\kappa_M R}{\kappa_M R + \kappa_T(1-R)} = \frac{kR}{kR+1-R} \equiv \rho(R) \tag{12}$$

where $k \equiv \kappa_M/\kappa_T$.

Integrating from $[B](0)=0$ yields $[B](t)=[B]_\infty(1-e^{-\lambda t})$, where

$$\lambda = \kappa_M R + \kappa_T(1-R) = \kappa_T\Big(kR+1-R\Big) \tag{13}$$

Defining the baseline text timescale $\tau \equiv 1/\kappa_T$, λt may be written as $u(t)=(kR+1-R)\frac{t}{\tau}$.

Substituting D_{cpt} for time t, and replacing the exponential decay $1-e^{-u}$ with a power-law form $1-(1+u)^{-\beta}$ to account for the multi-scale spectrum of neural network training (see, e.g., [16]), we directly recover the shift term. This analogy serves as a toy model to understand and motivate the formulation of the proposed scaling law.

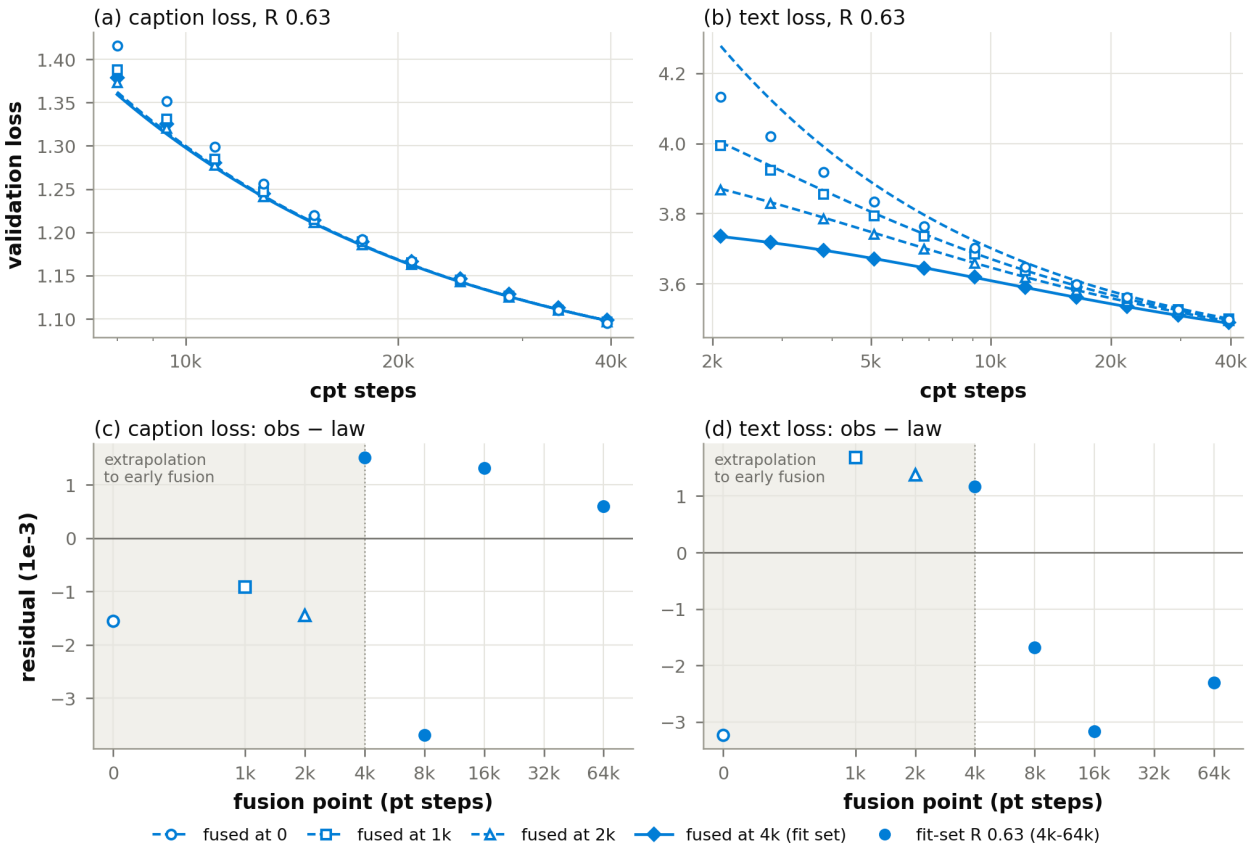


Figure 13: **Extrapolation of the scaling law to early fusion regimes ($D_{pt} = 0$).** The top panels map the predicted versus observed loss trajectories, while the bottom panels display the corresponding residuals to demonstrate accurate predictions even for early fusion.

E More on Base VLM Evaluation

E.1 Open-Ended Document Visual Question Answering

As mentioned in the main text, evaluating the reference answer's perplexity works partially for DocVQA. Let us provide a more detailed explanation on this.

For open-ended document visual question answering where candidate reformulation is impractical, we adapt DocVQA to zero-shot, teacher-forced reference perplexity. Conditioned on the document image and prompt, we compute the byte-normalized cross-entropy over the ground-truth answer continuation. This formulation provides a direct multimodal analogue to text language modeling perplexity.

This works partially: for caption-based pre-training, the score improves with training; when pre-trained with interleaved data from scratch, the visual grounding gap is negative (Figure 14). This suggests that DocVQA is sensitive to pre-training domain bias and serves as an imperfect evaluation proxy for base VLMs; nevertheless, we report it in the paper for completeness as a representative open-ended generation task.

E.2 Rewriting Strategy

In our rewriting pipeline, we use declarative templates that are generated deterministically from the question and candidate text to ensure that they share identical syntax to avoid confounders.

For MME, each image comes with a balanced Yes and No pair of questions that differ only in the probed attribute. We convert both questions to declarative statements using rule-based grammar transformations

Mixture shares				Caption residual		Text residual	
f	s_{txt}	s_{cap}	s_{int}	int. = text	$w_{\text{int}} = 0.18$	int. = caption	int. = text
Caption-free mixtures (text + interleaved)							
1.00	0.85	0.00	0.15	–	–	−0.012	+0.007
1.00	0.60	0.00	0.40	–	–	−0.040	+0.024
1.00	0.40	0.00	0.60	–	–	−0.077	+0.048
Low interleaved fraction ($f \leq 0.25$): no new parameter needed for caption							
0.15	0.00	0.85	0.15	−0.004	−0.002	–	–
0.22	0.38	0.48	0.14	−0.016	−0.011	−0.025	+0.023
High interleaved fraction ($f > 0.25$): one exchange rate recovers most caption runs							
0.27	0.45	0.40	0.15	−0.026	−0.019	−0.022	+0.020
0.33	0.40	0.40	0.20	−0.029	−0.020	−0.034	+0.027
0.40	0.00	0.60	0.40	−0.028	−0.019	–	–
0.50	0.20	0.40	0.40	−0.039	−0.022	−0.111	+0.073
0.67	0.40	0.20	0.40	−0.006	+0.046	−0.057	+0.042
0.75	0.20	0.20	0.60	−0.057	+0.014	−0.140	+0.082

Table 6: **Evaluation residuals (observed minus predicted in nats) for three-modality training mixtures (text, caption, interleaved) across varying interleaved multimodal shares (f).** Comparisons are made treating interleaved data as text or assigning an effective caption exchange rate ($w_{int} = 0.18$) on the caption loss axis, and treating interleaved data as text or caption on the text loss axis.

without inserting explicit negation. Discarding pairs that fail parsing or lacking matching syntax, we yield 1,158 pairs (2,316 statement) to this end.

For GQA's multi-choice questions, we extract a single statement and slot the each candidate option into it. We avoid negation in the statement, and hence exclude most binary questions containing negation. We also drop options that are uninformative like "*none of the above*". After grammatical verification, we are left with 6,405 question sets and 24,985 statements in total.

E.3 Evaluating Rewritten Benchmarks with Open Models

	MME (Acc.)			GQA (Acc.)		
Model	**Gen.**	**Ours**	**Gap (Δ)**	**Gen.**	**Ours**	**Gap (Δ)**
Qwen2.5-VL-3B	0.852	0.856	+0.244	0.774	0.618	+0.324
Gemma-3-4B-IT	0.786	0.844	+0.244	0.624	0.567	+0.285
SmolVLM2-2.2B	0.768	0.774	+0.172	0.697	0.607	+0.299

Table 7: **Comparison of standard zero-shot evalution versus log-likelihood based evaluation on rewritten benchmarks across open-weight VLMs.** "Generation" evaluates letter MCQ selection. "Ours" evaluates and ranks byte-normalized log-likelihood of rewritten statements; Gap (Δ) measures the accuracy difference of evaluation with-/without image prefix, confirming consistent visual grounding across models.

We further evaluate our rewritten benchmarks with open models, namely Gemma-3-4B-IT [43], Qwen2.5-VL-3B [3], and SmolVLM2-2.2B [4]. Additionally, we evaluate the original MME and GQA benchmarks with the same open models, by prompting them to answer the original questions directly. Note that these are instruction-tuned models where standard zero-shot generation is feasible.

Results are summarized in Tab. 7. We see that rewriting MME benchmarks can even improve the accuracy of open models, likely due to the more declarative and less ambiguous nature of the rewritten statements from binary Yes/No questions. Rewriting GQA benchmarks, however, leads to a drop in accuracy, yet still far above chance, reflecting the stricter criteria of full-sentence log-likelihood discrimination compared to predicting the correct option from a set of candidates. Furthermore, the relative performance ranking of the models remains the same across both formats although the scores differ. This indicates that the rewritten benchmarks still evaluate the same

[3] https://qwen.ai/blog?id=qwen2.5-vl
[4] https://huggingface.co/blog/smolvlm2

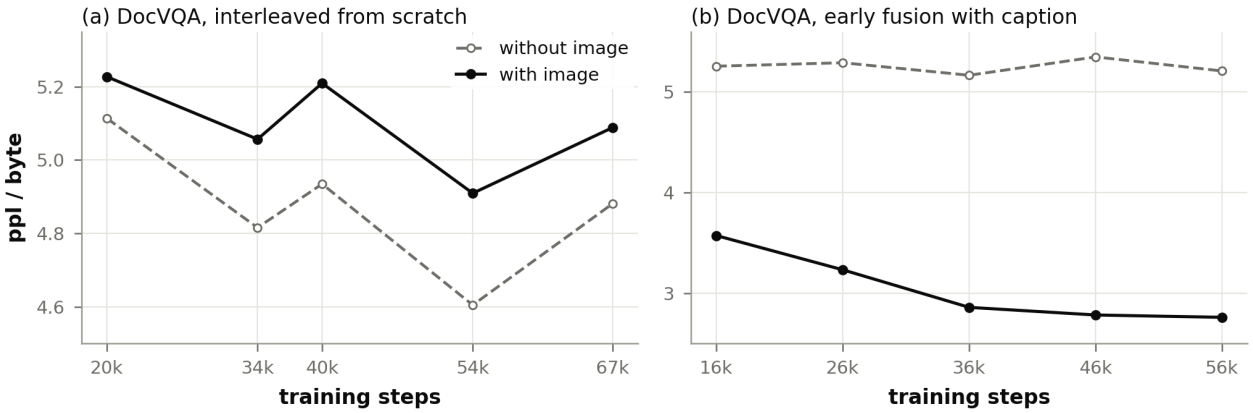


Figure 14: **DocVQA evaluation results via continuation perplexity per byte of the reference answers, with and without the image, for 1B model trained from with interleaved data only (left) and early fusion with caption (right).** We see that only early fusion with caption produces scores that improve with training and positive visual grounding gap. Training with interleaved data even has negative visual grounding gap, raising questions about the effectiveness of interleaved data in multimodal pre-training.

underlying visual capabilities. Importantly, all models maintain large, positive visual grounding gaps. These results validate the use of our rewritten benchmarks as a proxy for evaluating base VLMs.

F More on Fusion Timings

Derivation of the Optimal Fusion Timings. Under the budget constraint, D_{cpt} and R are completely determined by the variable D_{pt}:

$$D_{\text{cpt}} = (T_{\text{txt}} - D_{\text{pt}}) + T_{\text{cap}}, \qquad R = \frac{T_{\text{cap}}}{D_{\text{cpt}}}$$

The base term can be rewritten as

$$\begin{aligned} \mathcal{D}_{\text{base,cap}} &= w_{\text{cap}}\, D_{\text{pt}} + \big(R + w_{\text{cap}}(1-R)\big)\, D_{\text{cpt}} \\ &= w_{\text{cap}}\, T_{\text{txt}} + T_{\text{cap}} \end{aligned} \tag{14}$$

which holds similarly for the text base term. Therefore, the base term is constant and drops out of the optimization.

We are left with the shift term. Using the identity $(1-R)D_{\text{cpt}} = T_{\text{txt}} - D_{\text{pt}}$, we express the effective scale factor $X(D_{\text{pt}})$ directly as a linear function of D_{pt}:

$$X(D_{\text{pt}}) \equiv (k_{\text{cap}}R + 1 - R)D_{\text{cpt}} = k_{\text{cap}}\, T_{\text{cap}} + (T_{\text{txt}} - D_{\text{pt}})$$

Notice that $\rho(R) = \frac{k_{\text{cap}} R D_{cpt}}{X(D_{pt})} = \frac{k_{\text{cap}} T_{\text{cap}}}{X(D_{pt})}$. The caption shift term becomes:

$$\Delta L_{\text{cap}}(D_{\text{pt}}) = -B_{0,\text{cap}}(N) \cdot \frac{k_{\text{cap}} T_{\text{cap}}}{X(D_{\text{pt}})} \left[1 - \left(1 + \frac{X(D_{\text{pt}})}{\tau_{\text{cap}}}\right)^{-\beta_{\text{cap}}}\right]$$

Defining the normalized scale parameter $z(D_{\text{pt}}) = \frac{X(D_{\text{pt}})}{\tau_{\text{cap}}} > 0$, we then have

$$\Delta L_{cap}(D_{pt}) = -B_{0,\text{cap}}(N) \frac{k_{\text{cap}} T_{\text{cap}}}{\tau_{\text{cap}}} \cdot \phi(z(D_{pt})),$$

				val loss (nats)		text (lm-eval)				
Schedule	D_{pt}	D_{cpt}	R	caption	text	HSwag	PIQA	ARC-e	WGrande	LAMB
S0	0	56k	0.447	1.1488	3.4095	0.3691	0.6529	0.4933	0.5217	69.8
S1	4k	52k	0.482	1.1418	3.4194	0.3668	0.6436	0.4912	0.5146	73.0
S2	8k	48k	0.522	1.1277	3.4286	0.3678	0.6235	0.4928	0.5036	73.7
S3	16k	40k	0.626	1.0983	3.4574	0.3613	0.6349	0.4836	0.5367	89.3
S4	24k	32k	0.782	1.0486	3.5341	0.3591	0.6360	0.4726	0.5162	80.0
S5	32k	25k	1.000	0.9902	6.3135	0.2938	0.5871	0.3215	0.5130	2466.4
Spearman r_s						−0.94	−0.77	−0.94	−0.20	+0.94

	likelihood ranking (accuracy, with image)				generation ppl	
Schedule	MME	gap	GQA	gap	DocVQA	Δ
S0	0.6701	+0.0941	0.4103	+0.1280	3.88	+1.98
S1	0.6796	+0.1123	0.3694	+0.0924	3.73	+1.96
S2	0.6788	+0.1097	0.3975	+0.1198	3.79	+2.10
S3	0.6779	+0.0959	0.3891	+0.1085	4.13	+1.75
S4	0.6883	+0.1157	0.3969	+0.1226	3.78	+2.20
S5	0.6770	+0.0976	0.3869	+0.1027	3.95	+6.11
r_s	−0.20	−0.31	+0.37	+0.26	−0.31	−0.60

Table 8: **Downstream benchmark performance and final validation losses for the fusion timing experiments on the 300M model.** Vision benchmarks are evaluated using our rewriting strategy (MME and GQA) or via continuation perplexity (DocVQA), along with the gap between evaluation with and without image prefix. The bottom rows report Spearman rank correlation r_s between downstream metrics and their corresponding validation loss across schedules S0-S5.

where $\phi(z) = \frac{1-(1+z)^{-\beta}}{z}$.

One can show that $\phi(z)$ is strictly decreasing in z for $\beta > 0$, because

$$\phi'(z) = \frac{(1 + (\beta_{\text{cap}} + 1)z)\,(1+z)^{-\beta_{\text{cap}}-1} - 1}{z^2} < 0$$

Therefore the magnitude $|\Delta L_{cap}|$ is strictly increasing with D_{pt}. The optimal solutions thus lie at the edge:

$$\arg\min_{D_{\text{pt}}} L_{\text{cap}} = T_{\text{txt}} \implies D_{\text{cpt}} = T_{\text{cap}},\ R = 1 \tag{15}$$

$$\arg\min_{D_{\text{pt}}} L_{\text{txt}} = 0$$

$$\implies D_{\text{cpt}} = T_{\text{txt}} + T_{\text{cap}},\ R = \frac{T_{\text{cap}}}{T_{\text{txt}} + T_{\text{cap}}} \tag{16}$$

In other words, early fusion favors lower text loss while late fusion favors lower caption loss.

More on Fusion Timing Experimental Results. We give the full results of the fusion timing experiments in Tabs. 8 and 9. We note that S5 reuses our nearest available checkpoint at $D_{\text{pt}} = 32{,}000$ steps to match the target multimodal budget (25,000 caption steps). This leaves the text portion slightly over-budget (+3.3%, 5,7000 steps versus 56,000 steps for other runs). Given that text tokens contribute minimally to caption loss, and the text performance collapses catastrophically under pure-caption training, this has no measurable impact on the timing's conclusions. We further note that we do not run S2 for the 1B model to save computational resources, and that the difference with S1 and S3 is of noise-level significance.

				val loss (nats)		text (lm-eval)				
Schedule	D_{pt}	D_{cpt}	R	caption	text	HSwag	PIQA	ARC-e	WGrande	LAMB
S0	0	56k	0.447	0.9659	3.0258	0.4862	0.7029	0.5711	0.5264	27.0
S1	4k	52k	0.482	0.9600	3.0360	0.4795	0.6921	0.5720	0.5517	27.3
S3	16k	40k	0.626	0.9304	3.0791	0.4702	0.6997	0.5787	0.5438	24.8
S4	24k	32k	0.782	0.9029	3.1589	0.4569	0.6915	0.5673	0.5391	32.3
S5	32k	25k	1.000	0.8705	5.8537	0.3144	0.6262	0.3611	0.5130	557.6
Spearman r_s						−1.00	−0.90	−0.60	−0.40	+0.70

	likelihood ranking (accuracy, with image)				DocVQA ppl	
Schedule	MME	gap	GQA	gap	DocVQA	Δ
S0	0.7064	+0.1131	0.4134	+0.1386	3.10	+2.35
S1	0.7055	+0.1097	0.4236	+0.1429	2.91	+2.30
S3	0.7116	+0.1218	0.4158	+0.1374	2.94	+2.37
S4	0.7090	+0.1278	0.4159	+0.1326	2.67	+2.64
S5	0.6839	+0.1088	0.4209	+0.1318	2.80	+5.03
r_s	+0.20	+0.10	−0.40	+0.90	+0.80	−0.90

Table 9: **Downstream benchmark performance and final validation losses as in Table 8 but evaluated for the 1B model.**