

Beyond Repetition: Text Simplification and Curriculum Learning for Data-Constrained Pretraining

Matthew Theodore Roque* and Dan John Velasco*

Samsung R&D Institute Philippines
{roque.mt,dj.velasco}@samsung.com

*Equal Contribution

Abstract

Most studies on language model pretraining focus on large datasets, leaving open questions about optimization in data-constrained settings. In such settings, the effects of training data order and of including alternative versions of the same text remain underexplored. We address this by studying curriculum learning in pretraining, focusing on text-complexity ordering and data augmentation via simplification. We ask: (1) Does simplifying texts enhance representation quality more than reusing the original data? and (2) Does ordering data by text complexity yield better representations? To answer, we build on a pair of parallel corpora where human-written paragraphs are aligned with LLM-simplified variants, and test four data schedules: repeated exposure, low-to-high complexity, high-to-low, and interleaved. We analyze models' representation quality from a sample efficiency perspective via fine-tuning, as well as its zero-shot performance on linguistic knowledge, entity tracking, world knowledge, and commonsense reasoning. Our findings show that adding simplified data improves fine-tuning and zero-shot performance over a repeated-exposure baseline: smaller models benefit from low-to-high complexity, while larger models perform better with interleaved ordering.

1 Introduction

Scaling studies show that language model performance improves predictably with more data, parameters, and compute (Kaplan et al., 2020; Hoffmann et al., 2022). However, these studies typically assume that the amount of unique pretraining data is effectively unlimited (Muennighoff et al., 2025). In practice, pretraining often faces data-constrained settings where continued exposure to the same corpus is unavoidable. Under such conditions, two factors remain underexplored in modern decoder-only pretraining: (1) the order in which training




Figure 1: Average score on six language tasks by curriculum and fine-tuning data size. **(Top)** The 124M model benefits from SIMP → HW curriculum, suggesting smaller models gain from warming up on simpler text. **(Bottom)** The 256M model benefits from INTERLEAVED, favoring balanced exposure. All but HW → SIMP outperform BASELINE in sample efficiency.

data is presented, and (2) the use of simpler versions of the same text.

We examine these factors concretely. For training data order, we focus on **coarse-grained text complexity ordering**—presenting a simple corpus before a complex one—while keeping the core content constant, following the intuition of curriculum learning (Bengio et al., 2009). We define "simple" texts as those that use higher-frequency vocabulary and have shallower syntactic structures.

Our experiments use a high-quality English dataset where human-written paragraphs are paired with simplified counterparts produced by an LLM. This corpus was introduced and validated in concurrent work (Velasco and Roque, 2025), which demonstrates that simplification reduces surface-level complexity (sentence length, syntactic depth, lexical diversity) while preserving semantic content. Here, we do not revisit corpus construction in detail; instead, we leverage it to test how simplified data and ordering strategies interact under a fixed training budget.

We evaluate four data schedules:

- **BASELINE:** repeated exposure to human-written text.
- **INTERLEAVED:** human-written and simplified paragraphs uniformly mixed.
- **SIMP→HW:** training first on simplified, then human-written text (curriculum).
- **HW→SIMP:** training first on human-written, then simplified text (anti-curriculum).

All other training variables (architecture, tokenizer, context length, optimizer) are held constant across schedules, isolating the effect of text complexity and order.

Research questions. We ask two core questions:

- (1) In data-constrained settings, does replacing repeated exposure with simplified text improve representation quality?
- (2) Does ordering data by text complexity—simple to complex versus interleaved—yield better downstream and zero-shot performance?

Contributions. Our contributions are threefold:

- (1) We provide the first controlled study of how text simplification and curriculum scheduling interact in data-constrained pretraining.
- (2) We evaluate these schedules across fine-tuning and zero-shot tasks, covering linguistic knowledge, entity tracking, world knowledge, and commonsense reasoning.
- (3) We show that simplified data generally improves performance over repeated exposure, that smaller models benefit from simple-to-complex curricula, and that larger models favor balanced exposure via interleaving.

2 Related Work

Data-constrained Pretraining and Synthetic Data. Most scaling studies assume unlimited data. In data-constrained settings, Muennighoff et al. (2025) shows that training up to four epochs of repeated data is just as good as unique data, but further repetition offers no benefit. Recent works address the "data wall" either by using LLMs to generate synthetic data (Gunasekar et al., 2023; Ben Allal et al., 2024) or to rewrite existing data, broadly across general domains (Maini et al., 2024; Su et al., 2025; Nguyen et al., 2025; DatologyAI et al., 2025) and in specific domains such as math, code, and clinical text (Fujii et al., 2025; Liu and Nguyen, 2024). Rewriting into simpler forms remains underexplored; we investigate this setting.

Curriculum Learning. Humans learn better when examples follow a meaningful order (e.g., simple to complex). Bengio et al. (2009) formalized this as Curriculum Learning (CL), training neural networks on data of gradually increasing complexity. The opposite, anti-curriculum, has sometimes matched or outperformed CL (Kocmi and Bojar, 2017; Zhang et al., 2018). Difficulty metrics vary by application: in code language models, Nair et al. (2024) used software engineering metrics, while in machine translation, Zhou et al. (2020) measured sequence uncertainty via language model entropy. Recent studies applied CL to pretraining language models: Tsvetkov et al. (2016) examined linguistically inspired measures such as age of acquisition, while Oba et al. (2023) measured complexity via dependency tree depth. A recent large-scale study by Zhang et al. (2025) finds that using text complexity metrics such as Flesch Reading Ease, Lexical Diversity, and Compression Ratio can accelerate convergence and modestly outperform random shuffling.

However, most CL studies in pretraining conflate language and content complexity, failing to isolate the effect of ordering by language complexity. Most pretraining studies also assume large data volumes, leaving open the question of how to advance pretraining in data-constrained settings beyond repeated exposure. This work uniquely addresses this gap, testing whether LLM-based simplification and coarse-grained text complexity ordering improve representational quality beyond repeated exposure to the original data.

3 Methodology

3.1 Data

We reuse the two parallel corpora from [Velasco and Roque \(2025\)](#), derived from a 2B-token subset of FineWeb-Edu (ODC-By 1.0; [Penedo et al., 2024](#)). The human-written corpus (HW) and its simplified variant (SIMP) are aligned at the paragraph level: each paragraph in HW has a corresponding simplification in SIMP produced by an LLM, with paragraphs that fail basic length/formatting checks symmetrically removed from both sides to preserve one-to-one alignment. After filtering, token counts are approximately 2.00B (HW) and 1.71B (SIMP).

In brief, SIMP reduces surface-level complexity (shorter sentences, shallower syntax, more frequent vocabulary) while preserving core content; [Velasco and Roque \(2025\)](#) report readability and lexical/syntactic metrics alongside semantic-similarity checks that validate this property. Our study focuses on how to use these corpora under a fixed budget: repeated exposure vs. augmentation with simplifications and interleaving vs. ordered curricula. All other variables (architecture, tokenizer, context length, optimizer) are held constant across schedules.

3.2 Schedules

We compare four data schedules that differ only in the order of data presented to the model. Unlike most curriculum strategies that manipulate complexity per example (fine-grained), our schedules are coarse-grained, adjusting complexity at the corpus level rather than per example. The four schedules are as follows:

- BASELINE: two epochs of HW (simulating data-constrained scenarios).
- INTERLEAVED: HW and SIMP are uniformly interleaved, preserving each corpus’s within-source order (simulating random shuffling in a more balanced way).
- SIMP→HW: concatenation of SIMP and HW (simulating standard curriculum).
- HW→SIMP: concatenation of HW and SIMP (simulating anti-curriculum).

Each training example is a single paragraph, and both HW and SIMP corpora are perfectly parallel, containing the same number of paragraphs. Within each source, paragraph order is fixed across all

schedules, ensuring differences arise only from the sequence in which the two sources are presented. This design isolates the effect of data ordering by text complexity and the presence of LLM-rewritten text, while controlling for content coverage and total training steps.

3.3 Model and Training

We use 124M and 256M parameter causal language models based on the design principles of MobileLLM ([Liu et al., 2024](#)), adopting a deep-and-thin architecture with SwiGLU activations ([Shazeer, 2020](#)), grouped-query attention ([Ainslie et al., 2023](#)), and without embedding weight sharing for better comparability with contemporary decoder-only models. Each model has 30 transformer layers with 9 attention heads (3 key-value heads per layer) and embedding dimensions of 576 (124M) and 846 (256M). We refer to the configurations as 124M/256M by convention; the total parameter counts including embeddings are approximately 143M and 283M, respectively.

All corpora are tokenized using the LLaMA-2 BPE tokenizer ([Touvron et al., 2023](#)) with a 32,000-token vocabulary. Training examples are individual paragraphs, with no concatenation or sequence packing to control for total training steps. Inputs are right-padded to 512 tokens with the EOS token.

Optimization uses AdamW ([Loshchilov and Hutter, 2019](#)) with default hyperparameters, a peak learning rate of $3e-4$ linearly decayed over training, 5% warm-up, and no dropout. The effective batch size is 256 (8 examples per GPU \times 8 GPUs \times 4 gradient accumulation steps). Training is conducted in FP16 mixed precision on 8 \times NVIDIA P100 GPUs, without gradient checkpointing. All experiments use PyTorch with Hugging Face Transformers and Distributed Data Parallel (DDP), with a fixed random seed of 42 for data shuffling and parameter initialization.

3.4 Evaluation Setup

We evaluate each schedule under two complementary regimes to capture both transferable language understanding after fine-tuning and generalization without further task-specific training.

Fine-tuning on NLU tasks. We fine-tune the pre-trained models on a subset of the BabyLM evaluation pipeline’s natural language understanding (NLU) tasks, using the preprocessed training and validation splits provided therein ([Charpentier et al.,](#)

2025). From the original suite, we exclude:

- **MultiRC**, due to its substantially larger size and because preliminary experiments showed none of our pretraining setups outperformed training from scratch (58–59 points).
- **WSC**, due to its much smaller dataset size and the high variance (4–12 points) observed across seeds under different fine-tuning budgets.

The resulting set of tasks is: BoolQ, MNLI, MRPC, QQP, and RTE. Each task is framed as classification, with paired-input tasks (e.g., premise-hypothesis) concatenating the two sequences with a separator token. Fine-tuning is performed with all model parameters trainable. Batch sizes are set by memory constraints: 2 examples per GPU (effective batch size 16) for BoolQ, and 8 examples per GPU (effective batch size 64) for all other tasks. For each task, we search over learning rates $\{1e-4, 5e-5, 2e-5, 1e-5, 5e-6\}$ and epochs $\{1, 2, 3, 4, 5\}$. Model selection is based on the highest validation score according to the task’s standard metric (accuracy for BoolQ, MNLI, and RTE; F1 for MRPC and QQP). All results are averaged over three runs with different random seeds.

In addition to this **Full** fine-tuning setup, we introduce two smaller-scale, class-balanced variants to test whether having more downstream task examples diminishes the influence of pretraining differences. For each task, we identify the class with the fewest examples in the original training split, then construct:

1. **Small** dataset, in which all classes contain exactly half of the least-represented class size, and
2. **Tiny** dataset, in which all classes contain exactly one quarter of the least-represented class size.

This ensures that all classes are equally represented and that dataset size is systematically reduced across tasks.

These smaller datasets are fixed across runs so that each of the three seeds per setup uses the same subset of examples, with only model initialization differing. The same hyperparameter sweeps are applied to these reduced datasets as in the full-data setup.

| Task | Classes | Full | Small | Tiny |
|-------|---------|--------|-------------|-------------|
| MNLI | 3 | 10,000 | 4,911 (49%) | 2,454 (25%) |
| MRPC | 2 | 3,668 | 1,194 (33%) | 596 (16%) |
| BoolQ | 2 | 9,427 | 3,552 (38%) | 1,776 (19%) |
| RTE | 2 | 2,490 | 1,240 (50%) | 620 (25%) |
| QQP | 2 | 10,000 | 3,662 (37%) | 1,830 (18%) |

Table 1: Number of training examples per NLU task in the **Full**, **Small**, and **Tiny** fine-tuning setups. Small and tiny datasets are class-balanced subsets, with percentages shown relative to the full dataset. Percentages vary across tasks because subset sizes are determined relative to the least-represented class in each dataset.

Zero-shot evaluation. Using the LM Evaluation Harness (Gao et al., 2024), we assess models trained under different curricula via zero-shot performance on a suite of multiple-choice benchmarks, grouped into three capability areas:

- **Linguistic knowledge:** BLiMP (Warstadt et al., 2020) and BLiMP Supplement (Charpentier et al., 2025) for syntactic and morphological phenomena.
- **Discourse and world knowledge:** Entity Tracking (Kim and Schuster, 2023) for reference consistency, and MMLU (Hendrycks et al., 2021) and EWoK (Ivanova et al., 2024) for factual knowledge.
- **Commonsense reasoning:** ARC (Clark et al., 2018), HellaSwag (Zellers et al., 2019), PIQA (Bisk et al., 2020), Social IQa (Sap et al., 2019), and OpenBookQA (Mihaylov et al., 2018) for reasoning about everyday scenarios.

All zero-shot tasks are formatted according to their official specifications. For each candidate answer, we compute the sum of log-probabilities of its tokens given the prompt and select the option with the highest score. Accuracy is reported for all tasks. Evaluation is deterministic and performed on a single NVIDIA P100 GPU.

4 Results and Discussion

Across fine-tuning and zero-shot setups, outcomes hinge on two factors: augmenting with simplified data versus repeated exposure, and ordering of data from simplified to complex versus interleaving.

4.1 Simplification vs. Repetition

We ask whether replacing a second pass over the human-written corpus with LLM-based simplifications improves pretraining. The primary

| Model | Full | Small | Tiny |
|-------------|-------------|-------------|-------------|
| 124M | | | |
| BASELINE | 70.9 | 67.7 | 64.4 |
| INTERLEAVED | 71.1 | 68.2 | 64.8 |
| SIMP → HW | 71.4 | 68.3 | 65.2 |
| HW → SIMP | 71.1 | 67.5 | 64.4 |
| 256M | | | |
| BASELINE | 73.9 | 68.8 | 65.0 |
| INTERLEAVED | 75.7 | 70.7 | 66.9 |
| SIMP → HW | 75.3 | 69.5 | 65.6 |
| HW → SIMP | 75.5 | 69.1 | 64.8 |

Table 2: Macro-average across five NLU tasks (BoolQ, MNLI, MRPC-F1, QQP-F1, RTE) under three fine-tuning budgets: **Full**, **Small**, and **Tiny**. Best per model size per budget in bold. At 124M, SIMP→HW leads across all budgets (+0.5–0.8 vs. BASELINE); at 256M, INTERLEAVED leads across all budgets (+1.8–1.9), indicating size- and budget-dependent preferences. Per-task breakdowns with mean ± std. dev. are in Appendix A.

comparison is BASELINE (two epochs of HW) vs. INTERLEAVED, which uniformly mixes HW and SIMP.

Fine-tuning evaluation. At 124M, adding SIMP yields small, consistent gains over BASELINE when mixed (INTERLEAVED: +0.2, +0.5, +0.4 on **Full/Small/Tiny**). At 256M, the benefits of including SIMP are larger overall in the mixed setting (INTERLEAVED: +1.8/+1.9/+1.9). These advantages grow as fine-tuning data decreases: INTERLEAVED degrades less than BASELINE on Small and Tiny setups, suggesting simplification is most beneficial when data is scarce, improving sample efficiency.

Zero-shot evaluation. Across linguistic, discourse, and commonsense benchmarks, introducing SIMP tends to be neutral to positive at 124M and more evidently positive at 256M, with the largest single gain appearing on entity tracking when mixing (INTERLEAVED vs. BASELINE: \approx +4.9 points at 256M, Table 4). Other tasks show smaller, sometimes mixed, changes, so the aggregate trend favors simplification but with task-family variability (Table 3).

Summary. Compared to repeating HW, in-domain augmentation via LLM simplification yields modest gains at 124M and larger gains at 256M. While we do not claim causality, a plausible explanation is that the additional capacity of the 256M model may better absorb paraphrastic variety, making mixed (INTERLEAVED) exposure yield small but consistent improvements in both fine-tuned and zero-shot set-

tings. We consider this hypothesis specific to our training setup rather than a general rule.

While repetition underperforms simplification, it remains attractive given its zero generation cost and competitive results, consistent with findings by Muennighoff et al. (2025). **Our results complement rather than replace repeated exposure: since performance saturates beyond four epochs, text simplification offers a way to extend these gains further in data-constrained settings.**

4.2 Curriculum vs. Interleaving

We next ask whether ordering simplified and human-written text into two distinct phases (curriculum or anti-curriculum) provides advantages over mixing them uniformly. Here, INTERLEAVED serves as the natural baseline, representing random shuffling of the two corpora.

Fine-tuning evaluation. At 124M, the curriculum schedule SIMP→HW edges out INTERLEAVED across all budgets (+0.3, +0.1, +0.4), while HW→SIMP is equal or slightly worse (Table 2). The margins are small, but the consistent advantage for SIMP→HW suggests that a warm-up on simplified text may provide downstream benefits when model capacity is limited. At 256M, the picture changes: INTERLEAVED remains strongest across all budgets, while the ordered schedules are competitive but do not surpass it (SIMP→HW: −0.4, −1.2, −1.3 vs. INTERLEAVED; HW→SIMP: −0.2, −1.6, −2.1). In short, simple-to-complex ordering helps at 124M, especially in low-resource settings where it shows improved sample efficiency, but offers no advantage as model capacity increases.

Zero-shot evaluation. For zero-shot tasks, ordering effects are somewhat stronger than in the fine-tuning results. At 124M, SIMP→HW leads on BLiMP and BLiMP-Supplement, while HW→SIMP excels on entity tracking (Table 4). At 256M, however, INTERLEAVED typically matches or exceeds the ordered setups—for example, it leads on entity tracking (+4.3 vs. SIMP→HW) and holds a small edge on HellaSwag (Table 3). These patterns suggest that ordered exposure can steer smaller models toward specific strengths, but that random mixing is safer once capacity is sufficient.

Summary. Compared to interleaving, SIMP → HW yields modest gains at 124M—especially for NLU fine-tuning and linguistic probes—but these benefits diminish or reverse at 256M. This pattern

| Model | arc_chl | arc_e | hellaswag | openbookqa | piqa | social_iqa | Avg. |
|-------------|-------------|-------------|-------------|-------------|-------------|-------------|--------------|
| 124M | | | | | | | |
| BASELINE | 22.0 | 35.9 | 28.4 | 14.0 | 57.1 | 36.0 | 34.28 |
| INTERLEAVED | 23.4 | 34.6 | 28.4 | 15.4 | 58.4 | 35.7 | 34.50 |
| SIMP → HW | 23.1 | 33.7 | 28.0 | 17.0 | 57.3 | 36.1 | 34.42 |
| HW → SIMP | 20.4 | 33.5 | 28.1 | 14.8 | 57.6 | 36.9 | 34.18 |
| 256M | | | | | | | |
| BASELINE | 23.1 | 37.8 | 27.7 | 16.0 | 57.3 | 35.2 | 34.80 |
| INTERLEAVED | 22.7 | 36.9 | 28.6 | 16.2 | 56.8 | 36.6 | 35.02 |
| SIMP → HW | 21.8 | 35.5 | 29.1 | 18.2 | 58.4 | 36.5 | 35.54 |
| HW → SIMP | 24.7 | 34.7 | 28.2 | 16.6 | 56.4 | 35.5 | 34.28 |

Table 3: Zero-shot accuracy on commonsense reasoning benchmarks (ARC-Challenge, ARC-Easy, HellaSwag, OpenBookQA, PIQA, Social IQa). “Avg.” is the mean across tasks. At 124M, INTERLEAVED yields the highest average by a small margin; at 256M, SIMP → HW attains the best average, driven by gains on HellaSwag, OpenBookQA, and PIQA, while INTERLEAVED leads on Social IQa and HW → SIMP peaks on ARC-Challenge.

| Model | blimp | supp | ewok | entity | mmlu |
|-------------|-------------|-------------|-------------|-------------|-------------|
| 124M | | | | | |
| BASELINE | 71.8 | 61.9 | 53.9 | 22.4 | 24.7 |
| INTERLEAVED | 72.3 | 63.6 | 55.5 | 28.1 | 24.9 |
| SIMP → HW | 72.4 | 63.8 | 54.8 | 31.7 | 23.6 |
| HW → SIMP | 70.7 | 61.3 | 55.1 | 36.9 | 23.3 |
| 256M | | | | | |
| BASELINE | 73.8 | 65.6 | 55.0 | 30.1 | 23.5 |
| INTERLEAVED | 73.6 | 64.1 | 56.2 | 35.0 | 24.5 |
| SIMP → HW | 73.7 | 64.3 | 55.9 | 30.8 | 25.8 |
| HW → SIMP | 73.1 | 64.3 | 56.0 | 34.1 | 26.2 |

Table 4: Zero-shot evaluations on linguistic competence (BLiMP, BLiMP-Supplement), world knowledge (EWoK), discourse (Entity Tracking), and general-domain reasoning (MMLU). For 124M models, SIMP → HW leads on BLiMP, while HW → SIMP achieves the highest score on Entity Tracking. At 256M, performance differences are narrower, with small trade-offs across tasks.

suggests that smaller models benefit from a simple-to-complex progression in surface-level complexity, whereas larger models can internalize both variants without explicit ordering. **We view these as hypotheses under our training regime: ordering appears to matter more in low-capacity settings and is less critical when models scale up.**

5 Conclusion

This work examines whether augmenting pretraining data with LLM-based simplifications and ordering by text complexity improves representation quality in data-constrained settings. Through controlled data and training conditions, we isolated the effects of text complexity in curriculum design across two model sizes.

Our findings suggest three takeaways:

- (1) Adding simplified data outperforms repeating human-written text, yielding modest gains at 124M and clearer benefits at 256M.
- (2) Curriculum effects depend on model scale: smaller models benefit slightly more from a simple-to-complex curriculum, while larger models favor balanced exposure via interleaving.
- (3) Differences are most evident in zero-shot and low-resource fine-tuning scenarios, where schedule choice impacts representational quality, improving zero-shot performance and sample efficiency in small fine-tuning budgets.

Overall, our results show that in data-constrained settings, text simplification and curriculum learning can complement repeated exposure. Although the gains are modest, they suggest practical ways to extend the utility of pretraining without fresh data collection. Future work includes testing longer training horizons, exploring rewriting methods beyond simplification, and evaluating schedule effects at larger scales and in non-English domains.

Limitations

Our study is intentionally narrow in scope. We test only two model sizes (124M and 256M) with a single decoder-only architecture, leaving open how schedule effects scale to larger models or alternative designs. Training is constrained to a fixed budget of roughly 4B tokens (two full passes over a 2B-token human-written corpus), so results may differ under longer horizons or larger-scale training. We also study only one form of rewriting

(LLM-based simplification) while other transformation types (e.g., paraphrasing, elaboration, style transfer) may behave differently. Finally, the evaluation focuses on educational data in English and a limited set of NLU and zero-shot benchmarks, the results may not transfer directly to other domains, languages, or broader task families.

References

- Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebron, and Sumit Sanghai. 2023. [GQA: Training generalized multi-query transformer models from multi-head checkpoints](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4895–4901, Singapore. Association for Computational Linguistics.
- Loubna Ben Allal, Anton Lozhkov, Guilherme Penedo, Thomas Wolf, and Leandro von Werra. 2024. [Cosmopedia](#).
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. [Curriculum learning](#). In *Proceedings of the 26th International Conference on Machine Learning (ICML)*, pages 41–48. ACM.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. 2020. [Piqa: Reasoning about physical commonsense in natural language](#). In *Thirty-Fourth AAAI Conference on Artificial Intelligence*.
- Lucas Charpentier, Leshem Choshen, Ryan Cotterell, Mustafa Omer Gul, Michael Hu, Jaap Jumelet, Tal Linzen, Jing Liu, Aaron Mueller, Candace Ross, Raj Sanjay Shah, Alex Warstadt, Ethan Wilcox, and Adina Williams. 2025. [BabyLM turns 3: Call for papers for the 2025 babyLM workshop](#). *Preprint*, arXiv:2502.10645.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#). *arXiv:1803.05457v1*.
- DatologyAI, :, Pratyush Maini, Vineeth Dorna, Parth Doshi, Aldo Carranza, Fan Pan, Jack Urbanek, Paul Burstein, Alex Fang, Alvin Deng, Amro Abbas, Brett Larsen, Cody Blakeney, Charvi Bannur, Christina Baek, Darren Teh, David Schwab, Haakon Mongstad, and 12 others. 2025. [Beyondweb: Lessons from scaling synthetic data for trillion-scale pretraining](#). *Preprint*, arXiv:2508.10975.
- Kazuki Fujii, Yukito Tajima, Sakae Mizuki, Hinari Shimada, Taihei Shiotani, Koshiro Saito, Masanari Ohi, Masaki Kawamura, Taishi Nakamura, Takumi Okamoto, Shigeki Ishida, Kakeru Hattori, Youmi Ma, Hiroya Takamura, Rio Yokota, and Naoaki Okazaki. 2025. [Rewriting pre-training data boosts llm performance in math and code](#). *Preprint*, arXiv:2505.02881.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, and 5 others. 2024. [The language model evaluation harness](#).
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. 2023. [Textbooks are all you need](#). *Preprint*, arXiv:2306.11644.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). *Preprint*, arXiv:2009.03300.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Kathryn Millican, Bogdan Damoc, Arthur Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, and 2 others. 2022. [Training compute-optimal large language models](#). In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Anna A. Ivanova, Aalok Sathe, Benjamin Lipkin, Unnathi Kumar, Setayesh Radkani, Thomas H. Clark, Carina Kauf, Jennifer Hu, R. T. Pramod, Gabriel Grand, Vivian Paulun, Maria Ryskina, Ekin Akyurek, Ethan Wilcox, Nafisa Rashid, Leshem Choshen, Roger Levy, Evelina Fedorenko, Joshua Tenenbaum, and Jacob Andreas. 2024. [Elements of world knowledge \(ewok\): A cognition-inspired framework for evaluating basic world knowledge in language models](#). *arXiv preprint arXiv:2405.09605*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#). *arXiv preprint arXiv:2001.08361*.
- Najoung Kim and Sebastian Schuster. 2023. [Entity tracking in language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3835–3855, Toronto, Canada. Association for Computational Linguistics.
- Tom Kocmi and Ondřej Bojar. 2017. [Curriculum learning and minibatch bucketing in neural machine translation](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 379–386, Varna, Bulgaria. INCOMA Ltd.

- Jinghui Liu and Anthony Nguyen. 2024. [Rephrasing electronic health records for pretraining clinical language models](#). In *Proceedings of the 22nd Annual Workshop of the Australasian Language Technology Association*, pages 164–172, Canberra, Australia. Association for Computational Linguistics.
- Zechun Liu, Changsheng Zhao, Forrest Iandola, Chen Lai, Yuandong Tian, Igor Fedorov, Yunyang Xiong, Ernie Chang, Yangyang Shi, Raghuraman Krishnamoorthi, Liangzhen Lai, and Vikas Chandra. 2024. [Mobilellm: Optimizing sub-billion parameter language models for on-device use cases](#). *Preprint*, arXiv:2402.14905.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). *Preprint*, arXiv:1711.05101.
- Pratyush Maini, Skyler Seto, Richard Bai, David Granger, Yizhe Zhang, and Navdeep Jaitly. 2024. [Rephrasing the web: A recipe for compute and data-efficient language modeling](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14044–14072, Bangkok, Thailand. Association for Computational Linguistics.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. [Can a suit of armor conduct electricity? a new dataset for open book question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391, Brussels, Belgium. Association for Computational Linguistics.
- Niklas Muennighoff, Alexander M. Rush, Boaz Barak, Teven Le Scao, Aleksandra Piktus, Nouamane Tazi, Sampo Pyysalo, Thomas Wolf, and Colin Raffel. 2025. [Scaling data-constrained language models](#). *Preprint*, arXiv:2305.16264.
- Marwa Naïr, Kamel Yamani, Lynda Lhadj, and Riyadh Baghdadi. 2024. [Curriculum learning for small code language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 390–401, Bangkok, Thailand. Association for Computational Linguistics.
- Thao Nguyen, Yang Li, Olga Golovneva, Luke Zettlemoyer, Sewoong Oh, Ludwig Schmidt, and Xian Li. 2025. [Recycling the web: A method to enhance pre-training data quality and quantity for language models](#). *Preprint*, arXiv:2506.04689.
- Miyu Oba, Akari Haga, Akiyo Fukatsu, and Yohei Oseki. 2023. [BabyLM challenge: Curriculum learning based on sentence complexity approximating language acquisition](#). In *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pages 290–297, Singapore. Association for Computational Linguistics.
- Guilherme Penedo, Hynek Kydliček, Loubna Ben Alal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro von Werra, and Thomas Wolf. 2024. [The fineweb datasets: Decanting the web for the finest text data at scale](#). *arXiv preprint arXiv:2406.17557*.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. [Social IQa: Commonsense reasoning about social interactions](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4463–4473, Hong Kong, China. Association for Computational Linguistics.
- Noam Shazeer. 2020. [Glu variants improve transformer](#). *Preprint*, arXiv:2002.05202.
- Dan Su, Kezhi Kong, Ying Lin, Joseph Jennings, Brandon Norrick, Markus Kliegl, Mostofa Patwary, Mohammad Shoeybi, and Bryan Catanzaro. 2025. [Nemotron-CC: Transforming Common Crawl into a refined long-horizon pretraining dataset](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2459–2475, Vienna, Austria. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Yulia Tsvetkov, Manaal Faruqui, Wang Ling, Brian MacWhinney, and Chris Dyer. 2016. [Learning the curriculum with Bayesian optimization for task-specific word representation learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 130–139, Berlin, Germany. Association for Computational Linguistics.
- Dan John Velasco and Matthew Theodore Roque. 2025. [Rethinking the role of text complexity in language model pretraining](#). *arXiv preprint arXiv:2509.16551*.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [Hellaswag: Can a machine really finish your sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Xuan Zhang, Gaurav Kumar, Huda Khayrallah, Kenton Murray, Jeremy Gwinnup, Marianna J Martindale, Paul McNamee, Kevin Duh, and Marine

Carpuat. 2018. [An empirical exploration of curriculum learning for neural machine translation](#). *Preprint*, arXiv:1811.00739.

A Per-task Results of Fine-tuning Evaluation

Results are presented by fine-tuning budget: Full (Table 5), Small (Table 6), and Tiny (Table 7).

Yang Zhang, Amr Mohamed, Hadi Abdine, Guokan Shang, and Michalis Vazirgiannis. 2025. [Beyond random sampling: Efficient language model pretraining via curriculum learning](#). *Preprint*, arXiv:2506.11300.

Yikai Zhou, Baosong Yang, Derek F. Wong, Yu Wan, and Lidia S. Chao. 2020. [Uncertainty-aware curriculum learning for neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6934–6944, Online. Association for Computational Linguistics.

| Model | BoolQ | MNLI | MRPC | QQP | RTE | Avg. |
|-------------|------------|------------|------------|------------|------------|------|
| 124M | | | | | | |
| BASELINE | 70.6 ± 0.3 | 66.2 ± 0.7 | 82.2 ± 1.7 | 80.9 ± 0.3 | 66.4 ± 3.1 | 70.9 |
| INTERLEAVED | 69.5 ± 0.6 | 67.5 ± 0.2 | 80.6 ± 1.0 | 81.1 ± 0.3 | 69.2 ± 2.4 | 71.1 |
| SIMP → HW | 70.7 ± 0.6 | 67.1 ± 0.7 | 82.2 ± 2.2 | 81.3 ± 0.1 | 68.1 ± 2.8 | 71.4 |
| HW → SIMP | 71.2 ± 0.7 | 66.7 ± 0.7 | 83.3 ± 0.7 | 80.6 ± 0.1 | 66.2 ± 1.7 | 71.1 |
| 256M | | | | | | |
| BASELINE | 71.2 ± 0.1 | 69.2 ± 0.5 | 81.9 ± 1.1 | 81.5 ± 0.1 | 65.7 ± 3.1 | 73.9 |
| INTERLEAVED | 71.0 ± 0.3 | 69.8 ± 0.1 | 84.1 ± 1.7 | 82.2 ± 0.1 | 71.5 ± 1.8 | 75.7 |
| SIMP → HW | 71.0 ± 0.8 | 70.3 ± 0.3 | 82.9 ± 1.0 | 82.0 ± 0.1 | 70.1 ± 2.5 | 75.3 |
| HW → SIMP | 71.8 ± 0.2 | 69.5 ± 0.3 | 82.2 ± 1.0 | 82.0 ± 0.2 | 71.8 ± 2.6 | 75.5 |

Table 5: Per-task fine-tuning results under **Full** fine-tuning budget. “Avg.” is the mean across tasks. Results are grouped by model size (124M and 256M) and curriculum strategy.

| Model | BoolQ | MNLI | MRPC | QQP | RTE | Avg. |
|-------------|------------|------------|------------|------------|------------|------|
| 124M | | | | | | |
| BASELINE | 62.9 ± 1.0 | 63.1 ± 0.4 | 70.8 ± 1.0 | 78.1 ± 0.9 | 63.4 ± 4.5 | 67.7 |
| INTERLEAVED | 62.4 ± 0.3 | 64.0 ± 0.3 | 72.4 ± 1.0 | 78.7 ± 0.0 | 63.7 ± 1.6 | 68.2 |
| SIMP → HW | 64.1 ± 0.6 | 63.1 ± 0.3 | 71.5 ± 0.7 | 79.0 ± 0.2 | 63.9 ± 0.7 | 68.3 |
| HW → SIMP | 62.7 ± 0.7 | 63.1 ± 0.2 | 73.2 ± 0.3 | 78.3 ± 0.4 | 60.2 ± 3.3 | 67.5 |
| 256M | | | | | | |
| BASELINE | 63.9 ± 0.7 | 65.1 ± 0.2 | 72.8 ± 2.3 | 79.3 ± 0.1 | 63.2 ± 2.4 | 68.8 |
| INTERLEAVED | 63.7 ± 0.6 | 65.8 ± 0.5 | 73.1 ± 1.7 | 80.1 ± 0.2 | 71.1 ± 2.8 | 70.7 |
| SIMP → HW | 64.2 ± 0.3 | 66.1 ± 0.7 | 71.5 ± 1.7 | 79.7 ± 0.1 | 66.0 ± 0.7 | 69.5 |
| HW → SIMP | 63.7 ± 0.2 | 65.4 ± 0.7 | 71.6 ± 3.8 | 79.4 ± 0.2 | 65.3 ± 4.9 | 69.1 |

Table 6: Per-task fine-tuning results under **Small** fine-tuning budget. “Avg.” is the mean across tasks. Results are grouped by model size (124M and 256M) and curriculum strategy.

| Model | BoolQ | MNLI | MRPC | QQP | RTE | Avg. |
|-------------|------------|------------|------------|------------|------------|------|
| 124M | | | | | | |
| BASELINE | 59.5 ± 0.7 | 59.4 ± 1.0 | 67.9 ± 0.3 | 75.3 ± 0.4 | 59.7 ± 6.1 | 64.4 |
| INTERLEAVED | 58.4 ± 0.4 | 60.7 ± 0.3 | 70.4 ± 1.5 | 76.0 ± 0.2 | 58.3 ± 1.2 | 64.8 |
| SIMP → HW | 58.8 ± 1.0 | 59.8 ± 0.4 | 68.9 ± 2.6 | 76.5 ± 0.1 | 62.0 ± 0.8 | 65.2 |
| HW → SIMP | 59.5 ± 1.5 | 60.3 ± 0.1 | 71.6 ± 0.8 | 75.6 ± 0.5 | 55.1 ± 3.3 | 64.4 |
| 256M | | | | | | |
| BASELINE | 60.8 ± 1.3 | 60.7 ± 0.4 | 68.1 ± 1.2 | 76.5 ± 0.6 | 59.0 ± 1.8 | 65.0 |
| INTERLEAVED | 60.1 ± 0.7 | 61.6 ± 0.5 | 71.0 ± 1.9 | 77.5 ± 0.2 | 64.6 ± 0.7 | 66.9 |
| SIMP → HW | 60.5 ± 1.0 | 60.9 ± 0.9 | 68.6 ± 1.8 | 76.9 ± 0.3 | 61.1 ± 2.5 | 65.6 |
| HW → SIMP | 59.9 ± 0.6 | 61.4 ± 0.4 | 65.5 ± 4.1 | 76.2 ± 0.3 | 60.9 ± 4.2 | 64.8 |

Table 7: Per-task fine-tuning results under **Tiny** fine-tuning budget. “Avg.” is the mean across tasks. Results are grouped by model size (124M and 256M) and curriculum strategy.