

WRAP++: Web Discovery Amplified Pretraining

Jiang Zhou^{1,2} Yunhao Wang¹ Xing Wu^{1,†} Tinghao Yu^{1,†} Feng Zhang¹

¹Hunyuan Team, Tencent

²Georgia Institute of Technology

[†]Correspondence: {ucaswu,maxwellyu}@tencent.com

Abstract

Synthetic data rephrasing has emerged as a powerful technique for enhancing knowledge acquisition during large language model (LLM) pretraining. However, much of web-scale synthetic rephrasing operates at the *single-document* level, rewriting individual web pages in isolation. This confines synthesized examples to intra-document knowledge, missing cross-document relationships and leaving facts with limited associative context. We propose **WRAP++** (Web discover**R**y Amplified Pretraining), which *amplifies* the associative context of factual knowledge by *discovering* cross-document relationships from web hyperlinks and synthesizing joint QA over each discovered document pair. Concretely, **WRAP++** discovers high-confidence relational motifs including dual-links ($A \leftrightarrow B$) and co-mentions ($A \rightarrow E \leftarrow B$ with $A \rightarrow B$), and synthesizes QA that requires reasoning across both documents. This produces relational knowledge absent from either source document alone, creating diverse entry points to the same facts. Because the number of valid entity pairs grows combinatorially, this discovery-driven synthesis also amplifies data scale far beyond single-document rewriting. Instantiating **WRAP++** on Wikipedia, we amplify $\sim 8.4\text{B}$ tokens of raw text into 80B tokens of cross-document QA data. On SimpleQA, OLMo-based models at both 7B and 32B scales trained with **WRAP++** substantially outperform single-document approaches and exhibit sustained scaling gains, underscoring the advantage of cross-document knowledge discovery and amplification.

1 Introduction

Synthetic data has become an increasingly important component of large language model (LLM) pretraining. WRAP (Maini et al., 2024) showed that rephrasing noisy web text into QA format can improve pretraining, and later systems scaled this recipe substantially: Nemotron-CC (Su et al., 2025)

produced ~ 2 trillion synthetic tokens from Common Crawl, Phi-4 (Abdin et al., 2024) used 40% synthetic data in pretraining, and Qwen3 (Yang et al., 2025a) incorporated synthetic data into its training pipeline.

However, much of the web-scale synthetic pretraining literature has focused on intra-document factors such as rephrasing strategy, generator model, and source quality (Nguyen et al., 2025; Niklaus et al., 2026). These methods vary *how* a single document is rewritten while largely preserving *what* evidence is synthesized. Because many factual associations are distributed across multiple documents, this single-document paradigm can leave relational context underrepresented, limiting the ways in which a model can recover and connect related facts.

This limitation motivates a shift toward *cross-document synthesis*: bringing multiple facts into a shared context to learn relational knowledge jointly. However, moving to cross-document synthesis is non-trivial. A naïve approach of randomly pairing documents yields little improvement over single-document baselines (see §4.1), as forcing an LLM to synthesize joint QA from unrelated texts produces fabricated connections and low-quality data. Thus, the *document selection* mechanism is critical: cross-document synthesis only succeeds when the paired documents contain genuinely related facts.

Web hyperlinks provide a broad relevance signal by encoding human-curated judgments of importance (Zhou et al., 2022). For example, the Wikipedia pages of composers Hans Zimmer and Ludwig Göransson are topologically linked through shared collaborations with director Christopher Nolan. While single-document synthesis might only extract isolated facts (e.g., "Göransson won an Oscar for Oppenheimer"), **WRAP++** pairs these connected documents to synthesize multi-hop relational QA. As illustrated in Figure 1, the model is forced to explicitly reason across both

texts—deducing that Zimmer left Tenet to score Dune (his second Oscar), leading to Göransson’s hiring and subsequent second Oscar for Oppenheimer. This explicit relational supervision provides vital disambiguation cues and diverse retrieval paths absent from single-document synthesis, saving the LLM from having to implicitly infer these complex connections from massive unstructured text.

Based on this insight, we propose **WRAP++** (**Web discovRy Amplified Pretraining**), which effectively extends the synthetic data paradigm from single-document *rewriting* to cross-document *discovery and amplification*.

Amplifying Associative Context via Relation Discovery. We discover high-confidence relational motifs from web hyperlinks (Zhou et al., 2022)—specifically *dual-links* ($A \leftrightarrow B$) and *co-mentions* ($A \rightarrow E \leftarrow B$ with $A \rightarrow B$). Rather than simply concatenating the documents, we feed each discovered pair to an instruction-tuned LLM generator under three synthesis constraints: **Cross-Document Dependency** requires joint reasoning over both documents, **Explicit Factual Chaining** makes the underlying multi-hop reasoning explicit, and **Context-Free Internalization** presents the resulting QA as standalone knowledge rather than passage-grounded reading comprehension. This process produces relational knowledge—including comparisons, contrasts, and bridging facts—that creates diverse retrieval paths to the same facts. Moreover, because the relation space grows over document pairs rather than individual documents, this discovery-driven approach provides substantial headroom for scaling synthetic data.

We instantiate **WRAP++** on Wikipedia, achieving approximately $10\times$ data amplification by expanding $\sim 8.4\text{B}$ tokens of raw text into 80B tokens of cross-document QA data, compared with only $\sim 5.4\text{B}$ tokens produced by single-document WRAP. On SimpleQA (Wei et al., 2024), OLMo-based models (Olmo et al., 2025) at both 7B and 32B scales trained with **WRAP++** substantially outperform all single-document baselines and exhibit a more favorable scaling trajectory.

Our contributions are three-fold:

1. We propose **WRAP++**, a scalable framework for cross-document synthetic pretraining that uses naturally web hyperlink topology as a model-free signal for relation discovery.
2. We instantiate **WRAP++** on Wikipedia and

synthesize $80\text{B}+$ tokens of cross-document QA data from an 8.4B -token source corpus, showing that topology-guided combinatorial relation discovery enables data amplification far beyond single-document synthesis.

3. We show that **WRAP++** improves factual knowledge acquisition at scale: on SimpleQA, OLMo-based 7B and 32B models trained with **WRAP++** substantially outperform single-document baselines with sustained scaling gains, and **WRAP++** can be integrated into a realistic mid-training mixture without degrading general capabilities.

2 Method: WRAP++

WRAP++ is a framework that transitions synthetic pretraining data from single-document rewriting to topology-guided cross-document discovery and amplification. The framework consists of two core stages: Topological Relation Discovery (§2.2) and Joint QA Synthesis (§2.3). Figure 1 illustrates the overall pipeline.

2.1 Problem Formulation and Graph Abstraction

Let $\mathcal{D} = \{d_1, d_2, \dots, d_N\}$ denote a large-scale web corpus consisting of N documents. The inter-document references (e.g., hyperlinks) naturally induce a directed knowledge graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where each vertex $v_i \in \mathcal{V}$ corresponds to a document d_i , and a directed edge $e_{i,j} \in \mathcal{E}$ exists if d_i explicitly references d_j . Since each document in our corpus describes a single entity, we use “entity pair” and “document pair” interchangeably throughout this paper.

Conventional single-document synthesis operates solely on the local context of v_i , limiting the model’s exposure to isolated facts. In contrast, **WRAP++** leverages the topological structure of $\mathcal{G}$ to discover genuine semantic dependencies across documents, bringing related knowledge into a shared synthesis context to amplify the associative context of factual knowledge.

2.2 Topological Relation Discovery

A naïve approach of pairing random documents from $\mathcal{D}$ forces the synthesis model to hallucinate spurious connections. To ensure the semantic validity of cross-document synthesis, we discover high-confidence relational motifs directly from $\mathcal{G}$.




Figure 1: Overview of the **WRAP++** pipeline. Unlike single-document WRAP, which rewrites individual documents, **WRAP++** *discovers* cross-document entity relationships from web topology and *amplifies* them into pretraining data through joint QA synthesis.

We focus on two topological structures that provide strong inductive biases for relational reasoning:

Dual-link Motif. Two documents u and v form a dual-link relationship if they mutually reference each other. Formally, a dual-link pair (u, v) satisfies the bidirectional constraint:

$$e_{u,v} \in \mathcal{E} \wedge e_{v,u} \in \mathcal{E} \quad (1)$$

This mutual dependency typically indicates a strong, foundational semantic correlation (e.g., a notable director and their magnum opus, or a scientist and their core discovery). Discovering this motif ensures the underlying entity pair is highly coupled.

Co-mention Motif. Documents u and v share a co-mention relationship if they both reference a common structural hub E , while maintaining a direct link between themselves. Formally, the triplet (u, v, E) satisfies:

$$e_{u,E} \in \mathcal{E} \wedge e_{v,E} \in \mathcal{E} \wedge e_{u,v} \in \mathcal{E} \quad (2)$$

The shared structural context E imposes implicit analogical, hierarchical, or comparative relationships (e.g., two competing theories cited in the

same survey article). This motif explicitly encourages the subsequent synthesis model to generate relational knowledge that contrasts and compares the related entities, thereby amplifying their associative context.

2.3 Cross-Document Joint QA Synthesis

Given a discovered document pair (d_u, d_v) connected by a valid topological motif, we employ an instruction-tuned LLM generator $\mathcal{M}_\theta$ to synthesize a set of composite QA instances $\mathcal{S}_{u,v} = \{(q_i, c_i, a_i)\}_{i=1}^K$, where q_i is the question, c_i is the intermediate factual chain, and a_i is the final answer.

The generation process is conditioned on a structured prompt $\mathcal{P}$ and the concatenated document context:

$$\mathcal{S}_{u,v} \sim \mathcal{M}_\theta(d_u \oplus d_v, \mathcal{P}) \quad (3)$$

To amplify the associative context of the synthesized data and prevent the generator from degrading to shallow single-document summarization, $\mathcal{P}$ enforces three functional constraints on the output space:

- **Strict Cross-Document Dependency:** The generated question q_i should not be answer-

able from either document alone. Deriving the correct answer a_i must require complementary premises from both d_u and d_v , ensuring that the synthesis captures cross-document associations rather than merely rephrasing isolated single-document facts.

- **Explicit Factual Chaining:** Before outputting a_i , the generator must explicitly decode the traversal path c_i . By articulating the necessary facts extracted from both documents and linking them step-by-step, the pre-training model internalizes multi-hop knowledge structures, effectively creating diverse associative entry points to the underlying facts.
- **Context-Free Internalization:** The generator should not attribute facts to the provided passages (e.g., “According to Passage A”) or rely on local passage labels. Since the source documents are assumed to have already appeared in pretraining, the synthesized QA presents cross-document relations as standalone knowledge, reinforcing associations among known facts rather than repeating the original context as a reading-comprehension exercise.

3 Experiments

3.1 Experimental Setup

Synthesis Data. In principle, **WRAP++** is a general framework applicable to any text corpus containing hyperlinks. In this work, we instantiate it on Wikipedia because of its rich link structure and its widespread use in prior work on synthetic rewriting methods (Maini et al., 2024; Su et al., 2025). Specifically, we use the English subset of FineWiki (Penedo, 2025) as our base corpus ($\mathcal{D}$), which contains approximately 8.4B tokens. We parse the hyperlinks in FineWiki to construct the directed inter-document graph $\mathcal{G}$ used for topological relation discovery.

Synthesis Model. We use Qwen3-30B-A3B-Instruct-FP8 as our instruction-tuned generator $\mathcal{M}_\theta$. We select this mixture-of-experts model because its $\sim 3\text{B}$ active parameters per token make synthesis at the 82.7B-token scale practical; compared with Qwen3-235B-A22B, it uses $\sim 7.3\times$ fewer active parameters, with a 2.59-point pass@128 trade-off in our controlled 8B-token ablation (§4.4). The prompt is designed to enforce the strict cross-document dependency and

explicit factual chaining constraints described in §1, thereby encouraging high-quality relational QA generation. The full prompt template is provided in Appendix F.

Data Scale. Topological relation discovery substantially amplifies the data scale beyond individual documents. The dual-link motif yields highly coupled entity pairs that produce $\sim 3\text{B}$ tokens of cross-document QA data. Incorporating the co-mention motif broadens coverage, bringing the combined **WRAP++** dataset to $\sim 82.7\text{B}$ tokens.

Training Models. To assess the effect of cross-document synthetic data on parametric knowledge, we continue pretraining from the OLMo-3 stage-1 last checkpoint at both the 7B and 32B scales for one epoch. We choose OLMo because it is fully open-source and provides publicly released checkpoints throughout training, making it a suitable platform for controlled continued-pretraining experiments.

3.2 Evaluation Setup

Benchmark. We primarily evaluate on SimpleQA (Wei et al., 2024), a knowledge-intensive benchmark designed to measure short-form factual accuracy while minimizing sensitivity to formatting heuristics. Most SimpleQA questions can be answered directly from Wikipedia-derived knowledge, making it a natural testbed for studying factual knowledge acquisition under our setup. To assess transfer beyond short-form single-hop questions, we additionally evaluate the final token-matched OLMo-3-7B checkpoints on HotpotQA (Yang et al., 2018) and MuSiQue (Trivedi et al., 2022) in a closed-book setting (Appendix E).

Metric. We use pass@128 as our primary metric, defined as the probability that at least one of 128 sampled responses contains the correct fact. Our goal is to measure *knowledge recoverability* rather than only top-1 answer accuracy. In this setting, pass@128 is useful because it probes whether the correct factual association can be elicited from the model under sampling, even when it is not the single most likely surface form. We therefore treat it as a more sensitive indicator of parametric knowledge recoverability during continued pretraining than pass@1. This recoverability also represents useful headroom for subsequent post-training (Yang et al., 2026; Zhou et al., 2026): reinforcement learning can shift probability mass toward correct fac-

Table 1: SimpleQA results after 1-epoch continued training on OLMo-3-7B and OLMo-3-32B using different data recipes. We report pass@128 (%).

Data Recipe	OLMo-3-7B	OLMo-3-32B
<i>Pretrained Base</i>	34.76	42.35
+ WRAP	39.55	44.43
+ Extended WRAP	43.69	47.91
+ WRAP++	49.13	53.97

tual responses, potentially converting recoverable knowledge into higher pass@1 performance (Jiang et al., 2026a; Yuan et al., 2026; Fan et al., 2026; Jiang et al., 2026b).

Baselines. To isolate the benefit of cross-document synthesis, we compare **WRAP++** against two closely related single-document baselines derived from the same FineWiki corpus:

1. **WRAP (~5.4B tokens):** Standard single-document QA synthesis following the original WRAP recipe (Maini et al., 2024), representing the typical yield of intra-document fact extraction.
2. **Extended WRAP (~17.4B tokens):** An expanded single-document synthesis utilizing additional prompting strategies (e.g., exhaustive extraction) to push the limits of single-document scaling. For the token-matched OLMo-3-7B comparison, we cycle this corpus for approximately 4.6 passes to reach 80B consumed tokens, using the same batch size and learning-rate schedule as **WRAP++**. The full synthesis prompt is provided in Appendix F.

These corpus sizes illustrate the data-scaling constraint of single-document synthesis: on the same FineWiki corpus, standard WRAP yields only ~5.4B tokens and Extended WRAP reaches ~17.4B tokens, both well below the ~82.7B-token scale of **WRAP++**. Cycling Extended WRAP to 80B consumed tokens lets us separately test whether additional training on repeated single-document data can account for the scaling gains of **WRAP++**.

3.3 Main Results

Table 1 presents the main results of 1-epoch continued training across different synthesis recipes. We highlight two principal findings.

WRAP++ substantially outperforms single-document baselines. Across both model scales, continued pretraining with **WRAP++** yields substantially higher pass@128 on SimpleQA compared to all single-document approaches (+9.5 pp on 7B, +9.8 pp on 32B over WRAP; +5.4 pp on 7B, +6.1 pp on 32B over Extended WRAP). This advantage reflects two complementary factors. First, cross-document synthesis produces higher-quality relational knowledge *per token*: at the first token-matched 8B checkpoint, **WRAP++** already outperforms Extended WRAP by 1.95 points. Second, the combinatorial nature of relation discovery supplies a much larger space of distinct training examples. The token-matched OLMo-3-7B comparison shows that this advantage persists and generally widens through 80B consumed tokens despite identical token budgets and training schedules (Figure 2).

Surpassing the single-document scaling bottleneck. Single-document methods face a data-diversity bottleneck: the finite number of facts that can be extracted from each page. Extended WRAP pushes this limit through exhaustive extraction but produces only ~17.4B tokens. Our token-matched OLMo-3-7B comparison tests whether the remaining gap can be closed by training longer by cycling this corpus for approximately 4.6 passes. The repeated single-document data neither closes the gap nor produces the sustained improvement of **WRAP++**, indicating that additional exposure cannot substitute for a larger set of distinct cross-document associations. We analyze these scaling dynamics in detail next.

3.4 Scaling and Training Dynamics

Figure 2 plots pass@128 as a function of consumed training tokens. Most importantly, the token-matched OLMo-3-7B comparison equalizes the consumed-token budget, batch size, and learning-rate schedule through 80B tokens. Extended WRAP remains in a narrow 43.37–44.84 range as its 17.4B-token corpus is repeated, while **WRAP++** improves from 45.32 at 8B tokens to 49.18 at 80B tokens. **WRAP++** is better at every token-matched checkpoint, with the gap widening from 1.95 points at 8B to 5.53 points at 80B (and reaching 5.62 points at 56B). This controlled comparison rules out additional training tokens alone as the source of the gain and instead supports the value of the distinct cross-document associations enabled by




Figure 2: **SimpleQA pass@128 vs. consumed training tokens.** Solid and dashed lines denote OLMo-3-7B and OLMo-3-32B, respectively. WRAP and Extended WRAP contain ~ 5.4 B and ~ 17.4 B tokens. In the token-matched OLMo-3-7B comparison, we cycle Extended WRAP for approximately 4.6 passes to 80B consumed tokens under the same batch size and learning-rate schedule as **WRAP++**. **WRAP++** leads at all ten token-matched checkpoints by 1.95–5.62 points, showing that its scaling gain is not explained by training-token count alone.

relation discovery.

To further dissect how this scaling translates into knowledge recoverability, we track the evolution of the pass@ k curves throughout training. Figure 3 visualizes the unbiased SimpleQA pass@ k for both models. As training progresses (indicated by the light-to-dark blue gradient representing the 80B token influx), the curves exhibit a monotonic upward shift across all values of k .

Crucially, this improvement spans the entire logarithmic k -spectrum. The persistent lift at small k (the leftmost regions of the curves) shows that the model’s top-ranked answers increasingly contain the correct fact, reflecting higher precision. Simultaneously, the parallel gains at larger k indicate a broader and more robust set of associative retrieval paths to the same knowledge.

Table 2: Closed-book multi-hop factual QA at the final 80B-token checkpoint of the token-matched OLMo-3-7B comparison. We report pass@8 (%).

Data Setting	HotpotQA	MuSiQue
Extended WRAP	45.47	19.20
WRAP++	48.12	19.61

Robustness beyond SimpleQA pass@128. Table 2 shows that **WRAP++** also outperforms Extended WRAP on closed-book HotpotQA and MuSiQue by 2.65 and 0.41 pass@8 points, respectively, extending its advantage to multi-hop factual questions. Increasing the generation budget to 256 samples per SimpleQA question yields the same conclusion: **WRAP++** leads at all ten token-matched checkpoints on both pass@128 and pass@256, with a final pass@256 margin of 6.78 points (56.42 vs. 49.64). Appendix E reports the pretrained multi-hop baseline, complete pass@256 results, and evaluation details.

4 Ablations and Analysis

We conduct extensive ablations to validate each component of **WRAP++**. Due to experimental costs, unless otherwise noted, all ablation experiments use OLMo-3-7B continued pretraining with ~ 8 B tokens and report the results in terms of SimpleQA pass@128.

4.1 Necessity of Topological Relation Discovery

We explore whether the topological relation discovery is essential, or whether randomly pairing Wikipedia pages would suffice. Table 3 (Pairing




Figure 3: **Evolution of $\text{pass}@k$ performance during training.** The curves illustrate the unbiased $\text{pass}@k$ of OLMo-3-7B (a) and OLMo-3-32B (b) on SimpleQA ($k \in [1, 128]$ in log scale). The color gradient (from light to dark blue) tracks the accumulation of consumed **WRAP++** tokens (from 10B to 80B). The strictly monotonic upward shift across all values of k indicates robust, unsaturated knowledge internalization.

Table 3: Ablation results across different design choices (OLMo-3-7B, ~ 8 B tokens). We report SimpleQA $\text{pass}@128$. **WRAP++** (default) uses topological relation discovery, combined topologies, joint QA synthesis, and the Qwen3-30B-A3B synthesis model.

Ablation Axis	Variant	$\text{pass}@128$
Pairing strategy	Random entity pairing	43.46
	Topological discovery	45.11
Relation topology	Dual-link only	44.24
	Co-mention only	44.36
	Combined	45.11
Synthesis Method	Raw document concatenation	35.43
	QA with source docs prepended	38.93
	Cross-document joint QA	45.11
Synthesis model	Qwen3-30B-A3B	45.11
	Qwen3-235B-A22B	47.70

Strategy rows) shows a 1.65-point performance drop (from 45.11 to 43.46) when entities are paired randomly rather than via dual-link or co-mention relations. A question-level paired bootstrap test confirms that this difference is statistically significant (95% CI: [0.31, 2.99]; Appendix D). Qualitatively, random pairing forces the synthesis model to fabricate relationships between unrelated entities, producing factually incorrect comparisons and superficial connections. This confirms that principled relation discovery—specifically, dual-link and co-mention motif discovery—is important for high-quality cross-document synthesis.

4.2 Topology Comparison: Dual-Link vs. Co-Mention

We explore the contribution of each relation type at a matched token budget of ~ 8 B. As shown in

Table 3 (Relation Topology rows), both topologies provide strong relational signal at this budget, and their combination achieves the highest numerical performance (45.11). Paired bootstrap tests show that the combined topology significantly outperforms dual-link alone, whereas the combined-vs.-co-mention and dual-link-vs.-co-mention differences are not statistically conclusive (Appendix D). This pattern is consistent with the available data scale: the dual-link-only setting must reuse its ~ 3 B-token corpus within the 8B-token budget, whereas co-mention provides a much larger relation pool.

4.3 Necessity and Format of QA Synthesis

We explore the optimal data format for learning cross-document relationships by comparing our joint QA synthesis against two alternatives: (1) raw concatenation of related documents (no QA), and (2) prepending source documents to the synthesized QA pairs. Table 3 shows raw concatenation performs only marginally above the pretrained base (35.43 vs. 34.8), indicating that explicit *synthesis* is essential to convert document proximity into learnable relational knowledge. Moreover, prepending source documents to QA pairs causes a notable performance drop (38.93), likely by allowing the model to superficially copy answers rather than parametrically internalizing them. Thus, joint QA synthesis provides the most effective format for amplifying associative context.

Table 4: Performance comparison of **WRAP++** mixed with other single-document strategies. Each mixture contains $\sim 8\text{B}$ tokens total.

Single-Document Component	Proportion of WRAP++ in Mixture	
	0% (Baseline Only)	50% (1:1 Mix)
<i>Pretrained Model</i>	34.76	
Raw FineWiki	39.23	41.80
Distill	38.12	41.59
Nemotron-Extract	38.69	43.35
Knowledge List	38.60	42.03
WRAP++ (100%, Ours)	45.11	

4.4 Effect of Synthesis Model Scale

We explore the effect of synthesis model scale on **WRAP++** quality by comparing two synthesis models of different scales: Qwen3-30B-A3B and Qwen3-235B-A22B (22B active parameters). Table 3 (Synthesis Model Scale rows) shows that the larger model produces higher-quality cross-document QA, leading to better downstream pass@128. In practice, the choice between synthesis models involves a cost–quality tradeoff: the larger model is preferable when generation budget is not the bottleneck, while the smaller model enables broader coverage at lower compute cost.

4.5 Comparison with Other Single-Document Strategies

We contextualize the performance of **WRAP++** against other representative single-document rephrasing strategies Su et al. (2025) applied to the identical FineWiki source corpus, including: (a) **Distill**—rewriting into cleaner, more concise prose while preserving information; (b) **Nemotron-Extract**—extracting key factual statements and discarding redundancy using the Extract Knowledge strategy from Nemotron-CC; (c) **Knowledge List**—outputting structured knowledge in list format.

Table 4 presents both the isolated performance (0% and 100%) and the mixing dynamics (50% blending) at a strictly restricted budget of $\sim 8\text{B}$ tokens. When evaluated in isolation at the 8B-token budget, pure **WRAP++** (45.11) substantially outperforms all listed single-document baselines. This margin at a restricted data scale reveals an important dynamic in pretraining efficiency: discovery-driven synthesis already yields stronger knowledge recoverability, even before exploiting its data amplification headroom. The same pattern holds at the OLMo-3-32B scale under the same 8B-token budget: **WRAP++** reaches 51.24 pass@128, compared with 42.33 for Raw FineWiki and 42.79 for

Table 5: Integrating **WRAP++** into OLMo-3-7B under a fixed 100B-token mid-training budget. The **WRAP++** Mix replaces 6B tokens (6%) of the original mixture with **WRAP++** data. We report SimpleQA pass@128 and the average over 12 general benchmarks (see Appendix H).

Setting	SimpleQA	Gen. Avg
	pass@128	(12 tasks)
<i>Pretrained Base</i>	34.76	57.79
+ Midtrain (100B)	34.74	68.24
+ WRAP++ Mix (100B)	37.58	68.16

Nemotron-Extract (Appendix C). When blending **WRAP++** with other strategies in a 1:1 ratio, we observe a clear *uplift effect*: injecting **WRAP++** into any weaker baseline consistently improves upon its 0% counterpart (e.g., Knowledge List rises from 38.60 to 42.03), further highlighting the advantage of discovery-driven synthesis.

4.6 Integration with OLMo-3 Mid-Training Data

We further explore whether **WRAP++** data can be integrated into a realistic mid-training pipeline without harming general capabilities. We replace 6B tokens (6%) of OLMo-3’s original mid-training mixture with **WRAP++** data, keeping the total training budget fixed at 100B tokens. As a baseline, we train on the original mixture under otherwise identical conditions. Following Niklaus et al. (2026), we evaluate both SimpleQA and the average performance across 12 general tasks (detailed in Appendix H). The original OLMo-3 mixture spans diverse signals, including high-quality web data, code, mathematics, reasoning traces, PDFs, and explanatory, QA, and instruction data. As shown in Table 5, retaining 94% of this broad mixture and replacing 6% with knowledge-focused **WRAP++** data leaves the 12-task general average effectively unchanged (-0.08 points; 68.24 to 68.16), while improving SimpleQA by 2.84 points (34.74 to 37.58). The knowledge-intensive MMLU Redux benchmark also improves by 1.28 points (Appendix H). These results show that **WRAP++** strengthens factual knowledge within full-scale mid-training without meaningfully reducing general capabilities. Moreover, as the mid-training budgets of leading models grow toward the trillion-token regime (Yang et al., 2025a; Zeng et al., 2026), the 80B-token scale of **WRAP++** suggests strong

potential for integration into future large-scale training pipelines.

5 Related Work

Synthetic Data for Web-Scale LLM Pretraining. Synthetic data has been explored as a way to improve the quality and learnability of pretraining corpora, from generator-driven datasets such as TinyStories, Phi-style textbooks, and Cosmopedia (Eldan and Li, 2023; Gunasekar et al., 2023; Ben Allal et al., 2024) to web-grounded rephrasing methods. WRAP (Maini et al., 2024) showed that rewriting web documents into cleaner QA-style text can accelerate pretraining by $\sim 3\times$, motivating a line of work on synthetic web recycling. Subsequent methods improve the rephrasing recipe through richer target formats and quality-controlled rewriting (Su et al., 2025; Nguyen et al., 2025; Maini et al., 2025; Niklaus et al., 2026), study the effect of generator size (Maini et al., 2024, 2025), and analyze how source data quality affects synthetic data utility (Penedo et al., 2024; Li et al., 2024; Nguyen et al., 2025; Niklaus et al., 2026). Overall, this line of work primarily improves the quality, style, or density of individual documents, while leaving cross-document relational structure largely implicit.

Synthetic Data with Cross-Document Structure. Recent work has begun to make document or entity relations explicit during synthetic pretraining. EntiGraph (Yang et al., 2024) uses entity-centered synthesis to expand a small corpus for continued pretraining, while its synthesis is primarily driven by entity relations extracted from source documents, rather than explicit cross-document relation discovery. SoG (Ma et al., 2025) further emphasizes cross-document entity associations, using context-graph paths and entity-conditioned generation to improve multi-hop and domain-specific QA. SBP (Yang et al., 2025b) instead learns inter-document correlations implicitly by retrieving semantically similar document pairs and training a conditional synthesizer for pretraining from scratch. **WRAP++** is complementary to these efforts: it uses naturally available hyperlink topology as a model-free relational signal, avoiding costly entity extraction or embedding-based graph construction. This enables scalable cross-document synthesis over web-scale documents. On Wikipedia, **WRAP++** mines high-confidence relational motifs and converts them into joint QA data, amplifying an 8.4B-token source

corpus into 80B+ synthetic QA tokens.

6 Conclusion

We presented **WRAP++**, a scalable framework for cross-document synthetic pretraining. Instead of relying on costly entity extraction or embedding-based graph construction, **WRAP++** uses naturally available web hyperlink topology as a model-free signal for relation discovery. By mining high-confidence relational motifs such as dual-links and co-mentions from Wikipedia hyperlinks, **WRAP++** converts related document pairs into joint QA data that explicitly captures cross-document associations. This enables large-scale knowledge amplification, expanding an 8.4B-token source corpus into 80B+ synthetic tokens. Experiments on OLMo-3 7B and 32B models show that **WRAP++** substantially outperforms single-document baselines on SimpleQA, maintains favorable scaling up to 80B tokens, and can be integrated into a realistic mid-training mixture without degrading general capabilities.

Limitations

Our experiments instantiate **WRAP++** on Wikipedia, which is a clean and entity-centric corpus. Extending to noisier web corpora (e.g., Common Crawl), where hyperlinks include advertisements, navigation elements, and low-quality references, will require additional filtering heuristics; we are actively exploring this direction.

References

- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, and 8 others. 2024. [Phi-4 technical report](#). *Preprint*, arXiv:2412.08905.
- Loubna Ben Allal, Anton Lozhkov, Guilherme Penedo, Thomas Wolf, and Leandro von Werra. 2024. [Cosmopedia](#).
- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, and 1 others. 2020. Piqa: Reasoning about physical commonsense in natural language. *Proceedings of the AAAI conference on artificial intelligence*, 34(05):7432–7439.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg

- Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. [Evaluating large language models trained on code](#). *Preprint*, arXiv:2107.03374.
- 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 preprint arXiv:1803.05457*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, and 1 others. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Tri Dao. 2023. [Flashattention-2: Faster attention with better parallelism and work partitioning](#). *Preprint*, arXiv:2307.08691.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs. *arXiv preprint arXiv:1903.00161*.
- Ronen Eldan and Yuanzhi Li. 2023. Tinstories: How small can language models be and still speak coherent english? *arXiv preprint arXiv:2305.07759*.
- Zhiyuan Fan, Guanqiao Chen, Yanyi Huang, Mingkuan Zhao, Dadi Guo, and Yi R. Fung. 2026. [Learning diverse responses with prefix-conditioned supervised fine-tuning](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 247–276, San Diego, California, United States. Association for Computational Linguistics.
- Aryo Pradipta Gema and 1 others. 2024. Are we done with mmlu? *arXiv preprint arXiv:2407.03120*.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, and 1 others. 2023. Textbooks are all you need. *arXiv preprint arXiv:2306.11644*.
- Changjiang Jiang, Wenhui Dong, Zhonghao Zhang, Fengchang Yu, Wei Peng, Xinbin Yuan, Yifei Bi, Ming Zhao, Zian Zhou, Chenyang Si, and Caifeng Shan. 2026a. [Ivy-fake: A unified explainable framework and benchmark for image and video aigc detection](#). In *Proceedings of the 2026 International Conference on Multimedia Retrieval, ICMR '26*, page 2438–2447. Association for Computing Machinery.
- Changjiang Jiang, Qiannian Zhao, Lei Xin, Jinxiang Xie, Preslav Nakov, and Zhuohan Xie. 2026b. Think with structured grounding: Perceptual reinforcement learning for chart and visual-tabular understanding. *arXiv preprint arXiv:2608.22429*.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*.
- Jeffrey Li, Alex Fang, Georgios Smyrnis, Maor Ivgi, Matt Jordan, Samir Gadre, Hritik Bansal, Etash Guha, Sedrick Keh, Kushal Arora, and 1 others. 2024. Datacomp-lm: In search of the next generation of training sets for language models. *Advances in Neural Information Processing Systems*, 37:14200–14282.
- Shengjie Ma, Xuhui Jiang, Chengjin Xu, Cehao Yang, Liyu Zhang, and Jian Guo. 2025. Synthesize-on-graph: Knowledgeable synthetic data generation for continue pre-training of large language models. *arXiv preprint arXiv:2505.00979*.
- Pratyush Maini, Vineeth Dorna, Parth Doshi, Aldo Carranza, Fan Pan, Jack Urbanek, Paul Burstein, Alex Fang, Alvin Deng, Amro Abbas, and 1 others. 2025. Beyondweb: Lessons from scaling synthetic data for trillion-scale pretraining. *arXiv preprint arXiv:2508.10975*.
- Pratyush Maini, Skyler Seto, He Bai, David Grangier, Yizhe Zhang, and Navdeep Jaitly. 2024. [Rephrasing the web: A recipe for compute and data-efficient language modeling](#). *Preprint*, arXiv:2401.16380.
- 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. *arXiv preprint arXiv:1809.02789*.
- 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](#). In *Second Conference on Language Modeling*.
- Joel Niklaus, Guilherme Penedo, Hynek Kydlicek, Elie Bakouch, Lewis Tunstall, Ed Beeching, Thibaud Frere, Colin Raffel, Leandro von Werra, and Thomas Wolf. 2026. The synthetic data playbook: Generating trillions of the finest tokens.
- Team Olmo, :, Allyson Ettinger, Amanda Bertsch, Bailey Kuehl, David Graham, David Heineman, Dirk Groeneveld, Faeze Brahman, Finbarr Timbers, Hamish Ivison, Jacob Morrison, Jake Poznanski, Kyle Lo, Luca Soldaini, Matt Jordan, Mayee Chen, Michael Noukhovitch, Nathan Lambert, and 50 others. 2025. [Olmo 3](#). *Preprint*, arXiv:2512.13961.
- Panupong Pasupat and Percy Liang. 2015. Compositional semantic parsing on semi-structured tables. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1470–1480.

- Guilherme Penedo. 2025. [Finewiki](#). Source: Wikimedia Enterprise Snapshot API (<https://api.enterprise.wikimedia.com/v2/snapshots>). Text licensed under CC BY-SA 4.0 with attribution to Wikipedia contributors.
- Guilherme Penedo, Hynek Kydlíček, 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](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. *arXiv preprint arXiv:1806.03822*.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106.
- Dan Su, Kezhi Kong, Ying Lin, Joseph Jennings, Brandon Norick, 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.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [Commonsenseqa: A question answering challenge targeting commonsense knowledge](#). *Preprint*, arXiv:1811.00937.
- Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafford, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, and 21 others. 2024. [OLMo 2 Furious](#). *Preprint*, arXiv:2501.00656.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. [MuSiQue: Multi-hop questions via single-hop question composition](#). *Transactions of the Association for Computational Linguistics*, 10:539–554.
- Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. 2024. [Measuring short-form factuality in large language models](#). *Preprint*, arXiv:2411.04368.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025a. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Wanli Yang, Hongyu Zang, Junwei Zhang, Wenjie Shi, Du Su, Jingang Wang, Xueqi Cheng, and Fei Sun. 2026. [Beyond reasoning: Reinforcement learning unlocks parametric knowledge in llms](#). *Preprint*, arXiv:2605.07153.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [Hotpotqa: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Zitong Yang, Neil Band, Shuangping Li, Emmanuel Candes, and Tatsunori Hashimoto. 2024. Synthetic continued pretraining. *arXiv preprint arXiv:2409.07431*.
- Zitong Yang, Aonan Zhang, Hong Liu, Tatsunori Hashimoto, Emmanuel Candès, Chong Wang, and Ruoming Pang. 2025b. Synthetic bootstrapped pretraining. *arXiv preprint arXiv:2509.15248*.
- Xinbin Yuan, Jian Zhang, Kaixin Li, Zhuoxuan Cai, Lujian Yao, Jie Chen, Enguang Wang, Qibin Hou, Jinwei Chen, Peng-Tao Jiang, and 1 others. 2026. Se-gui: Enhancing visual grounding for gui agents via self-evolutionary reinforcement learning. *Advances in Neural Information Processing Systems*, 38:127658–127679.
- 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*, pages 3472–3483.
- Aohan Zeng, Xin Lv, Zhenyu Hou, Zhengxiao Du, Qinkai Zheng, Bin Chen, Da Yin, Chendi Ge, Chengxing Xie, Cunxiang Wang, and 1 others. 2026. Glm-5: from vibe coding to agentic engineering. *arXiv preprint arXiv:2602.15763*.
- Jiang Zhou, Xiaohu Zhao, Xinwei Wu, Tianyu Dong, Hao Wang, Yangyang Liu, Heng Liu, Linlong Xu, Longyue Wang, Weihua Luo, and 1 others. 2026. Incentivizing parametric knowledge via reinforcement learning with verifiable rewards for cross-cultural entity translation. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5616–5638.
- Jiawei Zhou, Xiaoguang Li, Lifeng Shang, Lan Luo, Ke Zhan, Enrui Hu, Xinyu Zhang, Hao Jiang, Zhao Cao, Fan Yu, Xin Jiang, Qun Liu, and Lei Chen. 2022. [Hyperlink-induced pre-training for passage retrieval in open-domain question answering](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7135–7146, Dublin, Ireland. Association for Computational Linguistics.

A Training Details

Architecture. We use the OLMo-3 architecture (Olmo et al., 2025), a decoder-only transformer with the Dolma-2 tokenizer (vocabulary size padded to a multiple of 128). Experiments are conducted at two scales: OLMo-3-7B (7 billion parameters) and OLMo-3-32B (32 billion parameters). We utilize the olmo-core training framework (Team OLMo et al., 2024), which provides a highly optimized and reproducible infrastructure for large-scale distributed training. All models use FlashAttention-2 (Dao, 2023) as the attention backend.

Continued Pretraining. All experiments initialize from official OLMo-3 pretrained checkpoints: step 1,413,814 for 7B and step 679,000 for 32B. We load only the model weights and optimizer state (no trainer state) and continue pretraining on synthetic data mixtures for 1 epoch. Table 6 summarizes the hyperparameters for each configuration.

The learning rate of 2.07×10^{-5} is inherited from the official OLMo-3 mid-training recipe. We use zero warmup steps because the optimizer state is loaded from the pretrained checkpoint, ensuring stable training from the first step. For the OLMo-3-7B scale, training on $\sim 8\text{B}$ tokens using 256 H20 GPUs takes approximately 4.5 hours with a Model FLOPs Utilization (MFU) of $\sim 65\%$. For the larger OLMo-3-32B scale, training on $\sim 80\text{B}$ tokens using 1,024 H20 GPUs requires approximately 1 day and 20 hours. These benchmarks demonstrate the efficiency and scalability of our training pipeline on modern hardware.

Compute and GPU Energy. Table 7 reports hardware, wall-clock time, GPU-hours, and estimated GPU electricity consumption for synthesis and continued pretraining. We compute GPU-hours as the number of GPUs multiplied by wall-clock hours and estimate electricity consumption using an average draw of 350 W per H20 GPU. These are GPU-only estimates: they exclude CPU, networking, storage, and data-center overhead such as cooling and power-delivery losses. We report energy rather than CO_2 emissions because the time- and location-specific carbon intensity of the cluster electricity supply is unavailable.

The additional synthesis cost can be offset by greater training-token efficiency. At the OLMo-3-7B scale, training on 8B WRAP++ tokens already outperforms training on 16B Extended WRAP to-

kens (45.32 vs. 44.45 SimpleQA pass@128). Assuming compute scales linearly with token count, synthesis and training together consume approximately 1.48 MWh for the 8B WRAP++ run, compared with 2.51 MWh for the 16B Extended WRAP run—a reduction of approximately 41%. At the same 8B-token training budget, WRAP++ has a higher one-time synthesis cost: total synthesis-and-training energy is 1.48 vs. 1.25 MWh for OLMo-3-7B and 2.65 vs. 2.43 MWh for OLMo-3-32B (WRAP++ vs. Extended WRAP). The relative contribution of this one-time synthesis overhead therefore decreases as model-training cost grows.

Data Format. All synthetic QA data is formatted as plain text with “Question:” and “Answer:” delimiters, consistent with prior WRAP work. For cross-document QA, the synthesized output directly states facts without referencing source passages, ensuring the model internalizes them as parametric knowledge rather than reading comprehension signals.

B Evaluation Details

In-Context Learning Setup. Because our continued pretraining experiments operate on base models (OLMo-3) that have not undergone instruction tuning, these models cannot reliably follow zero-shot formatting directives. To accurately probe their parametric knowledge, we adopt a 5-shot in-context learning protocol for all evaluations, including SimpleQA, HotpotQA, MuSiQue, and the 12 general benchmarks. Specifically, for each evaluation instance, we prepend the prompt with five demonstration question-answer pairs. For SimpleQA, these demonstrations are sampled directly from the SimpleQA dataset; to ensure strict evaluation integrity and prevent data contamination, any examples used as few-shot demonstrations are explicitly excluded from the active evaluation set during inference.

Decoding and Sampling Parameters. We adopt distinct decoding strategies tailored to the nature of each benchmark. For **factual QA benchmarks**, we employ nucleus sampling with a temperature of 0.6 and top- p of 0.95 to provide a diverse distribution for knowledge recoverability analysis. In contrast, for the **12 general benchmarks**, we use greedy decoding (temperature 0.0) to ensure deterministic and reproducible outputs across all model comparisons. These parameters are held constant

Table 6: Continued pretraining hyperparameters. All experiments use a linear decay schedule (no warmup, decaying to 0) and SkipStepAdamW (Olmo et al., 2025) with $\beta_1 = 0.9$, $\beta_2 = 0.95$, weight decay 0.1 (embedding weights excluded), max gradient norm 1.0, and auxiliary z-loss with multiplier 10^{-5} . Parameters are stored in bfloat16 and gradients are reduced in float32 via HSDP.

Hyperparameter	OLMo-3-7B (8B)	OLMo-3-7B (80B)	OLMo-3-32B
Sequence length	8,192	8,192	8,192
Global batch size (tokens)	$\sim 2\text{M}$ (2^{21})	$\sim 2\text{M}$ (2^{21})	$\sim 8\text{M}$ (2^{23})
Peak learning rate	2.07×10^{-5}	2.07×10^{-5}	2.07×10^{-5}
LR schedule	Linear $\rightarrow 0$	Linear $\rightarrow 0$	Linear $\rightarrow 0$
Warmup steps	0	0	0
Training steps	4,000	40,000	10,000
Training tokens	$\sim 8\text{B}$	$\sim 80\text{B}$	$\sim 80\text{B}$
Precision	bfloat16	bfloat16	bfloat16
Data parallel	HSDP (block wrap)	HSDP (block wrap)	HSDP (full wrap, shard 64)
Activation ckpt.	FFN-only	FFN-only	Budget (50%)
GPUs	$256 \times \text{H20}$	$256 \times \text{H20}$	$1024 \times \text{H20}$

Table 7: Compute and estimated GPU electricity consumption for data synthesis and continued pretraining. Energy is estimated as GPU-hours $\times$ 0.350 kW.

Job	Hardware and wall time	GPU-hours	GPU energy
Synthesize 8B Extended WRAP tokens	128 H20s $\times$ 19 h	2,432	0.85 MWh
Synthesize 8B WRAP++ tokens	128 H20s $\times$ 24 h	3,072	1.08 MWh
Synthesize 80B WRAP++ tokens	128 H20s $\times$ 240 h	30,720	10.75 MWh
Train OLMo-3-7B on 8B tokens	256 H20s $\times$ 4.5 h	1,152	0.40 MWh
Train OLMo-3-7B on 80B tokens	256 H20s $\times$ 45 h	11,520	4.03 MWh
Train OLMo-3-7B on 100B tokens (midtraining)	256 H20s $\times$ 56.25 h	14,400	5.04 MWh
Train OLMo-3-32B on 80B tokens	1,024 H20s $\times$ 44 h	45,056	15.77 MWh

across all model scales and data recipes to ensure a fair evaluation.

Unbiased Estimation of $\text{pass}@k$. While the $\text{pass}@k$ metric intuitively represents the probability of generating at least one correct answer within k attempts, empirically estimating this by drawing exactly k samples yields high variance. To achieve a more stable and unbiased estimate, we adopt the methodology introduced by Chen et al. (2021). For each evaluation instance, we generate n total samples ($n \geq k$) and determine the number of correct responses, c . The unbiased estimator for $\text{pass}@k$ is then computed as:

$$\text{pass}@k = 1 - \frac{\binom{n-c}{k}}{\binom{n}{k}} \quad (4)$$

where $\binom{\cdot}{\cdot}$ denotes the binomial coefficient. For our primary SimpleQA results, we generate $n = 128$ samples per question. In the larger-budget analysis, we generate $n = 256$ samples and estimate both $\text{pass}@128$ and $\text{pass}@256$ from the same sample pool. This formulation efficiently leverages all n generated samples to calculate the expected

pass rate for any evaluation budget $k \leq n$, thereby reducing variance without requiring repeated sampling passes.

C Cross-Scale Ablation Results

We repeat two representative single-document controls from the OLMo-3-7B ablation at the OLMo-3-32B scale under the same 8B-token continued-pretraining budget. Raw FineWiki trains directly on the original source documents without rewriting, whereas Nemotron-Extract applies the Extract Knowledge strategy from Nemotron-CC. Table 8 shows that the ordering is consistent across model scales: **WRAP++** substantially outperforms both controls at 7B and 32B. Together with the token-matched scaling curves in Figure 2, this result supports the cross-scale robustness of the main finding.

D Statistical Significance of Ablations

We assess the close results in the pairing-strategy and relation-topology ablations using question-level paired bootstrap tests with 10,000 resamples. For each comparison, Δ denotes $\text{pass}@128(\text{A})$ minus $\text{pass}@128(\text{B})$, measured in percentage points.

Table 8: Representative ablations across model scales under the same 8B-token continued-pretraining budget. We report SimpleQA pass@128.

Data Setting	7B	32B
Pretrained checkpoint	34.76	41.26
Raw FineWiki	39.23	42.33
Nemotron-Extract	38.69	42.79
WRAP++	45.11	51.24

Table 9 reports the corresponding 95% bootstrap confidence intervals. Topological discovery significantly outperforms random pairing, and the combined topology significantly outperforms dual-link alone; the other two within-topology differences are not statistically conclusive because their intervals include zero.

Table 9: Question-level paired bootstrap results for the 8B-token OLMo-3-7B ablations (10,000 resamples). Δ denotes A minus B in pass@128 percentage points.

Comparison (A vs. B)	Δ	95% CI
<i>Relation topology</i>		
Combined vs. dual-link	+0.87	[0.02, 1.79]
Combined vs. co-mention	+0.75	[-0.28, 1.60]
Dual-link vs. co-mention	-0.12	[-1.40, 1.86]
<i>Pairing strategy</i>		
Topological vs. random	+1.65	[0.31, 2.99]

The within-topology pattern is consistent with the corpus sizes at this budget. The dual-link-only setting must repeat its approximately 3B-token corpus to reach 8B training tokens, whereas co-mention provides a much larger relation pool. Consequently, adding co-mention to dual-link produces the clearest gain, while adding the smaller dual-link pool to co-mention provides a more modest incremental signal.

E Additional Factual QA Results

E.1 Closed-Book Multi-Hop Benchmarks

We evaluate the pretrained baseline and the final 80B-token checkpoints of the token-matched OLMo-3-7B comparison on HotpotQA (Yang et al., 2018) and MuSiQue (Trivedi et al., 2022). We use a closed-book setting aligned with our context-free internalization objective: the model receives no source documents or retrieved passages at inference time. Table 10 reports pass@8. **WRAP++** performs best on both benchmarks, improving over

Extended WRAP by 2.65 points on HotpotQA and 0.41 points on MuSiQue.

Table 10: Closed-book multi-hop factual QA at the final 80B-token checkpoint of the token-matched OLMo-3-7B comparison. We report pass@8 (%); Δ is **WRAP++** minus Extended WRAP.

Benchmark	Pretrained	Extended WRAP	WRAP++	Δ
HotpotQA	44.57	45.47	48.12	+2.65
MuSiQue	17.16	19.20	19.61	+0.41

E.2 Larger Sampling Budget on SimpleQA

To test whether the advantage persists when the OLMo-3-7B model is given more opportunities to recover the correct answer, we generate 256 responses per SimpleQA question for Extended WRAP and **WRAP++** at each token-matched checkpoint. We compute unbiased pass@128 and pass@256 from the same 256 samples. Because pass@128 is re-estimated from this larger sample pool, these values may differ slightly from the primary $n = 128$ evaluation. As shown in Table 11, **WRAP++** leads at every checkpoint on both metrics. Extended WRAP plateaus after its initial gain, whereas **WRAP++** increases from 45.32 to 49.18 pass@128 and from 52.88 to 56.42 pass@256. At 80B tokens, the margins are 5.53 and 6.78 points, respectively.

F Synthesis Prompt Templates

Single-Document WRAP Prompt (Baseline).

Following Maini et al. (2024), we use the standard QA-style prompt for single-document synthesis:

Convert the following paragraph into a conversational format with multiple tags of "Question:" followed by "Answer:".

Extended WRAP Prompt (Baseline). Extended WRAP instructs the generator to exhaustively extract distinct QA pairs from each source document:

Your goal is to exhaustively extract information and synthesize as many unique, non-repetitive question-and-answer pairs as possible from the provided text.

You must strictly follow this exact output format:

Question:

Answer:

Convert the following paragraph into this conversational format:

[text]

Table 11: SimpleQA with a 256-sample generation budget for the token-matched OLMo-3-7B comparison. Both pass@128 and pass@256 are estimated from the same 256 responses per question. Bold indicates the better method at each nonzero checkpoint.

Consumed Tokens	pass@128		pass@256	
	Extended WRAP	WRAP++	Extended WRAP	WRAP++
0B	35.91	35.91	41.31	41.31
8B	43.37	45.32	49.60	52.88
16B	44.45	46.49	50.98	54.13
24B	44.06	47.47	50.15	54.92
32B	44.84	47.68	50.80	54.89
40B	43.77	48.51	49.83	55.80
48B	44.35	48.10	50.24	55.20
56B	43.58	49.20	49.69	56.65
64B	44.19	49.49	50.13	56.84
72B	43.75	49.03	49.78	55.94
80B	43.65	49.18	49.64	56.42

WRAP++ Cross-Document Joint QA Prompt.

For cross-document synthesis, we provide the full text of both related entities and apply the following instruction. The prompt enforces three critical constraints: cross-document dependency, explicit factual chaining, and Context-Free Internalization (see §2.3).

You are an expert data generator for language model pretraining.

Below are two related Wikipedia passages:

[Passage A]
{text_a}

[Passage B]
{text_b}

Task:

1) Generate high-quality synthetic QA pairs that REQUIRE information from BOTH Passage A and Passage B to answer.

2) The Answer MUST begin with a step-by-step reasoning process. This reasoning must explicitly bridge facts from both passages.

3) Do not use external knowledge.

4) CRITICAL CONSTRAINT: The generated QA pair will be used to train a model WITHOUT these passages provided as context. Therefore, you MUST act as an omniscient AI stating absolute facts from your own inherent knowledge.

- DO NOT use any attribution phrases like ‘According to Passage A’, ‘Passage B mentions’, ‘As stated in the text’, or ‘Based on the provided documents’.

- State the facts directly and confidently.

Output format (strict):

Question: [Insert a complex question that bridges facts from both passages]

Answer: [Acting as an omniscient AI, directly state all necessary factual premises from both passages, and logically synthesize them to derive the conclusion.]

Therefore, [State the final, concise answer.]

Synthesis Hyperparameters. All QA synthesis uses Qwen3-30B-A3B-Instruct (FP8) as the generator with temperature 0.7, top- p = 0.8, and a maximum output length of 32,768 tokens. For the synthesis model scale ablation (§4.4), we additionally use Qwen3-235B-A22B-Instruct.

G Relation Discovery Statistics

Table 12: Statistics of the hyperlink relations discovered from FineWiki (English Wikipedia).

Statistic	Value
Source corpus	FineWiki (English)
Raw corpus tokens	~8.4B
Wikipedia articles processed	~6.7M
Dual-link pairs ($A \leftrightarrow B$)	~9.6M
Co-mention pairs ($A \rightarrow E \leftarrow B$, $A \rightarrow B$)	~232M
Dual-link synthesized tokens	~3B
Co-mention synthesized tokens	~79.7B
Total WRAP++ tokens	~82.7B
Single-doc WRAP tokens (baseline)	~5.4B
Extended WRAP tokens (baseline)	~17.4B

Table 12 reports the relation discovery and synthesis statistics. The dual-link motif yields a relatively small but high-precision set of ~9.6M entity pairs, while the co-mention motif provides a much

larger pool of $\sim 232\text{M}$ pairs, enabling substantial combinatorial expansion. The total synthesized corpus of $\sim 82.7\text{B}$ tokens is approximately $15\times$ larger than single-document WRAP on the same source, illustrating the amplification advantage of relation-driven cross-document synthesis.

H Mid-Training Evaluation Benchmarks

The 12 general tasks used for evaluation in Section 4.6 are categorized as follows:

- **General Knowledge:** ARC (Clark et al., 2018), MMLU Redux (Gema et al., 2024)
- **Reading Comprehension:** SQuAD v2 (Rajpurkar et al., 2018), DROP (Dua et al., 2019)
- **Reasoning:** OpenBookQA (Mihaylov et al., 2018), CSQA (Talmor et al., 2019)
- **Natural Language Understanding:** Winogrande (Sakaguchi et al., 2021), PIQA (Bisk et al., 2020), HellaSwag (Zellers et al., 2019)
- **Math:** GSM8K (Cobbe et al., 2021)
- **Table Understanding & Trivia:** WikiTableQuestions (Pasupat and Liang, 2015), TriviaQA (Joshi et al., 2017)

Table 13 provides per-benchmark results for the mid-training integration experiment described in Section 4.6.

I Additional Ablation: Co-Mention with Three Documents

In the main co-mention motif ($A \rightarrow E \leftarrow B$ while $A \rightarrow B$), we use the two target entities A and B as input to joint QA synthesis. A natural extension is to also include the bridging page E as a third input document, potentially providing additional shared context. We use a dedicated 3-document prompt template that instructs the synthesis model to generate QA requiring facts from all three passages.

Surprisingly, including the bridging page E does not improve—and slightly hurts—performance (Table 14). We hypothesize that the bridging page introduces distracting context: since E typically links to many entities, its content is broad and may divert the synthesis model from focusing on the specific relationship between A and B . The 2-document formulation used in WRAP++ strikes a better balance between relational grounding and synthesis focus.

J Synthesized QA Data Statistics

We report comprehensive statistics of the synthesized WRAP++ dataset to characterize the length distributions and data composition. Statistics are computed over the full corpus of 240,658,065 QA instances across 24,224 JSONL files, processed in parallel using 256 workers. Character-level and word-level lengths are measured on the raw synthesized text after extracting the “Question:” and “Answer:” fields from each record.

Dataset Composition. Table 15 summarizes the overall dataset composition by relation type.

Length Distributions. Table 16 reports the distributional statistics of the synthesized QA text. Questions are concise (median 203 characters, 32 words), while answers are substantially longer (median 1,386 characters, 212 words), reflecting the explicit factual chaining constraint that requires step-by-step reasoning before stating the final conclusion. The overall QA length is concentrated in the 1,000–2,000 character range (56.1% of all instances), with 27.6% in the 2,000–5,000 range and 16.0% in the 500–1,000 range. Fewer than 0.3% of instances fall outside the 500–5,000 character window, indicating a well-controlled generation process. Figure 5 visualizes the question and answer length distributions at both character and word levels.

Comparison Across Relation Types. Table 17 compares the QA length characteristics between the two relation motifs. Dual-link instances produce slightly shorter answers (median 1,210 vs. 1,394 characters), likely because mutual references tend to encode more focused bilateral relationships, whereas co-mention pairs often involve broader categorical or analogical connections that require more elaboration. Despite these differences, both subsets maintain similar question lengths and overall distributional shape. Figure 6 overlays the QA length histograms for both relation types, and Figure 7 provides a violin plot comparison across three length dimensions.

Source Document Lengths. The input Wikipedia passages exhibit substantial length variation. The first passage (text_a) has a median length of 4,458 characters (P5–P95: 519–27,963), while the second passage (text_b) is generally longer with a median of 9,575 characters (P5–P95: 943–50,004, where 50,004 indicates truncation at

Table 13: Per-benchmark breakdown for mid-training integration (OLMo-3-7B, full 100B-token schedule, 3-shot cloze format). Accuracy (%) is reported for all tasks except SQuAD v2 and DROP, which use token-level F1.

Category	Benchmark	Base	Midtrain	WRAP++ Mix
General Knowledge	ARC	77.47	84.98	85.07
	MMLU Redux	60.30	65.30	66.58
Reading Comprehension	SQuAD v2	42.27	48.59	49.99
	DROP	40.51	66.51	65.36
Reasoning	OpenBookQA	76.40	85.20	82.40
	CSQA	72.15	74.20	72.97
Language Understanding	WinoGrande	53.28	65.11	62.98
	PIQA	73.99	74.54	74.92
	HellaSwag	59.83	70.15	73.12
Math	GSM8K	38.36	79.53	77.79
Table & Trivia	WikiTableQ	36.55	43.74	44.46
	TriviaQA	62.42	61.01	62.23
Average		57.79	68.24	68.16

Table 14: Co-mention synthesis using 2 vs. 3 input documents (OLMo-3-7B, ~8B tokens, SimpleQA pass@8).

Co-mention Variant	SimpleQA pass@8
2-doc (entities A, B only)	15.74
3-doc (entities A, B + bridge E)	15.42

Table 15: Composition of the **WRAP++** synthesized QA dataset by relation type.

Relation Type	QA Instances	Proportion
Co-mention	231,292,954	96.1%
Dual-link	9,365,111	3.9%
Total	240,658,065	100.0%

the maximum context window). This asymmetry arises because co-mention pairs order documents by the directed edge $A \rightarrow B$, where B (the referenced entity) tends to be a more prominent article. Figure 8 visualizes this distributional asymmetry.

QA Length Bucketed Distribution. Table 18 and Figure 9 provide a bucketed view of the combined QA and answer length distributions. The synthesis process produces a unimodal distribution with the majority of instances in the 1,000–2,000 character range. No instances have empty answers, and fewer than 0.01% of answers are shorter than 200 characters, confirming that the explicit factual

chaining constraint effectively prevents degenerate outputs.

K Qualitative Analysis

Example: WRAP vs. WRAP++ Synthesis.

Source entity A: Ludwig Göransson (Swedish composer, film score artist)

Source entity B: Oppenheimer (2023 film directed by Christopher Nolan)

As shown in Table 19, the **WRAP++** output creates richer associative context: it contrasts with a commonly confused entity (Zimmer), provides cross-film comparisons, and generates reverse-direction queries—all contributing to more robust knowledge encoding.

L Responsible Use and Artifact Considerations

Risks, licenses, and intended use. Our experiments use public encyclopedic web content from huggingface dataset HuggingFaceFW/finewiki and publicly available models and benchmarks. Wikipedia text is distributed under the Creative Commons Attribution Share Alike 4.0 license and Wikimedia Terms of Use, and the other artifacts used in this work are subject to their respective licenses and terms. The generated QA data are intended for research on synthetic pretraining and factual knowledge acquisition, not as a verified factual database. As with other synthetic pretrain-

Per-Benchmark Score Trends During Mid-Training (WRAP++ Mix, 100B Tokens)




Figure 4: Per-benchmark score trajectories during mid-training with the **WRAP++** Mix (100B tokens) on OLMo-3-7B. Each subplot tracks one of the 12 evaluation tasks over the course of training. All benchmarks exhibit a clear upward trend, with knowledge-intensive tasks (e.g., HellaSwag, GSM8K, CSQA) showing the most pronounced gains.

Table 16: Length statistics of **WRAP++** synthesized QA data (aggregated over all 240.7M instances). “Chars” denotes character count; “Words” denotes whitespace-delimited token count.

Field	Mean	Std	Min	P5	P25	Median	P75	P95
<i>Character-level</i>								
Question	207	43	45	143	177	203	233	283
Answer	1,424	624	139	563	938	1,386	1,817	2,463
QA (combined)	1,651	637	242	766	1,161	1,610	2,051	2,716
<i>Word-level</i>								
Question	33	7	7	23	28	32	37	45
Answer	217	91	16	92	149	212	272	364

Table 17: Median QA lengths (characters) by relation type. Word-level medians are shown in parentheses.

Relation Type	Question	Answer	QA Total
Co-mention	204 (32 words)	1,394 (213 words)	1,618
Dual-link	197 (32 words)	1,210 (189 words)	1,428
Overall	203 (32 words)	1,386 (212 words)	1,610

ing corpora, the data may inherit factual errors or topical biases from the source corpus or synthesis model, so downstream use should include appropriate filtering and evaluation.

Personal or sensitive content. We do not collect private user data or conduct human-subject data collection. The source corpus is public encyclope-

Table 18: Bucketed length distributions for the combined QA text and answer text (character-level).

Length Range (chars)	QA (combined)		Answer only	
	Count	%	Count	%
[0, 200)	0	0.0	2,479	<0.01
[200, 500)	456,399	0.2	6,771,177	2.8
[500, 1,000)	38,609,776	16.0	61,411,377	25.5
[1,000, 2,000)	135,127,249	56.1	132,242,235	55.0
[2,000, 5,000)	66,336,557	27.6	40,114,788	16.7
[5,000, 10,000)	99,207	<0.1	89,225	<0.1
[10,000, ∞)	28,877	<0.1	26,784	<0.1

dic text and may naturally mention public figures, organizations, locations, or sensitive topics. Our generation process does not infer private attributes or intentionally introduce personal information beyond the provided source passages.

Question vs. Answer Length Distributions




Figure 5: Question vs. answer length distributions at character level (left) and word level (right). Questions are tightly concentrated around a median of 203 characters (32 words), while answers exhibit a broader, right-skewed distribution with a median of 1,386 characters (212 words), reflecting the explicit factual chaining required by the synthesis prompt.

Table 19: Qualitative comparison between single-document WRAP and cross-document **WRAP++** synthesis for the Ludwig Göransson and *Oppenheimer* example.

WRAP (single-document QA on entity A)	WRAP++ (cross-document QA on entities A+B)
Q: Who is Ludwig Göransson? A: Ludwig Göransson is a Swedish composer known for film scores. Q: What films has Göransson scored? A: He has scored <i>Black Panther</i> , <i>Tenet</i> , and <i>Oppenheimer</i> .	Q: What films earned Nolan’s original <i>Tenet</i> composer and his replacement their second Oscars? A: Hans Zimmer turned down <i>Tenet</i> for <i>Dune</i> , which earned his second Oscar (after <i>The Lion King</i>). His replacement, Ludwig Göransson, later won his second Oscar for <i>Oppenheimer</i> (after <i>Black Panther</i>). Therefore, the answer is <i>Dune</i> and <i>Oppenheimer</i> .




Figure 6: Synthesized QA length distribution by relation type. Both subsets exhibit a similar unimodal shape, with the co-mention distribution (blue) shifted slightly rightward relative to dual-link (red), consistent with the longer factual chains required to bridge co-mentioned entities.




Figure 7: Violin plot comparison of question length, answer length, and answer word count between co-mention and dual-link subsets. Black horizontal lines indicate medians. Both motifs produce similarly distributed questions, while co-mention answers are moderately longer, reflecting the additional elaboration needed to bridge indirectly related entities.




Figure 8: Source document length distributions for the two input passages. Passage A (the referencing entity) is typically shorter (median $\approx 4,700$ chars), while Passage B (the referenced entity) tends to be longer and more prominent (median $\approx 9,500$ chars), with a visible mass accumulation at the 50K truncation boundary.

Length Distribution (Bucketed, Full Dataset: 240.7M instances)




Figure 9: Bucketed length distributions computed over the full dataset of 240.7M instances. Left: combined QA length; right: answer-only length. The 1K–2K character bucket dominates both distributions (>55%), with a secondary concentration in the 2K–5K range for QA and the 500–1K range for answers. Extreme lengths (>5K or <200) are negligible.