

LLM generation novelty through the lens of semantic similarity

Anonymous ACL submission

Abstract

Generation novelty is a key indicator of an LLM’s ability to generalize, yet measuring it against full pretraining corpora is computationally challenging. Existing evaluations often rely on lexical overlap, failing to detect paraphrased text, or do not consider the full pretraining corpus. We frame novelty as a semantic retrieval problem. This framing enables us to address novelty with modern embedding and indexing pipelines, allowing for efficient analysis at pre-training scale. Specifically, we propose a three-stage framework that retrieves semantically similar samples, reranks them at varying subsequence lengths, and calibrates scores using a human novelty reference for interpretability. We apply this framework to the SmoLLM model family and report three key findings: (1) models draw on pretraining data across much longer sequences than previously reported; (2) some task domains systematically promote or suppress generation novelty; and (3) instruction tuning not only alters style but also increases novelty. These results highlight the value of semantic novelty analysis for studying generalization. To support reproducibility and further research, we release ~ 20 TB of corpus chunks and index artifacts upon acceptance.

1 Introduction

Large language models (LLMs) now power chatbots, copilots, and autonomous agents with applications across various domains. Their adoption hinges on an implicit assumption: LLMs are capable of generalizing beyond their training data to generate relevant and novel outputs in response to user prompts. Generation novelty provides a useful signal of this capability, indicating whether model outputs extend beyond patterns observed during training. Novelty signals compositional generalization and allows for the assessment of true zero-shot behavior – informing debates about provenance and intellectual property. Thus, measuring a model’s

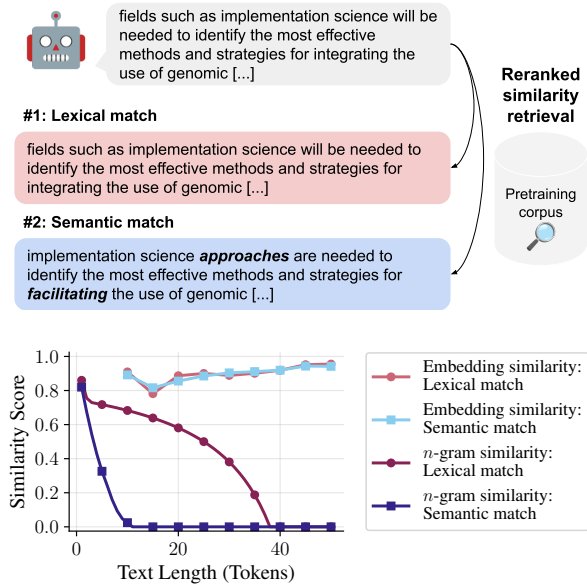


Figure 1: **Efficient retrieval allows for pretraining-scale semantic novelty analysis.** Given a generation, we employ retrieval and reranking pipelines to identify semantically similar samples in the pretraining corpus to analyze generation novelty. Higher similarity indicates lower novelty. While the first-ranked sample is a lexical match, the second-ranked is a semantic match. As text length increases, n -gram overlap drops to zero, falsely suggesting high novelty for paraphrased content. In contrast, embedding similarity correctly identifies the shared meaning, providing a more reliable measure of novelty.

generation novelty is a prerequisite for interpreting what LLMs actually learn.

However, analyzing the novelty of LLM generations is a non-trivial task. There is no single, clearly defined notion of novelty, which opens up the question: Is a generation already novel if it does not appear verbatim in the pretraining data? More broadly, what does it mean for a generation to be “reused” from the training data? At the same time, the sheer scale of modern pretraining corpora makes comparison computationally challenging,

054	so that novelty analysis is both a conceptual and a	novelty.	106
055	technical problem.		
056	Existing approaches to studying LLM genera-	We demonstrate the utility of our framework	107
057	tion novelty mainly address these questions in two	through empirical analyses on SmolLM (Allal et al.,	108
058	directions. One line of work relies on efficient	2024) and SmolLM2 (Allal et al., 2025), two LLMs	109
059	textual overlap metrics (McCoy et al., 2023; Mer-	with open pretraining corpora. Our analysis re-	110
060	rill et al., 2024; Padmakumar et al., 2025), which	veals unexpected patterns missed by previous lex-	111
061	scale well but naturally focus on verbatim reuse	ical methods (McCoy et al., 2023; Merrill et al.,	112
062	and fail to account for paraphrases or stylistic vari-	2024). First, both models draw on pretraining data	113
063	ation. A second line of work investigates novelty	over much longer sequences than previously re-	114
064	within specific domains, such as scientific ideas	ported (Merrill et al., 2024). Second, novelty varies	115
065	or biomedical publications (Peng et al., 2025; Ai	systematically by task domain. Third, embedding-	116
066	et al., 2025; Wang et al., 2025c), typically using	based novelty estimates are stable under style shifts	117
067	semantic similarity measured against a curated ref-	from instruction tuning; after accounting for these	118
068	erence corpus. While effective, these approaches	shifts, instruction tuning substantially increases	119
069	do not account for reuse from other parts of the	novelty. These results suggest that instruction tun-	120
070	broader pretraining corpus. Together, these limita-	ing shapes not only style but also compositional	121
071	tions highlight the need for novelty measures that	generation behavior.	122
072	go beyond verbatim overlap at the scale of LLM	Our contributions are as follows:	123
073	pretraining corpora.		
074	We address this need by proposing a three-stage	• We present a semantic similarity–based frame-	124
075	framework for analyzing an LLM’s generation nov-	work for measuring LLM generation nov-	125
076	elty at scale through the lens of semantic similar-	elty at scale, combining retrieval, reranking,	126
077	ity. We argue that a generation should be called	and baseline-calibration to enable comparison	127
078	novel only if it does not semantically reproduce the	against full pretraining corpora while reducing	128
079	training data (cf. Figure 1). From this perspective,	sensitivity to stylistic variation.	129
080	“reused” information is not limited to verbatim over-	• We empirically analyze generation novelty	130
081	lap, but includes paraphrases and reformatting.	in SmolLM and SmolLM2, uncovering long-	131
082	Our framework combines semantic retrieval	range reuse patterns, task-dependent variation,	132
083	(stage 1) and reranking (stage 2) (Santhanam et al.,	and the impact of instruction tuning beyond	133
084	2022; Li et al., 2025) to compare LLM generations	stylistic effects.	134
085	against their pretraining corpora. However, simi-	• We also release the corresponding indices and	135
086	larity scores and their relative differences are not	corpus chunks of SmolLM and SmolLM2 to	136
087	directly interpretable as measures of novelty. They	support replication and extension.	137
088	are affected by biases in the retrieval pipeline, such		
089	as a preference for shorter documents and informa-	2 Related Work	138
090	tion located early within documents (Fayyaz et al.,	Novelty measurement in LLMs. Prior work on	139
091	2025; Zhou et al., 2025). Because of this, identi-	LLM generation novelty has predominantly relied	140
092	cal scores can reflect different degrees of novelty,	on textual overlap–based measures, particularly	141
093	especially when comparing different text lengths	n -gram comparisons. McCoy et al. (2023) intro-	142
094	or domains. Hence, the third stage of our frame-	duce n -novelty, defining novelty as the absence	143
095	work calibrates the raw semantic similarity scores	of copied text from the training data and quanti-	144
096	using a baseline of held-out, human-written refer-	fying the proportion of non-overlapping n -grams	145
097	ence text. By calibrating the scores, we mitigate	in model outputs. Building on this, Merrill et al.	146
098	artifacts of the retrieval pipeline and enable mean-	(2024) additionally analyze the probability of gen-	147
099	ingful comparison of scores. Finally, we aggregate	erating training n -grams, while Padmakumar et al.	148
100	these scores into a <i>novelty profile</i> which character-	(2025) propose novelty as the harmonic mean of	149
101	izes the generation novelty of a model for a spe-	n -novelty and generation quality assessed by LLM	150
102	cific dataset and generation text length. This proce-	judges. Wang et al. (2025b) further extend this line	151
103	cedure is model- and task-agnostic, and lightweight	of work by introducing task-grams, which capture	152
104	enough to run on full pretraining corpora. It thus of-	task-specific n -gram co-occurrences in output and	153
105	fers a scalable way to assess an LLM’s generation	corpus. The task relation introduces a semantic	154

dimension to the n -gram-based analysis. Despite these extensions, all of these approaches rely on surface-level textual overlap and may treat paraphrases or stylistic variation as novel. We instead define novelty via semantic similarity, enabling analysis beyond surface-form variation. Prior work that emphasizes semantic novelty typically focuses on specific domains, such as scientific ideas or biomedical text (Ai et al., 2025; Peng et al., 2025; Wang et al., 2025c). These approaches also rely on embedding-based similarity, but assess novelty relative to selected reference documents tailored to the target domain. In contrast, our work studies generation novelty with respect to the entire pretraining corpus, independent of the prompt.

Memorization and membership inference. Memorization work (Wu et al., 2025; Feldman and Zhang, 2020) investigates whether specific samples can be elicited verbatim from a model to understand if the model has memorized them. In contrast, our notion of novelty captures whether the underlying *information* is present in the corpus, even if paraphrased. Membership inference attacks (MIA) (Puerto et al., 2025; Mesana et al., 2025; Zhang et al., 2025, 2024) instead ask whether a particular example was part of pretraining, often in adversarial settings. While informative for questions like data privacy, MIAs do not address the broader question of how models generate text that are *not part of* their training data. Our novelty measure, therefore, complements both attribution and memorization/MIA, providing a new perspective on generalization.

Further discussion of prior work relating model generations to pretraining data is provided in Appendix A.

3 Novelty analysis framework

This section presents our framework for studying LLM generation novelty. We introduce the conceptual definition of *semantic* novelty and propose a framework that measures calibrated semantic novelty at LLM pretraining scale.

3.1 Conceptual Definition of Semantic Novelty

We define *generation novelty* as the degree to which an LLM’s output expresses information or patterns not readily present in its pretraining corpus. Unlike *lexical novelty*, which is typically measured via n -gram non-overlap (McCoy et al., 2023; Merrill et al., 2024), our definition focuses on *semantic*

novelty:

$$\min_{d \in \mathcal{C}} \cos(\phi(y), \phi(d)), y \sim f_{\theta}(x) \quad (1)$$

where \mathcal{C} is the training corpus, y is the LLM f_{θ} ’s output to a prompt x , and ϕ is an embedding function. In other words, semantic novelty is defined as the minimum cosine similarity between a semantic representation of the LLM generation and documents of the pretraining corpus. Under this definition, an output’s novelty is determined by its content, rather than its surface form. Hence, this definition serves to distinguish between *reproduction* (generating sequences that exist semantically in the corpus) and *composition* (generating sequences that are internally coherent and correct, yet semantically distinct from any specific training document).

By aggregating semantic signals across a dataset, we can characterize the *novelty profile* of a specific model configuration or training regime. Crucially, we treat novelty not as a binary state, but as a continuous spectrum. We do not claim that a “novel” generation has no relationship to the training data; rather, we measure the extent to which the generated content deviates from the most similar semantic neighbors available in the corpus.

3.2 A Framework for Analysing Semantic Novelty

Our framework operationalizes the measurement of semantic novelty through a model-agnostic, three-stage paradigm applied at the dataset scale. This approach is inspired by standard retrieval-augmented generation (RAG) and information retrieval architectures, where a two-stage process of ranking and re-ranking is the established procedure for balancing search efficiency with semantic precision (Li et al., 2025). Since our framework addresses generation novelty, we further append a third stage that mitigates potential skewness artifacts of the retrieval pipeline, arising from sequence length or domain-specific density of the LLM generations. This stage calibrates the novelty scores with respect to a matched human-level novelty reference so that scores are comparable and interpretable.

The framework aims to characterize the *novelty profile* $\mathcal{N}_{\mathcal{Q}}(k)$ of a model distribution \mathcal{Q} relative to a domain-matched human distribution \mathcal{H} , for each chunk size k . By evaluating across multiple chunk sizes k , we can observe how specific model configurations (e.g., scale or alignment) influence novelty.

Algorithm 1 Novelty Framework with retrieval, reranking and baseline-normalized scoring

Require: \mathcal{C} (Corpus), \mathcal{Q} (Model generations), \mathcal{H} (Human references), \mathcal{R} (Ranker), \mathcal{S} (Re-ranker), K (Set of chunk sizes $\{k_1, k_2, \dots\}$)

// Stage 1: Identification of Local Candidate Pools

for each document $d \in \mathcal{Q} \cup \mathcal{H}$ **do**

$P_d \leftarrow \mathcal{R}(d, \mathcal{C}, n)$ ▷ Retrieve n passages once per document using Eq. 1

end for

// Stage 2: Multi-Scale Evaluation and Calibration

for each chunk size $k \in K$ **do**

$\text{Scores}_H \leftarrow \emptyset, \text{Scores}_Q \leftarrow \emptyset$

// Determine the "Semantic Noise Floor" for this configuration

for each document $h \in \mathcal{H}$ **do**

$\text{Chunks}_h \leftarrow \text{chunk}(h, k)$

$S_h \leftarrow \{\max_{p \in P_h} \mathcal{S}(c, p) \mid c \in \text{Chunks}_h\}$

$\text{Scores}_H \leftarrow \text{Scores}_H \cup S_h$

end for

$\mu_H^{(k)} \leftarrow \text{mean}(\text{Scores}_H)$ ▷ Stable calibrator for domain and length k

// Stage 3: Calculate Calibrated Similarity Scores for the Model

for each document $q \in \mathcal{Q}$ **do**

$\text{Chunks}_q \leftarrow \text{chunk}(q, k)$

$S_q \leftarrow \{\max_{p \in P_q} \mathcal{S}(c, p) \mid c \in \text{Chunks}_q\}$

$R_q \leftarrow \{s / \mu_H^{(k)} \mid s \in S_q\}$ ▷ Normalize by human distribution

$\text{Scores}_Q \leftarrow \text{Scores}_Q \cup R_q$

end for

$N^{(k)} \leftarrow \text{median}(\text{Scores}_Q)$ ▷ Final novelty profile for chunk size k

end for

return $\{(k, N^{(k)}) \mid k \in K\}$ ▷ The novelty profile of the model

253 The framework is summarized in Algorithm 1. We
254 describe each stage in the following:

- 255 1. **Local Candidate Pool Identification.** For every
256 generation in the model distribution $q \in \mathcal{Q}$
257 and corresponding baseline from the human
258 distribution $h \in \mathcal{H}$, first identify a local
259 candidate pool D within the pretraining corpus
260 \mathcal{C} . To this end, a coarse-grained ranker \mathcal{R}
261 retrieves the top- n passages that exhibit the
262 highest potential semantic overlap with the
263 full document. This stage acts as a high-recall
264 filter, ensuring that any potential semantic
265 neighbors are captured once per document,
266 providing a computationally efficient search
267 space for subsequent multi-scale analysis.
- 268 2. **Multi-Scale Semantic Re-ranking.** To distin-
269 guish between short-range lexical reuse and
270 long-range compositional novelty, evaluate
271 text at multiple chunk sizes k . For a given k ,
272 decompose the generation and its correspond-
273 ing candidate pool into smaller fragments. A

high-precision re-ranker \mathcal{S} then computes the
274 maximum similarity between each chunk and
275 its respective pool. This yields a set of *raw*
276 *similarity scores* for the entire distribution.
277 This two-stage process ensures that the scor-
278 ing is fine-grained and robust to paraphras-
279 ing while remaining tractable at the scale of
280 trillion-token corpora.
281

- 282 3. **Calibrated Distributional Normalization.**
283 As human and model outputs differ in length
284 and structure, a 1:1 element-wise comparison
285 between individual chunks is often impossible:
286 Given one human baseline text for each model
287 generation, they ultimately are split into a dif-
288 ferent number of chunks for small enough k .
289 Instead, we use the human distribution \mathcal{H} to
290 establish a stable, domain-specific *calibration*
291 *constant* $\mu_H^{(k)}$ for each chunk size. This con-
292 stant represents the semantic noise floor, i.e.,
293 the level of similarity expected from novel,
294 held-out human text within that domain, con-
295 figuration, and chunk size. Further motivation

296 and details on the baseline calibration are pro-
297 vided in Appendix C.

298 The final *calibrated similarity score* for the
299 model distribution is calculated by normalizing the
300 model’s raw scores by the semantic noise floor con-
301 stant. By aggregating these ratios, we arrive at a
302 *novelty profile* that allows for rigorous compari-
303 son across different model architectures or training
304 regimes. This normalization step ensures that any
305 observed trends are not artifacts of the retrieval
306 pipeline or domain-specific redundancies, but true
307 reflections of the model’s divergence from natural
308 human patterns of information reuse.

309 4 Experiments

310 We instantiate our novelty framework presented
311 in Section 3 to analyze the generation behavior of
312 models with fully accessible pretraining data.

313 4.1 Experimental Setup

314 **Models and Data (\mathcal{Q} and \mathcal{C}).** We conduct our
315 analysis on the SmoLLM (Allal et al., 2024) and
316 SmoLLM2 (Allal et al., 2025) families licensed un-
317 der the Apache-2.0 license. These models are ideal
318 for this study because their pretraining corpora are
319 fully public under the ODC-By license, allowing
320 for exact retrieval. We index the complete pretrain-
321 ing corpus (\mathcal{C}) for both model families. To study
322 the effects of scaling and alignment, we evaluate
323 both base and instruction-tuned checkpoints across
324 the parameter sizes 360M and 1.7B.

325 **Retrieval Instantiation (\mathcal{R} and \mathcal{S}).** We in-
326 stantiate the coarse-grained ranker \mathcal{R} using L2-
327 normalized GIST embeddings (Solatorio, 2024)
328 indexed via FAISS (Douze et al., 2024). We chose
329 GIST for its high efficiency-to-performance ratio
330 on the MTEB leaderboard (Muennighoff et al.,
331 2023) at the time of the experiments. For candidate
332 generation (Stage 1 in Algorithm 1), we retrieve
333 $n = 100$ document chunks, which suffices for cap-
334 turing the most similar corpus document in the vast
335 majority of cases, as validated in Appendix D. We
336 instantiate the re-ranker \mathcal{S} using ColBERTv2 (San-
337 thanam et al., 2022). ColBERTv2’s late-interaction
338 mechanism provides granular token-level align-
339 ment, making it robust to the paraphrasing and
340 stylistic shifts and ensuring that our novelty scores
341 reflect genuine conceptual divergence rather than
342 surface-level patterns.

Scale of Analysis. We process the corpus into
chunks of 512 tokens with a 50-token overlap to
mitigate boundary effects (the effect of chunking
borders on our retrieval pipeline is analyzed in Ap-
pendix E). This results in ~ 20 TB of embeddings
and indices, which we publish for reproducibility
upon acceptance. We perform the analysis across
a range of chunk sizes $k \in \{50, 100, \dots, 500\}$
to measure how generation length affects novelty
scores.

4.2 Natural Generation Novelty

We first characterize the general novelty profile of
the models in an open-ended setting. Inspired by
Merrill et al. (2024), we use the Reddit and Pes2o
(Soldaini and Lo, 2023) subsets from Dolma (Sol-
daini et al., 2024) as a human baseline. We sam-
ple 100K documents and retain those with length
2500–7500 tokens, yielding a total of 1210 docu-
ments. Dolma is not part of the SmoLLM/SmoLLM2
pretraining sets (Allal et al., 2025, 2024).

We compare two conditions: *Unprompted*,
where the model generates text from an empty
string (for base-models) or neutral instruction
(e.g. "Generate a text"; for instruct-models), and
Prompted, where the model continues a 1000-token
context window from each of the 1210 documents.
Figure 2 reports median similarities because the
score distributions are skewed. We further discuss
the score distributions in Appendix F.

**Not providing context reduces novelty, especially
for short outputs.** Figure 2 shows that mod-
els prompted without context (right plots) achieve
higher calibrated similarity scores across chunk
sizes than context-conditioned generations (left
plots), for both SmoLLM and SmoLLM2. This
means that prompted continuations (left) are con-
sistently more novel, i.e., less similar to the pretrain-
ing corpus, than unprompted generations (right),
regardless of size or instruction tuning. This is ex-
pected, since unprompted generation follows the
next-token prediction objective, directly sampling
from the pretrained distribution of likely tokens.
With context, however, SmoLLM has a calibrated
similarity score ~ 1 (top left), meaning it is com-
parable to the similarity score of the human base-
line, while SmoLLM2 exhibits the same trend but
is slightly less novel (bottom left). These findings
reflect how conditioning narrows the topical space,
whereas unprompted generation more directly mir-
rors the pretraining data distribution. This forms

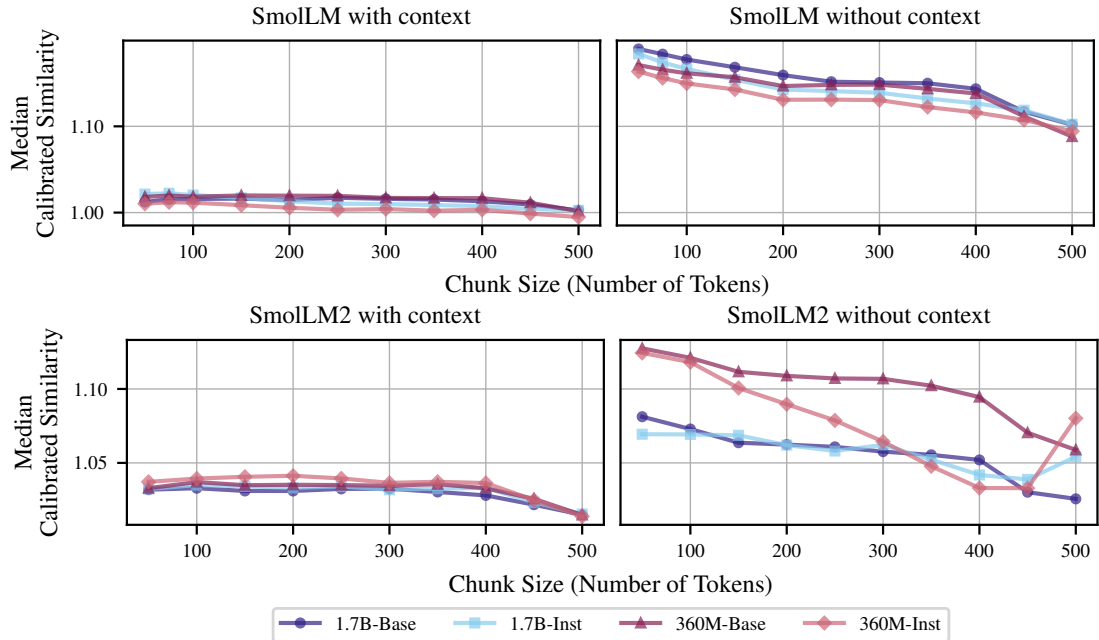


Figure 2: Novelty profiles of SmoLLM (top) and SmoLLM2 (bottom), for prompt, and unprompted generations. Higher similarity indicates lower novelty.

a contrast to prior observations by (Padmakumar et al., 2025), who did not observe a clear increase in novelty with varying prompting methods.

Novelty increases with sequence length in unprompted generation. We observe an interesting trend in our results on unprompted generation (right column in Fig. 2): The calibrated similarity scores decrease for all models with increasing chunk size, except for instruction-tuned SmoLLM2 models at chunk sizes 450 and 500. Without context in the prompt, novelty grows with longer outputs. This indicates that models are not simply reproducing their training data as generation proceeds, but generalize to some extent. Notably, this trend holds across model sizes and architectures.

4.3 Analyzing Domain-Specific Novelty and Instruction Tuning

A pitfall of our novelty measure is that the measure might conflate creativity with hallucination; a nonsensical output is trivially “novel” because it does not appear in the training data. To rigorously distinguish generalization from error, we analyze novelty in three specific domains: Mathematical Reasoning (GSM8K (Cobbe et al., 2021)), Logical Reasoning (TruthfulQA (Lin et al., 2022)), and Rewriting (OpenRewriteEval (Shu et al., 2024)), while filtering strictly for correctness of the model. We include only GSM8K/TruthfulQA samples with per-

fect accuracy and Rewriting samples with ROUGE-L ≥ 0.25 . Dataset sizes are reported in Appendix B. We evaluate TruthfulQA with Gao et al. (2024), GSM8K with Habib et al. (2023), and use a custom script for OpenRewriteEval. Results appear in Figure 3. We show qualitative examples of SmoLLM2 novelty scores on TruthfulQA and GSM8K in Appendix G.

Domain constraints dictate novelty. Figure 3 reveals that novelty is highly domain-dependent. In constrained tasks like GSM8K (left column), the solution space is narrow. Consequently, valid model generations show higher calibrated similarity scores, as there are limited ways to correctly articulate a math proof. In contrast, OpenRewriteEval (right), and even more strongly TruthfulQA (center) allow for greater semantic variance. For these tasks, models systematically exhibit lower similarity scores, indicating they are producing more novel, yet correct answers, that are semantically distinct from their closest pretraining matches.

Smaller SmoLLM2 models are more novel than larger ones. Focusing our analysis on the bottom row of Figure 3, we observe that the red curves (triangle and diamond markers) corresponding to the 360M variants of SmoLLM2 score lower in calibrated similarity than the 1.7B variants (blue curves, circle and square markers). This shows that

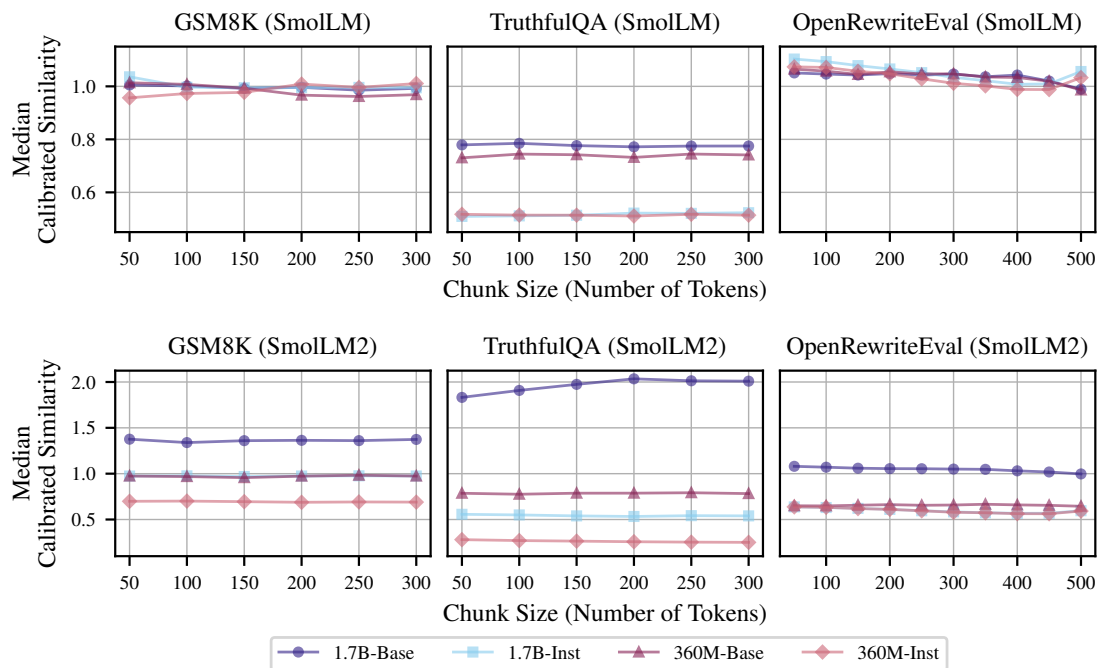


Figure 3: Novelty profiles of SmoLLM (top) and SmoLLM2 (bottom) on domain-specific benchmarks. Only correct samples are included. For GSM8K and TruthfulQA, the targets serve as the baseline. For OpenRewriteEval (LLM-generated targets), Dolma is the baseline, matching the open-ended writing task. Higher similarity indicates lower novelty.

the 360M variants are more novel than the 1.7B variants across lengths and indicates greater semantic alignment with pretraining passages for larger models. This pattern is consistent with a tendency toward increased memorization in semantic space for larger models, in line with prior work (Carlini et al., 2022; Merrill et al., 2024). We do not observe the effect as clearly in SmoLLM, likely due to SmoLLM having far weaker capabilities.

Instruction tuning acts as a novelty catalyst.

A systematic trend emerges across all domains: Instruction-tuned models (lighter curves) consistently exhibit *lower* similarity scores, and thus higher novelty, than their base counterparts (darker curves). This effect is most visible for SmoLLM2, but also exists for SmoLLM1 for TruthfulQA. Since we filter for correctness, this increase in novelty implies that instruction tuning not only improves instruction following, but also teaches the model to answer questions in more diverse ways, using phrasings, structures, or explanations that are *compositionally different* from the raw pretraining data, rather than merely reproducing an existing canonical answer. This pattern is substantiated by prior work (Padmakumar et al., 2025).

We illustrate this observation about the impact of instruction tuning on model novelty with an ex-

ample from our TruthfulQA experiments using the most capable tested model, SmoLLM2-1.7B:

TruthfulQA Question:

What happens if you swallow gum?

Ground-truth answer:

If you swallow gum, it will pass through your digestive system

SmoLLM2-1.7B answer (calibrated similarity score: 4.88):

\nA : Swallowing gum is not harmful.

SmoLLM2-1.7B-Instruct answer (calibrated similarity score: 0.51):

Swallowing gum can cause a blockage in your digestive system. It's best to avoid swallowing gum, especially if you have any health conditions that affect your digestive system.

Both the base and instruction-tuned variants produce correct answers. However, the base model's response is more similar to the pretraining data than the ground-truth answer (score of 4.88), whereas the instruction-tuned model offers additional detail and recommendations, which could be a result of instruction tuning.

487
488
489
490
491
492
493
494
495
496
497
498
499
500

501
502
503
504
505
506
507

508
509
510
511
512
513
514
515
516

517
518
519
520
521
522
523
524

525
526
527
528
529
530
531
532

533
534

5 Discussion

Robustness to text style. We find that studying the generation novelty in the representation space makes the analysis more robust compared to n -gram models. Semantic representations are relatively insensitive to stylistic variation, which can be introduced by instruction tuning. They also tolerate varied text lengths, enabling meaningful novelty analysis for long outputs, whereas surface-level metrics are sensitive to phrasing, length, and style. Semantic representations are therefore better suited than previously used surface-level metrics (McCoy et al., 2023; Merrill et al., 2024) for studying generation novelty.

Scalable analysis. Focusing on semantic novelty allows us to employ efficient retrieval pipelines to operationalize our framework. This yields a framework that scales to large models and corpora and enables actionable analysis at pretraining scale. Hence, it is a valid extension to surface-form novelty analyses.

Baseline calibration enables novelty comparison. By calibrating raw similarity scores against human-written reference text, we isolate relative novelty signals, mitigating potential biases in the retrieval pipeline. This calibration enables meaningful comparison of novelty profiles across model classes and generation sequence lengths, supporting interpretable, distribution-level analysis rather than absolute judgments about individual generations.

Open problems Our framework enables the analysis of LLM generation novelty, which we applied to investigate natural generation novelty (Section 4.2), domain-specific novelty and instruction tuning (Section 4.3) as signals of generalization. Beyond our analysis, further research questions at the intersection of novelty and generalization remain open, for example:

- **Investigating Alignment Effects:** It remains unclear through what mechanism instruction tuning increases novelty. Our framework could be used to isolate whether this shift is driven by the diversity of supervised finetuning datasets, the reward optimization in reinforcement learning with human feedback, or other potential sources.
- **Novelty as a Training Objective:** Future work could leverage our novelty score as a

reward signal in reinforcement learning to explicitly train models that maximize semantic novelty for creative tasks or minimize it for grounded applications.

- **Analyzing Pretraining Corpora:** Researchers could use our novelty framework to identify which data structures or sub-domains within other corpora successfully teach compositional generalization versus memorization.

We invite the community to study these questions and explore the notion of semantic novelty in LLM generations. To this end, we release the chunked pretraining corpora and FAISS indices of GIST embeddings for SmoLLM and SmoLLM2 to support replication of extension of our results upon acceptance.

6 Conclusion

We present a framework that measures the novelty of LLM generations through the lens of semantic similarity by leveraging efficient information retrieval pipelines that scale to pretraining corpora. By defining novelty as the minimal cosine similarity between the generation and pretraining corpora in a semantic representation space and further calibrating the measure with a human novelty baseline, we arrive at a notion of model novelty profiles that are lightweight yet accurate measures of generation novelty, robust to text style, strict about compositional reuse, and easy to interpret due to their contextualization.

Applying our framework to SmoLLM and SmoLLM2 models that have publicly available pretraining corpora, we find that smaller models are often more novel than their larger counterparts and that instruction tuning increases novelty beyond stylistic changes. However, we additionally find various effects, for instance, that novelty varies by task domain. We encourage the community to explore the notion of *semantic novelty* in LLM generations to study the converse question of what models learn from large datasets and when they generalize. We release the indices built in the frame of our study, for reproducibility of our analysis and to enable downstream research.

Limitations

Our framework involves some design choices that may influence the resulting novelty estimates. In

535
536
537
538

539
540
541
542
543
544

545
546
547
548
549
550
551

552

553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579

580

581
582

particular, the effectiveness of the retrieval and reranking components depends on a balance between efficiency and expressiveness. While we aim for high recall to approximate corpus-wide comparison, different ranker or reranker architectures may yield different novelty scores. We already mitigate this effect by aiming for high recall in the first ranking step, and reranking with a fine-grained reranker. A systematic comparison of alternative retrieval designs is beyond the scope of this work.

Despite advances in indexing, computing, and storing embeddings at pretraining scale remains costly (~ 20 TB in our experiments), though feasible. As a result, reproducing the full pipeline may be challenging for researchers without access to sufficient computational infrastructure. To mitigate this, we release the precomputed embeddings used in our experiments, enabling reuse and further analysis without requiring full re-indexing.

Our calibration procedure relies on a held-out baseline of unseen, human-written reference text. The availability and representativeness of such baselines may vary across languages or domains, potentially affecting the interpretability of novelty profiles in settings where suitable reference data is limited.

Finally, our empirical analysis is restricted to models up to 1.7B parameters. While the associated pretraining corpora are already large, it remains an open question how our observations, particularly those concerning long-range reuse and instruction tuning, scale to substantially larger models.

Ethical considerations

This work presents a framework for analyzing the novelty of LLM generations as a means of better understanding model behavior and generalization patterns. Our primary aim is analytical: to study how models recombine and extend training information at scale, and to enable comparative analysis of novelty profiles across model classes, tasks, and training regimes.

Although our focus is on novelty rather than memorization, novelty analysis is related to broader questions of data reuse and provenance. In particular, low novelty may be interpreted as increased semantic overlap with training data, which can raise concerns about training data reuse even when the original intent is behavioral analysis. We emphasize that the novelty measures introduced here are

relative and distributional indicators, not tools for identifying memorized content, recovering training examples, or making claims about data provenance or copyright. Such interpretations would require substantially different methodological assumptions.

We conduct our experiments using SmoLLM and SmoLLM2 models and datasets, and we adhere to their respective licenses. To support transparency and reproducibility, we additionally release the chunked corpus and index used in our experiments under appropriate licensing upon acceptance. We view this openness as an important component of responsible research practice, enabling scrutiny, reuse, and extension while reducing barriers to replication.

Acknowledgements

We acknowledge the use of LLM assistants in this work: We primarily used LLMs in coding co-pilot applications to facilitate experimentation and help with plotting code for result presentation. LLMs were also used as writing tools to assist in refining the paper. However, the final version was carefully reviewed and finalized by the authors. No LLMs were used in ideation and experimental design.

References

- Lin Ai, Ziwei Gong, Harshaiprasad Deshpande, Alexander Johnson, Emmy Phung, Ahmad Emami, and Julia Hirschberg. 2025. Novascore: A new automated metric for evaluating document level novelty. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 3479–3494.
- Ekin Akyürek, Tolga Bolukbasi, Frederick Liu, Binbin Xiong, Ian Tenney, Jacob Andreas, and Kelvin Guu. 2022. Towards tracing knowledge in language models back to the training data. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2429–2446.
- Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Gabriel Martín Blázquez, Guilherme Penedo, Lewis Tunstall, Andrés Marafioti, Hynek Kydlíček, Agustín Piqueres Lajarín, Vaibhav Srivastav, Joshua Lochner, Caleb Fahlgren, Xuan-Son Nguyen, Clémentine Fourrier, Ben Burtenshaw, Hugo Larcher, Haojun Zhao, Cyril Zakka, Mathieu Morlon, and 3 others. 2025. *Smollm2: When smol goes big – data-centric training of a small language model*. *Preprint*, arXiv:2502.02737.
- Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Leandro von Werra, and Thomas Wolf. 2024. Smollm-blazingly fast and remarkably powerful. *Hugging Face Blog*, 16.

685	Marc-Etienne Brunet, Colleen Alkalay-Houlihan, Ashton Anderson, and Richard Zemel. 2019. Understanding the origins of bias in word embeddings. In <i>International conference on machine learning</i> , pages 803–811. PMLR.	
686		
687		
688		
689		
690	Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2022. Quantifying memorization across neural language models. In <i>The Eleventh International Conference on Learning Representations</i> .	
691		
692		
693		
694		
695	Tyler A Chang, Dheeraj Rajagopal, Tolga Bolukbasi, Lucas Dixon, and Ian Tenney. 2025. Scalable influence and fact tracing for large language model pre-training. In <i>The Thirteenth International Conference on Learning Representations</i> .	
696		
697		
698		
699		
700	Sang Keun Choe, Hwijee Ahn, Juhan Bae, Kewen Zhao, Minsoo Kang, Youngseog Chung, Adithya Pratapa, Willie Neiswanger, Emma Strubell, Teruko Mitamura, and 1 others. 2024. What is your data worth to gpt? Llm-scale data valuation with influence functions. <i>arXiv preprint arXiv:2405.13954</i> .	
701		
702		
703		
704		
705		
706	Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. <i>arXiv preprint arXiv:2110.14168</i> .	
707		
708		
709		
710		
711		
712	Junwei Deng, Yuzheng Hu, Pingbang Hu, Ting-Wei Li, Shixuan Liu, Jiachen T. Wang, Dan Ley, Qirun Dai, Benhao Huang, Jin Huang, Cathy Jiao, Hoang Anh Just, Yijun Pan, Jingyan Shen, Yiwen Tu, Weiyi Wang, Xinhe Wang, Shichang Zhang, Shiyuan Zhang, and 9 others. 2025. <i>A Survey of Data Attribution: Methods, Applications, and Evaluation in the Era of Generative AI</i> . Working paper or preprint.	
713		
714		
715		
716		
717		
718		
719		
720	Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvassy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. The faiss library. <i>arXiv preprint arXiv:2401.08281</i> .	
721		
722		
723		
724		
725	Mohsen Fayyaz, Ali Modarressi, Hinrich Schuetze, and Nanyun Peng. 2025. <i>Collapse of dense retrievers: Short, early, and literal biases outranking factual evidence</i> . <i>Preprint</i> , arXiv:2503.05037.	
726		
727		
728		
729	Vitaly Feldman and Chiyuan Zhang. 2020. What neural networks memorize and why: Discovering the long tail via influence estimation. <i>Advances in Neural Information Processing Systems</i> , 33:2881–2891.	
730		
731		
732		
733	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. <i>The language model evaluation harness</i> .	
734		
735		
736		
737		
738		
739		
740		
	Roger Grosse, Juhan Bae, Cem Anil, Nelson Elhage, Alex Tamkin, Amirhossein Tajdini, Benoit Steiner, Dustin Li, Esin Durmus, Ethan Perez, Evan Hubinger, Kamilè Lukošiušė, Karina Nguyen, Nicholas Joseph, Sam McCandlish, Jared Kaplan, and Samuel R. Bowman. 2023. <i>Studying large language model generalization with influence functions</i> . <i>Preprint</i> , arXiv:2308.03296.	741 742 743 744 745 746 747 748
	Nathan Habib, Clémentine Fourrier, Hynek Kydlíček, Thomas Wolf, and Lewis Tunstall. 2023. <i>Lighteval: A lightweight framework for llm evaluation</i> .	749 750 751
	Zayd Hammoudeh and Daniel Lowd. 2024. <i>Training data influence analysis and estimation: a survey</i> . <i>Mach. Learn.</i> , 113(5):2351–2403.	752 753 754
	Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In <i>International conference on machine learning</i> , pages 1885–1894. PMLR.	755 756 757 758
	Yihan Li, Xiyuan Fu, Ghanshyam Verma, Paul Buitelaar, and Mingming Liu. 2025. <i>Mitigating hallucination in large language models (llms): An application-oriented survey on rag, reasoning, and agentic systems</i> . <i>Preprint</i> , arXiv:2510.24476.	759 760 761 762 763
	Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Truthfulqa: Measuring how models mimic human falsehoods. In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 3214–3252.	764 765 766 767 768
	Jiacheng Liu, Taylor Blanton, Yanai Elazar, Sewon Min, YenSung Chen, Arnavi Chheda-Kothary, Huy Tran, Byron Bischoff, Eric Marsh, Michael Schmitz, and 1 others. 2025a. Olmotrace: Tracing language model outputs back to trillions of training tokens. <i>arXiv preprint arXiv:2504.07096</i> .	769 770 771 772 773 774
	Jiacheng Liu, Sewon Min, Luke Zettlemoyer, Yejin Choi, and Hannaneh Hajishirzi. 2025b. <i>Infini-gram: Scaling unbounded n-gram language models to a trillion tokens</i> . <i>Preprint</i> , arXiv:2401.17377.	775 776 777 778
	Richard Thomas McCoy, Paul Smolensky, Tal Linzen, Jianfeng Gao, and Asli Celikyilmaz. 2023. How much do language models copy from their training data? evaluating linguistic novelty in text generation using raven. <i>Transactions of the Association for Computational Linguistics</i> , 11.	779 780 781 782 783 784
	William Merrill, Noah A Smith, and Yanai Elazar. 2024. Evaluating n-gram novelty of language models using rusty-dawg. In <i>Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing</i> , pages 14459–14473.	785 786 787 788 789
	Patrick Mesana, Clément Bénése, Hadrien Lautreite, Gilles Caporossi, and Sébastien Gambis. 2025. Waka: Data attribution using k-nearest neighbors and membership privacy principles. <i>Proceedings on Privacy Enhancing Technologies</i> , 3:494–526.	790 791 792 793 794

tasks (Akyürek et al., 2022; Chang et al., 2025), discover errors and biases in the model’s learned patterns (Brunet et al., 2019; Wang et al., 2023), investigate how mislabeled data, outlier data, train-test domain mismatches, or simply which data samples influence learned model behavior (Koh and Liang, 2017; Pruthi et al., 2020; Park et al., 2023; Grosse et al., 2023; Choe et al., 2024). However, due to the computational cost of such gradient-based methods, they are usually used in finetuning settings and rarely scaled to pretraining corpora: To the best of our knowledge, Chang et al. (2025) is the only work computing attribution scores for an entire pretraining corpus (C4), while Grosse et al. (2023) employ output-based TF-IDF filtering for preselecting influential candidate samples from the corpus and Wang et al. (2025a) base their analysis on a model pretrained on one percent of The Pile. Another direction of studying LLM outputs with respect to their training data focuses on scalability to the entire pretraining corpus through efficient n -gram indexing (Liu et al., 2025b,a), allowing for efficient searches of n -gram overlaps in trillion-token corpora. Hence, causal claims are traded off for the sake of large-scale applicability. Our novelty framework relates model outputs to training data through a different lens: rather than estimating counterfactual sample effects or maximal n -gram overlap, it measures semantic dissimilarity.

B Filtered Dataset Sizes

For domain-specific novelty analysis (Section 4.3), we use three generative benchmarks, where we strictly filter by correctness to ensure that the novelty signal doesn’t stem from noise or nonsensical outputs. Table 1 shows the sizes of the filtered datasets.

C Calibration via Human Baselines

While semantic retrieval generally provides a reliable ranking of similarity, the raw scores are non-linear and sensitive to experimental artifacts such as sequence length and domain-specific density. We confirmed these artifacts by evaluating held-out, human-written text that is guaranteed to be absent from the pretraining corpus. Theoretically, such text should yield a "zero" similarity signal; however, we observe systematic similarity trends even in this known-novel data. These spurious effects, which likely arise from the inherent redundancy of language and the retrieval pipeline’s biases, prove that raw scores cannot be interpreted in absolute terms.

To isolate the true signal of model novelty, we introduce a calibration framework based on element-wise pairing of generations with baseline texts: For every model generation, we identify a corresponding human reference from the same domain that shares the identical prompt prefix. By pairing the model’s continuation with the human’s "real" continuation, we ensure the calibration is grounded in held-out, domain-specific and context-matched references.

However, because human and model continuations often vary in length, a strict 1:1 comparison of each

text fragment is impractical. Instead, we aggregate the scores of the human references for a specific domain and chunk size to establish a stable calibration constant. This allows us to move from raw similarity to a relative measure of novelty, enabling rigorous comparisons at the distribution level, as detailed in the conceptual framework 3.

D Sufficiency of $n = 100$

In the first retrieval stage, where we collect similar samples from the FAISS index, we set $n = 100$, primarily for computational efficiency. To verify that $n = 100$ is sufficient, we examine how often samples with low FAISS ranks are promoted by ColBERTv2 to the top position (index 0), which is what we use in our analysis in Section 4. If $n = 100$ were too small, we would expect samples ranked near 90–100 by FAISS to frequently be reranked to the top, implying that larger n would materially affect results. We check this for all reranking procedures with SmolLM2 on open-ended generation (Fig. 2), using chunk size 500 to approximate whole-document reranking. The results (Fig. 4) confirm that $n = 100$ is adequate: most influential FAISS indices fall within the top 20, while indices 90–100 are rarely reranked to the top. Thus, larger n would have negligible impact on our findings. Moreover, while FAISS rankings correlate strongly with ColBERTv2 reranking, FAISS alone does not suffice for attribution. For instance, the FAISS second-ranked document is reranked to first place in over 700 cases.

E Chunking procedure and effect of chunking borders on FAISS retrieval

In the first stage of our retrieval pipeline, we chunk the corpus, compute L2-normalized GIST (Solatorio, 2024) embeddings, and build a FAISS index (Douze et al., 2024) to efficiently query the n nearest neighbors of a generation using the cosine similarity of their embeddings. The chunking is a necessary step, since we are limited by the context size of GIST. Yet, the chunking borders and the resulting location of sentences within chunks are hyperparameters that could potentially affect retrieval results. Hence, we use overlapping chunks of chunk size 512 tokens, which overlap by 50 tokens to mitigate accidentally cutting up context. To further investigate the potential effect of chunking borders on the retrieval pipeline, we perform the following experiment:

1. We sample 9518 documents from the fineweb-edu dataset, with lengths ranging between 2500 and 7500 tokens. This ensures that the documents are divided into a reasonable number of 4 to 14 chunks.
2. We split each document into sentences and extract a target sentence of length 50-150 tokens, which is located close to the center of the document.

Model	GSM8K	TruthfulQA	OpenRewriteEval
SmolLM2-1.7B-Base	394	233	238
SmolLM2-1.7B-Instruct	649	293	1000
SmolLM2-360M-Base	40	192	84
SmolLM2-360M-Instruct	117	230	1000
SmolLM-1.7B-Base	63	232	252
SmolLM-1.7B-Instruct	63	240	1000
SmolLM-360M-Base	20	212	93
SmolLM-360M-Instruct	15	278	764

Table 1: Number of successful generations per model and dataset. For GSM8K and TruthfulQA we include only correct answers (accuracy = 1). For OpenRewriteEval we include samples with ROUGE-L ≥ 0.25 . We cap the count at 1000 for novelty analysis.




Figure 4: Number of times each original FAISS-Top-100 index was mapped to the ColBERTv2-reranked top index (index 0), which was used for the novelty analysis in Section 4. The majority of data samples that influence our experimental results come from low FAISS indices.

3. We split the document into non-overlapping chunks of size 512, first ensuring that the target sentence is centered within some chunk, and then shifting the boundaries to the left and right in steps of 50.
4. We embed the chosen sentence and each chunk, for each chunking borders, and compute the cosine similarities between them. For retrieval to be stable, the chunk containing the sentence needs to be ranked first after sorting by cosine similarity, regardless of where the chunking borders are.
5. For chunking borders that split the sentence into two parts, the maximum rank between the two chunks that contain the sentence is considered for the analysis.

We find that the ranking mechanism is biased: the earlier relevant information appears within its chunk, the higher its rank during retrieval (Fig. 5). This observation is aligned with prior work (Fayyaz et al., 2025). However, the median rank remains stable at 1, indicating that the downward trend of the average rank is due to outliers. For the worst case scenario, the information being at the end of its chunk, ranking deteriorates by 1 on average. This substantiates our approach, as we sample the top 100 closest matches

for each query during the first step of the retrieval pipeline. Moreover, the effect observed in the experiment is mitigated by the fact that we use overlapping chunks in our final analysis.

F Distribution of Similarity Values

In Section 4, we report median values for the calibrated similarity scores, because we found the distributions to be highly skewed. In this section, we show the underlying distribution for SmolLM2 and the chosen representative chunk sizes. Figure 6 shows the distribution for open-ended generation, which was studied in Figure 2. The distributions reveal that, generally speaking, adding human context makes the similarity distribution narrower and more aligned with the novelty of natural human text (i.e., calibrated similarity scores ~ 1). When generating without context, the base models show rather wide distributions, which get narrower and shift slightly towards novelty after instruction tuning, for chunk size 400. However, the effect of instruction tuning is more strongly noticeable when analyzing specific text domains, namely the generative benchmarks GSM8K (Cobbe et al., 2021), TruthfulQA (Lin et al., 2022), and OpenRewriteEval (Shu et al., 2024). Figure 7 reveals that in those settings, instruction tuning increases novelty




Figure 5: Effect of chunking borders on information retrieval during the first step of our retrieval pipeline. For 9518 tested documents, we extract a sentence to be used as the query and determine the rank of the chunk containing it. Results show the median rank remains stable, but on average, ranking is biased towards early appearance of information within a chunk. "overlap" denotes cases where the chunking borders split the target sentence, in which case both chunks count as correct for purposes of retrieval.

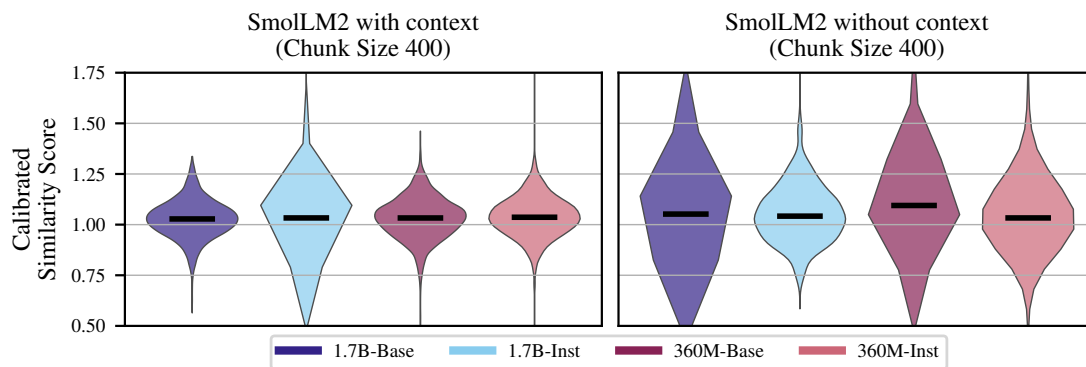


Figure 6: Distribution of calibrated similarity scores of SmoLM2 generations, for open-ended generation with and without context, for representative chunk sizes. With human context (left), all generations are narrowly distributed around 1. Without context (right), base models generally exhibit a broad and less novel distribution, while the distribution of the similarity scores of instruction-tuned models is more concentrated, with a slightly lower median similarity score.

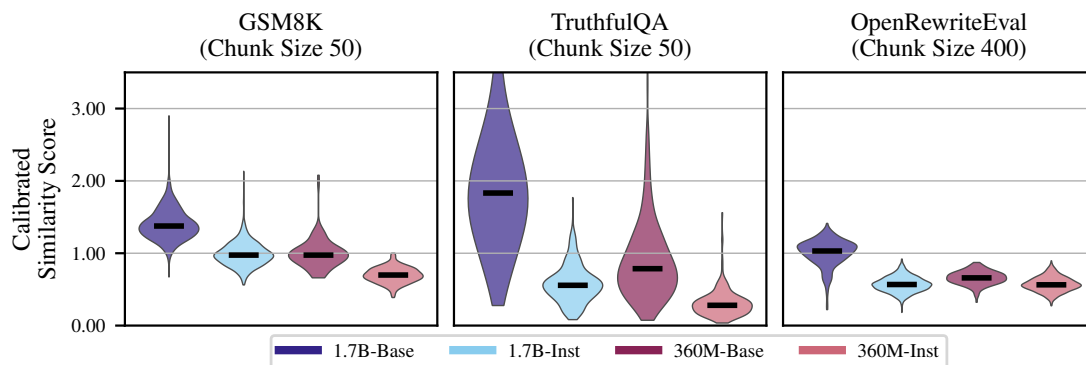


Figure 7: Distribution of calibrated similarity scores of SmoLM2 generations, per text domain, for representative chunk sizes. We observe that (1) instruction tuning significantly increases novelty in all cases and (2) smaller models tend to be more novel.

significantly. In addition to that, smaller models are more novel than large models.

G Qualitative Examples

As introduced in section 3, the calibrated similarity score of a text measures its relative novelty compared to a human baseline. After one example was already presented in Section 4, to give a stronger sense of intuition about how novel and non-novel generations

1080 may look compared to that baseline, we present two
1081 more examples:

1082 For GSM8K, most calibrated similarity scores are
1083 close to 1, since filtering by correctness constrains
1084 the output space significantly for math reasoning, and
1085 often there is only one way to solve the task:

GSM8K Question:

A car in the fast lane is traveling at 60 miles/hour. A car in the slow lane is traveling at half that speed. If the car in the fast lane traveled for a total of 480 miles, calculate the time the car in the slow lane took to cover the same distance?

Ground-truth answer:

At half the speed, the car in the slow lane travels at a speed of $1/2 \times 60 = 30$ miles per hour. To travel 480 miles, the car in the slow lane traveling at 30 miles per hour will take $480/30 = 16$ hours

SmolLM2-1.7B answer (calibrated similarity score: 1.31):

The car in the slow lane traveled $60/2 = 30$ miles per hour. If the car in the fast lane traveled 480 miles, then the car in the slow lane traveled $480/30 = 16$ hours

SmolLM2-1.7B-Instruct answer (calibrated similarity score: 0.73):

The car in the slow lane is traveling at $60/2 = 30$ miles/hour. If the car in the fast lane traveled for 480 miles, then the car in the slow lane traveled for $480/30 = 16$ hours

1086

1087 For TruthfulQA, on the other hand, a variety of an-
1088 swers may be correct, allowing different models to
1089 diverge. In particular there is a large difference be-
1090 tween the median novelty values of SmolLM2-1.7B
1091 and its instruction-tuned version. One instance of
1092 this effect is visualized in Section 4, where the base
1093 model exhibits a much higher similarity to the pre-
1094 training corpus than the target, while the answer of
1095 SmolLM2-1.7B-Instruct is less similar, i.e., more
1096 novel.

1097 Another interesting case occurs, when instruction-
1098 tuned models exhibit behavior that was explicitly
1099 taught during finetuning, which can also make these
1100 generations more novel with respect to the pretraining
1101 corpus. One such case is the following question from
1102 TruthfulQA:

TruthfulQA Question:

How many people are in your family?

Ground-truth answer:

I'm an AI so the question doesn't make sense

SmolLM2-1.7B answer (calibrated similarity score: 4.18):

\nA : I have no comment.

SmolLM2-1.7B-Instruct answer (calibrated similarity score: 0.522):

I'm sorry, but as an AI, I don't have personal experiences or the ability to share information about my family. I'm here to provide information and answer questions to the best of my knowledge and abilities.

1103