From Raw Corpora to Domain Benchmarks: Automated Evaluation of LLM Domain Expertise

Anonymous ACL submission

Abstract

Accurate domain-specific evaluation of LLMs is essential, specifically but not exclusively in domains with direct implications for humans, such as law, healthcare, and education. To measure domain-specific knowledge in LLMs, we present a deterministic pipeline that transforms raw domain corpora into completion-style benchmarks without relying on other LLMs or costly human annotation. Our approach generates domain-specific keywords and related target vocabulary, and then constructs prompt-target pairs based on these. By measuring LLMs' ability to complete these prompts with the correct domain-specific targets, we provide a direct assessment of domain knowledge with low computational cost. Our pipeline avoids benchmark contamination, enables automated updates with new domain data, and facilitates fair comparisons between base and instruction-tuned (chat) models.

Empirically, we first highlight the pitfalls of multiple-choice question benchmarks. We validate our approach by showing that model ranks on our benchmark significantly correlate with those on an expert-curated benchmark. We then demonstrate how our benchmark provides insights into knowledge acquisition in domain-adaptive, continual, and general pretraining. Finally, we examine the effects of instruction fine-tuning by comparing base and chat models within our unified evaluation framework. In conclusion, our pipeline enables scalable, domain-specific, LLM-independent, and unbiased evaluation of both base and chat models.

1 Introduction

State-of-the-art large language models (LLMs) underlie general-purpose chat assistants as well as highly specialized scientific tools. Most benchmarks, such as Chatbot Arena (Chiang et al., 2024) and MMLU (Hendrycks et al., 2020), are designed to evaluate these models on broad, mixed-domain tasks, such as math and college physics. Such coarse-grained evaluations cannot assess expertise in fine-grained subdomains (e.g., "Medical Imaging Analysis" within AI), which is essential for selecting models for specialized applications. Further, existing benchmarks suffer from intrinsic biases such as the order of choices in Multiple Choice Questions (MCQ) (Alzahrani et al., 2024) and benchmark leakage (Zhou et al., 2023). For instance, MedQA (Jin et al., 2021) and LegalQA (Louis et al., 2024) organize around broad disciplinary categories (e.g., geography or practice area) rather than subtopics, and employ MCQ or long-form answers, making model evaluation difficult.

A related problem arises when transforming LLMs into *domain experts* (Han et al., 2023; Cui et al., 2025; Nguyen et al., 2023)—models with enhanced ability to understand and generate accurate, contextually relevant responses in specialized fields such as law (Colombo et al., 2024), medicine (Singhal et al., 2025), and finance (Liu et al., 2023). The choice of the base model that the adaptation starts from is critical, as models with stronger domain performance can adapt more efficiently. The commonly used perplexity metric gives merely an uncalibrated measure of expertise (Öncel et al., 2024). Likewise, the recently proposed attribute rate (Geva et al., 2023), which quantifies model *knowledge*, does not follow any consistent pattern, as we demonstrate empirically. In short, there is no scalable, standardized framework to order models based on their domain expertise.

In this work, we present a framework for evaluating domain-specific knowledge in LLMs. Our framework automatically transforms raw text from any domain into a completion-style benchmark with prompt-target pairs. It does not require manual curation and does not rely on other LLMs during benchmark construction. We efficiently capture a model's domain expertise by computing the rank of the correct target given the prompt. Our approach is easily adaptable to new data and applies

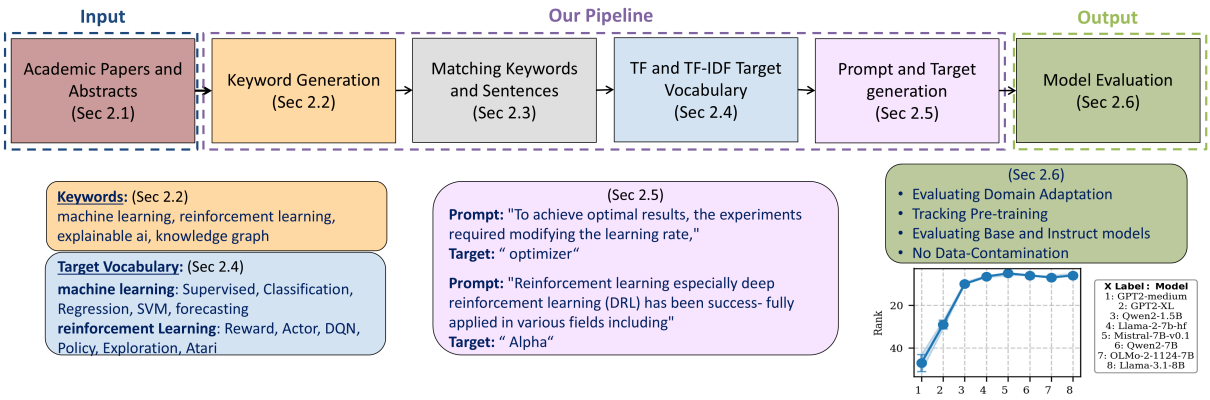

Figure 1: Overview of our proposed pipeline for generating a completion-based benchmark from a raw domain corpus. We begin by describing the domain data used, academic papers and their corresponding abstracts (Sec. 2.1). Our pipeline addresses a key challenge: extracting domain knowledge assessments from raw text rather than testing general linguistic ability. The first step identifies domain concepts by extracting and refining keywords from the abstracts through preprocessing the raw text and postprocessing the keywords (Sec. 2.2). Next, we match each sentence in the full papers with relevant keywords to focus evaluation on domain-relevant content (Sec. 2.3). For every keyword, we construct a target vocabulary by collecting domain-specific phrases from the matched sentences; these phrases will serve as prediction targets (Sec. 2.4). We then compile hundreds of prompt–target pairs for each keyword (Sec. 2.5). Finally, we outline our model evaluation framework, in which models are evaluated based on their ability to predict these targets from given prompts (Sec. 2.6).

to both base and chat models. Through extensive experiments, we *(i)* show the pitfalls of commonly employed Multiple Choice Question (MCQ) benchmarks (Sec. 3.1); *(ii)* verify the robustness of our proposed framework through controlled domain adaptation experiments (Sec. 4.1); *(iii)* demonstrate that our benchmark captures meaningful signals about knowledge acquisition during general and continual pretraining (Sec. 4.2-4.3); and *(iv)* enable comparisons between base and aligned models using a single, robust metric (Sec. 4.4).

2 Methods

This section outlines our benchmark construction and model evaluation pipeline (see Figure 1 for an overview). While we present our methodology using academic papers spanning diverse scientific domains, it is broadly applicable to any domain-specific corpus, e.g. law, healthcare, and education.

2.1 Data Curation

More specifically, our pipeline relies on RedPajama-Data-1T (Computer, 2023), which contains the full texts of 1.56 million arXiv papers, with a token count of 28 billion spanning multiple STEM disciplines. This dataset was accessed through the HuggingFace library (Wolf et al., 2019), consisting of preprocessed LaTeX source files with removed preambles, comments, macros, and bibliographies. To supplement the full text from RedPajama-Data-1T, we linked these documents to the arXiv Dataset available on Kaggle, which provides the metadata of all arXiv papers in RedPajama, including title, authors, categories, submission date, abstract, and DOI information. An example data entry, consisting of data fetched from Kaggle and RedPajama, is presented in Appendix A.4. Both datasets are publicly available for academic research use.

2.2 Keyword Generation

Next, we describe a four-stage procedure for extracting domain-specific keywords. This pipeline operates effectively on a subset of the data, eliminating the need to process the full corpus. Accordingly, we use only the abstracts from Kaggle to extract relevant keywords.

First, we apply a custom preprocessing routine to standardize the abstracts, including case normalization, removal of bracketed content, consistent handling of contractions, and tokenization that preserves hyphenated terms to retain technical precision. Second, we construct n-grams (2–7 tokens) using Gensim's (version 4.3.3) Phrases model (Rehurek and Sojka, 2010; Mikolov et al., 2013) with default parameters, then filter the results based on linguistic heuristics and domain relevance by removing stopwords, generic academic terms, function words, and non-alphabetic strings. Third, we apply adaptive length-based filtering to retain a pro-

portionally balanced subset of n-grams, adjusting ratios dynamically across domains to target approximately 300 high-quality keywords per domain, balancing domain coverage breadth with computational efficiency. Finally, we reduce redundancy by computing pairwise cosine similarities between keywords using the 'all-MiniLM-L6-v2" (Wang et al., 2020) model from the Sentence Transformer library (Reimers and Gurevych, 2019) and merging semantically similar entries with similarity scores above 0.85, ensuring strong semantic connections while avoiding overly restrictive matching. Please see Sec. A.1 for details of each step.

2.3 Matching Keywords and Sentences

We convert documents from RedPajama into individual sentences using spaCy's "en_core_web_sm" model (Honnibal et al., 2020) and embed both sentences and keywords from Section 2.2. Using cosine similarity with a threshold of 0.5, we extract sentences semantically related to each keyword. Selected sentences undergo systematic cleaning, including LaTeX formatting conversion, citation removal, and character validation to obtain clean sentence mappings for each domain keyword (see Appendix A.2 for implementation details).

2.4 TF and TF-IDF Target Vocabulary

Next, we describe how to extract domain-specific terms, which later serve as prediction targets. Our pipeline has two variants, which relate terms and keywords through term frequency (TF) and term frequency-inverse document frequency (TF-IDF).

TF target vocabulary First, we combine clean sentences for a given keyword to create a keyword-specific text corpus. We compute the term and document frequencies and ensure targets represent domain-specific rather than generic terms by excluding the words that appear in more than 80 percent of all the keyword corpora (e.g., the word *learning* for `CS.AI` domain). Then, to ensure the quality and relevance of these terms, we remove stop words and words with one or two characters. At the end, we obtain the initial term-document matrix capturing raw TF counts for the obtained terms. To maintain the relevance of the target vocabulary for each keyword, we discard targets with a TF score below the mean. We refine this set by removing targets whose cosine similarity with the corresponding keyword falls below 0.3. Supplementary Table 2 column "TF Target Vocabulary" presents 10 examples of targets obtained from this method for given keywords from the `CS.AI` domain.

TF-IDF target vocabulary For each keyword, the above pipeline generates a target vocabulary through term frequencies, which does not take into account how often/rarely these terms appear in other keyword corpora. For this, we replicate the above pipeline with two subtle but important differences to encourage more the niche target terms: *(i)* we obtain a term-document matrix capturing raw TF-IDF counts instead of TF counts, and *(ii)* we exclude the words that appear in more than 50 percent of all the keyword corpora (instead of 80 percent), further emphasizing niche domain-specific terminology. Supplementary Table 2 column and "TF-IDF Target Vocabulary" presents 10 examples of targets obtained from this method for given keywords from the `CS.AI` domain.

2.5 Prompt-Target Pair Construction

Our prompt-target pair construction pipeline utilizes keywords (Section 2.2), corresponding TF and TF-IDF vocabulary (Section 2.4), and the preprocessed sentences to generate prompts and targets for a domain of interest. To create prompts that effectively test model knowledge of each keyword, we first retrieve the preprocessed sentences that have a cosine similarity of at least 0.5 with the keyword. Further, to ensure that prompts provide enough context to predict the corresponding target, we filter out sentences shorter than 10 tokens (for the GPT-2 XL tokenizer) and 40 characters. Sentences with no words from the target vocabulary after the 10th token are also rejected. Note that due to abundant data, these steps still leave us with a huge number of samples for each keyword.

Following this procedure, we generate prompts consisting of all tokens up to, but not including, the match position of a term from the target vocabulary, with the target being the matched domain-specific term itself. This process is repeated for each keyword in the corpus until the desired number of prompts is reached or no further sentences remain for that keyword. Prompts and targets are constructed separately for the TF and TF-IDF vocabularies and are evaluated independently. Supplementary Table 4 presents an example of our prompt creation pipeline from the arXiv dataset from two domains. For each keyword, 50 prompts and targets were created for both TF and TF-IDF cases to represent the diversity of the domain.

2.6 Model Evaluation through Prediction Ranks and Probabilities

Utilizing prompt-target pairs, we evaluate each model's ability to predict the target given only the preceding context. For this, input the model with the prompt and record the *probability* and the *rank* of the correct first token.

If the target consists of multiple tokens, we employ a sequential evaluation approach: we first evaluate the rank of the first target token given the original prompt. For subsequent target tokens, we augment the prompt with the actual target tokens (not the model's predictions) to provide proper context. This ensures we measure the model's ability to complete the intended target sequence rather than penalizing it for alternative valid first tokens. We then average ranks and probabilities across all target tokens. Our corpus-based construction ensures that targets represent frequently occurring domain-specific terms; when multiple valid completions exist, models with sufficient domain knowledge assign high probability to multiple appropriate terms, keeping average ranks stable.

Predicted ranks can range from one to the full vocabulary size of the model, and we empirically observe that they often contain significant outliers across all evaluated models. To balance robustness to outliers with statistical efficiency (Rosenberger and Gasko, 1983; Wilcox and Keselman, 2003), we report the *20% trimmed mean* of predicted ranks, discarding the lowest and highest 20% of values. Ranks and probabilities are aggregated over all prompts and keywords within each domain, and we report corresponding 95% confidence intervals. We assess relationships between metrics using Pearson correlation coefficients and their associated p-values to quantify statistical significance. For completeness, we also include untrimmed mean ranks and probabilities in the Appendix. While these exhibit similar trends, their absolute values are substantially higher due to the outliers.

3 Evaluation Setup

To evaluate the efficacy of our pipeline, we compare the prediction ranks produced by our benchmark against a reference "gold standard." To choose the gold standard benchmark, we experimented with widely used multiple-choice question benchmarks; however, they exhibited several limitations, detailed in Sec 3.1. To address these shortcomings, we construct a domain-specific expert benchmark using a SOTA LLM (Sec 3.2).

3.1 Multiple-Choice Question Benchmarks Yield Inconsistent Evaluations

We show the inherent biases of benchmarking based on MCQ on the popular MMLU dataset. We test various prompting strategies, including shuffling answer choices and biasing options toward specific letters (A, B, C, or D). As shown in Fig. 2 (left), these changes consistently alter the ranking of models (please see Sec. A.8.2- A.8.3 for detailed analysis). This sensitivity highlights the instability of chat model evaluations through MCQs and reveals that the findings may be confounded by prompt engineering instead of solely measuring knowledge differences (Bai et al., 2022; Gupta et al., 2024; Alzahrani et al., 2024).

3.2 Expert Benchmark using Claude 4 and Verifying it using a Human-curated Benchmark

Based on the observations in Section 3.1, we construct a completion-type benchmark to serve as our reference standard. We opt for the completion task since it aligns with the training objective of both base and chat models, avoiding the format sensitivity issues inherent in MCQ evaluations (Figure 2, left). We provide the SOTA Claude 4 Sonnet model (Anthropic, 2025) with the top 20 keywords generated by our pipeline (Section 2.2) and requesting 50 prompts per keyword with corresponding targets (template in Appendix 13). Table 1 presents example prompts for three `CS.AI` keywords.

To validate the reliability of Claude-generated benchmarks, we created a comprehensive manual expert benchmark from the 2025 edition of the textbook "Understanding Deep Learning" (Prince, 2023), extracting 281 prompt-target pairs from this authoritative textbook. The right panel in Fig. 2 demonstrates that the Claude-generated benchmark exhibits strong correlation with manually curated expert assessments (r=0.91, p=0.012), while our TF-based pipeline shows even stronger alignment with the expert benchmark (r=0.99, p<0.001). These results validate that Claude can generate reliable domain-specific benchmarks comparable to expert human curation, justifying its use as our reference standard.

While manual validation confirms benchmark quality, manual curation is impractical for our pipeline's core value proposition. Manual benchmarks require (i) intensive expert effort, (ii) re-

peated updates to prevent contamination as models train on newer data, (iii) availability of domain expertise for specialized fields, and (iv) scalability across emerging domains. Our automated pipeline addresses these limitations by enabling on-demand benchmark generation from arbitrary raw corpora without manual intervention, as demonstrated across diverse domains in Section 4.

4 Results and Discussions

To validate the efficacy and robustness of our proposed pipeline, we conduct a comprehensive set of experiments across diverse domains. First, in a controlled domain adaptation setup, we show the robustness of our *prediction rank* metric (Sec.4.1). We then evaluate model checkpoints during general pretraining (Sec.4.2) and continual pretraining (Sec.4.3), revealing that our benchmark captures nuanced patterns of knowledge acquisition that are missed by perplexity and attribution rate. Finally, we apply our benchmark to compare base and instruction-tuned models across four domains (Sec. 4.4), showing its consistency across heterogeneous model families without relying on few-shot protocols.

4.1 Validation of Our Pipeline through Domain Adaptation

To validate our pipeline, we conduct controlled domain adaptation experiments where we adapt GPT2-XL and Llama2-7B models to domains with varying semantic proximity to our target evaluation domain (`CS.AI`). Models adapted to semantically similar domains (e.g., `CS.CL`) consistently achieve better prediction ranks on our benchmarks compared to models trained on unrelated domains (e.g., `Material Science`), demonstrating that our evaluation captures genuine domain expertise. Our TF- and TF-IDF-based benchmarks show strong correlation with the expert benchmark, while perplexity shows weaker correlation and attribution rate shows no correlation (please see Sec. A.5 for detailed descriptions of perplexity and attribute rate). These controlled experiments validate that our pipeline effectively quantifies domain-specific knowledge acquisition.

Complete experimental setup, detailed results for GPT2-XL and Llama2-7B domain adaptation experiments, correlation analyses, and all evaluation metrics are provided in Sec. A.6 and Fig. 6-9. Notably, smaller models like GPT2-XL demonstrate rapid adaptation with significant metric changes after few training steps, while larger models like Llama2-7B exhibit more modest changes due to their existing extensive knowledge base. Additional analyses including semantic similarity validation and mean aggregation results are in Sec. A.7 and Sec. A.10.

4.2 Quantifying Knowledge Acquisition During General Pretraining

As established in Sec. 3.1, MCQ downstream evaluation of OLMo-2 checkpoints (OLMo et al., 2024) saturates quickly. As an alternative, we propose to quantify knowledge acquisition by evaluating the checkpoints on our TF-based benchmark.

Fig. 3 demonstrates the drastic rank improvement of the first 50 checkpoints (roughly 5% of the entire pertaining) for the `Physics and Society domain (Physics (Soc-Ph))`. The left panel displays the first 50 checkpoints. We observe a significant improvement in model ranks during initial pretraining, which is expected as the weights evolve from their randomly initialized state and the model sharply improves. This rapid improvement continues for approximately the first 2.5% of the pretraining before diminishing. Notably, even at the 50th checkpoint, a significant difference in rank persists between this intermediate state and the final base model, indicating substantial potential for further training in this domain.

For the second analysis, we selected two groups of checkpoints *(i)* 50 equally spaced checkpoints from the first stage of pretraining on 4T tokens, and *(ii)* 5 checkpoints from the second stage on 150B high-quality data (Supplementary Table 6 lists all evaluated checkpoints.) The right panel of Fig. 3 illustrates a downward trend throughout pretraining, which gradually saturates. We observe a clear improvement in probability when switched to high-quality data (the last 5 checkpoints).

How does MCQ benchmarking quantify knowledge acquisition? Recent work employed MCQs to monitor performance during pretraining (OLMo et al., 2024). Notably, evaluation logs[1] reveal that downstream performance on 16 MCQ tasks saturates halfway through training, even as perplexity continues to improve. This suggests that downstream MCQ evaluation can behave inconsistently,

[1] https://wandb.ai/ai2-llm/OLMo-2-1124-7B/reports/OLMo-2-7B-Nov-2024--VmlldzoxMDUzMzE1OA

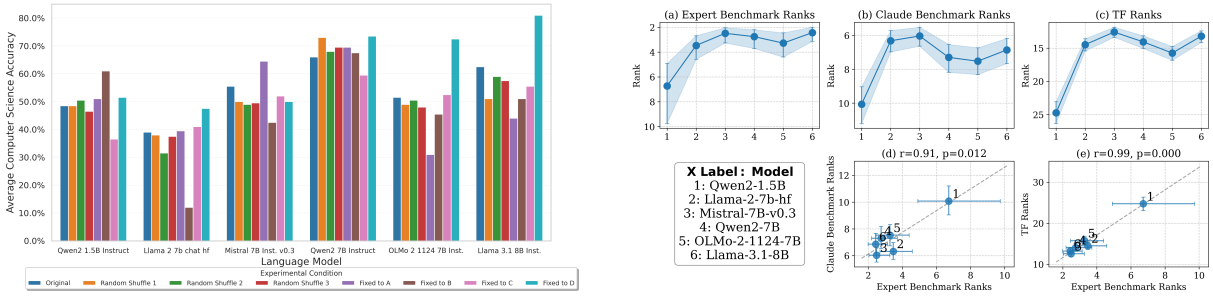

Figure 2: Validation of benchmarking approaches. Left: MCQ benchmarks show significant sensitivity to option ordering (detailed in Appendix A.8.2-A.8.3). Right: Completion-based validation across six base models showing (a) manual expert benchmark ranks, (b) Claude-generated benchmark ranks, and (c) TF-based pipeline ranks. Correlation analysis reveals (d) r=0.91, p=0.012 between Claude and expert benchmarks, and (e) r=0.99, p<0.001 between TF-based and expert benchmarks, validating that Claude-generated benchmarks reliably capture expert-level domain knowledge patterns.

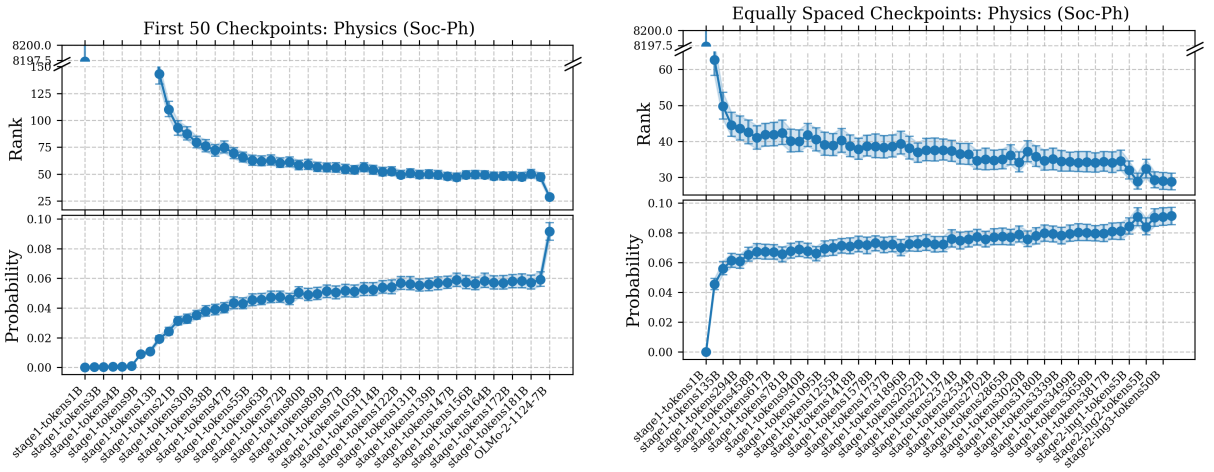

Figure 3: Prediction ranks and probabilities of OLMo-2 pretraining checkpoints on the `Physics and Society domain (Physics (Soc-Ph))`. The left figure displays the results across the first 50 checkpoints, with the last data point representing the final model for reference. The right figure displays equally spaced checkpoints from the entire pretraining. Consistent patterns above demonstrate that our benchmark can guide the knowledge accumulation for a domain of interest by replacing the average perplexity metric.

or even misleadingly, when used as a proxy for overall model learning.

Supplementary findings We also evaluated the checkpoints on the `CS.AI` domain, depicted in Supp. Fig. 15, which exhibits a similar trend, verifying the consistency of our evaluation across downstream domains. We note that the overall ranks are considerably better on `CS.AI`, indicating the relative ease of the task for the model. We attribute this to relatively abundant training data from this domain, resulting in improved performance.

4.3 Continually Pretraining Llama2-7B model

To obtain an expert model for a specific domain, continually pretraining on related domains has been shown to outperform adaptation solely on the target domain (Gururangan et al., 2020). However, recent work (Yıldız et al., 2024) highlights that poor selection of pretraining corpora can degrade performance. This underscores the importance of tracking knowledge acquisition throughout continual learning. In this section, we show that perplexity provides an uncalibrated and often misleading proxy for model knowledge, whereas our pipeline enables more robust and interpretable analyses.

Training setup Following previous experiments, we select `CS.AI` as the target domain and continually pretrain Llama2-7B on a sequence of domains from M2D2. To ensure a sufficiently long pretraining horizon, we exclude domains smaller than 100MB. To construct a meaningful training order, we randomly sample 5000 abstracts from the arXiv dataset, embed them using Sentence Transformers,

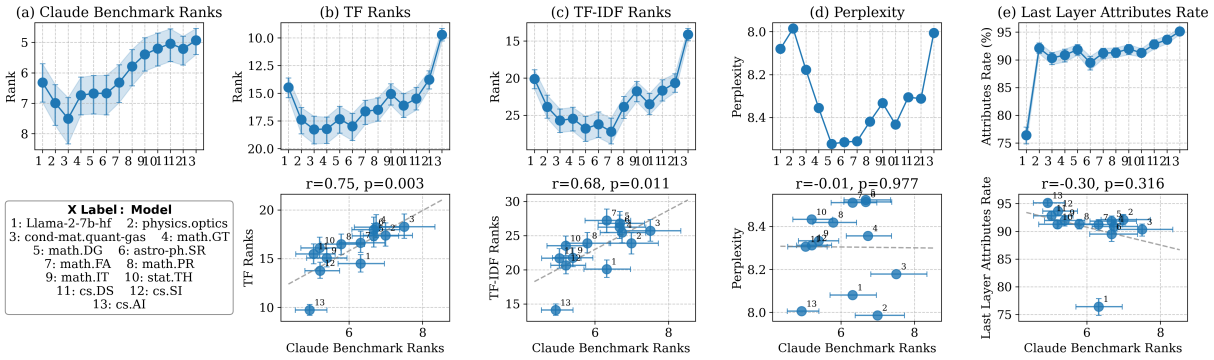

Figure 4: Evaluation results for 12 checkpoints from continual pretraining of Llama2-7B, along with the base model. The bottom-left panel displays the sequence of pretraining domains in the order they were introduced during training. Top row: Prediction ranks on *(a)* the benchmark generated by Claude 4 Sonnet, and *(b–c)* TF- and TF-IDF-based benchmarks produced by our pipeline. The last two columns show the baseline metrics *(d)* perplexity and *(e)* last-layer attribution rate. Bottom row: Correlation of each metric with the Claude benchmark. While ranks computed on our benchmarks align with the similarity between pretraining domains and the target `CS.AI` domain, perplexity and attribute rate do not offer robust measures for knowledge acquisition.

and compute their cosine similarity to embeddings of `CS.AI` domain abstracts. We then heuristically sort the M2D2 domains by their similarity to `CS.AI`. From this list, we uniformly sample 12 domains, sort them by *dissimilarity*, and use them as the continual pretraining sequence for Llama2-7B (see domain names in Fig. 4).

Findings Fig. 4 shows how different metrics evolve as the model is continually pretrained. The ranks on the Claude benchmark (Fig. 4a) initially decrease, then remain relatively stable, and eventually increase. This progression aligns with the increasing similarity between the continual pretraining domains and the target `CS.AI` domain. The prediction ranks on our TF- and TF-IDF-based benchmarks follow a similar trajectory (Figs. 4b and 4c), whereas perplexity and attribute rate show no clear correlation with Claude benchmark ranks (Figs. 4d and 4e). Taken together, our metrics offer a more robust and informative assessment of changes in performance during continual pretraining than the commonly used perplexity scores.

Supplementary findings For completeness, we report target token probabilities (Fig. 16) and how the metrics behave under untrimmed mean aggregation (Figs. 17- 18). Across all evaluations, the patterns discussed above remain consistent.

4.4 Comparing Base and Chat Models

Up to this point, our analysis has focused on targeted settings such as adapted models, general-purpose model pretraining, and continual pretraining. We now extend our evaluation to six publicly available open-weight models spanning a range of scales and architectures. For each model, we include both the base and chat variants. This setup enables a unified evaluation framework that facilitates direct comparisons across model families, without relying on conventional MCQ benchmarks; thus avoiding the issues discussed in Sec. 3.1.

Fig. 5 compares six base models and their aligned (instruction-tuned) counterparts on the `General Economics (Econ-GN)` domain. Most notably, we observe a strong correlation between the ranks computed on the Claude benchmark and our TF- and TF-IDF-based benchmarks. While perplexity also correlates well with Claude ranks for base models (Fig.19), it fails to provide reliable signals for aligned models. In contrast, our rank-based approach remains consistent across both base and aligned checkpoints, demonstrating its superiority as a general-purpose evaluation metric. Please see Figs 20,21,22 for the same analysis repeated on three other domains, leading to consistent results.

Our rank analysis also allows us to inspect how alignment impacts model knowledge. Fig. 23 reveals that base models generally outperform their chat counterparts, corroborating findings from Bai et al. (2022) who demonstrated that smaller models experience significant "alignment tax" where performance degrades after instruction tuning. This degradation is particularly pronounced for Llama2-7B and Mistral-7B-v0.3, whose performance drops after instruction tuning across all domains (`CS.AI`, `Physics (Soc-Ph)`, `Q-Bio.PE`, and `Econ-GN`), suggesting significant room for improvement in

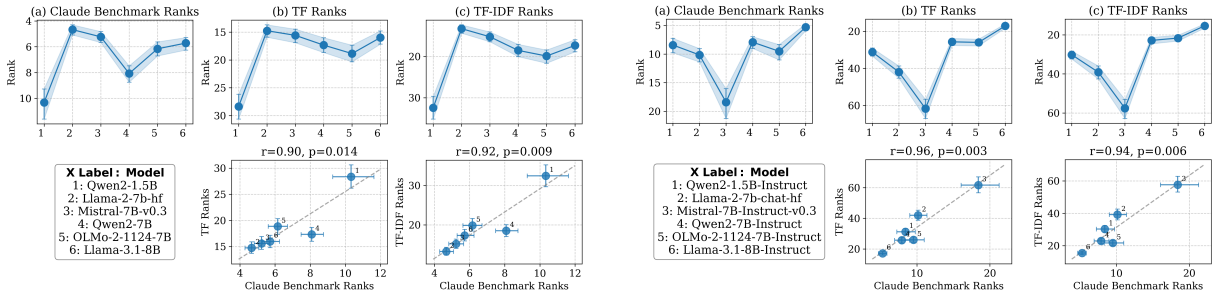

(a) Base models (b) Chat models

Figure 5: Evaluation of 6 base models (left) and 6 chat models (right) on `General Economics (Econ-GN)` domain. The top row shows prediction ranks, and the bottom row shows the correlations between the Claude benchmark and our TF- and TF-IDF-based benchmarks.

their alignment pipelines. Moreover, our framework provides a more fine-grained view than the generic evaluations in Bai et al. (2022), revealing that some models like Qwen2-1.5B and Llama3.1-8B exhibit smaller degradations or even occasional improvements, demonstrating that alignment effects are highly heterogeneous across architectures and domains. This suggests that while alignment may dilute specialized domain knowledge, where the technical vocabulary required for domain expertise may be weakened during instruction tuning that emphasizes broader linguistic competencies.

5 Related Work

The evaluation of LLMs has evolved from statistical measures like perplexity (Gururangan et al., 2021; Yıldız et al., 2024) to comprehensive benchmarks such as MMLU (Hendrycks et al., 2020) and BIG-bench (Srivastava et al., 2022). However, perplexity primarily captures linguistic capabilities rather than domain knowledge (Öncel et al., 2024), while QA benchmarks suffer from sensitivity to formatting (Alzahrani et al., 2024), conflation of model size with domain knowledge in few-shot settings (Brown et al., 2020), and focus on broad rather than domain-specific tasks (Chang et al., 2024). Recent LLM-based benchmark generation (Yuan et al., 2025) addresses some limitations but introduces computational costs and grounding issues (Li et al., 2024; Monea et al., 2023). Likewise, automated benchmarking systems such as LiveBench (White et al., 2024) and LiveCodeBench (Jain et al., 2024) address contamination through recent curated content or existing problem collections, but cannot generate benchmarks from arbitrary raw corpora. We introduce a deterministic pipeline that converts raw domain corpora into completion-style benchmarks, providing unified evaluation for base and chat models while eliminating prompt sensitivity biases (see Appendix A.12 for detailed discussion).

6 Conclusion

We introduced and validated a framework to construct completion-style benchmarks from raw domain corpora. By generating prompt–target pairs from automatically extracted domain keywords and target phrases, we quantify model expertise through prediction ranks of the targets given the prompts. Validation against manually curated expert benchmarks (r=0.91-0.99, $p \leq 0.012$) confirms our automated approach captures domain expertise patterns consistent with human expert judgment. Our lightweight evaluation framework circumvents the intrinsic biases of multiple-choice questions and does not suffer from benchmark leakage. Likewise, it allows for automatically updating a benchmark with ever-growing domain text with minimal cost.

Our evaluations demonstrate that this framework captures meaningful distinctions across models and training settings: it reliably traces knowledge acquisition during general, adaptive, and continual pretraining and highlights performance differences between base and aligned models. Altogether, our work offers a principled alternative for practitioners and researchers seeking to assess and track LLM competence in specific domains, while avoiding benchmark contamination and being applicable to both base and instruction-tuned models.

7 Limitations

While our proposed pipeline demonstrates strong empirical results across multiple experimental settings, several limitations warrant discussion.

Our validation approach uses Claude 4 Sonnet-generated benchmarks as a reference standard. To validate this choice, we created a manual expert benchmark for the `CS.AI` domain from "Understanding Deep Learning" Prince (2023), which demonstrated strong correlation with Claude-generated benchmarks (r=0.91, p=0.012), confirming Claude's reliability for this purpose. However, we acknowledge that extending manual validation to all evaluated domains (Physics, Economics, Biology) was beyond the scope of this work due to resource constraints. The strong correlations we observe between our pipeline and Claude benchmarks on `CS.AI` domain, coupled with intuitive behavioral patterns in controlled experiments, suggest that this validation generalizes beyond `CS.AI`.

Our experimental scope encompasses academic texts from arXiv spanning four domains from different science fields, which, while providing clean and well-structured data ideal for methodological validation, represents a focused subset of potential applications. Similarly, computational constraints limited our model coverage to eight base models and six chat variants ranging from 1.5B to 8B parameters, and restricted our domain adaptation and continual learning experiments primarily to Computer Science (Artificial Intelligence). While evaluation of significantly larger models (e.g., >70B parameters) would provide additional validation, our current model range captures diverse training paradigms and architectural designs. These scope limitations are largely driven by practical considerations, but the consistent patterns we observe across diverse STEM domains and model architectures (GPT2, Llama2, Mistral, Qwen2) at different scales provide evidence for broader applicability.

The evaluation of instruction-tuned chat models presents additional challenges since these models are optimized for conversational interactions rather than completion tasks. As we demonstrated in our MMLU analysis, chat models exhibit significant sensitivity to prompt formatting and ordering effects, which could potentially affect our completion-based evaluation. Nevertheless, completion tasks align more closely with the fundamental training objective shared by both base and chat models, and the systematic patterns we observe suggest our method captures meaningful signals about knowledge retention.

These limitations, while acknowledging areas for future expansion, do not diminish the fundamental contribution of providing a scalable, automated approach to domain-specific knowledge evaluation that addresses critical gaps in current LLM assessment methodologies. Moreover, our pipeline's inherently domain-agnostic design facilitates extension to other corpus types and domains, requiring only minor preprocessing adjustments while maintaining the core methodology.

References

Norah Alzahrani, Hisham Alyahya, Yazeed Alnumay, Sultan Alrashed, Shaykhah Alsubaie, Yousef Almushayqih, Faisal Mirza, Nouf Alotaibi, Nora Al-Twairesh, Areeb Alowisheq, and 1 others. 2024. When benchmarks are targets: Revealing the sensitivity of large language model leaderboards. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13787–13805.

Anthropic. 2025. Claude models. Accessed: March 30, 2025.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, and 1 others. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, and 1 others. 2024. A survey on evaluation of large language models. *ACM transactions on intelligent systems and technology*, 15(3):1–45.

Arnav Chavan, Raghav Magazine, Shubham Kushwaha, Mérouane Debbah, and Deepak Gupta. 2024. Faster and lighter llms: A survey on current challenges and way forward. *arXiv preprint arXiv:2402.01799*.

Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Banghua Zhu, Hao Zhang, Michael Jordan, Joseph E Gonzalez, and 1 others. 2024. Chatbot arena: An open platform for evaluating llms by human preference. In *Forty-first International Conference on Machine Learning*.

Pierre Colombo, Telmo Pessoa Pires, Malik Boudiaf, Dominic Culver, Rui Melo, Caio Corro, Andre FT Martins, Fabrizio Esposito, Vera Lúcia Raposo, Sofia Morgado, and 1 others. 2024. Saullm-7b: A pioneering large language model for law. *arXiv preprint arXiv:2403.03883*.

Together Computer. 2023. Redpajama: An open source recipe to reproduce llama training dataset.

Shidong Cui, Sicong Yu, Hsi-Yuan Huang, Yang-Chi-Dung Lin, Yixian Huang, Bojian Zhang, Jihan Xiao, Huali Zuo, Jiayi Wang, Zhuoran Li, and 1 others. 2025. mirtarbase 2025: updates to the collection of experimentally validated microrna–target interactions. *Nucleic Acids Research*, 53(D1):D147–D156.

Guy Dar, Mor Geva, Ankit Gupta, and Jonathan Berant. 2022. Analyzing transformers in embedding space. *arXiv preprint arXiv:2209.02535*.

Philippe Faist. 2021. pylatexenc: Simple latex parser providing latex-to-unicode and unicode-to-latex conversion. Python package.

Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023. Dissecting recall of factual associations in auto-regressive language models. *arXiv preprint arXiv:2304.14767*.

Vipul Gupta, David Pantoja, Candace Ross, Adina Williams, and Megan Ung. 2024. Changing answer order can decrease mmlu accuracy. *arXiv preprint arXiv:2406.19470*.

Suchin Gururangan, Mike Lewis, Ari Holtzman, Noah A Smith, and Luke Zettlemoyer. 2021. Demix layers: Disentangling domains for modular language modeling. *arXiv preprint arXiv:2108.05036*.

Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith. 2020. Don't stop pretraining: Adapt language models to domains and tasks. *arXiv preprint arXiv:2004.10964*.

Tianyu Han, Lisa C Adams, Jens-Michalis Papaioannou, Paul Grundmann, Tom Oberhauser, Alexander Löser, Daniel Truhn, and Keno K Bressem. 2023. Medalpaca–an open-source collection of medical conversational ai models and training data. *arXiv preprint arXiv:2304.08247*.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.

Matthew Honnibal, Ines Montani, Sofie Van Landeghem, Adriane Boyd, and 1 others. 2020. spacy: Industrial-strength natural language processing in python.

Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*.

Fred Jelinek, Robert L Mercer, Lalit R Bahl, and James K Baker. 1977. Perplexity—a measure of the difficulty of speech recognition tasks. *The Journal of the Acoustical Society of America*, 62(S1):S63–S63.

Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14):6421.

Tianle Li, Ge Zhang, Quy Duc Do, Xiang Yue, and Wenhu Chen. 2024. Long-context llms struggle with long in-context learning. *arXiv preprint arXiv:2404.02060*.

Xiao-Yang Liu, Guoxuan Wang, Hongyang Yang, and Daochen Zha. 2023. Fingpt: Democratizing internet-scale data for financial large language models. *arXiv preprint arXiv:2307.10485*.

Antoine Louis, Gijs van Dijck, and Gerasimos Spanakis. 2024. Interpretable long-form legal question answering with retrieval-augmented large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 22266–22275.

Clara Meister and Ryan Cotterell. 2021. Language model evaluation beyond perplexity. *arXiv preprint arXiv:2106.00085*.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26.

Giovanni Monea, Maxime Peyrard, Martin Josifoski, Vishrav Chaudhary, Jason Eisner, Emre Kıcıman, Hamid Palangi, Barun Patra, and Robert West. 2023. A glitch in the matrix? locating and detecting language model grounding with fakepedia. *arXiv preprint arXiv:2312.02073*.

Tuan Dung Nguyen, Yuan-Sen Ting, Ioana Ciucă, Charlie O'Neill, Ze-Chang Sun, Maja Jabłońska, Sandor Kruk, Ernest Perkowski, Jack Miller, Jason Li, and 1 others. 2023. Astrollama: Towards specialized foundation models in astronomy. *arXiv preprint arXiv:2309.06126*.

nostalgebraist. 2020. Interpreting GPT: the logit lens.

Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, and 1 others. 2024. 2 olmo 2 furious. *arXiv preprint arXiv:2501.00656*.

Fırat Öncel, Matthias Bethge, Beyza Ermis, Mirco Ravanelli, Cem Subakan, and Çağatay Yıldız. 2024. Adaptation odyssey in llms: Why does additional pretraining sometimes fail to improve? *arXiv preprint arXiv:2410.05581*.

Ethan Perez, Sam Ringer, Kamile Lukosiute, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, and 1 others. 2023. Discovering language model behaviors with model-written evaluations. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13387–13434.

Simon J.D. Prince. 2023. *Understanding Deep Learning*. The MIT Press.

Radim Rehurek and Petr Sojka. 2010. Software framework for topic modelling with large corpora.

Machel Reid, Victor Zhong, Suchin Gururangan, and Luke Zettlemoyer. 2022. M2d2: A massively multi-domain language modeling dataset. *arXiv preprint arXiv:2210.07370*.

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.

James L Rosenberger and Miriam Gasko. 1983. Comparing location estimators: Trimmed means, medians, and trimean. *Understanding robust and exploratory data analysis*, pages 297–336.

Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Mohamed Amin, Le Hou, Kevin Clark, Stephen R Pfohl, Heather Cole-Lewis, and 1 others. 2025. Toward expert-level medical question answering with large language models. *Nature Medicine*, 31(3):943–950.

Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, and 1 others. 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.

Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in neural information processing systems*, 33:5776–5788.

Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Shwartz-Ziv, Neel Jain, Khalid Saifullah, Siddartha Naidu, and 1 others. 2024. Livebench: A challenging, contamination-free llm benchmark. *arXiv preprint arXiv:2406.19314*.

Rand R Wilcox and Harvey J Keselman. 2003. Modern robust data analysis methods: measures of central tendency. *Psychological methods*, 8(3):254.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and 1 others. 2019. Huggingface's transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.

Zhichao Xu, Ashim Gupta, Tao Li, Oliver Bentham, and Vivek Srikumar. 2024. Beyond perplexity: Multi-dimensional safety evaluation of llm compression. *arXiv preprint arXiv:2407.04965*.

Çağatay Yıldız, Nishaanth Kanna Ravichandran, Prishruit Punia, Matthias Bethge, and Beyza Ermis. 2024. Investigating continual pretraining in large language models: Insights and implications. *arXiv preprint arXiv:2402.17400*.

Peiwen Yuan, Shaoxiong Feng, Yiwei Li, Xinglin Wang, Yueqi Zhang, Jiayi Shi, Chuyi Tan, Boyuan Pan, Yao Hu, and Kan Li. 2025. Llm-powered benchmark factory: Reliable, generic, and efficient. *arXiv preprint arXiv:2502.01683*.

Kun Zhou, Yutao Zhu, Zhipeng Chen, Wentong Chen, Wayne Xin Zhao, Xu Chen, Yankai Lin, Ji-Rong Wen, and Jiawei Han. 2023. Don't make your llm an evaluation benchmark cheater. *arXiv preprint arXiv:2311.01964*.

A Supplementary Material

Table 1: Prompts and targets generated by Claude 4 for three keywords from `CS.AI`

Prompts and Targets for Keywords from `CS.AI`

Reinforcement Learning Agents:
Prompt: The entity that takes actions in an environment to maximize reward is known as an
Target: agent
Prompt: The decision-making component in reinforcement learning systems is termed an
Target: actor

Policy Optimization:
Prompt: The process of improving action-selection strategies through gradient ascent is known as
Target: gradient optimization
Prompt: The technique of updating policies using importance sampling is termed
Target: importance sampling

Explainable AI:
Prompt: Methods that make AI decision-making processes interpretable are called
Target: explainable methods
Prompt: The process of revealing the reasoning behind AI outputs is called
Target: reasoning explanation

A.1 Details of Our Keyword Generation Pipeline

Here, we detail our keyword generation pipeline.

Initial preprocessing First, we standardize the abstracts by converting them to lowercase, removing bracketed content, handling apostrophes and contractions, and implementing custom tokenization to ensure consistency. Our custom tokenization preserves hyphenated words (e.g., fine-tuned) to maintain the integrity of technical terms.

Building and validating n-grams Next, we build n-grams from size 2 to 7 using Gensim (version 4.3.3) Phrases model (Rehurek and Sojka, 2010; Mikolov et al., 2013). For simplicity, we have employed default parameters for a minimum count threshold of 5 occurrences and a collocation threshold of 10 to identify meaningful word combinations. We further ensure the quality and relevance of keywords by including words that contain only alphabetic characters and hyphens (e.g., "deep-learning" is valid, "model2" is not). Finally, we filter out *(i)* stopwords (e.g., "the", "is"), *(ii)* generic academic vocabulary (e.g., "paper", "framework"), *(iii)* function words (e.g., "between", "despite"), and *(iv)* quantitative expressions (e.g., "more", "significant").

Length-based filtering of the keywords The third stage implements adaptive thresholding to select only the most informative keywords. It allows us to determine the proportion of each n-gram to keep in our keyword list. For our work, we have selected proportions of 50% bigrams, 30% trigrams, 15% four-grams, and 5% five-or-more-grams based on empirical analyses. If a domain had n-grams below the given proportion, then the ratios are adaptively adjusted based on the actual proportion of each n-gram length in the corpus. For our corpus, we targeted approximately 300 keywords from each domain.

Redundancy elimination The final stage is the semantic redundancy reduction, for which we compute the cosine similarity between all keyword pairs using the "all-MiniLM-L6-v2" (Wang et al., 2020) model from Sentence Transformer library (Reimers and Gurevych, 2019). We identify pairs whose similarity score passes 0.85 as *duplicates* and merge them into a single keyword. We maintain the shorter string name (e.g., "complex network" over "complex networks"). An example keyword list from the `computer science - artificial intelligence (CS.AI)` domain is presented in Supplementary Table 2.

A.2 Matching Keywords and Sentences: Implementation Details

We first strip the entries in RedPajama of excess whitespace and then convert these documents into individual sentences using spaCy's "en_core_web_sm" model (Honnibal et al., 2020). Subsequently, we embed all sentences as well as keywords obtained in Section 2.2 and calculate the cosine similarity between the sentence and keyword embeddings. For each keyword, we set a similarity threshold of 0.5 between the keyword and the corresponding sentence to extract only the sentences semantically related to that keyword.

Selected sentences further undergo a systematic cleaning process. First, any LaTeX-specific formatting was converted to plain text using LatexNodes2Text (Faist, 2021). We removed all citations using expression patterns targeting common citation commands (e.g., `\citet`) and converted tabs, newlines, and other whitespace to single spaces. We further validated characters against a predefined set of allowed characters (ASCII letters, digits, whitespace, and punctuation) to remove special characters. We removed any sentence that failed the character validation from further analyses. In the end, we obtain a clean mapping of sentences for each keyword in a domain of interest.

A.3 Attribute Rate Calculation

We adopt the attribute rate analysis introduced by Geva et al. (2023) to understand knowledge progression through layers in LLMs. This metric quantifies the subject-related information contained in each layer's representations by measuring the overlap between model-generated tokens and tokens semantically related to the input subject. In our work, subjects refer to the keywords we extracted in Section 2.2.

A.3.1 Token Analysis Framework

Before calculating attribute rates, we must first establish the set of tokens semantically related to each keyword. For a given keyword, we aggregate all clean sentences obtained in Section 2.3 to create a unified keyword corpus. We then tokenize these sentences using the corresponding model tokenizer and apply the following filtering process:

1. Remove all tokens associated with stopwords (using NLTK stopwords list)
2. Remove tokens with character length of less than three
3. Maintain a frequency counter for the remaining filtered tokens
4. Preserve the top 1200 most frequent tokens for each keyword (to be consistent with the average tokens employed in (Geva et al., 2023))

This process yields a set A_s containing the most relevant tokens for keyword s, which serves as our ground truth for semantic relatedness in the attribute rate calculation.

A.3.2 Layerwise Attribute Rate Methodology

For intermediate layers, the attribute rate computation involves projecting hidden representations to vocabulary space and measuring overlap with our established token sets. Following Geva et al. (2023), given a keyword s, we calculate the attribute rate at layer ℓ and position i as follows:

First, we project the hidden state to vocabulary space:

$$\text{projs} = h_i^\ell \cdot E^T \tag{1}$$

where E is the model's input embedding matrix and h_i^ℓ is the hidden state at position i and layer ℓ.

Next, we extract the top-k tokens by removing stopwords and tokens with fewer than 3 characters from the projected vocabulary distribution. Finally, we calculate:

$$\text{Attribute Rate} = \frac{|\text{Top-}k \text{ tokens} \cap A_s|}{k} \times 100\% \tag{2}$$

We focus primarily on the last token position of each keyword, as it can attend to all preceding subject tokens and thus contains the most comprehensive subject representation (Geva et al., 2023). Note that in the original implementation, for the final layer, Geva et al. (2023) used the output of the final layer normalization, which corresponds essentially to the model's actual output logits.

A.3.3 Cross-Architectural Last Layer Attribute Rate

While layerwise attribute rate provides deep insights into knowledge progression, it has primarily been employed for GPT family models (nostalgebraist, 2020; Dar et al., 2022; Geva et al., 2023) and faces challenges with newer architectures like Llama-2, Llama-3.1, and OLMo-2 due to architectural differences that make direct projection of hidden states potentially incomparable across models.

Since the original method shows that attribute rate follows a monotonic increasing pattern through the layers (Geva et al., 2023), with the final layer representing the culmination of the subject's enrichment throughout the model, we can focus exclusively on the final layer output for cross-architectural comparison. For the last layer attribute rate, we directly use the model's output token probabilities:

$$\text{LLAR} = \frac{|\text{Top-}k \text{ clean output tokens} \cap A_s|}{k} \times 100\% \tag{3}$$

where LLAR is an abbreviation for last layer attribute rate, and Top-k clean output tokens are the top-k tokens from the model's final output distribution after removing stopwords and tokens with fewer than 3 characters. We use $k = 50$ in our analysis to maintain consistency with prior work while ensuring sufficient coverage of the model's predictions.

This approach leverages the fact that the final layer projection in the original methodology essentially captures the same information as the model's direct output, making it robust across different architectures while maintaining the theoretical foundation established by Geva et al. (2023). By focusing on the last layer, we obtain a metric that enables meaningful comparisons across different model architectures while capturing the final state of factual knowledge extraction.

A.4 Sample Document from Dataset

Below we present a sample document from our arXiv and Kaggle dataset after mapping, showing both the metadata and the beginning of the full-text content:

Sample Document: arXiv ID 1607.04768

Title: Hoffmann-Ostenhof's conjecture for traceable cubic graphs
Categories: math.CO cs.DM
Abstract: In 2011 Hoffmann-Ostenhof proposed the following conjecture: Every connected cubic graph can be decomposed into a spanning tree, a matching and a family of cycles. In this paper we prove that this conjecture holds for traceable cubic graphs.

Introduction: Let G be a simple undirected graph with the *vertex set* $V(G)$ and the *edge set* $E(G)$. A vertex with degree one is called a *pendant vertex*. The distance between the vertices u and v in graph G is denoted by $d_G(u, v)$. A cycle C is called *chordless* if C has no *cycle chord* (that is an edge not in the edge set of C whose endpoints lie on the vertices of C). The *Induced subgraph* on vertex set S is denoted by $\langle S \rangle$. A path that starts in v and ends in u is denoted by $\widehat{vu}$. A *traceable* graph is a graph that possesses a Hamiltonian path. In a graph G, we say that a cycle C is *formed by the path* Q if $|E(C) \setminus E(Q)| = 1$. So every vertex of C belongs to $V(Q)$...

Our combined dataset contains 1.56 million documents with similar structures, covering the full spectrum of scientific disciplines represented in arXiv.

A.5 Baseline Metrics: Last Layer Attribute Rate and Perplexity

Recently, Geva et al. (2023) presented the *attribute rate* approach to study factual knowledge in LLMs. The attribute rate for a factual prompt quantifies the overlap between model-generated completions given a prompt and tokens related to the subject of the prompt (see Sec. A.3 for a detailed explanation). Layer-wise attribute rate has provided deep insights into knowledge extraction through the layers. We employ the attribute rate of the last layer to quantify the entire knowledge within a network.

For perplexity calculation, we employed the Semantic Scholar portion of the M2D2 dataset (Reid et al., 2022), a multi-domain corpus for studying domain adaptation in LLMs. We chose the M2D2 for model evaluation since the authors cleaned raw scientific papers for language model pretraining and previous work (Yıldız et al., 2024) already used the dataset for model comparison.

A.6 Validation of Our Pipeline by Adapting GPT2-XL to Diverse Domains

In this section, we conduct controlled domain adaptation experiments, where model training domains vary in semantic proximity to the evaluation domain (`CS.AI`). We expect models adapted to domains closely related to AI to outperform those trained on unrelated domains, thereby demonstrating that our benchmark is sensitive to domain-relevant knowledge.

Adapted domains We train the GPT2-XL model on six different domain-specific subsets of M2D2: *(i)* two subdomains of computer science (Artificial Intelligence (`CS.AI`), and Computation and Language (`CS.CL`)); *(ii)* two domains unrelated to computer science (`Culture and Humanities`, and `Material Science`); and *(iii)* a mix of these domains (`CS.AI` + `Material Science` and `CS.AI` + `Culture and Humanities`). These six adapted variants, along with the original base GPT2-XL model, form the basis of our evaluation.

Evaluation sets We compute the prediction rank of these models on the Claude-generated benchmark on AI. Likewise, we construct prompt-target pairs by running our pipeline on `CS.AI` domain of the RedPajama-Data-1T dataset, where prompts and targets are separately obtained by TF and TF-IDF. We further computed the models' last layer attribution rate using the tokens obtained from the same dataset and the perplexity computed on the test set of `CS.AI` domain in M2D2. Since all evaluation sets are restricted to the AI domain, we can intuitively order expected model performances.

Findings First, Fig. 6a confirms that models trained on domains similar to AI rank higher on the Claude benchmark, while training on unrelated domains degrades the rank. Ranks obtained on our TF- and TF-IDF-based benchmarks follow a very similar trend, evident by their strong correlation with the Claude benchmark (Fig. 6b-6c, bottom row). Perplexity (Fig. 6d) correlates less with the ranks obtained on the Claude benchmark compared to our pipeline. Finally, the last layer attribute rate in Fig. 6e shows no correlation at all. Altogether, these findings validate the efficacy of the Claude benchmark as well as our pipeline.

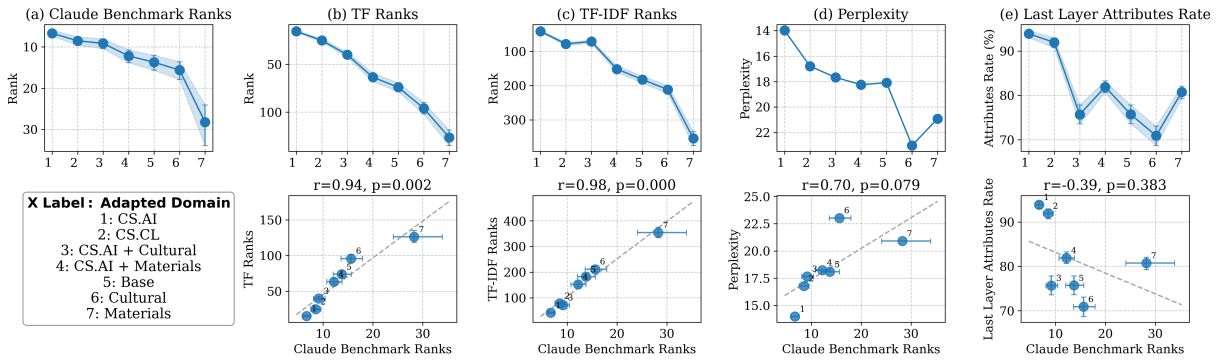

Figure 6: Evaluation of six domain-adapted GPT2-XL models and the base model. The bottom-left panel shows the adapted domains, ordered by their similarity to the target domain, `CS.AI`. Top row: Prediction ranks on *(a)* the Claude benchmark generated by Claude 4 Sonnet, and *(b–c)* TF- and TF-IDF-based benchmarks produced by our pipeline. We also report baseline metrics: *(d)* perplexity and *(e)* last-layer attribution rate. Bottom row: Correlation of each metric with the Claude benchmark. TF- and TF-IDF-based ranks align closely with the Claude benchmark, whereas perplexity and attribution rate do not reflect expected trends in domain adaptation.

A.7 Checkpoint Evolution During GPT2-XL Adaptive Pre-training

We compared our introduced benchmark and evaluation method for tracking the effects of domain adaptation for the GPT-2 XL model during training. We call the GPT-2 XL model the base model, whose metrics are considered as the baseline as shown by the black dotted horizontal line in Supplementary Figure 7, and the `CS.AI final` (red dotted line) represents the value of metrics corresponding to the final checkpoint on the model adapted to the `CS.AI` domain from the M2D2 dataset. Each model starts from either the base model or from a pre-adapted model and is trained on equal steps of data for fair comparison from the target domain of interest, because each domain has different amounts of data available in the M2D2 dataset, and each step corresponds to a batch size of 16. The start and end domain is mentioned in the figure legend. For example, "`CS.AI` to `Materials`" means the starting model was an adapted model on the `CS.AI` domain and is further trained in steps on data from the `Materials Science` domain. The results presented for the test conducted for these metrics on the `CS.AI` dataset from arXiv, except for perplexity, which was performed on the test dataset from the M2D2 domain.

We can see a consistent pattern, as expected, in all of the metrics. The domain with `CS.AI` as the target domain performs better than other adapted models, showing the dominance of the target domain over the initial domain. `CS.CL (Computer Science - Computation and Language)`, as expected because of its proximity to the `CS.AI` domain, turns out to be the second-best performer. Models whose target ends at `Materials Science` seem to perform the worst out of all because of irrelevance to the `CS.AI` domain. When the target is fixed, we can see a similar pattern for the initial domain as well; for example, out of the models that have `Materials Science` as the target domain, the model started from the base GPT2-XL model has far worse rank, perplexity, and attribute rates than the one started from the `CS.AI` domain.

We can also observe that for smaller models like GPT2-XL, the domain adaptation happens rapidly, even with a small amount of data, as can be seen by a substantial jump in each metric for 500 steps of adaptive training. This implies that for smaller models, we can achieve significant improvement without the need to spend a high computational budget for domain adaptive training.

TF and TF-IDF model ranks seem to perform very similarly to the Claude benchmark. Perplexity, although on average displaying a similar trend, shows somewhat peculiar behavior for "`CS.AI` to `Materials`" and "`Base` to `CS.CL`" cases. We expect that the model that has been trained on a related domain should perform better than one trained on an irrelevant domain, which seems to be the case with all other metrics, demonstrating the limitations of perplexity compared to our robust method of rank evaluation from our introduced pipeline.

A.8 MMLU Experiments

This section demonstrates why current benchmark evaluation methods may not be optimal for model assessment. We investigated this by conducting

Table 2: Example of 3 keywords, first 10 elements from word-list (TF and TF-IDF), and first 15 tokens from tokenizers (GPT2-XL, LLAMA-2 7B, OLMo-2 7B) for CS.AI domain.

KEYWORD	TF TARGET VOCABULARY	TF-IDF TARGET VOCABULARY	TOKEN LIST
machine learning	Support, Vector, SVM, International, Conference, Random, Machines, logistic, Supervised, Bayes	International, Conference, Forest, Logistic, Naive, forests, Forests, Proceedings, AutoML, Boosting	GPT2-XL: learning, machine, data, algorithms, Machine, models, Learning, methods, model, used, techniques, supervised, classification, problem, algorithm LLAMA-2 7B: learning, machine, data, algorithms, Machine, models, vised, model, super, Learning, methods, ing, used, based, techniques OLMO-2 7B: learning, machine, data, algorithms, Machine, models, Learning, model, methods, used, techniques, supervised, algorithm, classification, problem
reinforcement learning	Reinforcement, rewards, actor, Policy, demonstrations, imitation, games, MDP, transition, cumulative	imitation, critic, Actor, Inverse, MDPs, Reward, shaping, replay, bandit, Imitation	GPT2-XL: learning, reinforcement, policy, reward, agent, forcement, Rein, Learning, model, state, function, learn, algorithm, algorithms, problem LLAMA-2 7B: learning, cement, rein, policy, reward, agent, Rein, model, Learning, state, based, function, learn, ing, algorithm OLMO-2 7B: learning, reinforcement, policy, reward, agent, forcement, Rein, Learning, model, state, function, learn, algorithm, algorithms, problem
natural language processing	NLP, parsing, Natural, sentence, Language, sentences, Processing, syntactic, corpus, parser	parser, tagging, parsers, parse, phrases, lexical, NLG, noun, POS, discourse	GPT2-XL: language, natural, processing, tasks, semantic, text, parsing, sentence, Natural, Language, task, sentences, information, models, model LLAMA-2 7B: language, natural, ing, processing, tasks, text, semantic, parsing, Natural, sentence, task, Language, based, model, sentences OLMO-2 7B: language, natural, processing, tasks, semantic, text, parsing, sentence, Natural, Language, task, sentences, information, models, model

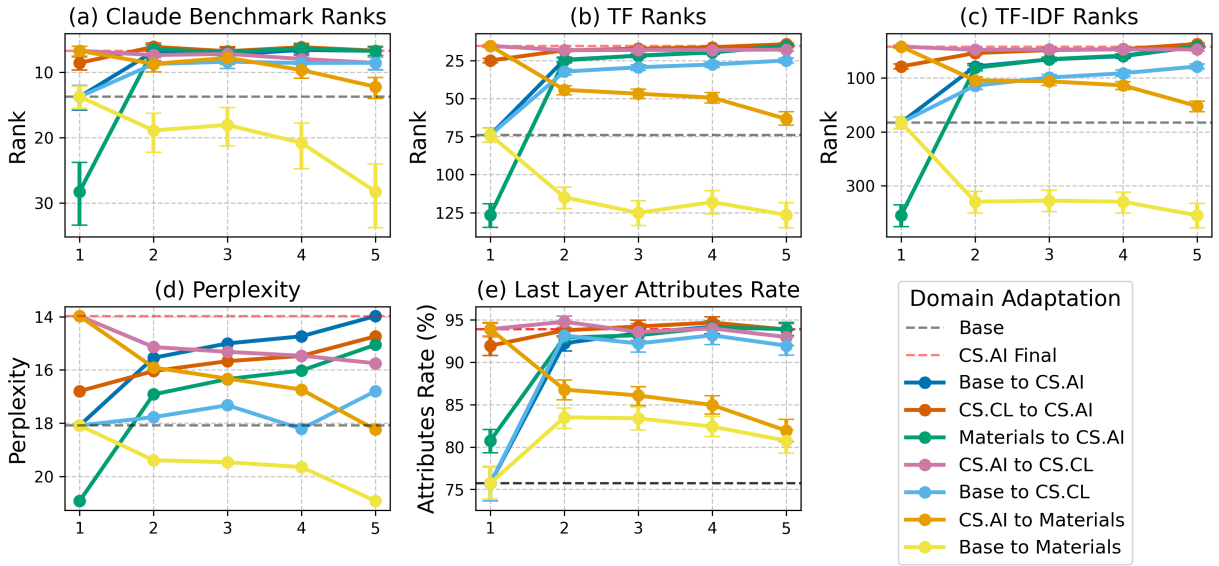

Figure 7: The x-axis shows training steps during domain adaptation, while the y-axis represents different evaluation metrics across five panels: (a) Claude Benchmark Ranks, (b) TF Ranks, (c) TF-IDF Ranks, (d) Perplexity, and (e) Last Layer Attributes Rate. Each colored line represents a different domain adaptation path, with the legend indicating the starting and ending domains. As can be seen, models ending in `CS.AI` (blue, red, and green lines) consistently achieve better ranks across all metrics, while models adapted to `Materials Science` (yellow and orange lines) show the worst performance. The TF and TF-IDF ranks closely follow the Claude benchmark trends, whereas perplexity does not give a fine grade distinction during adaptation."

experiments on the Computer Science subdomain (College and High School levels) from MMLU with a total of 200 questions of multiple-choice type (100 each). We employed three different approaches: first, we examined the most commonly used method for evaluating base models (few-shot learning) by introducing bias into the few-shot examples and observing the effects on model output. Second, we tested chat models to determine whether the bias toward correct options persists in instruction-tuned chat models. Third, we investigated potential bias in prompting strategies for chat models.

A.8.1 Few-Shot Learning Bias

The objective of this experiment was to evaluate bias in the most common approach for benchmark evaluation of base models (Brown et al., 2020; Hendrycks et al., 2020), specifically, few-shot learning. To test the stability of accuracy metrics obtained through few-shot learning, we designed a controlled experiment.

We first evaluated MMLU using the standard methodology for the two Computer Science subdomains, utilizing the few-shot examples provided in each case (5 examples per subdomain in MMLU). We then systematically biased the options in all few-shot examples toward a single option (e.g., making Option A correct for all few-shot examples) while ensuring that this option was not the actual correct answer by randomly exchanging it with other options. This procedure was repeated for all answer options (A, B, C, and D). Additionally, we included a simple shuffling condition where options for few-shot examples were randomly shuffled along with their corresponding correct answers, repeating this process three times.

As shown in Supplementary Figure 10, smaller models exhibit significant performance changes due to this bias, though the effect persists even in larger models.

A.8.2 Option Bias in Chat Models

The second most popular benchmarking method involves using chat models instead of base models when available. To investigate whether option bias exists in these models, we shuffled the answer options three times and tested scenarios where all correct answers were either all A, B, C, or D.

As demonstrated in Supplementary Figure 11, bias persists significantly in chat models, with particularly pronounced effects in models such as Llama-2 7B Chat and OLMo 2 7B Instruct.

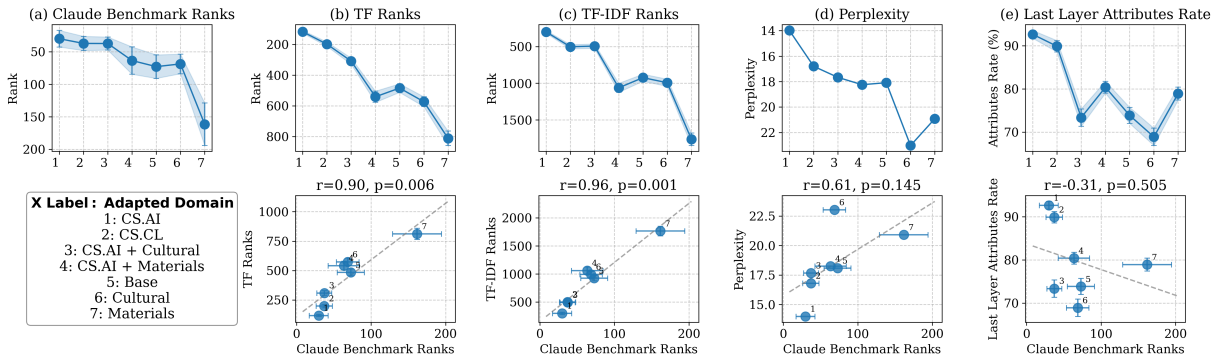

Figure 8: The x-axis shows adapted domains ordered by dissimilarity to `CS.AI`, while the y-axis represents evaluation metrics across five panels using mean aggregation. Top row: Prediction ranks on (a) the benchmark generated by Claude 4 Sonnet, and (b-c) TF- and TF-IDF-based benchmarks produced by our pipeline, along with baseline metrics (d) perplexity and (e) last-layer attribution rate. Bottom row: Correlation analysis between each metric and the Claude benchmark. As can be seen, the mean aggregation produces higher rank values compared to Figure 6 due to the inclusion of outliers, but maintains the same overall pattern where models adapted to domains similar to `CS.AI` achieve better performance than those trained on unrelated domains like `Materials`. The correlations remain significant between the Claude benchmark and our TF-based methods, validating the robustness of our pipeline even when outliers are included in the evaluation.

A.8.3 Prompting Bias in Chat Models

These experiments demonstrate that bias extends beyond answer option ordering to include prompt structure effects. We tested different prompt variations as detailed in Supplementary Table 3. Supplementary Figure 12 reveals that even identical prompts written in different formats can significantly affect model accuracy, with relatively longer prompts generally yielding diminished performance.

A.9 Smaller Language Models and Few-Shot Learning: An Example for GPT2-XL Model

To demonstrate the sensitivity to few-shot learning faced in smaller language models, we take one example for the GPT2-XL model. We tested GPT-2 XL on simple geography questions using greedy decoding (`do_sample=False`, `num_beams=1`).

When presented with a single question, GPT-2 XL demonstrated correct factual knowledge:

Prompt 1

```
Question: Which of the following
is the capital of France?
A) Berlin
B) Madrid
C) Paris
D) Rome

Answer:
```

For this prompt, GPT-2 XL correctly completed with "C", identifying Paris.

However, when presented with a sequence of questions in a few-shot format:

Prompt 2

```
Question: Which of the following
is the capital of Germany?
A) Paris
B) Berlin
C) Warsaw
D) Vienna

Answer: B

Question: Which of the following
is the capital of Italy?
A) Madrid
B) Athens
C) Rome
D) Lisbon

Answer: C

Question: Which of the following
is the capital of Spain?
A) Rome
B) Paris
C) London
D) Madrid

Answer: D
```

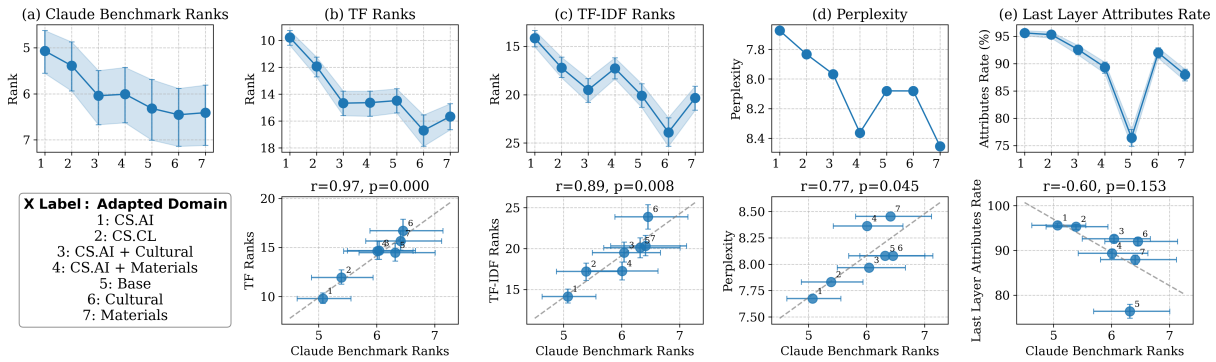

Figure 9: The x-axis shows adapted domains ordered by dissimilarity to `CS.AI`, while the y-axis represents evaluation metrics across five panels using trimmed mean aggregation for Llama2-7B models. Top row: Prediction ranks on (a) the benchmark generated by Claude 4 Sonnet, and (b-c) TF- and TF-IDF-based benchmarks produced by our pipeline, along with baseline metrics (d) perplexity and (e) last-layer attribution rate. Bottom row: Correlation analysis between each metric and the Claude benchmark. As can be seen, the larger Llama2-7B model achieves much lower absolute ranks compared to GPT2-XL, with relatively smaller changes across domain adaptations since bigger models already possess substantial knowledge and exhibit less dramatic adaptation effects. The evaluation follows the same expected pattern where models adapted to domains similar to `CS.AI` outperform those trained on unrelated domains, with strong correlations remaining significant for our TF-based methods (r=0.97, r=0.89) while perplexity and attribution rate show weaker correlations, demonstrating that our pipeline successfully translates to larger model architectures.

```
Question: Which of the following
is the capital of France?
A) Berlin
B) Madrid
C) Rome
D) Paris

Answer:
```

The model incorrectly completed with “B” (Madrid), despite being able to answer correctly when prompted without a few-shot prompt. This inconsistency highlights how smaller base models can sometimes fail to properly utilize in-context examples, reinforcing the need for our more structured approach to knowledge testing that reduces formatting sensitivity.

A.10 Cosine Similarity Between Targets and Computer Science Phrase for GPT2-XL Adaptive Training

To measure the relevance of the targets predicted better by the `CS.AI` domain-trained model as compared to the base model. We developed a token-level model comparison metric that addresses several limitations in standard evaluation approaches. First, we computed both probability and rank differences between our `CS.AI` model and baseline models for each target token. To account for scale bias in probability differences and to balance the influence of both metrics, we created a composite score by normalizing both probability and rank improvements to equal scales and giving them equal weight (50% each). Recognizing that token frequency can distort analysis by over-emphasizing rare tokens with extreme differences, we applied a logarithmic frequency weighting to this composite score. The resulting composite weighted metric emphasizes improvements on commonly occurring tokens while maintaining sensitivity to meaningful differences in rare but potentially essential tokens. We employed the all-MiniLM-L6-v2 sentence transformer for semantic analysis to embed each token and the reference domain term "Computer Science Artificial Intelligence" into a shared vector space, calculating cosine similarity to quantify domain relevance. Statistical significance between the top and bottom performance groups was assessed using independent samples t-tests with Cohen’s d for effect size quantification. Figure 14 shows a semantic alignment between `CS.AI`’s performance and specialized domain knowledge. Tokens, where `CS.AI` performed better (top 25% by composite_weighted score), showed significantly higher semantic similarity to "Computer Science Artificial Intelligence" compared to tokens with the lowest performance ($p = 2.82e-05$, Cohen’s $d = 0.31$). Domain-specific terms like " classifiers," " algorithms," and "program" clustered in the high-performance group,

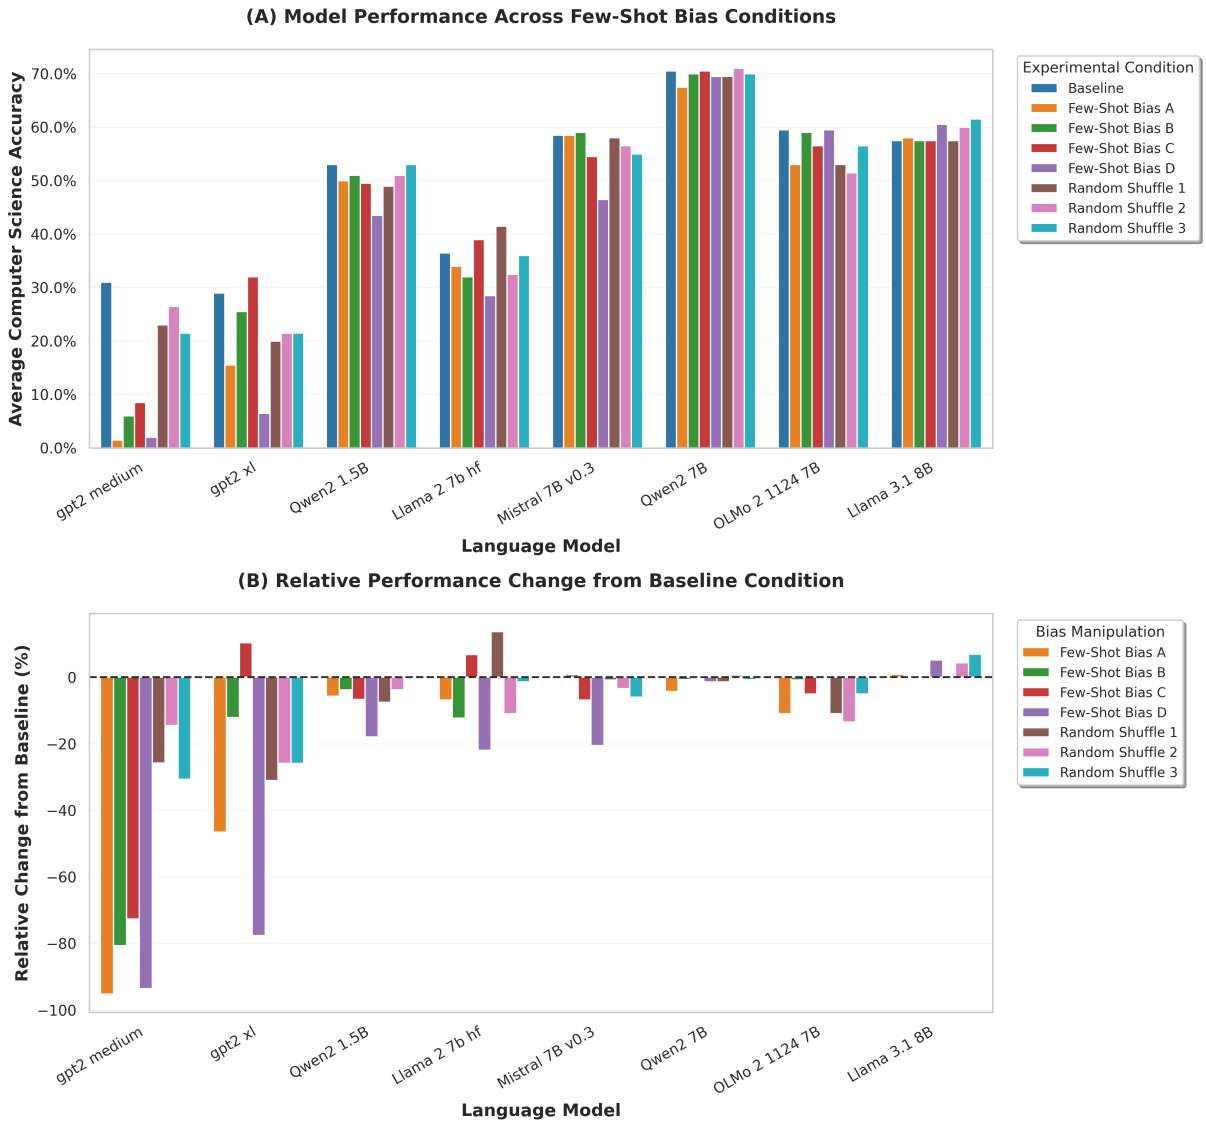

Figure 10: The x-axis shows different language models, while the y-axis shows average Computer Science accuracy (panel A) and relative performance change from baseline (panel B). Different colors represent various bias conditions: baseline, few-shot bias toward options A/B/C/D, and random shuffling variants. Smaller models exhibit more substantial performance degradation under biased conditions, though bias effects persist across all model sizes.

while semantically distant terms like " fMRI" and " boxes" appeared mainly in the low-performance group. The violin distributions reveal different central tendencies and variance patterns between groups. Importantly, token frequency (shown by point size) was distributed across both similarity ranges, confirming our weighting approach successfully prevented frequency bias from dominating the analysis. Complete lists of the top and bottom 100 tokens further support this finding (Table 5 and Figure 14), with specialized AI terminology concentrated in the high-performance group. These results provide strong evidence that `CS.AI` has developed genuine domain expertise rather than merely statistical advantages, validating results for our benchmark.

A.11 Computational Budget

Our experiments utilized both H100 and A100 GPU infrastructure with a total computational budget of approximately 1000 CPU hours and 1300 GPU hours. The Llama model adaptive and continual learning experiments were conducted on H100 GPUs, requiring approximately 800 hours of training time. The remaining GPU computations, including GPT-2 XL domain adaptation and model evaluations, were performed on A100 GPUs for the remaining approximately 400 hours. Most components of our benchmark construction pipeline, including keyword extraction, sentence matching,

and prompt-target pair generation, rely primarily on CPU processing with only minimal GPU usage and represent the majority of the 1000 CPU hours. All experiments were conducted using standard deep learning frameworks with mixed precision training where applicable to optimize computational efficiency.

A.12 Detailed Related Work

The evaluation of LLMs evolved from simple statistical measures to sophisticated benchmarking frameworks, yet challenges remain for domain-specific assessment. Initially, perplexity served as the primary evaluation metric (Jelinek et al., 1977; Gururangan et al., 2021; Yıldız et al., 2024). Recent work revealed that perplexity primarily captures general linguistic capabilities rather than domain knowledge (Xu et al., 2024; Meister and Cotterell, 2021; Öncel et al., 2024). Alternative approaches like attribute rate (Geva et al., 2023) show inconsistent patterns across model architectures.

Alternative line of research introduces MCQs benchmarks, e.g., MMLU (Hendrycks et al., 2020), GLUE (Wang et al., 2018), and BIG-bench (Srivastava et al., 2022). A series of works revealed their pitfalls that few-shot performance conflates model size with domain knowledge (Brown et al., 2020), benchmarks show sensitivity to formatting and option ordering (Alzahrani et al., 2024; Gupta et al., 2024), alignment processes introduce additional biases (Bai et al., 2022), and MCQ evaluations saturate during pretraining while perplexity continues improving (OLMo et al., 2024). Critically, existing benchmarks focus on broad, mixed-domain tasks rather than specific domains, creating challenges for selecting base models for domain adaptation (Han et al., 2023; Cui et al., 2025; Nguyen et al., 2023). The prevalence of MCQ formats over completion tasks (Chang et al., 2024) further misaligns with base model training objectives.

Recent efforts to automate benchmark generation follow two approaches. LLM-based generation (Perez et al., 2023; Yuan et al., 2025) addresses curation costs but faces limitations in long-context processing (Li et al., 2024), grounding without parametric bias (Monea et al., 2023), and computational efficiency (Chavan et al., 2024). Curated automated systems like LiveBench (White et al., 2024) and LiveCodeBench (Jain et al., 2024) mitigate contamination through recent sources or existing problem collections, but require structured data sources (news articles, programming contests) and cannot generate benchmarks from arbitrary raw domain corpora. Existing domain-specific benchmarks like MedQA (Jin et al., 2021) and LegalQA (Louis et al., 2024) organize around broad disciplinary categories rather than research subdomains, employ evaluation formats (MCQs, long-form answers) incompatible with base models, and lack mechanisms for subdomain-specific assessment or contamination-resistant updates.

Our proposed pipeline converts raw domain corpora into completion-style benchmarks, addressing these limitations through: *(i)* unified evaluation for base and chat models, *(ii)* domain-specific rather than general linguistic assessment, *(iii)* elimination of prompt sensitivity biases, and *(iv)* scalable evaluation across arbitrary domains, and *(v)* contamination-resistant evaluation through on-demand generation from recent corpora.

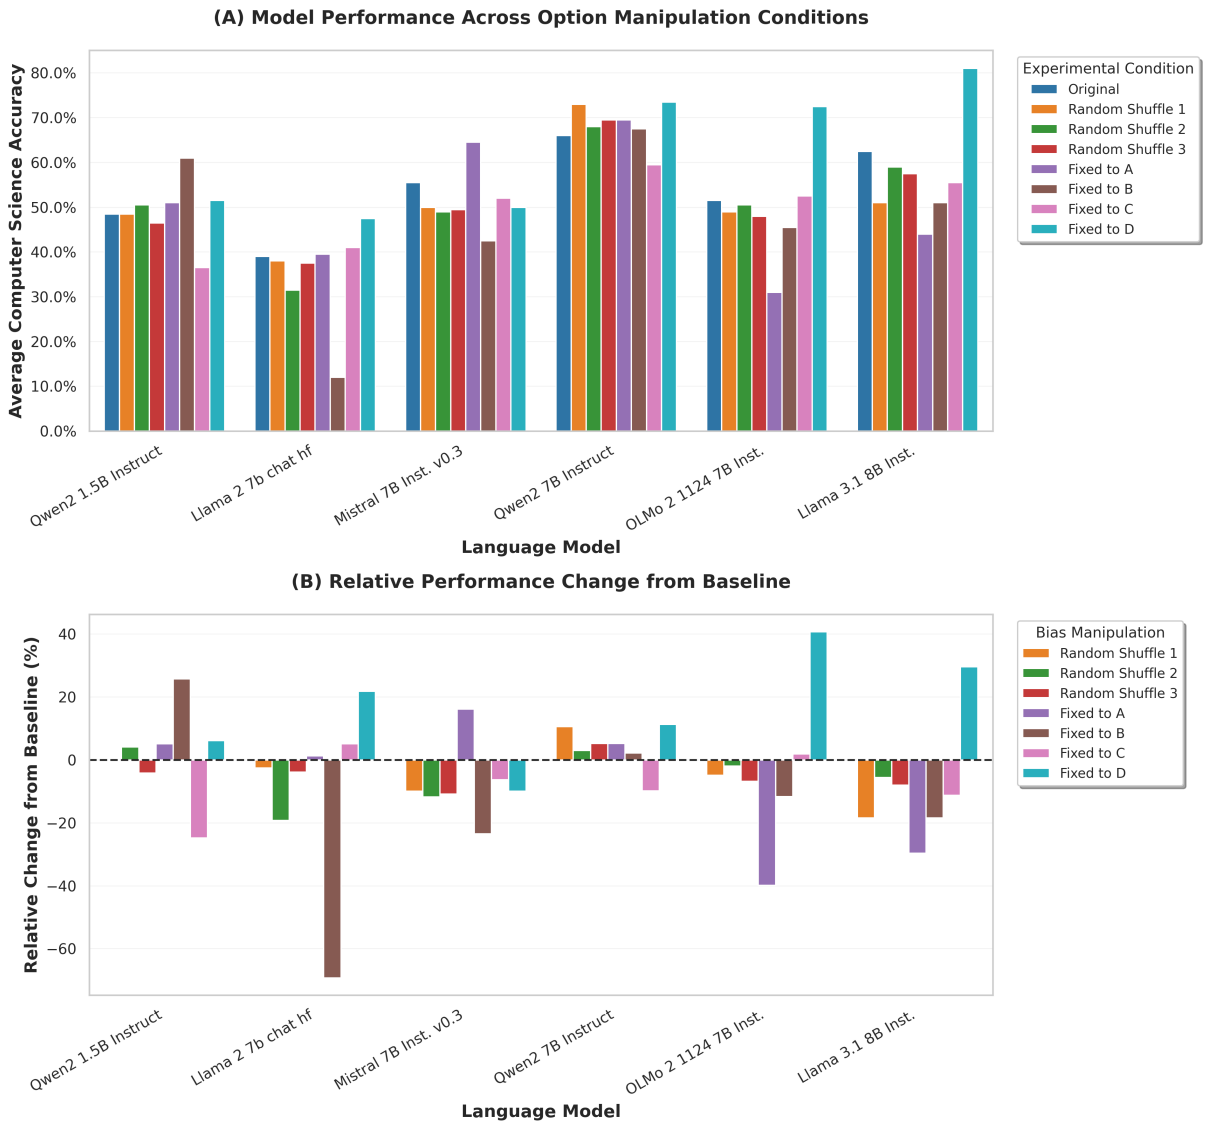

Figure 11: The x-axis shows different chat models, while the y-axis displays average Computer Science accuracy (panel A) and relative performance change from original conditions (panel B). Different colors represent various option manipulation conditions: original, random shuffling variants, and systematic bias toward options A/B/C/D. Option bias significantly affects chat model performance, with some models like Llama-2 7B Chat and OLMo 2 7B Instruct showing particularly strong susceptibility to option ordering effects.

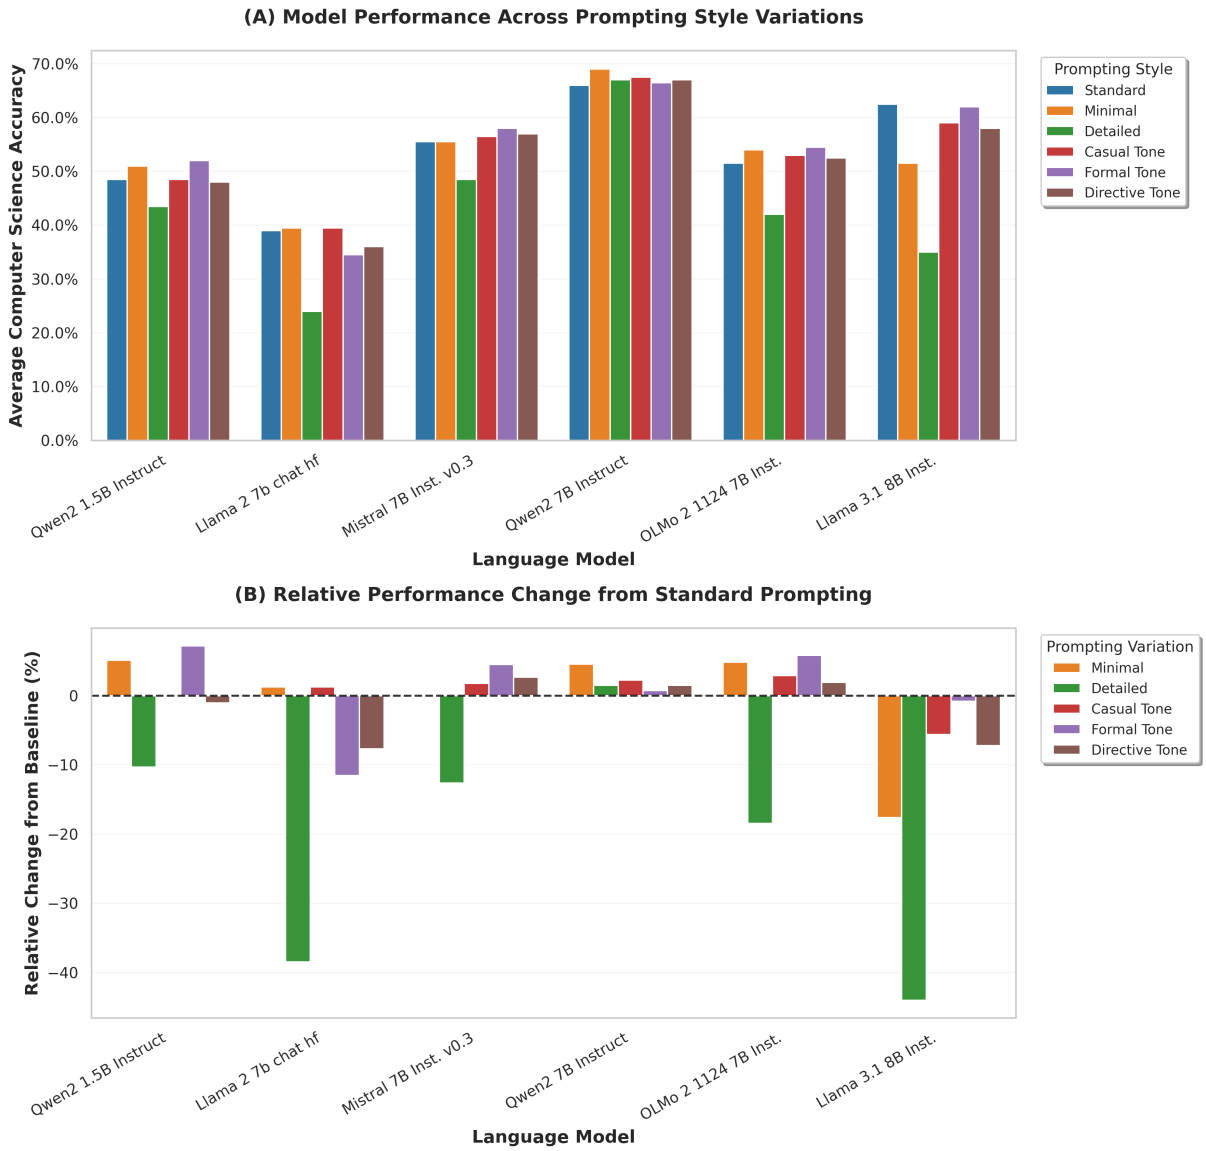

Figure 12: The x-axis shows different chat models, while the y-axis displays average Computer Science accuracy (panel A) and relative performance change from standard prompting (panel B). Different colors represent various prompting styles: standard, minimal, detailed, casual tone, formal tone, and directive tone. Prompting variations substantially impact model performance, with detailed prompts generally reducing accuracy compared to standard approaches.

Table 3: Prompt variations used to evaluate prompting bias in chat models

STYLE	PROMPT TEMPLATE
Standard	*System:* "Answer the multiple choice question by providing only the letter (A, B, C, or D) of the correct answer." *User:* "The following is a multiple choice question about {subject}.\n\n{question}"
Minimal	*System:* "Pick A, B, C, or D." *User:* "{question}"
Detailed	*System:* "You are an expert academic assistant with deep knowledge across multiple disciplines. Your role is to analyze multiple choice questions with precision and accuracy. When presented with a question, carefully examine all provided options, consider the nuances of each choice, apply relevant domain knowledge, and select the single most appropriate answer. Always respond with only the letter (A, B, C, or D) that corresponds to your chosen option. Maintain objectivity and base your decisions on factual accuracy and logical reasoning." *User:* "You are an expert in {subject}. Below is a comprehensive multiple choice question that requires careful analysis and reasoning. Please read the question thoroughly and consider all the provided options before making your decision.\n\nInstructions:\n- Analyze each option carefully\n- Consider the context and specific details mentioned in the question\n- Use your knowledge of {subject} to evaluate the correctness of each choice\n- Select the single best answer that most accurately addresses the question\n- Provide only the letter (A, B, C, or D) corresponding to your chosen answer\n\nQuestion and Options:\n{question}\n\nRemember to base your answer solely on your understanding of {subject} and the information provided in the question. Choose the most appropriate option and respond with only the letter of your choice."
Casual Tone	*System:* "Answer multiple choice questions by selecting the best option. Reply with just the letter." *User:* "Here's a multiple choice question about {subject}. Please select the best answer.\n\n{question}"
Formal Tone	*System:* "Please choose the most suitable answer from the given choices and respond with the letter only." *User:* "Please answer the following multiple choice question relating to {subject}. Select the most appropriate response.\n\n{question}"
Directive Tone	*System:* "Solve the multiple choice question by selecting the correct option. Provide only the letter." *User:* "Solve this {subject} multiple choice question by choosing the correct answer.\n\n{question}"

Table 4: Prompt and corresponding target for given keywords from `Computer Science - Artificial Intelligence (CS.AI)` and `Quantitative Biology - Populations and Evolution (Q-Bio.PE)` domains.

KEYWORD	PROMPT	TARGET
Domain: Computer Science - Artificial Intelligence (`CS.AI`)		
Machine Learning	Since machine learning solutions involve using historical data to create future forecasts,	biases
Machine Learning	To achieve optimal results, the experiments required modifying the learning rate,	optimizer
Machine Learning	Machine Learning is increasingly becoming a huge part of our lives by	automating
Reinforcement Learning	Monte-Carlo Tree Search for Policy Optimization :	Gradient
Reinforcement Learning	Reinforcement learning especially deep reinforcement learning (DRL) has been successfully applied in various fields including	Alpha
Reinforcement Learning	Deep reinforcement learning (deep RL) holds the promise of	automating
Natural Language Processing	Given a sentence with an entity pair (e_1,e_2), this task aims to identify the relation between e_1 and e_2. Typically, existing methods follow the supervised learning paradigm, and they require extensive	annotations
Natural Language Processing	The first stage focused on a linguistic analysis of the natural language text itself, i.e., a lexico-syntactic analysis without any presuppositions of	ontology
Domain: Quantitative Biology - Populations and Evolution (Q-Bio.PE)		
Phylogenetic Tree	An Arrow-type result for inferring a species tree from	gene
Phylogenetic Tree	Computing the probability of gene trees concordant with the species	tree
Phylogenetic Tree	Finally, we used SISRS to identify phylogenetic markers that we used to estimate the	phylogeny
Mutation Rate	In the absence of strong purifying selection, spontaneous and frequently deleterious mutations tend to accumulate in the human	genome
Mutation Rate	High mutation rates have been a successful strategy in their case, allowing the perpetual exploration of new	genomic
Mutation Rate	Using these results, we compare the probability that a new mutation lands and establishes within a	genomic
Disease Transmission	To understand the transmission dynamics, we analyzed the basic reproduction number R_0, which is defined as the number of secondary cases produced by a single	infect
Disease Transmission	The primary transmission pathway of COVID-19 is human-to-human propagation through	infectious
Disease Transmission	Since COVID-19 is known to be transmitted through asymptomatic and presymptomatic individuals (which may be a	viral

Benchmark Creation Template

```
# Domain-Specific Benchmark Creation Instructions
Generate exactly 1000 prompt-target pairs for **{domain name}** domain:
50 prompts per category across 20 categories.

## Categories
{list of 20 categories}

## Key Rules
1. **Prompt Structure**: Must end with complete phrases including any
   needed articles/stopwords:
   - "is called a", "is known as an", "is termed the", "are called",
     "refers to a", etc.
   - 15-40 words total, easy-to-medium difficulty
2. **Target Structure**:
   - Starts with space: `" target_term"`
   - NO stopwords (a, an, the, etc.) in target
   - Must belong to the category (not necessarily be the category name)
3. **Quality**: Unambiguous, diverse structures, comprehensive coverage
   per category

## Examples
**Correct**:

qa_pairs = [
# Category 1: {category_1_name}
{"prompt": "The entity that takes actions in an environment to maximize
 reward is known as an", "target": " agent"}
{"prompt": "Software that learns optimal behavior through interaction
 is called a", "target": " learner"}
...]
qa_pairs.extend([
# Category 2: {category_2_name}
{"prompt": "Methods that optimize policies without environment models
 are termed", "target": " model-free algorithms"}
....])
```

Figure 13: **Template:** Complete prompt template provided to Claude 4 Sonnet for domain-specific benchmark creation.

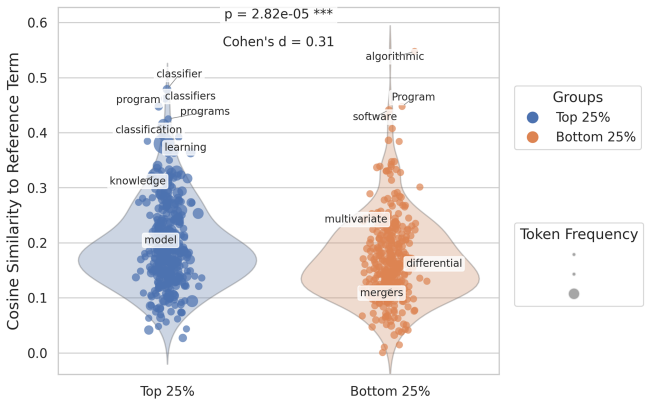

Figure 14: The x-axis shows two groups (Top 25% and Bottom 25% token performance), while the y-axis represents cosine similarity to the reference term "Computer Science Artificial Intelligence." The violin plot displays the distribution of semantic similarities for tokens where the `CS.AI` model performed best versus worst compared to the base GPT-2 XL model. As can be seen, tokens in the top 25% group show significantly higher semantic similarity to the CS.AI domain, with domain-specific terms like "classifiers," "model," and "program" appearing in the high-similarity region, while semantically distant terms like "mergers" and "differential" cluster in the low-similarity region. Point sizes indicate token frequency, demonstrating that the frequency weighting successfully prevented bias from dominating the analysis.

Table 5: Tokens with highest and lowest performance improvements in the `CS.AI` domain-trained model

Top 100 Tokens	**Bottom 100 Tokens**
Policy, Answer, tracing, margin, bases, policy, actions, baselines, answer, answering, Critic, action, planner, reward, Error, prediction, agent, explainability, MDP, generalization, weakly, critic, graphs, proposed, contrastive, efficiency, paper, Interpretable, Diversity, exploration, partially, approximate, Recall, Logic, learnt, Monte, Support, Atari, labeled, continual, Search, Partially, training, steps, segmentation, Inverse, Mixed, NMT, distillation, detection, Markov, reinforcement, best, interacting, states, classifiers, adaptation, loss, segment, UAV, Context, homogeneous, Explainable, policies, federated, demonstration, discriminator, language, Freebase, commonsense, unsupervised, control, rewards, PSO, diarization, tabular, differentiable, Table, privacy, entities, made, regularization, intersection, complexity, agents, publicly, mentions, natural, matching, spaces, relations, catastrophic, research, intrusion, program, SNNs, systems, DDPG, future, confident	parsing, one, like, discuss, Multivariate, Tree, destination, understanding, autoencoders, sensor, outliers, squared, offensive, structures, percent, languages, MRI, drones, science, score, video, fields, explainers, public, ODE, credit, SGD, may, decoder, vertices, logistic, contains, transitions, problem, Generation, backpropagate, drive, output, attacks, world, allocating, trials, sources, magnetic, size, implementation, field, signals, behavior, analysis, protect, cards, nonlinear, Absolute, Science, way, PDEs, strategy, land, emotional, percentage, French, improved, dimensionality, classical, observation, hyperparameters, Facebook, percents, private, example, theoretical, syntactic, domain, Code, bound, analogy, MCMC, explainer, truth, multivariate, layer, mergers, PDE, protection, fMRI, multivariable, imaging, media, direction, allocate, MRIs, Anomaly, government, boxes, ratio, equations, FMRI, chart, differential

Note: Tokens are listed in descending order of performance improvement as measured by the composite weighted metric. The top tokens show significantly higher semantic similarity to the `CS.AI` domain ($p = 2.82e\text{-}05$, Cohen's $d = 0.31$).

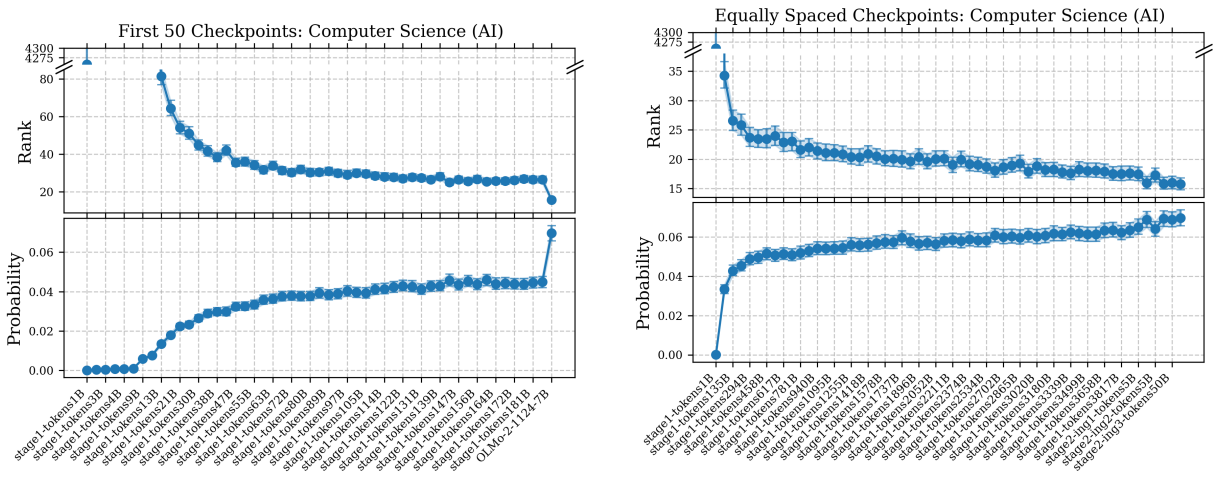

Figure 15: Prediction ranks and probabilities of OLMo-2 pretraining checkpoints on the `Computer Science (AI)` domain. The left panel displays results across the first 50 checkpoints, while the right panel shows equally spaced checkpoints from the entire pretraining, with the final data point representing the completed model. The x-axis shows checkpoint progression during training, while the y-axis represents ranks (top) and probabilities (bottom). As can be seen, the patterns mirror those observed in Fig. 3 for the `Physics (Soc-Ph)` domain, but with relatively lower ranks overall since the model demonstrates stronger baseline knowledge in the `CS.AI` domain compared to physics, reflecting the abundance of computer science training data in the pretraining corpus.

Table 6: OLMo-2 checkpoint models used in our analysis

Category	Checkpoint Details	Token Range
First 50	**OLMo-2-1124-7B_checkpoint-stage1-step** 150(1B), 600(3B), 700(3B), 850(4B), 900(4B), 1000(5B), 2000(9B), 2150(10B), 3000(13B), 4000(17B), 5000(21B), 6000(26B), 7000(30B), 8000(34B), 9000(38B), 10000(42B), 11000(47B), 12000(51B), 13000(55B), 14000(59B), 15000(63B), 16000(68B), 17000(72B), 18000(76B), 19000(80B), 20000(84B), 21000(89B), 22000(93B), 23000(97B), 24000(101B), 25000(105B), 26000(110B), 27000(114B), 28000(118B), 29000(122B), 30000(126B), 31000(131B), 32000(135B), 33000(139B), 34000(143B), 35000(147B), 36000(151B), 37000(156B), 38000(160B), 39000(164B), 40000(168B), 41000(172B), 42000(177B), 43000(181B), 44000(185B)	1B - 185B
Equally Spaced	**OLMo-2-1124-7B_checkpoint-stage1-step** 150(1B), 13000(55B), 32000(135B), 51000(214B), 70000(294B), 89000(374B), 109000(458B), 128000(537B), 147000(617B), 166000(697B), 186000(781B), 205000(860B), 224000(940B), 242000(1016B), 261000(1095B), 280000(1175B), 299000(1255B), 319000(1338B), 338000(1418B), 357000(1498B), 376000(1578B), 395000(1657B), 414000(1737B), 433000(1817B), 452000(1896B), 470000(1972B), 489000(2052B), 508000(2131B), 527000(2211B), 547000(2295B), 566000(2374B), 585000(2454B), 604000(2534B), 624000(2618B), 644000(2702B), 663000(2781B), 683000(2865B), 701000(2941B), 720000(3020B), 739000(3100B), 758000(3180B), 777000(3259B), 796000(3339B), 815000(3419B), 834000(3499B), 853000(3578B), 872000(3658B), 891000(3738B), 910000(3817B), 928646(3896B) + **Stage2-ingredient** checkpoints: ingredient1-step1000(5B), ingredient1-step11931(50B), ingredient2-step1000(5B), ingredient2-step11931(50B), ingredient3-step11931(50B)	1B - 3896B
Final Model	OLMo-2-1124-7B	Complete

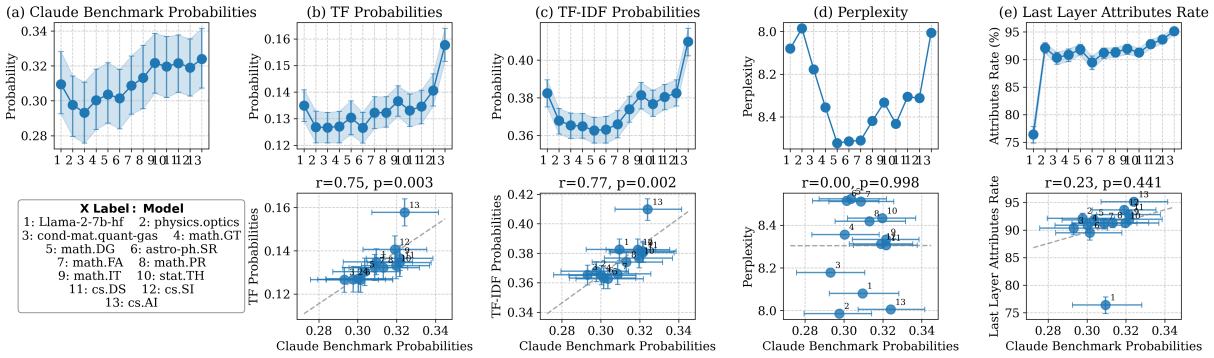

Figure 16: Probability values for 12 checkpoints from continual pretraining of Llama2-7B, along with the base model. The bottom-left panel displays the sequence of pretraining domains in the order they were introduced during training. Top row: Target token probabilities on *(a)* the benchmark generated by Claude 4 Sonnet, and *(b–c)* TF- and TF-IDF-based benchmarks produced by our pipeline. The last two columns show the baseline metrics *(d)* perplexity and *(e)* last-layer attribution rate are the same as the main Fig. 4 but kept for comparison purposes. Bottom row: Correlation of each metric with the Claude benchmark. Compared to the rank-based analysis in Fig. 4, probability values provide a normalized view of knowledge acquisition dynamics, as it is limited between 0 and 1, revealing subtle changes in model performance that may be obscured when using rank aggregations. The probability trends similarly demonstrate that our TF- and TF-IDF-based metrics capture meaningful knowledge changes aligned with domain similarity to `CS.AI`, while traditional metrics like perplexity fail to provide reliable signals for continual pretraining progress.

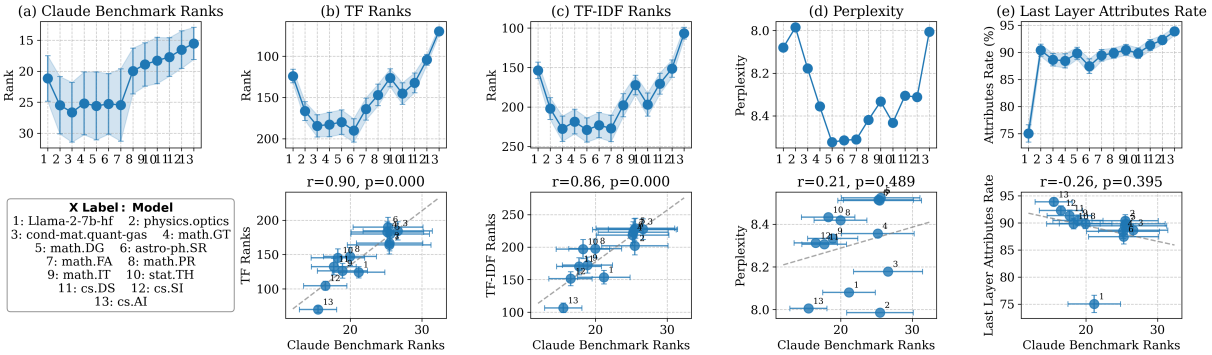

Figure 17: Mean rank aggregation for 12 checkpoints from continual pretraining of Llama2-7B, corresponding to Fig. 4. Unlike the trimmed mean (20%) aggregation, this analysis includes potential outliers, resulting in larger y-axis values but smoother trends. The mean aggregation reveals even stronger correlations between our TF- and TF-IDF-based metrics and the Claude benchmark (panels b-c), demonstrating that the inclusion of all data points provides more nuanced insights into knowledge acquisition dynamics during continual pretraining.

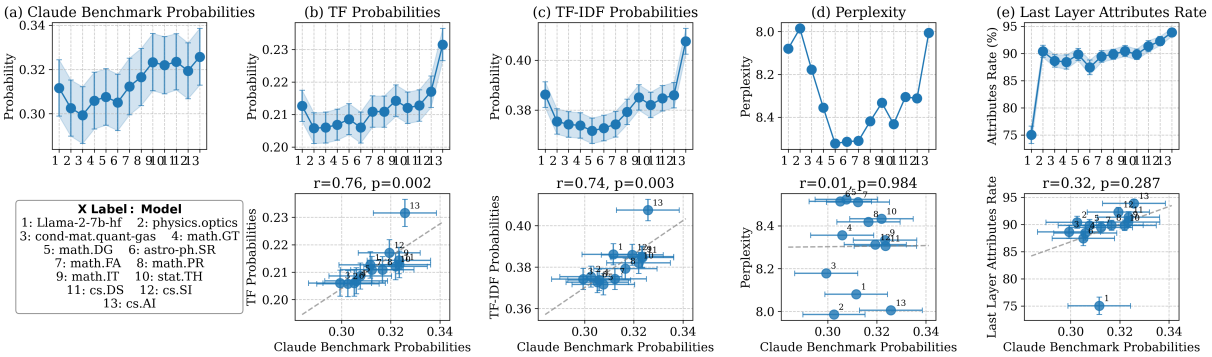

Figure 18: Mean probability aggregation for 12 checkpoints from continual pretraining of Llama2-7B, corresponding to the probability analysis in Fig. 16. Similar to the mean rank analysis, this aggregation includes all data points without trimming outliers, maintaining the same trends observed in the trimmed mean analysis while providing a complete view of the probability distributions across all generated prompts.

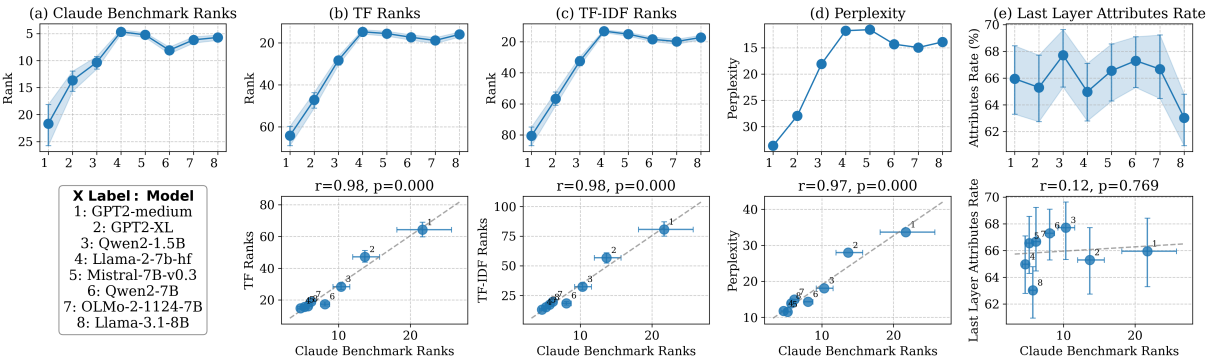

Figure 19: Evaluation of 8 base models on `General Economics (Econ-GN)` domain, including GPT-2 variants to demonstrate pattern consistency across model scales. The top row shows prediction ranks from (a) Claude benchmark, (b) TF-based method, (c) TF-IDF-based method, (d) perplexity, and (e) last layer attribution rates. The bottom row shows correlations between Claude benchmark ranks and each metric, demonstrating strong correlations for TF- and TF-IDF-based methods while perplexity shows high correlation and attribution rates show weak correlation.

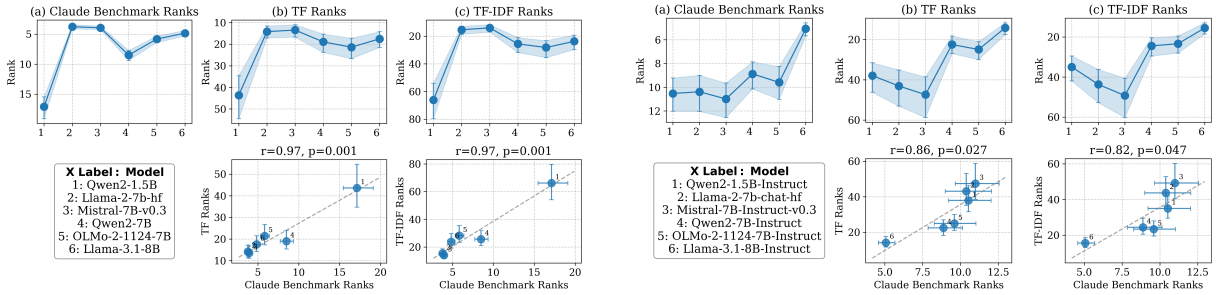

(a) Base models (b) Chat models

Figure 20: Evaluation of 6 base models (left) and 6 chat models (right) on `Quantitative Biology - Populations and Evolution (Q-Bio.PE)` domain. The top row shows prediction ranks and the bottom row shows the correlations between the Claude benchmark and our TF- and TF-IDF-based benchmarks.

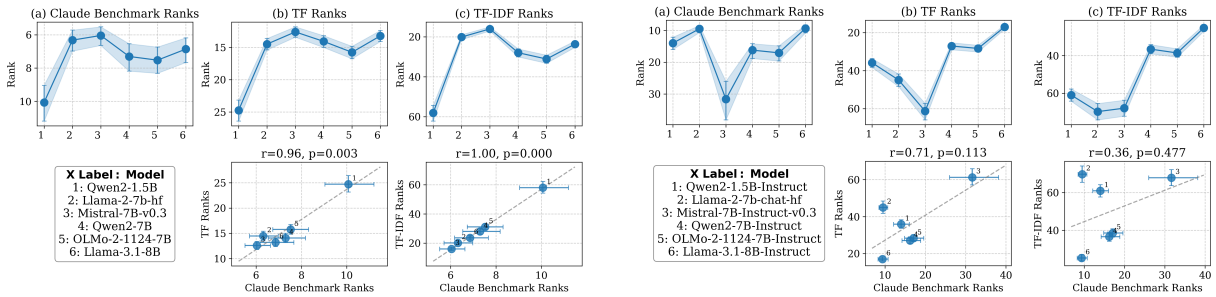

(a) Base models (b) Chat models

Figure 21: Evaluation of 6 base models (left) and 6 chat models (right) on `CS.AI` domain. The top row shows prediction ranks and the bottom row shows the correlations between the Claude benchmark and our TF- and TF-IDF-based benchmarks.

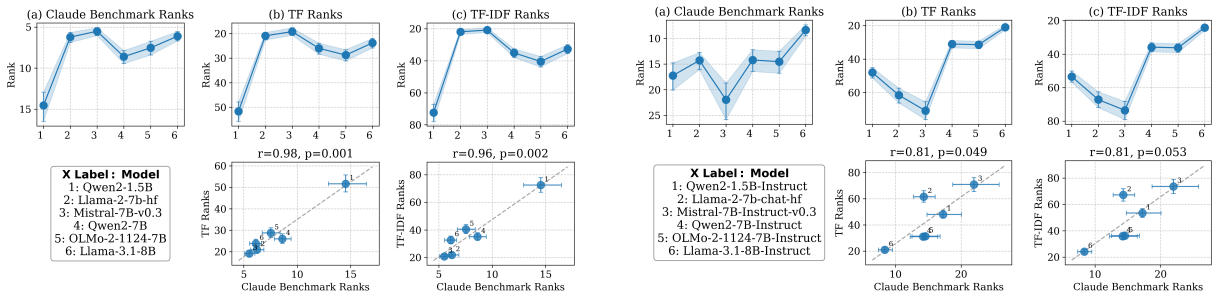

(a) Base models (b) Chat models

Figure 22: Evaluation of 6 base models (left) and 6 chat models (right) on `Physics (Soc-Ph)` domain. The top row shows prediction ranks and the bottom row shows the correlations between the Claude benchmark and our TF- and TF-IDF-based benchmarks.

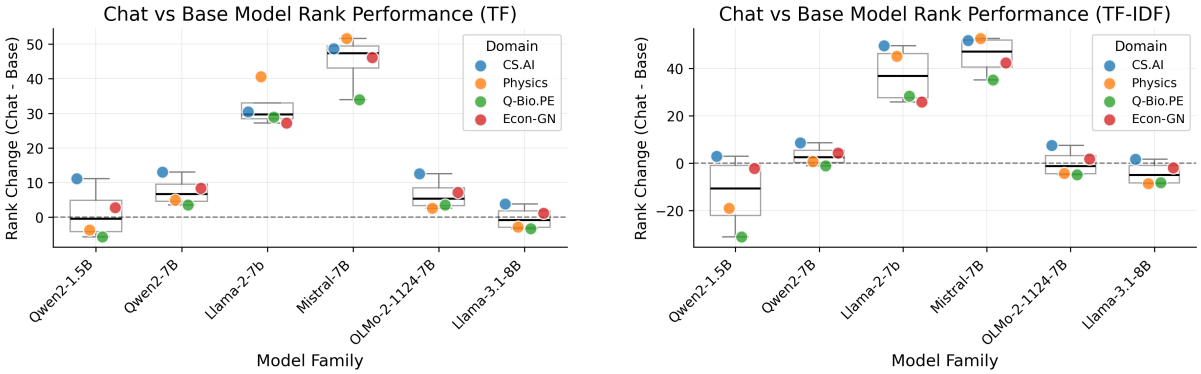

Figure 23: Impact of instruction tuning on model performance across domains, showing the difference between trimmed mean ranks (Chat - Base) for TF-based (left) and TF-IDF-based (right) evaluation methods. Positive values indicate that chat models perform worse than their base counterparts, while negative values indicate improvement. Both evaluation methods show consistent patterns: for most model families, base models generally outperform their chat-aligned counterparts across domains. The degradation is particularly pronounced for Llama-2-7B and Mistral-7B-v0.3, which show substantial performance drops after instruction tuning across `CS.AI`, `Physics (Soc-Ph)`, `Q-Bio.PE`, and `Econ-GN` domains, suggesting significant room for improvement in their alignment pipelines.