Technical Appendices of AC-ODM

Anonymous Author(s)
Affiliation
Address
email

A Evolution of Domain Weights During Training

Figure 1a–1d illustrate the evolution of domain weights across 22 distinct domains in The Pile dataset during training 1B Pythia model under AC-ODM algorithm. AC-ODM initializes from the original domain weights of The Pile and undergoes dynamic updates during the warmup phase. After approximately 15,000 training steps, the domain weights stabilize. Afterward, minor fluctuations are observed, which correspond to the evolving state of the LLM. The adaptive nature of AC-ODM's domain weight generation during this critical phase allows it to better align with the evolving model state, thereby facilitating faster reductions in both training loss and perplexity compared to prior methods.

Both AC-ODM and ODM algorithms eventually converge to stable domain weights. However, AC-ODM exhibits more substantial adjustments in domain weights during the first third of training, while ODM Albalak et al. [2023] stabilizes after only the first fifth of the total training steps. Notably, even after reaching stability, AC-ODM continues to experience slight fluctuations in domain weights, enabling dynamic adaptation to evolving LLM state. In contrast, domain weights in ODM remain nearly constant in the later stages of training, indicating a lack of flexibility in response to parameter updates in later stage.

A comparison of domains with the large magnitudes of increases or decreases in weights across Figure 1a–1d reveals consistent patterns. Regardless of the token proportion, domains characterized by high-quality and general-purpose texts tend to experience weight increases during training. Examples include HackerNews in Figure 1a, Gutenberg (PG-19) and BookCorpus2 in Figure 1b, StackExchange and USPTO Backgrounds in Figure 1c, and Book3 in Figure 1d. In contrast, domains containing noisier texts or highly domain-specific contents exhibit significant weight reductions, such as Enron Emails in Figure 1a, DM Mathematics and Wikipedia (en) in Figure 1b, Github and FreeLaw in Figure 1c, and PubMed Central in Figure 1d. These observations align with human intuitive expectations: during LLM pretraining, data domains rich in high-quality, generalizable content are more effective at driving model convergence in the early stages of training.

B Analysis of Results of MMLU Tasks

Table 1: Zero-shot accuracy of AC-ODMs among different groups in MMLU.

Algorithm	STEM	Social Sciences	Humanities	Other	Average
AC-ODM	0.24213	0.30433	0.25381	0.24626	0.25146
AC-ODM-410M	0.28219	0.38231	0.29908	0.28924	0.29980

We evaluate the performance of AC-ODM across four domain-specific groups in the MMLU benchmark, along with the overall average accuracy. As shown in Table 1, AC-ODM achieves better accuracy in the *Social Sciences* group, achieving approximately 21% higher than average. This indicates that AC-ODM effectively adapts to domain shifts in this group, likely benefiting from the

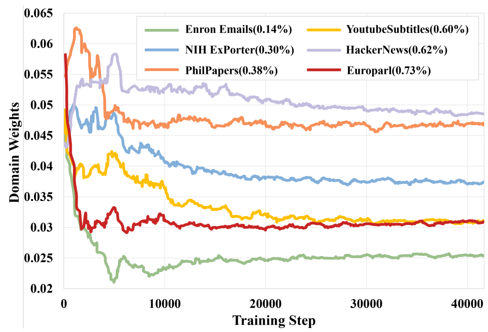

(a)

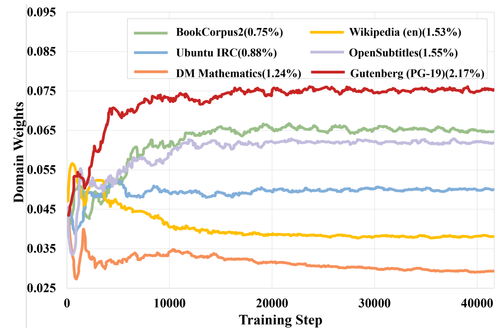

(b)

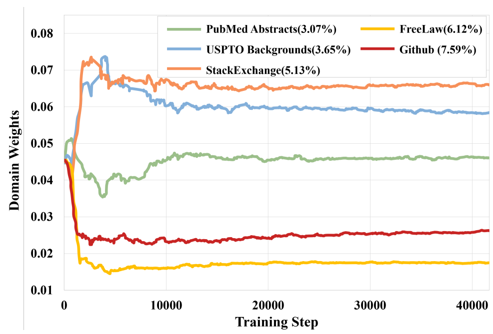

(c)

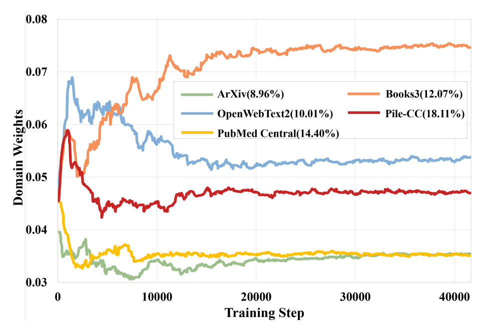

(d)

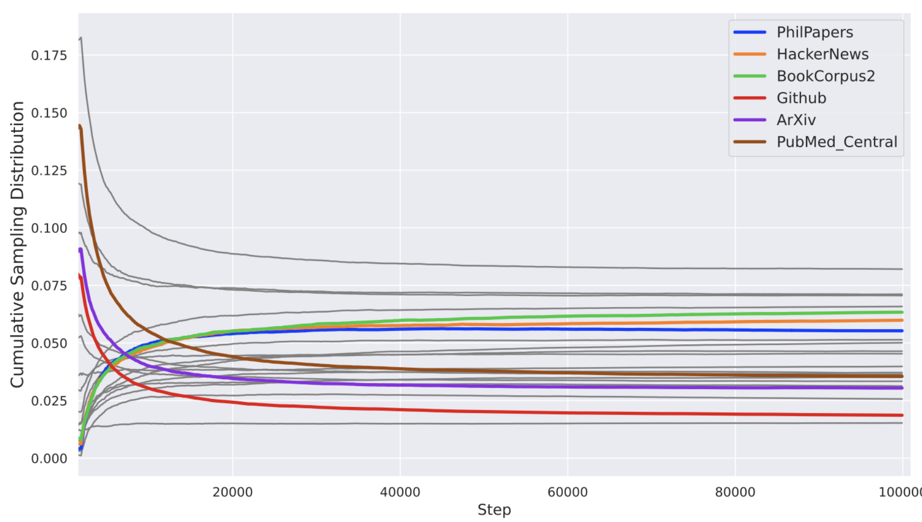

(e)

Figure 1: Evolution of domain weights during training. The legend indicates the proportion of tokens of each domain (in percentage). (a) Six domains with smallest token proportions; (b) Six domains with token proportions below 3%; (c) Five domains with token proportions below 8%; (d) Five domains with highest token proportions; (e) The cumulative sampling distribution of ODM Albalak et al. [2023].

alignment between Social Sciences content and the training distribution in The Pile. In contrast, AC-ODM underperforms in the *STEM* and *Other* groups, where accuracy falls slightly below the overall average. The *Humanities* group yields performance close to the average. These observations suggest that AC-ODM facilitates the LLM's ability to better acquire and generalize semantic patterns related to humanities and social science domains from The Pile.

Compared to the direct application of AC-ODM, the proxy-based AC-ODM-410M variant consistently improves performance across all groups, yielding an overall 19% increase in average accuracy. The most notable gains occur in the *Social Sciences* and *Other* groups, with improvements of 26% and 17%, respectively. These results indicate that AC-ODM trained on a 410M-parameter proxy model can effectively capture the underlying domain relationships present in The Pile, which are transferable to larger models and particularly beneficial for tasks involving humanities, social sciences, and general knowledge. However, the relatively limited gains in STEM-related domains also suggest that AC-ODM pays less attention to exploring domain-specific features relevant to science and engineering. This limitation may stem from the relatively low proportion of STEM-related content in The Pile dataset itself, which we would like to investigate in the future.

Figure 2 illustrates the task-level accuracy of AC-ODMs across different groups within the MMLU benchmark. In the *STEM* group, AC-ODMs achieve strong performance on tasks such as *Electrical Engineering* and *Computer Security*. Within the *Social Sciences* group, notable improvements are observed in *US Foreign Policy*, *Professional Psychology*, *High School Psychology*, and *Econometrics*. For the *Humanities* group, AC-ODMs perform well on *World Religions*, *Logical Fallacies*, and *Jurisprudence*. In the *Other* group, tasks such as *Marketing*, *Human Aging*, *College Medicine*, and *Clinical Knowledge* benefit significantly from AC-ODMs. These results suggest that AC-ODM's domain weight optimization strategy effectively guides the LLMs to acquire semantic information associated with general-purpose knowledge domains.

Compared to AC-ODM, the proxy-based AC-ODM-410M consistently improves performance across all tasks. Notably, for particularly challenging tasks such as *High School Statistics*, *Elementary Mathematics*, and *Management*, AC-ODM-410M achieves non-zero accuracy where AC-ODM fails completely (0% accuracy). These findings highlight that the use of a well-trained proxy model during training enables AC-ODM to capture meaningful domain relationships, ultimately enhancing LLM performance. Proxy-based training allows the model to better infer the latent structure of domain-specific knowledge while fulfilling difficult tasks, thereby leading to more effective adaptation and improved generalization.

References

Alon Albalak, Liangming Pan, Colin Raffel, and William Yang Wang. Efficient online data mixing for language model pre-training, 2023. URL `https://arxiv.org/abs/2312.02406`.

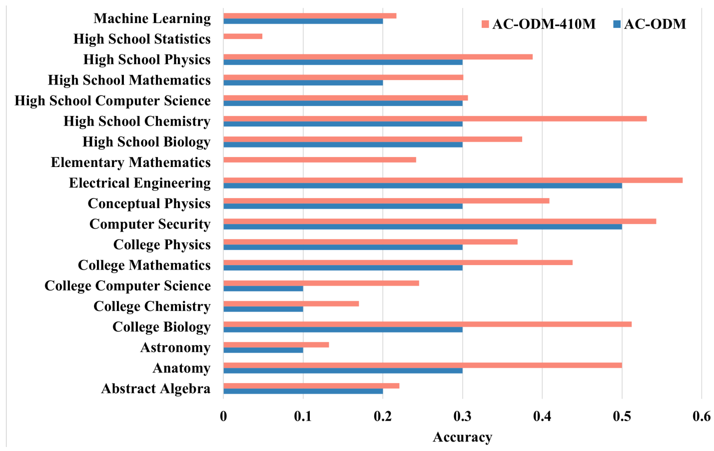

(a)

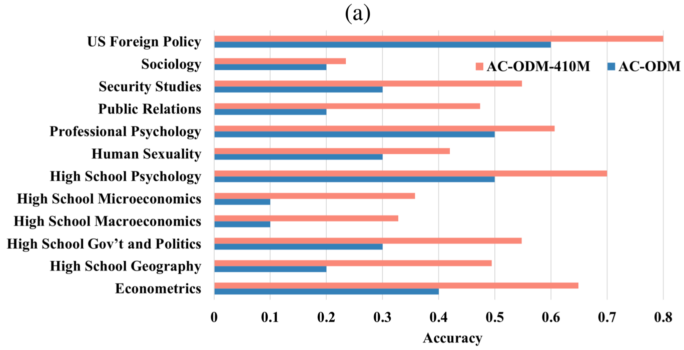

(b)

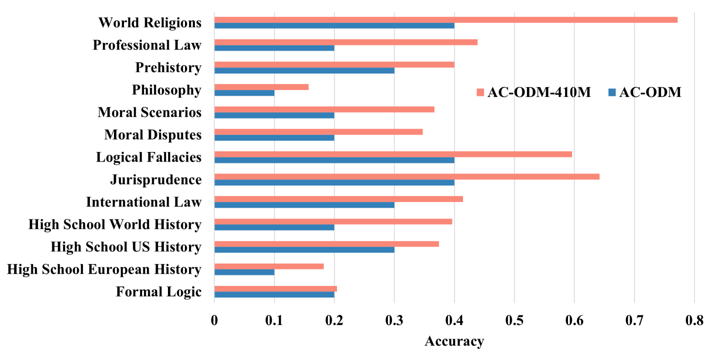

(c)

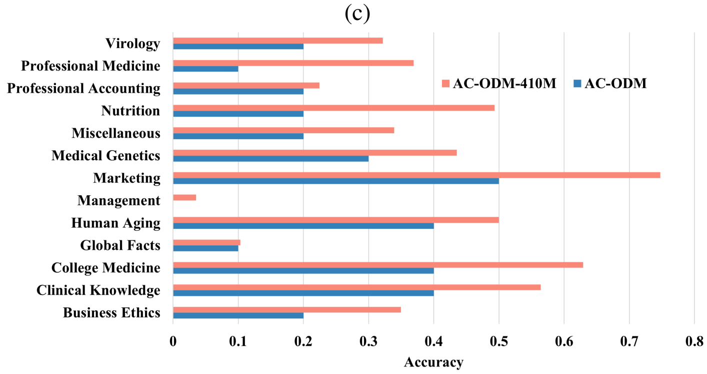

(d)

Figure 2: Zero-shot accuracy of AC-ODMs across MMLU tasks, grouped by subject category. (a) STEM; (b) Social Sciences; (c) Humanities; (d) Other.