

Robust Uncertainty Quantification for Self-Evolving Large Language Models via Continual Domain Pretraining

Xiaofan Zhou, Lu Cheng
University of Illinois at Chicago
Correspondence: xzhou77@uic.edu

Abstract

Continual Learning (CL) is essential for enabling self-evolving large language models (LLMs) to adapt and remain effective amid rapid knowledge growth. Yet, despite its importance, little attention has been given to establishing statistical reliability guarantees for LLMs under CL, particularly in the setting of continual domain pretraining (CDP). Conformal Prediction (CP) has shown promise in offering correctness guarantees for LLMs, but it faces major challenges in CDP: testing data often stems from unknown or shifting domain distributions, under which CP may no longer provide valid guarantees. Moreover, when high coverage is required, CP can yield excessively large prediction sets for unanswerable queries, reducing informativeness. To address these challenges, we introduce an adaptive rejection and non-exchangeable CP framework. Our method first estimates the distribution of questions across domains in the test set using transformer-based clustering, then reweights or resamples the calibration data accordingly. Building on this, adaptive rejection CP allows the LLM to selectively abstain from answering when its confidence or competence shifts significantly. Extensive experiments demonstrate that our framework enhances both the effectiveness and reliability of CP under CDP scenarios. Our code is available at: <https://github.com/AlearZhou/CPCL>

1 Introduction

LLMs have demonstrated impressive performance across diverse domains, from personal assistants (Achiam et al., 2023) and robotics (Wang et al., 2025) to biology (Jung et al., 2024) and scientific discovery (Rane et al., 2023). Yet, in the face of rapid knowledge growth, new and private information emerges at a pace that quickly renders their static, pre-trained knowledge outdated. This dynamic landscape underscores the need for self-evolving LLMs that can continuously adapt

to novel environments. CL has therefore become a critical capability, allowing LLMs to incorporate new knowledge while preserving previously acquired skills (Shi et al., 2024; Gururangan et al., 2021; Li et al., 2023; Lu et al., 2025).

As research on continual learning for LLMs advances — particularly in high-stakes domains where hallucinations can have severe consequences (Alansari and Luqman, 2025; Xu et al., 2026) — robustly quantifying uncertainty throughout the model’s self-evolving process becomes essential. A central paradigm in uncertainty quantification (UQ) is ensuring that model outputs cover the ground-truth answer with high probability. Conformal Prediction (CP) (Zhou et al., 2025; Shafer and Vovk, 2008; Angelopoulos and Bates, 2021) offers such a framework: a model-agnostic method that provides formal statistical guarantees that the correct answer is included with a user-specified probability. This kind of method calculates uncertainty based on the model’s internal signals and can therefore be applied to LLMs, where uncertainty can be revealed through internal states (Huang et al., 2026), sampling frequency (Su et al., 2024; Li et al., 2024), or token entropy (Moskvoretskii et al., 2025). In recent years, CP has been successfully applied across diverse LLM tasks, including multiple-choice QA (Kumar et al., 2023; Ye et al., 2024), open-ended QA (Su et al., 2024; Li et al., 2024), long-form answering (Mohri and Hashimoto, 2024; Rubin-Toles et al., 2025), and even LLM-as-a-judge evaluations (Sheng et al., 2025), substantially enhancing reliability.

In this work, we study open-ended QA under continual domain pretraining (CDP), a CL setting where LLMs are continually fine-tuned on knowledge from diverse domains. Applying CP to CDP introduces two major challenges: (1) *Violation of Exchangeability*. CP’s coverage guarantees rely on exchangeability (e.g., I.I.D. data). In CDP, selecting fine-tuning data across domains can break




Figure 1: Different quantiles across domains (i.e., “sc”, “ht”, “sp”, “en”, “gg”) in TriviaQA for NE (left), indicating uncertainty about whether the LLM can answer, and the answer uncertainty (right), measuring the likelihood that the LLM will produce the ground-truth answer. x -axis is quantile and y -axis is quantile value for the uncertainty measurement. “whole” represents average quantile. We see that quantile values vary among different domains. Results are from Llama-3.1-8B-Instruct.

this assumption, leading to uneven coverage across domains (Figure 1). To address this, we propose a non-exchangeable CP framework that leverages a buffer dataset and calibrates using the covariate distribution of the test data. (2) *Uninformative Prediction Sets*. CP may produce excessively large sets when LLMs face uncertainty in new domains, making predictions uninformative. While adding “Can’t answer” labels can help (Li et al., 2024), naive thresholding fails under catastrophic forgetting (Toneva et al., 2018). We develop an adaptive CP method based on label-conditional CP (Löfström et al., 2015), which dynamically adjusts to shifting model capabilities and achieves optimal efficiency by producing the smallest prediction sets.

Overall, our main contributions are:

- To the best of our knowledge, this work is the first to investigate CP in the context of CDP for LLMs. We introduce a model-agnostic uncertainty quantification framework tailored to CDP, enabling reliable adaptation of LLMs across evolving domains.
- We address two key challenges of robust UQ for CDP via CP: non-exchangeability and adaptive rejection under catastrophic forgetting, and theoretically recover finite sample exact coverage under domain distribution shift. In addition, we provide a theoretical analysis for guaranteeing better efficiency of our approach (Appendix A).
- Extensive experiments demonstrate that our approach –AR-NECP–effectively mitigates the distribution shift between calibration and testing data, enabling LLMs to reject unanswerable

questions while maintaining reliable correctness coverage.

2 Related work

2.1 CDP in LLM

CDP fine-tunes LLMs sequentially across domains (Shi et al., 2024). Gururangan et al. (2020) show that domain-adaptive and task-adaptive pretraining (DAPT, TAPT) consistently enhance downstream performance. To mitigate domain interference, Gururangan et al. (2021) proposes DEMix layers with modular experts that disentangle representations. Qin et al. (2022) present ELLE, which expands model capacity using domain prompts for efficient lifelong adaptation, while recyclable tuning (Qin et al., 2023) reuses prior weights via distillation to speed convergence and retain knowledge. Han et al. (2020) introduces ECONET for continual pretraining on temporal and event reasoning. Lu et al. (2025) propose a novel Lora adaptor for dealing with catastrophic forgetting. Domain-specific variants such as PMC-LLaMA (Wu et al., 2024), SaulLM-7B (Colombo et al., 2024), Lawyer LLaMA (Huang et al., 2023), EcomGPT-CT (Ma et al., 2023), CFGPT (Li et al., 2023), and AF Adapter (Yan et al., 2023) further realize vertical continuity, tailoring the pretraining–fine-tuning pipeline to specialized domains while balancing transfer and forgetting.

2.2 CP for LLM

CP enables LLMs to produce reliable prediction sets or calibrated outputs (Zhou et al., 2025; Campos et al., 2024). Existing CP methods for LLMs fall into three categories: (1) sampling-based CP, which samples multiple responses and selects those exceeding a calibrated uncertainty threshold (Su et al., 2024; Li et al., 2024; Gui et al., 2024), though its achievable coverage is limited by model accuracy; (2) adaptive-sampling CP, which continues sampling until reaching the desired coverage (Quach et al., 2023) but at high computational cost (Su et al., 2024); and (3) abstractive CP, which replaces uncertain content with generic statements in generative tasks (Mohri and Hashimoto, 2024; Rubin-Toles et al., 2025). Other applications include LLM-as-a judge (Sheng et al., 2025), classification, Multi-choice QA (Kumar et al., 2023; Ye et al., 2024). However, all these methods assume exchangeability between calibration and test data, which often breaks under domain shifts or CL.

As the first CP method for CDP, our method provides approximate finite-sample or *global* asymptotic coverage guarantees across domains, reduces computation by calibrating once per dataset instead of per instance, and extends CP to LLMs’ CDP—achieving scalable, domain-aware, and more reliable calibration under domain shifts. The most related work is (Gao and Liu) that proposes a nearest-neighbor calibration achieving *local* asymptotic coverage under traditional CL settings in each domain, while neglecting the non-exchangeable problem caused by different domain distribution in test and calibration data.

3 Preliminary

3.1 Background in CP

CP is a model-agnostic framework that provides distribution-free uncertainty quantification with guaranteed coverage (Shafer and Vovk, 2008; Angelopoulos and Bates, 2021). Given a data set $\{(x_i, y_i)\}_{i=1}^n$, it first divides it into training data $D_{train}\{(x_i, y_i)\}_{i=1}^{n_{train}}$, $D_{cal}\{(x_i, y_i)\}_{i=1}^{n_{cal}}$ and $D_{test}\{(x_i, y_i)\}_{i=1}^{n_{test}}$, where x_i is the question, y_i is the corresponding answer to x_i , n_* is the number of samples in the corresponding dataset, $n = n_{train} + n_{cal} + n_{test}$.

Let $S(f, (x, y))$ denote a nonconformity score that measures how unusual label y is for input x under model f . For example, in classification, it is usually calculated by:

$$S(f, (x, y)) = 1 - f_y(x), \quad (1)$$

where $f_y(x)$ is probability of f predict x as y .

Using the calibration set $\{(x_i, y_i)\}_{i=1}^{n_{cal}}$, it compute non-conformity scores as:

$$s_i = S(f, (x_i, y_i)), \quad i = 1, \dots, n_{cal}. \quad (2)$$

Based on the above equation, it calculates a threshold quantile for including labels by:

$$q_\alpha = \text{Quantile}\left(\{s_i\}_{i=1}^{n_{cal}}; \frac{(n_{cal} + 1)(1 - \alpha)}{n_{cal}}\right). \quad (3)$$

Where α is the max error rate required to be achieved. For a new test input x_{n+1} and candidate labels $y \in Y$, we compute the non-conformity score for y as $s_{n+1}(y) = S(f, (x_{n+1}, y))$. Then the prediction set is defined as:

$$C_\alpha(x_{n+1}) = \{y : s_{n+1}(y) < q_\alpha\}, \quad (4)$$

which guarantees finite-sample coverage:

$$\mathbb{P}(y_{n+1} \in C_\alpha(x_{n+1})) \geq 1 - \alpha. \quad (5)$$

Besides the coverage guarantee, CP also needs to ensure a small prediction set, i.e., high prediction efficiency, to be informative

3.2 Sampling-based CP in LLMs

One common strategy of applying CP in LLMs is to use sampling-based methods, especially for open-ended QA tasks (Su et al., 2024; Li et al., 2024). Given an LLM, one would only need two datasets, calibration data $D_{cal}\{(X_i, Y_i)\}_{i=1}^{n_{cal}}$ and $D_{test}\{(X_i, Y_i)\}_{i=1}^{n_{test}}$.

It involves two stages, the calibration stage and the testing stage. In the calibration stage, for each data point (X_i, Y_i) , it first asks the LLM to generate the answer for M times, and uses a similarity-based method, e.g., Rouge-L score, cosine similarity, to cluster the answers to get a clustered answer set:

$$C_i = \{c_{ij} | j \in [1, N]\}, \quad (6)$$

where c_{ij} is the representative answer in the cluster j for question i and N is the number of clusters. Then it uses the number of answers in each cluster as the uncertainty score, as the clustering may not be perfect, and each question may have various answers. It then calculates the non-conformity score as:

$$S(f, (x, y)) = \max_{j \in [1, M], c_j = y} \left(1 - \frac{\text{freq}_j}{M}\right), \quad \text{if such } j \text{ exists,} \\ S(f, (x, y)) = 0, \quad \text{otherwise.} \quad (7)$$

The quantile is calculated by Eq. 3. In the inference stage, the prediction set is built by:

$$C_\alpha(x_{n+1}) = \{c_j : 1 - \frac{\text{freq}_j}{M} < q_\alpha\}. \quad (8)$$

The limitation of this method is that the target coverage is upper-bounded by the model’s ability.

4 Proposed Approach

4.1 Problem Formulation

We consider K QA datasets $D^k, k \in [1, \dots, K]$, each from different domains, and split them into $D_{train}^k, D_{buffer}^k$, where each dataset D_{train}^k is used to finetune a corresponding LLM ϕ_k, D_{buffer}^k is used to calibrate CP method. At each training step,

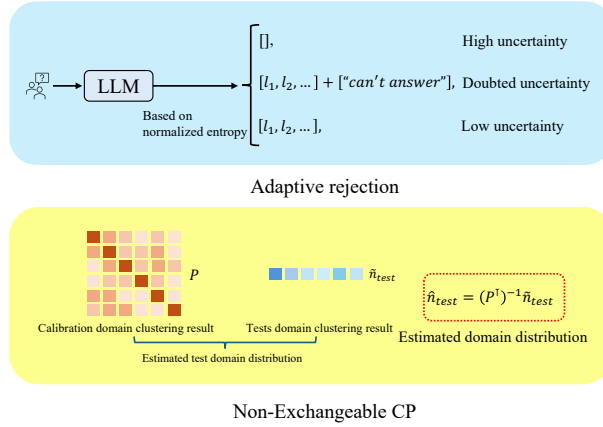


Figure 2: Our two major contributions: Adaptive rejection module and non-exchangeable CP

ϕ_k has only limited access to data from previous tasks, which is a common setting in CDP. For illustration simplicity, we describe our proposed CP method on the final finetuned model ϕ_K . At this stage, a union dataset D_{test} containing all domains from previous steps is provided, without domain labels, and an unknown domain distribution. **Our goal is to enable LLM ϕ_K to generate prediction sets that marginally guarantee to cover the right answer with $1-\alpha$ probability.**

Exchangeable Assumption within Domains.

Our problem setting assumes that data within each domain is exchangeable following (Li et al., 2024), i.e., data from calibration and testing sets in the same domain are drawn independently and identically from the same data distribution D .

Overview. Our framework follows protocols in sampling-based CP and comprises two key components: the *Adaptive Rejection CP* and the *Non-Exchangeable CP* modules, and is named as ARNECP. The Adaptive Rejection CP introduces a dynamic rejection mechanism that adaptively determines the threshold for abstaining from answering while maintaining the ground truth within the prediction set, thereby ensuring the desired coverage guarantee. The Non-Exchangeable CP component first estimates the domain composition of the testing data by analyzing grouping statistics between the calibration and testing sets. It then either resamples or reweights the calibration samples based on the clustering results from both sets, so that the distribution of nonconformity scores in the calibration data closely aligns with that of the testing data. This adjustment statistically restores exchangeability between the two distributions, enabling reliable coverage under domain shift. The overall procedure

of the proposed algorithm is summarized in Algorithm 1. A theoretical analysis of our method for guaranteeing better CP performance can be seen in Appendix A.

4.2 Adaptive Rejection CP

A simple abstention baseline assigns a ‘can’t answer’ label using a lightweight MLP’s predicted probability (Li et al., 2024), but it lacks a coverage-aware, principled rejection mechanism. We enable refusal via a **label-conditional CP** scheme calibrated to maintain coverage $1 - \alpha$ on answerable questions; items labeled ‘can’t answer’ are then rejected. Concretely, we fix the error rate on answerable instances at (α) —matching the global CP level—to yield well-calibrated prediction sets. To mitigate catastrophic forgetting across domains, we incorporate the estimated *answerability rate* from calibration data to adaptively set the rejection threshold. We detail the procedure next.

Rejection Mechanism To begin with, each sample $x_i \in D_{\text{cal}}$ in the dataset is given an unanswerable score s_i . In our project, we propose to use normalized entropy (NE) as the unanswerable score, which is a statistical metric to predict the LLM’s ability to answer a question. It is calculated as:

$$NE(x_i) = -\frac{1}{\log M} \sum_{j=1}^K f(c_j) \log f(c_j), \quad (9)$$

where M is total number of samples, $f(c_j)$ is the number of answers in cluster c_j divided by M , and K is the total number of clusters that x_i has, . $NE \in [0, 1]$, where higher NE indicates lower confidence to predict the question. We further formulate the probability that the LLM cannot

(p^0) /can (p^1) answer the question as:

$$p^0(x_i) = NE(x_i); p^1(x_i) = 1 - NE(x_i). \quad (10)$$

Setting the error rate of unanswerable questions as α_0 and the error rate of answerable questions as α_1 , we can get the overall error rate as:

$$\alpha = (1 - r_{correct})\alpha_0 + r_{correct}(\alpha_1 + \alpha - \alpha\alpha_1), \quad (11)$$

where $r_{correct}$ is the ratio of answerable questions in the calibration dataset. Under the exchangeability assumption within domains, Eq. 11 remains valid for the test set. We consider a question as answerable if any sampled question’s answer regex matches the ground truth. Thus, we can get the relation between α_0 and α_1 as:

$$\alpha_1 = \frac{(1 - r_{correct})(\alpha - \alpha_0)}{r_{correct}(1 - \alpha)}. \quad (12)$$

After this, we denote the answerable question as x^1 and the unanswerable question as x^0 , with n_{cal}^0 as the number of questions that cannot be answered and n_{cal}^1 as the number of questions that can be answered in the calibration data. We calculate the quantile for adding the ‘can’t answer’ label as:

$$q_{\alpha_0}^0 = \text{Quantile}\left(\{p^0(x_i^0)\}_{i=1}^{n_{cal}^0}; \frac{(n_{cal}^0 + 1)(1 - \alpha_0)}{n_{cal}^0}\right), \quad (13)$$

and the quantile for rejecting to answer as:

$$q_{\alpha_1}^1 = \text{Quantile}\left(\{p^1(x_i^1)\}_{i=1}^{n_{cal}^1}; \frac{(n_{cal}^1 + 1)(1 - \alpha_1)}{n_{cal}^1}\right). \quad (14)$$

In the testing stage, LLM refuses to answer the question when $p^0(x_i) < q_{\alpha_0}^0$ and $p^1(x_i) > q_{\alpha_1}^1$, where the NE score shows that the LLM has a low probability of answering the question but a high probability of being unanswerable. We add the ‘can’t answer’ label to the prediction set when $p^0(x_i) < q_{\alpha_0}^0$ and $p^1(x_i) < q_{\alpha_1}^1$ to ensure a guaranteed coverage. For the remaining cases, where the LLM has a low probability of being unanswerable, we use the standard CP method.

Prediction Set for Answerable Questions After labeling a portion of questions as unanswerable, some truly answerable questions may be mistakenly classified as unanswerable. This misclassification introduces a distribution bias between the set

of all answerable questions and the subset predicted as answerable. To mitigate this bias, we recompute the quantile to ensure that the prediction set properly includes the retained answerable questions, as follows:

$$\hat{q}_{\alpha}^{text} = \text{Quantile}\left(\{\hat{s}_i\}_{i=1}^{\hat{n}_{cal}}; \frac{(\hat{n}_{cal} + 1)(1 - \alpha)}{\hat{n}_{cal}}\right), \quad (15)$$

where the frequency-based non-conformity score \hat{s}_i corresponds to the answerable questions with $p^1(x_i) < q_{\alpha_1}^1$, and \hat{n}_{cal} is the number of such questions. For answerable questions, clusters with higher cluster scores than \hat{q}_{α}^{text} are included in the prediction set c_i , if not rejected. This progress can be formulated as:

$$C(x_i) = \begin{cases} \text{Refuse to answer,} \\ \quad \text{if } p^0(x_i) < q_{\alpha_0}^0, p^1(x_i) > q_{\alpha_1}^1 \\ \text{Add a “can’t answer” label,} \\ \quad \text{if } p^0(x_i) < q_{\alpha_0}^0, p^1(x_i) < q_{\alpha_1}^1 \\ \text{Use standard CP, otherwise.} \end{cases} \quad (16)$$

Due to the concentration issue caused by using frequency alone (Su et al., 2024), we propose to set the cluster scores as frequency minus NE based on the method in (Su et al., 2024).

Grid Search for Best Prediction Efficiency With the above quantiles, we use a grid search to achieve the best efficiency (i.e., smallest prediction set), which is especially helpful when the answerable rate is high or the coverage requirement is large, making it hard to find unanswerable questions. Specifically, we choose the pair of (α_0, α_1) according to Eq. 12. This can be formulated as:

$$\alpha_0, \alpha_1 = \arg \min_{\alpha_0, \alpha_1} E(|C_{\alpha}(x_i)| \mid x_i \in D_{cal}), \quad (17)$$

where $|C_{\alpha}(x_i)|$ is the number of answers in the prediction set for x_i .

4.3 Non-Exchangeable CP Across Domains

In CDP, data from different domains can be non-exchangeable. To achieve reliable coverage across different domains, after each domain’s pretraining, we draw a small subset from D^k as D_{buffer}^k . We further split it into two parts, $D_{cluster}^k$ and D_{cal}^k , with sizes $n_{cluster}^k$ and n_{cal}^k , respectively. $D_{cluster}^k$ will be used for constructing a cluster centroid for each

domain, and D_{cal}^k will be used to calibrate for CP. Given a joint test question distribution D_{test} , the domain information is unknown in CDP. Moreover, the exchangeability assumption no longer holds in this setting, as testing data and calibration data may have different domain distributions. To address these challenges, we first approximate the domain distribution in the test data and then resample the calibration data accordingly. Our method mitigates the domain distribution shift under the assumption that within each domain, data is exchangeable.

Domain Cluster Centroids via Transformer Encoders. We use a Transformer encoder $T(\cdot)$ to obtain semantic embeddings for all questions, representing each domain in a shared embedding space. For domain k with n_{cluster}^k questions $\{x_i^k\}_{i=1}^{n_{\text{cluster}}^k}$, each question is encoded as

$$\mathbf{h}_i^k = T(x_i^k) \in \mathbb{R}^d. \quad (18)$$

The domain centroid is then the mean embedding:

$$\mathbf{c}_k = \frac{1}{n_{\text{cluster}}^k} \sum_{i=1}^{n_{\text{cluster}}^k} \mathbf{h}_i^k, \quad (19)$$

which is further ℓ_2 -normalized for comparability:

$$\tilde{\mathbf{c}}_k = \frac{\mathbf{c}_k}{\|\mathbf{c}_k\|_2}. \quad (20)$$

A new question x_i is assigned to the domain with the highest cosine similarity:

$$c(x_i) = \arg \max_{k \leq K} (\tilde{\mathbf{c}}_k^\top \mathbf{h}_i). \quad (21)$$

Testing data domain distribution estimation from the Buffer. After obtaining the cluster centroid for each domain, we group both calibration and testing data according to their nearest domain centroids. Under the I.I.D. assumption within each domain, the statistics of this clustering on the calibration can be used to estimate the proportion of domains in the testing set.

In the calibration data, after clustering, we obtain each domain's cluster assignment result. We denote P_{ij} , where $i, j \in \{1, \dots, K\}$, as the ratio of questions from domain i that are clustered into domain j . Each domain in the test set will follow a similar transition pattern:

$$\tilde{n}_{\text{test}}^j = \sum_{i=1}^K \hat{n}_{\text{test}}^i P_{ij}, \quad (22)$$

where $\tilde{n}_{\text{test}}^j$ is the number of test samples clustered into domain j . Let $P = [P_{ij}]_{i,j=1}^K$, the number of questions in the testing data clustered into domain k as $\tilde{\mathbf{n}}_{\text{test}} = [\tilde{n}_{\text{test}}^1, \dots, \tilde{n}_{\text{test}}^K]^\top$, and estimate true number of questions in each domain k in $\hat{\mathbf{n}}_{\text{test}} = [\hat{n}_{\text{test}}^1, \dots, \hat{n}_{\text{test}}^K]^\top$. The system can then be written compactly as:

$$P^\top \hat{\mathbf{n}}_{\text{test}} = \tilde{\mathbf{n}}_{\text{test}}. \quad (23)$$

As P^\top is a diagonally dominant matrix, it is invertible (Horn and Johnson, 2012). We can then estimate the true domain counts by:

$$\hat{\mathbf{n}}_{\text{test}} = (P^\top)^{-1} \tilde{\mathbf{n}}_{\text{test}}. \quad (24)$$

Since the transformer-based clustering method demonstrates strong separability performance (Wu and Varshney, 2024), Eq. 24 holds in practice.

Resampling CP After estimating the number of questions in each domain of the testing data, we resample the calibration data from each domain to form \hat{D}_{cal}^k . The union $\hat{D}_{\text{cal}}^{\text{union}}$ thus approximates the domain distribution of the testing dataset.

Reweighting CP Another approach for alleviating distribution shift in CP is Weighted CP (Tibshirani et al., 2019), which reweights calibration samples based on the estimated domain proportions. Let $\{s_i, w_i\}_{i=1}^{n_{\text{cal}}}$ denote the nonconformity scores and their associated weights, where for each calibration sample $x_i \in D_{\text{cal}}^k$, the weight is defined as

$$w_i = \frac{\hat{n}_{\text{test}}^k}{\sum_{k=1}^K \hat{n}_{\text{test}}^k}. \quad (25)$$

The weighted quantile threshold \hat{q}_α is then

$$\hat{q}_\alpha = \inf \left\{ q \in \mathbb{R} : \frac{\sum_{i=1}^{n_{\text{cal}}} w_i \mathbf{1}(s_i \leq q)}{\sum_{i=1}^{n_{\text{cal}}} w_i} \geq 1 - \alpha \right\}, \quad (26)$$

where $\mathbf{1}(\cdot)$ denotes the indicator function. In the Adaptive Rejection CP framework, all thresholds $q_{\alpha_0}^0$, $q_{\alpha_1}^1$, and $\hat{q}_\alpha^{\text{ext}}$ are recalculated accordingly based on the weighted calibration scores. Compared with resampling-based strategies, Weighted CP provides a deterministic adjustment and generally yields more stable coverage when certain domains in the test set are severely underrepresented, as resampling may fail to include sufficient representative samples from minority domains.

Algorithm 1 Adaptive Rejection and Non-Exchangeable CP (AR-NECP)

Require: Split D_{buffer}^k into $\{D_{cluster}^k, D_{cal}^k\}_{k=1}^K$, x_{test} in D_{test} , LLM ϕ_K , significance α , transformer model T

Ensure: Prediction sets $C_\alpha(x)$

- 1: **Step 1: Domain Centroids**
 - 2: **for** $k = 1, \dots, K$ **do**
 - 3: Get c_k from $D_{cluster}^k$ using Eq. 19.
 - 4: **end for**
 - 5: Assign $x \in D_{cal} \cup D_{test}$ using Eq. 21.
 - 6: **Step 2: Non-exchangeable CP**
 - 7: Compute P_{ij} and \tilde{n}_{test} from clustering
 - 8: Estimate $\hat{n}_{test} = (P^\top)^{-1} \tilde{n}_{test}$
 - 9: Resample or reweight \hat{D}_{cal}^k according to \hat{n}_{test}
 - 10: **Step 3: Adaptive Rejection CP Calibration**
 - 11: Calculate $q_{\alpha_0}^0, q_{\alpha_1}^1, \hat{q}_\alpha^{text}$ for each domain clustered dataset \hat{D}_{cal}^{union} by method in Section 4.2.
 - 12: **Step 4: Inference**
 - 13: **for** $x_i \in D_{test}$ **do**
 - 14: Build $C_\alpha(x_i)$ with $q_{\alpha_0}^0, q_{\alpha_1}^1, \hat{q}_\alpha^{text}$ by algorithm described in Section 4.2.
 - 15: **end for**
 - 16:
 - 17: **return** $C_\alpha(x)$
-

5 Experiment

For a comprehensive evaluation, we focus on the following research questions:

- **RQ. 1.** How effective is our AR-NECP method in addressing the distribution shift in CDP?
- **RQ. 2.** How does the adaptive rejection CP improve the prediction efficiency of CP?
- **RQ. 3.** How do different UQ methods for LLM answer rejection influence CP’s prediction efficiency?
- **RQ. 4.** How well does \hat{n}_{test}^k (estimated number of test questions per domain) work in practice?

5.1 Experiment Setting

We evaluate **AR-NECP** on three widely used open QA datasets: TriviaQA (Joshi et al., 2017), HotPotQA (Yang et al., 2018), and MMLU (Hendrycks et al., 2021b,a). For each dataset, we first split it into multiple domains. We divide **TriviaQA** into five domains using keyword-based search. The corresponding keywords are: *science & art (sc)*, *history (ht)*, *sports & politics (sp)*, *entertainment (en)*, and *geography (gg)*. For **HotPotQA**, we divide

the data into three domains using keyword-based search: *science & art (sc)*, *history (ht)*, *entertainment & geography (eg)*. For **MMLU**, we follow Liang (2024) and divide the dataset into four domains: *STEM (st)*, *Humanities (hu)*, *Social Science (ss)*, and *Other (ot)*.

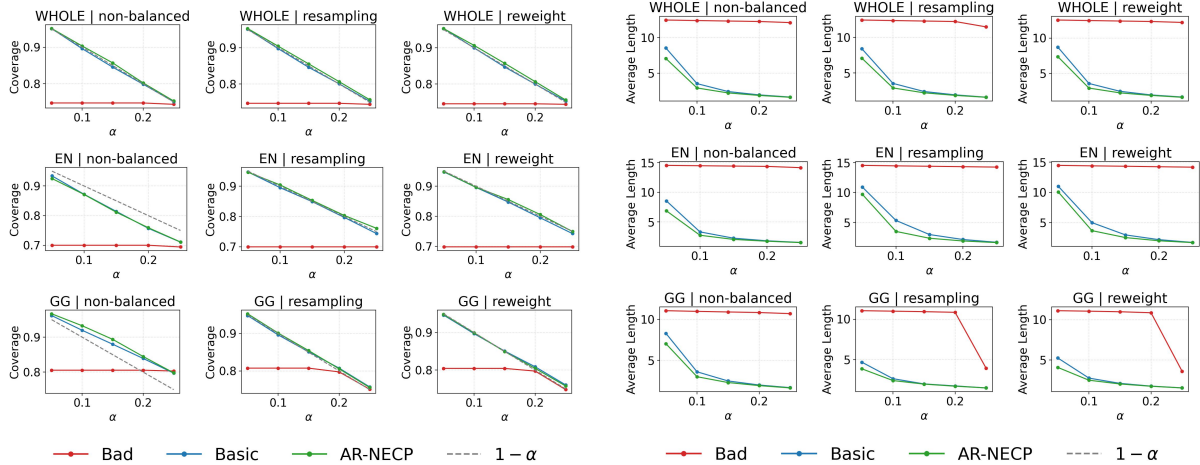
For each domain, we set the number of training samples $n_{train}^k = 10,000$, buffer samples $n_{buffer}^k = 2,000$, and both clustering and calibration samples $n_{cluster}^k = n_{cal}^k = 1,000$, with a unified calibration size $n_{cal}^{union} = 5,000$. The test size n_{test}^k will be defined later when we discuss distribution shifts and how our method addresses them. Three LLMs are used: Mistral-7B-Instruct-v0.3 (Mistral7B) (Jiang et al., 2023), gemma-7b-it (gemma7b) (Team et al., 2024), and Llama-3.1-8B-Instruct (Llama8B) (Touvron et al., 2023). Our discussion will be focused on results for Llama8B on TriviaQA, and results for the other two models can be found in the Appendix. We fine-tune each model using the hyperparameters provided in the Appendix B.2. We vary error levels α among [0.05, 0.10, 0.15, 0.20, 0.25].

We compare with two **baselines**: frequency-based CP (Bad) and basic “can’t answer” CP (Basic), adapting state-of-the-art methods from (Su et al., 2024) and (Li et al., 2024) to the CDP setting. The former ignores LLMs’ ability to abstain, while the latter introduces a “can’t answer” label but still fails to reject responses when the LLM cannot provide a valid answer. We apply the same data balancing techniques (for addressing non-exchangeability across domains) to all compared methods. Results are averaged from five runs.

Metrics. We adopt three metrics to assess performance: (1) Coverage: the ratio of prediction sets that contain the ground truth. (2) Efficiency: the average size of the prediction sets. (3) Unanswerable efficiency: efficiency measured specifically on unanswerable questions, used to verify the effectiveness of our proposed methods. The first two metrics are standard in CP.

5.2 Main Results (RQ1 & RQ2)

We evaluate coverage and efficiency across different values of α under three testing distribution settings: (1) When easy questions dominate the test data, calibration causes over-coverage; (2) When hard questions dominate, CP yields under-coverage; and (3) When calibration and test distributions are similar, the target coverage is achieved. For all datasets and all different domain distribution settings, we set 3000 questions for the major do-



(a) Coverage of the three CP methods under different quantiles, different distribution shifts, and using different balancing methods (resampling, reweighting, and non-balancing). Results that lie closest to the $1 - \alpha$ line indicate that the corresponding methods achieve the desired performance, neither exceeding it (overcoverage) nor falling short (undercoverage).

(b) Average prediction set size of the three CP methods under different quantiles, under different distribution shifts, and using different balancing methods. A smaller value means better performance.

Figure 3: Main results

main, with 200–500 questions randomly sampled for the remaining domains. In the balanced case, each domain contains 800–1500 randomly sampled questions. Results for TriviaQA are shown in Figure 3a and Figure 3b. We have the following observations:

- The basic non-balanced CP method is sensitive to distribution shift. When the easier domain “GG” dominates the testing data, it tends to produce over-coverage; whereas when the harder domain “EN” dominates, it fails to reach the target coverage (the dotted line in Figure 3a). This confirms that domain distribution indeed affects the coverage and highlights the necessity of addressing such distribution shifts.
- Both resampling and reweighting methods produce prediction sets whose coverage closely matches the target level, empirically supporting that our approach successfully mitigates non-exchangeability under distribution shift. This also indicates that our estimation of domain proportions performs well in practice.
- Our adaptive rejection CP consistently yields tighter prediction sets through its rejection mechanism, especially under high coverage requirements, demonstrating the robustness and effectiveness of the grid-search calibration. When the coverage level is low, the improvement becomes limited because the high inclu-

sion threshold already filters out uncertain or unanswerable cases.

- The Bad CP baseline consistently shows lower efficiency than both the Basic CP and AR-NECP, and fails to achieve the target coverage in most scenarios, revealing the limitations of previous CP settings.
- On the MMLU datasets, the coverage findings remain consistent. However, efficiency gains are limited, as MMLU is a 4-class classification task that inherently avoids the oversized prediction set problem. For HotpotQA, the domain split exhibits high overlap, resulting in minimal distribution shift. Our method achieves larger efficiency improvements on this dataset compared to TriviaQA, as it is inherently more challenging, while maintaining the coverage guarantee across all settings.

5.3 How do different UQ methods for LLM answer rejection influence prediction efficiency? (RQ3)

We compare two UQ approaches for LLM answer rejection: our proposed method, which uses NE to measure uncertainty, and a commonly used alternative that trains a small MLP classifier (Li et al., 2024) on $D_{\text{cluster}}^{\text{union}}$ —the union of all cluster sets $D_{\text{cluster}}^k, k \in [1, 2, \dots, K]$ —to predict the LLM’s confidence in answering each question. For illustrative purposes, we evaluated both methods at five

Table 1: Comparison of prediction efficiency for AR-NECP (MLP) and AR-NECP (NE) under reweighting across different domains distribution and α values.

Distribution	Method	$\alpha=0.05$	$\alpha=0.10$	$\alpha=0.15$	$\alpha=0.20$	$\alpha=0.25$
whole	AR-NECP (MLP)	8.66	3.60	2.47	1.92	1.65
	AR-NECP (NE)	7.95	3.38	2.37	1.88	1.62
en-d	AR-NECP (MLP)	11.07	4.51	3.06	2.24	1.79
	AR-NECP (NE)	10.49	3.72	2.66	2.10	1.72
gg-d	AR-NECP (MLP)	4.58	2.59	2.06	1.76	1.55
	AR-NECP (NE)	4.29	2.54	2.04	1.75	1.54

different quantile levels using the reweighting strategy. The comparison results, presented in Table 1, show that the NE-based AR-NECP consistently outperforms the MLP-based variant, demonstrating the superior ability of NE to distinguish between answerable and unanswerable questions. The smaller performance gap between AR-NECP (MLP) and AR-NECP (NE) in *gg-d* domain is mainly because *gg* is the most recent domain used to finetune LLMs in CDP, making its questions relatively easier for the LLM to answer; hence, the rejection mechanism has a marginal impact. In the dataset HotpotQA, it also shows such a phenomenon, NE consistently outperforms MLP on the efficiency.

5.4 How well does \hat{n}_{test}^k work in practice? (RQ4)

We examine whether estimating the number of test questions per domain is useful in applications. For each testing distribution, we compare the estimated count \hat{n}_{test}^k with the true count n_{test}^k in every domain k and report the mean relative error defined as: $\Delta^k = \frac{|n_{\text{test}}^k - \hat{n}_{\text{test}}^k|}{n_{\text{test}}^k}$. Results are presented in Table 2. Overall, the errors are small across domains, which shows the estimator is accurate and useful in practice. The small bias comes from sampling variability; with more samples, for instance, en in en-d, gg in gg-d, the domain-wise I.I.D. assumption holds more closely and the estimates improve. Similar findings are observed on HotpotQA and MMLU. Even though the domains in HotpotQA have substantial overlap, our method still achieves considerable error reduction, demonstrating its effectiveness.

Shift	sc	ht	sp	en	gg
en-d	0.08	0.04	0.05	0.01	0.10
gg-d	0.10	0.11	0.11	0.06	0.02
whole	0.04	0.02	0.03	0.02	0.04

Table 2: Mean relative error Δ^k between true and estimated domain counts under different test-time shifts. *en-d* and *gg-d* denote test sets dominated by the *en* and *gg* domains, respectively. Lower values are better.

6 Conclusion

This work takes an important step toward ensuring the reliability of self-evolving LLMs. We are the first to CDP where LLMs must continuously adapt to shifting knowledge domains. Our proposed AR-NECP framework addresses two fundamental barriers: handling unknown domain shifts between calibration and testing data, and enabling LLMs to selectively abstain when their competence varies across domains. By reweighting calibration data according to estimated domain distributions and incorporating a label-conditional rejection mechanism, AR-NECP delivers stronger statistical guarantees and more informative predictions. Extensive experiments verify its robustness and efficiency. Looking forward, this approach opens a pathway toward trustworthy CL, laying the foundation for LLMs that can evolve while maintaining reliability in real-world, dynamically changing environments.

Limitation

There are several limitations in this work that warrant future research. (1) Although our estimated number of test questions per domain is theoretically justified, in practice, it can still exhibit noticeable variance. How to accurately estimate and control this variance is beyond the scope of this paper and can be explored in future work. (2) Another limitation lies in the need to resample calibration answers after each fine-tuning stage on a new dataset, which incurs additional computational cost. Ideally, calibration answers could be sampled once and reused in later stages to improve efficiency. However, this is not feasible because fine-tuning across domains introduces non-exchangeability among models trained at different steps. A potential direction for future research is to estimate and correct this bias by modeling how the number of training steps or the order of domain adaptation affects the calibration consistency.

Acknowledgments

This work is supported by the National Science Foundation (NSF) Grant #2312862, NSF-Simons SkAI Institute, NSF CAREER #2440542, NSF #2533996, National Institutes of Health (NIH) #R01AG091762, NSF ACCESS Computing Resources, NAIRR, a Google Research Scholar Award, and Cisco gift grant.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Aisha Alansari and Hamzah Luqman. 2025. Large language models hallucination: A comprehensive survey. *arXiv preprint arXiv:2510.06265*.
- Anastasios N Angelopoulos and Stephen Bates. 2021. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511*.
- Margarida Campos, António Farinhas, Chrysoula Zerva, Mário AT Figueiredo, and André FT Martins. 2024. Conformal prediction for natural language processing: A survey. *Transactions of the Association for Computational Linguistics*, 12:1497–1516.
- 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*.
- Rui Gao and Weiwei Liu. Model uncertainty quantification by conformal prediction in continual learning. In *Forty-second International Conference on Machine Learning*.
- Yu Gui, Ying Jin, and Zhimei Ren. 2024. Conformal alignment: Knowing when to trust foundation models with guarantees. *Advances in Neural Information Processing Systems*, 37:73884–73919.
- 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**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Rujun Han, Xiang Ren, and Nanyun Peng. 2020. Econet: Effective continual pretraining of language models for event temporal reasoning. *arXiv preprint arXiv:2012.15283*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. 2021a. Aligning ai with shared human values. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021b. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Roger A Horn and Charles R Johnson. 2012. *Matrix analysis*. Cambridge university press.
- Fanding Huang, Guanbo Huang, Xiao Fan, Yi He, Xiao Liang, Xiao Chen, Qinting Jiang, Faisal Nadeem Khan, Jingyan Jiang, and Zhi Wang. 2026. **Semantic-space exploration and exploitation in rlvr for llm reasoning**. *Preprint*, arXiv:2509.23808.
- Quzhe Huang, Mingxu Tao, Chen Zhang, Zhenwei An, Cong Jiang, Zhibin Chen, Zirui Wu, and Yansong Feng. 2023. Lawyer llama technical report. *arXiv preprint arXiv:2305.15062*.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. **Mistral 7b**. *Preprint*, arXiv:2310.06825.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. **triviaqa: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension**. *arXiv e-prints*, arXiv:1705.03551.
- Sung Jae Jung, Hajung Kim, and Kyoung Sang Jang. 2024. Llm based biological named entity recognition from scientific literature. In *2024 IEEE International Conference on Big Data and Smart Computing (Big-Comp)*, pages 433–435. IEEE.
- Bhawesh Kumar, Charlie Lu, Gauri Gupta, Anil Palepu, David Bellamy, Ramesh Raskar, and Andrew Beam. 2023. Conformal prediction with large language models for multi-choice question answering. *arXiv preprint arXiv:2305.18404*.
- Jiangtong Li, Yuxuan Bian, Guoxuan Wang, Yang Lei, Dawei Cheng, Zhijun Ding, and Changjun Jiang. 2023. Cfgpt: Chinese financial assistant with large language model. *arXiv preprint arXiv:2309.10654*.
- Shuo Li, Sangdon Park, Insup Lee, and Osbert Bastani. 2024. **TRAQ: Trustworthy retrieval augmented question answering via conformal prediction**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3799–3821, Mexico City, Mexico. Association for Computational Linguistics.
- Kaizhao Liang. 2024. MMLU Auxiliary Trained Set Labelled by e5-mistral-7b-instruct. <https://huggingface.co/datasets/kz919/mmlu-auxiliary-train-e5-mistral-7b-instruct>. Accessed: Date of Access.

- Tuve L fstr m, Henrik Bostr m, Henrik Linusson, and Ulf Johansson. 2015. Bias reduction through conditional conformal prediction. *Intelligent Data Analysis*, 19(6):1355–1375.
- Yuheng Lu, Bingshuo Qian, Caixia Yuan, Huixing Jiang, and Xiaojie Wang. 2025. [Controlled low-rank adaptation with subspace regularization for continued training on large language models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 19165–19181, Vienna, Austria. Association for Computational Linguistics.
- Shirong Ma, Shen Huang, Shulin Huang, Xiaobin Wang, Yangning Li, Hai-Tao Zheng, Pengjun Xie, Fei Huang, and Yong Jiang. 2023. Ecomgpt-ct: Continual pre-training of e-commerce large language models with semi-structured data. *arXiv preprint arXiv:2312.15696*.
- Christopher Mohri and Tatsunori Hashimoto. 2024. Language models with conformal factuality guarantees. *arXiv preprint arXiv:2402.10978*.
- Viktor Moskvoretskii, Maria Marina, Mikhail Salnikov, Nikolay Ivanov, Sergey Pletenev, Daria Galimzianova, Nikita Krayko, Vasily Konovalov, Irina Nikishina, and Alexander Panchenko. 2025. Adaptive retrieval without self-knowledge? bringing uncertainty back home. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6355–6384.
- Yujia Qin, Cheng Qian, Xu Han, Yankai Lin, Huadong Wang, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2023. Recyclable tuning for continual pre-training. *arXiv preprint arXiv:2305.08702*.
- Yujia Qin, Jiajie Zhang, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2022. Elle: Efficient lifelong pre-training for emerging data. *arXiv preprint arXiv:2203.06311*.
- Victor Quach, Adam Fisch, Tal Schuster, Adam Yala, Jae Ho Sohn, Tommi S Jaakkola, and Regina Barzilay. 2023. Conformal language modeling. *arXiv preprint arXiv:2306.10193*.
- Nitin Liladhar Rane, Abhijeet Tawde, Saurabh P Choudhary, and Jayesh Rane. 2023. Contribution and performance of chatgpt and other large language models (llm) for scientific and research advancements: a double-edged sword. *International Research Journal of Modernization in Engineering Technology and Science*, 5(10):875–899.
- Maxon Rubin-Toles, Maya Gambhir, Keshav Ramji, Aaron Roth, and Surbhi Goel. 2025. Conformal language model reasoning with coherent factuality. *arXiv preprint arXiv:2505.17126*.
- Glenn Shafer and Vladimir Vovk. 2008. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3).
- Huanxin Sheng, Xinyi Liu, Hangfeng He, Jieyu Zhao, and Jian Kang. 2025. [Analyzing uncertainty of llm-as-a-judge: Interval evaluations with conformal prediction](#). *Preprint*, arXiv:2509.18658.
- Haizhou Shi, Zihao Xu, Hengyi Wang, Weiyi Qin, Wenyuan Wang, Yibin Wang, Zifeng Wang, Sayna Ebrahimi, and Hao Wang. 2024. Continual learning of large language models: A comprehensive survey. *ACM Computing Surveys*.
- Jiayuan Su, Jing Luo, Hongwei Wang, and Lu Cheng. 2024. [API is enough: Conformal prediction for large language models without logit-access](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 979–995, Miami, Florida, USA. Association for Computational Linguistics.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Riviere, Mihir Sanjay Kale, Juliette Love, and 1 others. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Ryan J Tibshirani, Rina Foygel Barber, Emmanuel Candès, and Aaditya Ramdas. 2019. Conformal prediction under covariate shift. *Advances in neural information processing systems*, 32.
- Mariya Toneva, Alessandro Sordani, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. 2018. An empirical study of example forgetting during deep neural network learning. *arXiv preprint arXiv:1812.05159*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timoth e Lacroix, Baptiste Rozi re, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Jiaqi Wang, Enze Shi, Huawen Hu, Chong Ma, Yiheng Liu, Xuhui Wang, Yincheng Yao, Xuan Liu, Bao Ge, and Shu Zhang. 2025. Large language models for robotics: Opportunities, challenges, and perspectives. *Journal of Automation and Intelligence*, 4(1):52–64.
- Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Weidi Xie, and Yanfeng Wang. 2024. Pmc-llama: toward building open-source language models for medicine. *Journal of the American Medical Informatics Association*, 31(9):1833–1843.
- Xinbo Wu and Lav R Varshney. 2024. Transformer-based causal language models perform clustering. *arXiv preprint arXiv:2402.12151*.
- Shijia Xu, Zhou Wu, Xiaolong Jia, Yu Wang, Kai Liu, and April Xiaowen Dong. 2026. [Self-correcting rag: Enhancing faithfulness via mmkp context selection and nli-guided mcts](#). *Preprint*, arXiv:2604.10734.

Yongyu Yan, Kui Xue, Xiaoming Shi, Qi Ye, Jingping Liu, and Tong Ruan. 2023. Af adapter: Continual pretraining for building chinese biomedical language model. In *2023 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 953–957. IEEE.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Fanghua Ye, Mingming Yang, Jianhui Pang, Longyue Wang, Derek F. Wong, Emine Yilmaz, Shuming Shi, and Zhaopeng Tu. 2024. [Benchmarking LLMs via uncertainty quantification](#). In *The Thirty-eight Conference on NIPS Datasets and Benchmarks Track*.

Xiaofan Zhou, Baiting Chen, Yu Gui, and Lu Cheng. 2025. Conformal prediction: A data perspective. *ACM Computing Surveys*.

A Theoretical Proof

In this section, we provide theoretical proofs for the following statements:

1. Our Adaptive Rejection CP can achieve better efficiency than standard CP with the ‘Can’t answer’ label, when the coverage requirement is high, and the uncertainty score for identifying unanswerable questions is reasonable.
2. Our Adaptive Rejection CP can achieve the required coverage.

For our Non-exchangeable CP, we additionally prove:

1. Weighted CP achieves finite-sample exact coverage.
2. Resampling CP achieves asymptotic (finite-sample) coverage.

Efficiency of Adaptive Rejection CP. Our Adaptive Rejection CP (AR-CP) achieves higher efficiency than standard CP, particularly when the target coverage $1 - \alpha$ is large. Since AR-CP employs a grid search to find the pair (α_0, α_1) that minimizes the expected prediction set size, we have the least-efficient baseline:

$$\begin{aligned} E_{\alpha_0, \alpha_1}(|C_{\alpha}^{\text{ARCP}}(x_i)| \mid x_i \in D_{\text{test}}) \\ = E_{\alpha}(|C_{\alpha}(x_i)| \mid x_i \in D_{\text{test}}), \end{aligned} \quad (27)$$

when $\alpha_1 = 0$ and $\alpha_0 = \alpha$. Thus, it suffices to show that for small α , there exist (α_0, α_1) satisfying Eq. 12 such that

$$\begin{aligned} E_{\alpha_0, \alpha_1}(|C_{\alpha}^{\text{ARCP}}(x_i)| \mid x_i \in D_{\text{test}}) \\ < E(|C_{\alpha}(x_i)| \mid x_i \in D_{\text{test}}). \end{aligned} \quad (28)$$

We can express the expected prediction set size under AR-CP and $E(|C_{\alpha}(x_i)| \mid x_i \in D_{\text{test}})$ as:

$$\begin{aligned} E_{\alpha_0, \alpha_1}(|C_{\alpha}^{\text{ARCP}}(x_i)| \mid x_i \in D_{\text{test}}) \\ = P(p^0(x_i) < q_{\alpha_0}^0, p^1(x_i) < q_{\alpha_1}^1) \\ + (E(|C_{\alpha}^{\text{ans}}(x_i)|) + \epsilon)P(p^1(x_i) < q_{\alpha_1}^1) \\ < E(|C_{\alpha}^{\text{ans}}(x_i)|) + P(p^0(x_i) < q_{\alpha}^0), \end{aligned} \quad (29)$$

where $\epsilon \in \mathbb{R}$ is a small correction term used to mitigate calibration shift in q_{α}^{ext} . $E(|C_{\alpha}^{\text{ans}}(x_i)|)$ is the average prediction set size of answers without the ‘Can’t answer’ label.

set $t = P(p^0(x_i) < q_{\alpha_0}^0) - P(p^0(x_i) < q_{\alpha}^0) > 0$ which will be close to $(\alpha - \alpha_0)r_{\text{correct}}$ when uncertainty measurement performs well and close to $\alpha - \alpha_0$ when uncertainty measurement performs bad. we can get Eq. 28 stand when:

$$\begin{aligned} t + (E(|C_{\alpha}(x_i)| \mid x_i \in D_{\text{test}}) + \epsilon)P(p^1(x_i) < q_{\alpha_1}^1) \\ < E(|C_{\alpha}(x_i)| \mid x_i \in D_{\text{test}}) \end{aligned} \quad (30)$$

Let $A := \mathbb{E}(|C_{\alpha}(x_i)| \mid x_i \in D_{\text{test}})$, $p := \mathbb{P}(p^1(x_i) < q_{\alpha_1}^1)$, and $t := \mathbb{P}(p^0(x_i) < q_{\alpha_0}^0) - \mathbb{P}(p^0(x_i) < q_{\alpha}^0)$. Then inequality Eq. 28 holds if and only if

$$\frac{t + \epsilon p}{1 - p} < A \iff p < \frac{A - t}{A + \epsilon}, \quad (31)$$

provided that $0 \leq p < 1$ and $t < A$. Since p is monotone in α_1 , there exists

$$\alpha_1^* := \inf \{ \alpha_1 : \mathbb{P}(p^1(x_i) < q_{\alpha_1}^1) \leq (A - t)/(A + \epsilon) \},$$

such that Eq. 28 holds for all $\alpha_1 \geq \alpha_1^*$.

Here $t \approx (\alpha - \alpha_0)r$, $r \in (r_{\text{correct}}, 1)$, so a smaller α_0 (i.e., stricter rejection threshold) increases t and tightens the bound $(A - t)/(A + \epsilon)$, requiring a smaller p to satisfy (31). Since $p = \mathbb{P}(p^1(x_i) < q_{\alpha_1}^1)$ increases with larger α_1 , There exists a critical value α_1^* beyond which The inequality naturally holds. In practice, $t < A$ almost always holds when α is small because A —the expected prediction set size—is typically much larger than $\alpha - \alpha_0$.

Coverage Guarantee of Weighted CP. Under the domain I.I.D. assumption and the explanation in Section 4.3, our estimated \hat{n}_{test}^k provides a valid approximation of the domain proportions in the testing data (Tibshirani et al., 2019). The coverage guarantee of Weighted CP under covariate shift follows directly from prior work (Tibshirani et al., 2019). By reweighting calibration samples according to the density ratio between the test and calibration distributions,

$$w(x) = \frac{p_{\text{test}}(x)}{p_{\text{cal}}(x)}, \quad (32)$$

The resulting weighted conformal predictor satisfies the asymptotic coverage guarantee:

$$\mathbb{P}_{(x,y) \sim P_{\text{test}}}(y \in C_{\alpha}(x)) \geq 1 - \alpha. \quad (33)$$

Coverage Guarantee of Resampled CP. Under the domain I.I.D. assumption, when the number of resampled calibration samples is sufficiently large, the resampled calibration data approximates the test domain distribution. Hence, calibration and test sets can be regarded as I.I.D., guaranteeing that the prediction set achieves the desired coverage.

B Experiment setting

B.1 Keywords for dividing datasets into different domains

We categorize the TriviaQA dataset into five domains according to distinctive lexical cues. Science & Art includes a broad range of scientific and artistic terminology, encompassing words such as physics, chemistry, biology, scientist, astronomy, mathematics, theory, planet, galaxy, telescope, microscope, quantum, molecule, enzyme, neuron, genome, DNA, equation, calculus, genetics, and biochemistry, along with artistic and literary terms like poet, poem, novel, novelist, painter, composer, symphony, art, gallery, museum, playwright, theatre, literature, ballet, opera, sonnet, sculpture, poetry, and sculptor. History includes keywords referring to historical eras, figures, and events, including dynasty, emperor, treaty, revolution, war, ancient, medieval, empire, king, queen, monarch, colonial, renaissance, pharaoh, roman, greek, battle, republic, treaties, and monarchy. Sports & Politics merges athletic and political expressions, including olympics, fifa, uefa, nba, mlb, nfl, nhl, footballer, coach, tournament, world cup, premier league, grand slam, stadium, athlete, league, match, finals, medal, goal, club, referee, umpire, arena,

president, prime minister, parliament, election, constitution, senate, cabinet, ministry, chancellor, congress, referendum, party, coalition, diplomacy, treaties, regime, governor, and minister. Entertainment focuses on mass media and performing arts, including film, movie, actor, actress, TV, television, oscars, academy award, grammy, emmy, director, screenplay, soundtrack, series, episode, band, album, song, box office, festival, orchestra, opera, musical, and concert. Finally, Geography covers geographic and geopolitical terminology such as capital, river, mount, mountain, province, city, country, island, lake, ocean, sea, peninsula, desert, continent, bay, harbor, border, valley, delta, latitude, longitude, archipelago, and coast.

We categorize the HotpotQA dataset into three semantic domains based on representative keywords. Science & Art contains terms that reflect both scientific and artistic concepts, including physics, chemistry, biology, scientist, astronomy, mathematics, theory, planet, galaxy, telescope, microscope, quantum, enzyme, neuron, genome, DNA, equation, calculus, genetics, biochemistry, poet, poem, painter, composer, symphony, gallery, playwright, ballet, opera, sculpture, poetry, and sculptor.

History focuses on terminology related to historical events, figures, and periods, including dynasty, emperor, treaty, revolution, war, ancient, medieval, empire, king, queen, monarch, colonial, renaissance, pharaoh, roman, greek, battle, republic, treaties, and monarchy. Entertainment & Geography combines cultural and locational entities, covering entertainment-related words such as film, movie, actor, actress, TV, television, oscars, academy award, grammy, emmy, director, screenplay, soundtrack, series, episode, band, album, song, box office, festival, orchestra, opera, musical, concert, and novel, as well as geographical expressions like capital, river, mount, mountain, province, city, country, island, lake, ocean, sea, peninsula, desert, continent, bay, harbor, border, valley, delta, latitude, and longitude.

B.2 Training and sampling detail

The learning rate is set to 1×10^{-4} , and the model is trained for three epochs at each step. During answer sampling, we use a temperature of 1.2, top- $p = 0.9$, top- $k = 100$, and a repetition penalty of 1.0.

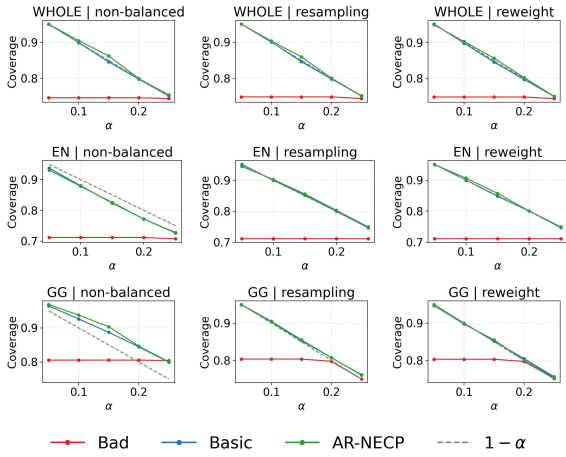


Figure 4: Coverage of the three CP methods under different quantiles under different distribution shifts and using different balancing methods. Results for Mistral7B on TriviaQA.

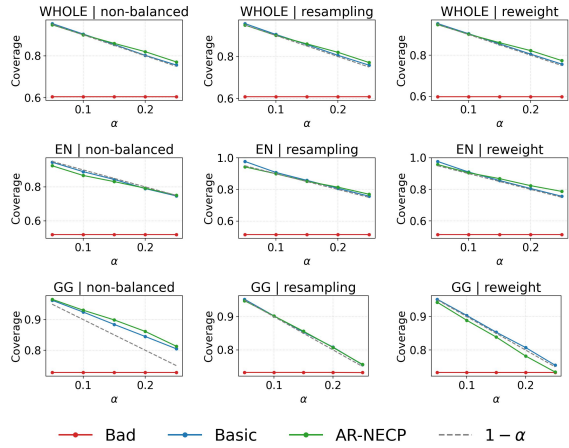


Figure 6: Coverage of the three CP methods under different quantiles under different distribution shifts and using different balancing methods. Results for Gemma7B on TriviaQA.

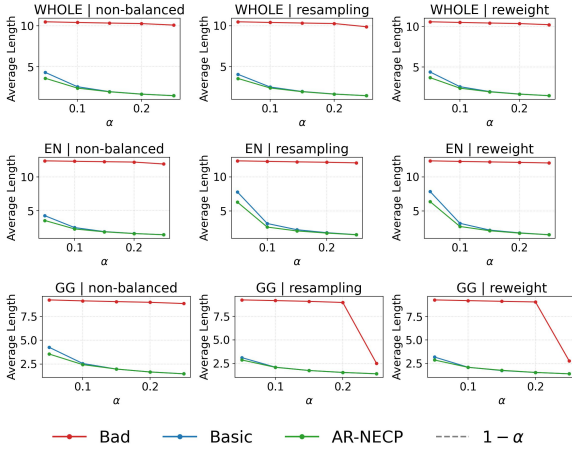


Figure 5: Average prediction set size of the three CP methods under different quantiles, under different distribution shifts, and using different balancing methods. Results for Mistral7B on TriviaQA.

C Additional experiments

C.1 Additional experiment on Main (RQ1 & RQ2)

The results of Gemma7B and Mistral7B are shown in Figure 4, 7, 5, 6. They show the same finding as the main text. The results of Llama8B, Gemma7B, and Mistral7B on HotpotQA are shown in Figures 8, 11, 9, 12, 13, and 10. It is worth noting that HotpotQA is a multi-hop reasoning dataset, which makes it challenging for LLMs to fine-tune effectively. This harder learning problem tends to make the model more resistant to catastrophic forgetting, resulting in minor shifts in the non-conformity score distribution. Therefore, in

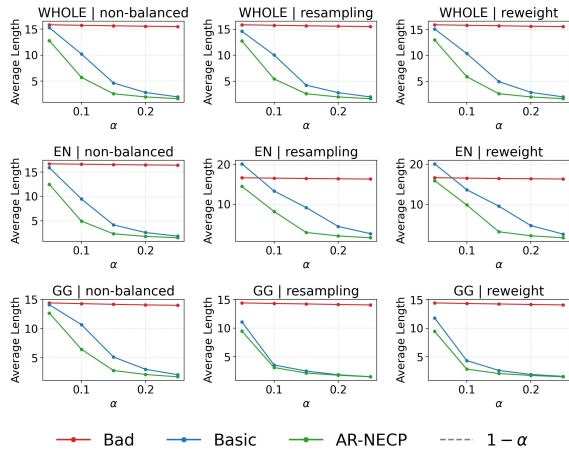


Figure 7: Average prediction set size of the three CP methods under different quantiles, under different distribution shifts, and using different balancing methods. Results for Gemma7B on TriviaQA.

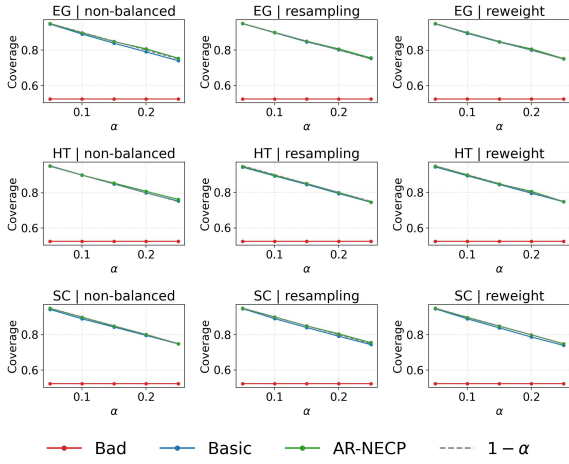


Figure 8: Coverage of the three CP methods under different quantiles under different distribution shifts and using different balancing methods. Results for Mistral7B on HotpotQA.

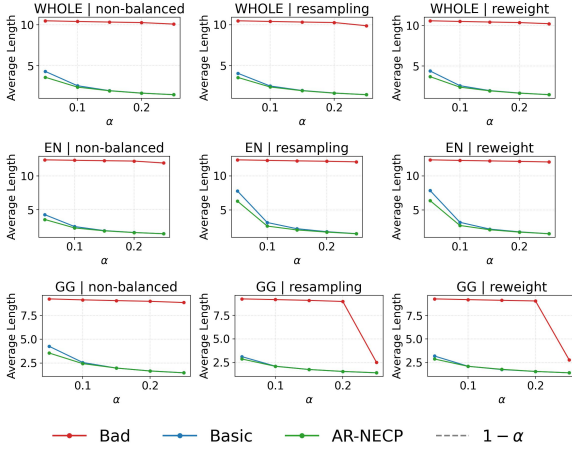


Figure 9: Average prediction set size of the three CP methods under different quantiles, under different distribution shifts, and using different balancing methods. Results for Mistral7B on HotpotQA.

this dataset, we mainly focus on evaluating the performance of our Adaptive Rejection CP.

In the MMLU dataset, which is a multiple-choice QA benchmark, it is relatively easier for LLMs to obtain correct answers through sampling. Therefore, we mainly focus on evaluating our method’s effectiveness in addressing distribution shifts. For efficiency analysis, the coverage results are presented in Figure 14, 16, 18, 15, 17, 19. From these results, we observe that our method effectively mitigates the distribution shift, while also improving efficiency. However, the improvement in efficiency is limited, as the task itself is relatively simple—the model only needs to choose from four possible answers.

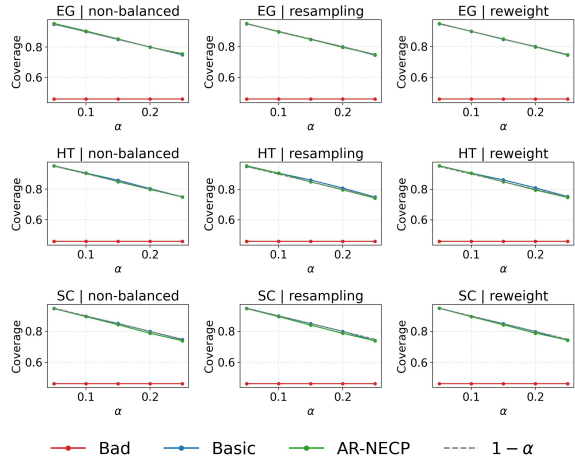


Figure 10: Coverage of the three CP methods under different quantiles under different distribution shifts and using different balancing methods. Results for Gemma7B on HotpotQA.

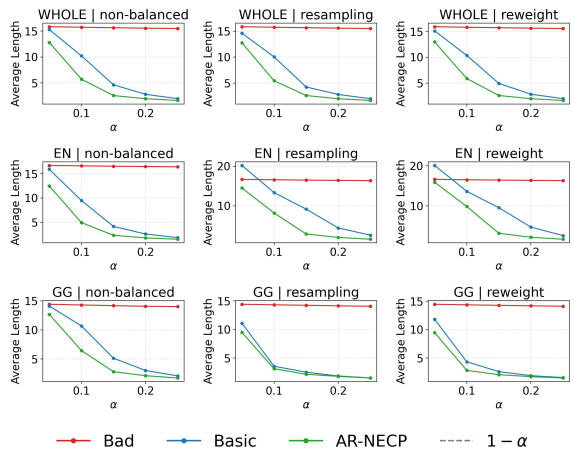


Figure 11: Average prediction set size of the three CP methods under different quantiles, under different distribution shifts, and using different balancing methods. Results for Gemma7B on HotpotQA.

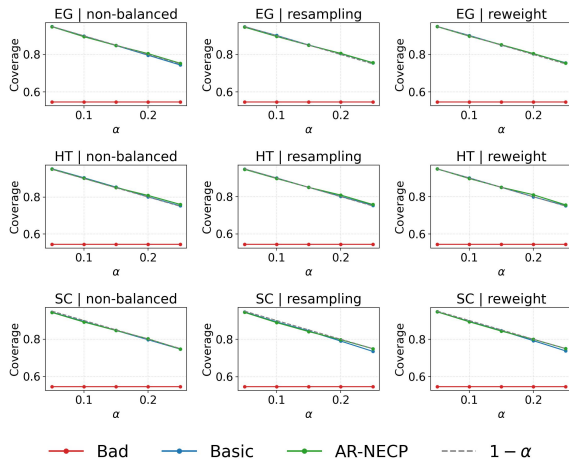


Figure 12: Coverage of the three CP methods under different quantiles under different distribution shifts and using different balancing methods. Results for Llama8B on HotpotQA.

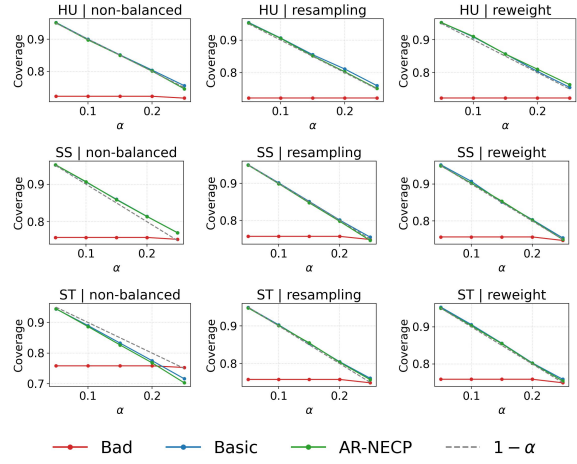


Figure 14: Coverage of the three CP methods under different quantiles under different distribution shifts and using different balancing methods. Results for Mistral7B on MMLU.

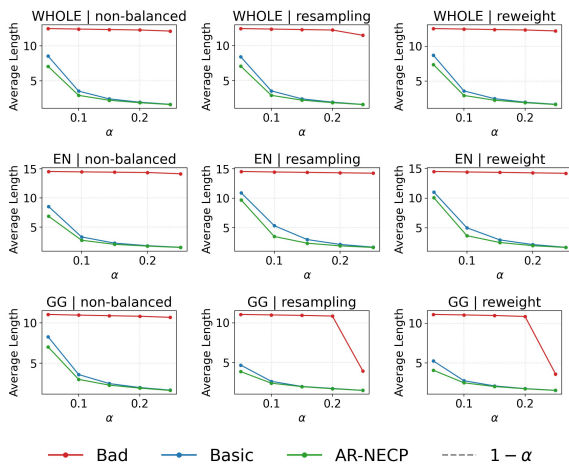


Figure 13: Average prediction set size of the three CP methods under different quantiles, under different distribution shifts, and using different balancing methods. Results for Llama8B on HotpotQA.

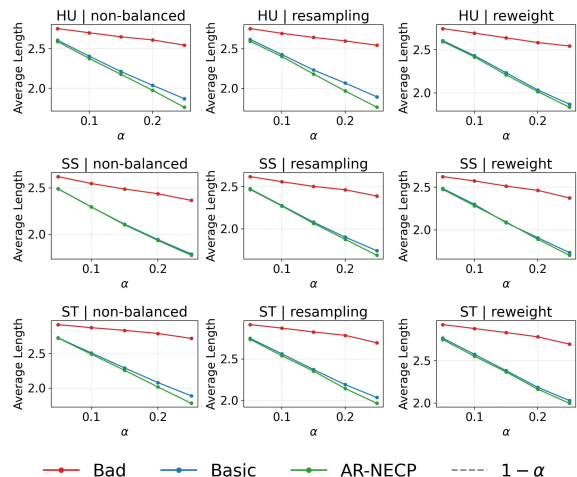


Figure 15: Average prediction set size of the three CP methods under different quantiles, under different distribution shifts, and using different balancing methods. Results for Mistral7B on MMLU.

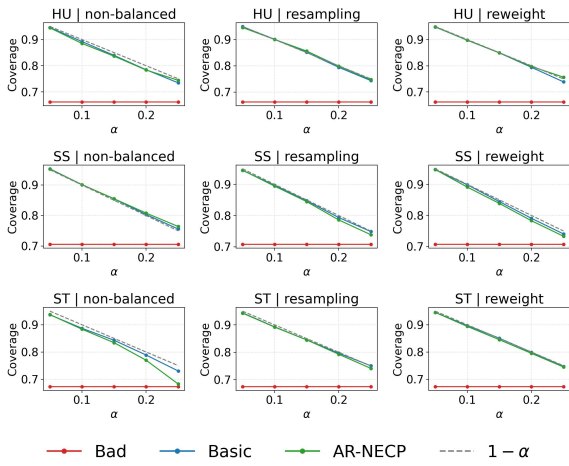


Figure 16: Coverage of the three CP methods under different quantiles under different distribution shifts and using different balancing methods. Results for Gemma7B on MMLU.

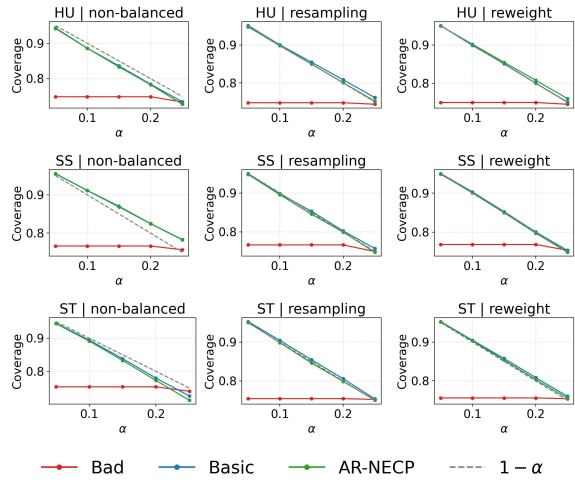


Figure 18: Coverage of the three CP methods under different quantiles under different distribution shifts and using different balancing methods. Results for Llama8B on MMLU.

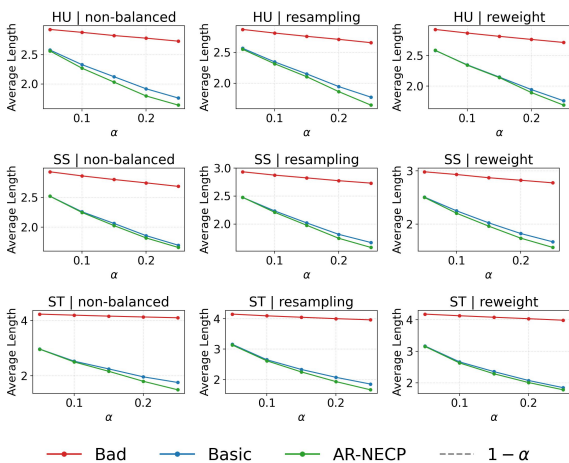


Figure 17: Average prediction set size of the three CP methods under different quantiles, under different distribution shifts, and using different balancing methods. Results for Gemma7B on MMLU.

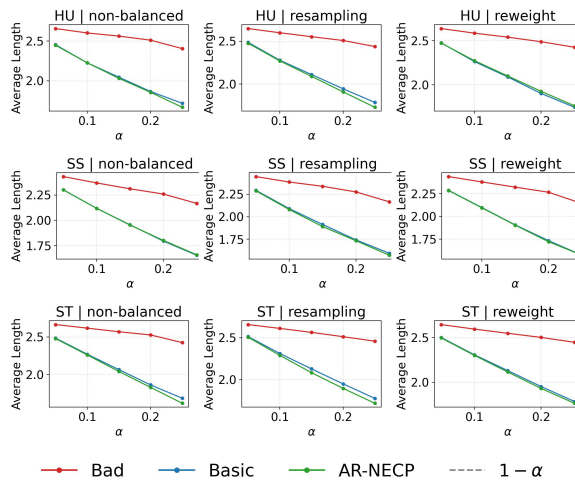


Figure 19: Average prediction set size of the three CP methods under different quantiles, under different distribution shifts, and using different balancing methods. Results for Llama8B on MMLU.

C.2 Additional experiment on RQ3

Here are the results from TriviaQA on the model Mistral7B and gemma-7b-it. Under most settings, NE performs better than MLP. Tables 6, 7, 5 are

Table 3: Comparison of prediction efficiency for AR-NECP (MLP) and AR-NECP (NE) under reweighting across different domains distribution and α values. These results come from Mistral7B.

Distribution	Method	$\alpha=0.05$	$\alpha=0.10$	$\alpha=0.15$	$\alpha=0.20$	$\alpha=0.25$
whole	AR-NECP (MLP)	4.16	2.45	1.91	1.66	1.46
	AR-NECP (NE)	3.67	2.35	1.89	1.63	1.45
en-d	AR-NECP (MLP)	8.06	3.05	2.14	1.71	1.47
	AR-NECP (NE)	6.95	2.78	2.09	1.70	1.47
gg-d	AR-NECP (MLP)	2.89	2.06	1.85	1.64	1.41
	AR-NECP (NE)	2.73	2.02	1.80	1.63	1.40

Table 4: Comparison of prediction efficiency for AR-NECP (MLP) and AR-NECP (NE) under reweighting across different domains distribution and α values. These results come from Gemma7B.

Distribution	Method	$\alpha=0.05$	$\alpha=0.10$	$\alpha=0.15$	$\alpha=0.20$	$\alpha=0.25$
whole	AR-NECP (MLP)	13.9	8.97	3.38	2.46	1.94
	AR-NECP (NE)	13.25	6.36	2.91	2.13	1.76
en-d	AR-NECP (MLP)	19.9	12.91	8.28	3.54	2.53
	AR-NECP (NE)	16.44	11.11	5.19	2.74	2.0
gg-d	AR-NECP (MLP)	10.66	3.42	2.43	1.92	1.73
	AR-NECP (NE)	8.82	2.72	2.08	1.78	1.61

the results from HotpoQA. Under most settings, NE performs better than MLP.

Table 5: Comparison of prediction efficiency for AR-NECP (NE on top) vs. AR-NECP (MLP) under reweighting across distributions and α for Llama8B (HotpotQA).

Distribution	Method	$\alpha=0.05$	0.10	0.15	0.20	0.25
sc-d	AR-NECP (NE)	9.48	4.41	2.84	2.14	1.84
	AR-NECP (MLP)	9.79	5.47	3.10	2.25	1.84
ht-d	AR-NECP (NE)	9.31	4.28	2.77	2.15	1.84
	AR-NECP (MLP)	9.14	5.04	3.06	2.20	1.73
eg-d	AR-NECP (NE)	10.09	5.70	3.11	2.30	1.87
	AR-NECP (MLP)	9.72	6.03	3.17	2.34	1.84

We do not present experiments on RQ3 for MMLU, as the improvement brought by our method is limited. Therefore, it would be meaningless to compare the performance between NE and MLP on this dataset.

C.3 Additional experiment on RQ4

Each model in triviaQA shares the same \hat{n}_{train}^k and n_{test}^k . Below are for HotpotQA in Table 8

Below are for MMLU in Table 9

D AI assistant usage

We used ChatGPT to assist with paper revision and coding.

Table 6: Comparison of prediction efficiency for AR-NECP (NE on top) vs. AR-NECP (MLP) under reweighting across distributions and α for Gemma7B (HotpotQA).

Distribution	Method	$\alpha=0.05$	0.10	0.15	0.20	0.25
sc-d	AR-NECP (NE)	9.00	4.58	2.86	2.12	1.75
	AR-NECP (MLP)	10.85	6.19	3.23	2.19	1.72
ht-d	AR-NECP (NE)	9.14	5.02	2.90	2.10	1.72
	AR-NECP (MLP)	9.75	6.05	3.05	2.16	1.74
eg-d	AR-NECP (NE)	8.66	4.38	2.69	1.98	1.68
	AR-NECP (MLP)	9.89	5.92	3.10	2.17	1.68

Table 7: Comparison of prediction efficiency for AR-NECP (NE on top) vs. AR-NECP (MLP) under reweighting across distributions and α for Mistral7B (HotpotQA).

Distribution	Method	$\alpha=0.05$	0.10	0.15	0.20	0.25
sc-d	AR-NECP (NE)	9.00	4.58	2.86	2.12	1.71
	AR-NECP (MLP)	10.85	6.19	3.23	2.19	1.72
ht-d	AR-NECP (NE)	9.14	5.02	2.90	2.10	1.72
	AR-NECP (MLP)	9.75	6.05	3.05	2.16	1.74
eg-d	AR-NECP (NE)	8.66	4.38	2.69	1.98	1.68
	AR-NECP (MLP)	9.89	5.92	3.10	2.17	1.68

Shift	sa	hd	eg
sa-d	0.08	0.15	0.19
hd-d	0.24	0.08	0.11
eg-d	0.19	0.22	0.03

Table 8: Mean relative error Δ^k between true and estimated domain counts under different test-time shifts. *sc-d*, *ht-d* and *eg-d* denote test sets dominated by the *sc*, *ht* and *eg* domains, respectively. Lower values are better.

Shift	st	hu	ss	ot
st-d	0.01	0.03	0.10	0.08
hu-d	0.05	0.01	0.09	0.13
ss-d	0.06	0.02	0.04	0.17
ot-d	0.09	0.06	0.17	0.05
whole-d	0.02	0.03	0.07	0.06

Table 9: Mean relative error Δ^k between true and estimated domain counts under different test-time shifts. *hu-d*, *ot-d*, *ss-d* and *st-d* denote test sets dominated by the *hu*, *ot*, *ss* and *st* domains, respectively. Lower values are better.