

# Multi-Source Test-Time Adaptation as Dueling Bandits for Extractive Question Answering

Hai Ye Qizhe Xie Hwee Tou Ng

Department of Computer Science, National University of Singapore  
{yehai, qizhex, nght}@comp.nus.edu.sg

## Abstract

In this work, we study multi-source test-time model adaptation from user feedback, where  $K$  distinct models are established for adaptation. To allow efficient adaptation, we cast the problem as a stochastic decision-making process, aiming to determine the best adapted model after adaptation. We discuss two frameworks: multi-armed bandit learning and multi-armed dueling bandits. Compared to multi-armed bandit learning, the dueling framework allows pairwise collaboration among  $K$  models, which is solved by a novel method named Co-UCB proposed in this work. Experiments on six datasets of extractive question answering (QA) show that the dueling framework using Co-UCB is more effective than other strong baselines for our studied problem<sup>1</sup>.

## 1 Introduction

Large language models (LLMs) can be fine-tuned or prompted with texts to achieve good performance in NLP tasks (Devlin et al., 2019; Brown et al., 2020; Ouyang et al., 2022). However, because of the unexpected distribution shift at test time, the effectiveness of LLMs can degenerate (Wang et al., 2021c). They may also generate outputs that are untrustworthy or toxic and fail to meet user expectations (Ouyang et al., 2022). One critical issue that we need to address is to improve the generalization ability of LLMs. Recent research on test-time adaptation (TTA) suggests a possible way to do this, by continually updating the deployed model with target data from an arbitrary test distribution (Wang et al., 2021a).

Interacting with users is important during test-time adaptation. First, user feedback allows the model to better align with humans (Stiennon et al., 2020; Ouyang et al., 2022). Users can directly teach the model to learn by interaction so that the

<sup>1</sup>Code of the paper is available at <https://github.com/oceanpy/Multi-source-TTA>.

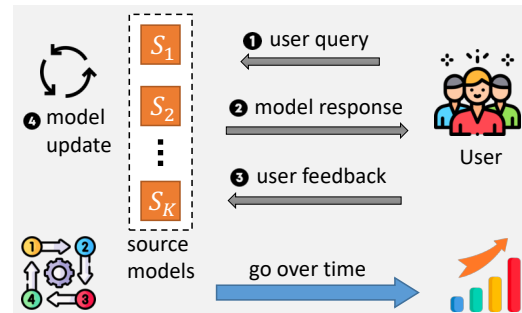


Figure 1: The illustration of multi-source test-time adaptation from user feedback studied in this work. Each model is trained from a distinct source domain. With unlabeled test data, models are adapted online from user feedback.

model can be better trained to follow human instructions and reduce the generation of toxic and harmful content. Besides, obtaining feedback from users can also reduce the cost of data annotation by experts, and the collected data will be more in line with the distribution of the users (Nguyen et al., 2017; Gao et al., 2022), which makes the adaptation more economical and effective.

Leveraging multiple learned models of tasks is also important for TTA. As in previous work, utilizing multiple known tasks helps the model better learn new tasks (or distributions), such as meta-learning (Hospedales et al., 2022) and multi-source domain adaptation (Ramponi and Plank, 2020). To take advantage of known tasks, compared to reusing task data, directly using their learned models has gained popularity recently (Pfeiffer et al., 2021; Wang et al., 2021b), which is much cheaper for online adaptation and has better data privacy protection (Kundu et al., 2020). Recent work on lightweight tuning empowers LLMs to store knowledge of a large number of tasks cheaply (Houlsby et al., 2019; Liu et al., 2021). Platforms like Huggingface (Wolf et al., 2019) also allow users to share locally trained models, promoting a large amount of knowledge stored as models in the cloud.

So, it has become more critical for TTA to adapt from multiple learned models of tasks.

Based on the above discussion, we propose to study an important but under-explored problem – multi-source test-time adaptation from user feedback – where  $K$  source models are given, each trained from a distinct source domain, to adapt to a new target domain (Figure 1). Previous work on leveraging multiple knowledge sources is to learn an ensemble (Guo et al., 2018; Ahmed et al., 2021), which means jointly accessing all the models is needed for training. Due to its high cost, it is not suitable for real-time updates required by TTA. In order to adapt efficiently, we turn this problem into a stochastic decision-making process that trades off model exploration and exploitation. We aim to determine the best adapted model that can perform well in the target domain.

We formulate the problem in two frameworks: multi-armed bandit learning and multi-armed dueling bandits (Kuleshov and Precup, 2014). Bandit learning samples one source model each time to receive binary feedback (👍 or 👎) (§4). However, it lacks collaboration among sources and can result in a sub-optimal adapted model. In order not to introduce too much cost, pairwise collaboration between models is explored in dueling bandits (Yue et al., 2009), where two distinct source models are chosen each time for dueling with user preference feedback (e.g., 🗳️/🗳️) (§5). A novel method, Co-UCB, is proposed to allow collaborative updates.

We choose to study the task of extractive question answering (QA), since there are large datasets in different domains that can be used (Fisch et al., 2019). More importantly, extractive QA is suitable for eliciting users to leave feedback, since the surrounding context around predicted answer spans can help users to verify the answers. Gao et al. (2022) has simulated user feedback for TTA in extractive QA, but not in the multi-source scenario. Following previous work (Gao et al., 2022), we simulate user feedback with the annotated answer spans. We conduct our simulation experiments on the MRQA benchmark (Fisch et al., 2019), where six domains of extractive QA are studied. We compare the two proposed frameworks to assess their effectiveness and reveal the differences. We also look into the effect of noisy preference feedback.

Our contributions in this work are as follows:

- We are the first to study multi-source test-time adaptation from user feedback;

- We propose a novel formulation of the problem as dueling bandits and solve it by a new method;
- Preference feedback is discussed for extractive QA for the first time; and
- Extensive experiments and analysis are conducted to verify our method.

## 2 Related Work

**Domain Adaptation.** Adapting from source domain(s) to target domain(s) is important for generalized machine learning (Ramponi and Plank, 2020). Test-time adaptation (TTA) attracts much attention recently, which adapts with the test data on the fly (Sun et al., 2020; Iwasawa and Matsuo, 2021; Wang et al., 2021a; Ye et al., 2022). TTA is more suitable for domain generalization since it needs no source data, which is different from unsupervised domain adaptation (UDA) (Ramponi and Plank, 2020; Ye et al., 2020). Multi-source DA is a more challenging problem than UDA, since it needs to determine suitable source knowledge for adaptation (Guo et al., 2018, 2020). Multi-source TTA has not been explored in NLP. Different from multi-source DA, multi-source TTA has no access to the source training data, and only the source models are given, which makes it more challenging to exploit useful sources for adaptation.

**Learning from Human Feedback.** Human feedback is a useful signal to refine the model outputs and adapt to new domains (Gao et al., 2022), follow human instructions (Ouyang et al., 2022), etc. Human feedback has been explored for different NLP tasks such as machine translation (Nguyen et al., 2017; Kreutzer and Riezler, 2019; Mendonça et al., 2021), semantic parsing (Lawrence and Riezler, 2018; Yao et al., 2020; Elgohary et al., 2021), document summarization (Gao et al., 2018; Stiennon et al., 2020), question answering (Kratzwald et al., 2020; Gao et al., 2022), and dialogue systems (Shuster et al., 2022). In particular, learning from human feedback has gained a lot of interests recently in the context of alignment of large language models (LLMs) (Stiennon et al., 2020; Ouyang et al., 2022; OpenAI, 2023). Fundamentally, alignment research is necessary and appealing from two aspects: (1) Alignment enables a model to go beyond supervised learning (Stiennon et al., 2020) (2) Alignment leads to a safer system (OpenAI, 2023). The proposed co-UCB could potentially be used for alignment in future work.

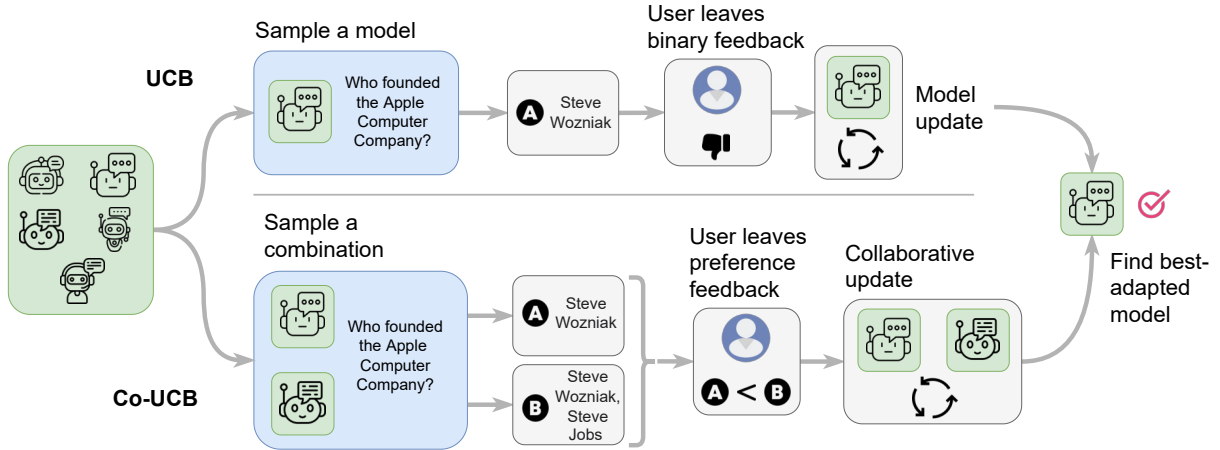


Figure 2: The illustration of UCB and Co-UCB for multi-armed bandit learning and dueling bandits respectively. For UCB, one model is sampled and receives binary feedback. For Co-UCB, two models are chosen each time and preference feedback is made by the user, then the models are updated.

### 3 Preliminaries

#### 3.1 Problem Definition

**Multi-source TTA.** We study multi-source test-time domain adaptation from  $K$  source models by interacting with users, where each model is trained from a distinct source domain. Test-time data is from a target domain  $X$  which is unlabeled. We focus on online adaptation where the test data  $\mathbf{x} \sim X$  comes as a stream and the model is continually updated with newly emerged test data at each time  $t^2$ . The parameters of each model at time  $t$  are inherited from the last time  $t - 1$ . At each time  $t$ , we obtain the test data  $\mathbf{x}_t \sim X$  which is the concatenation of the question  $\mathbf{q}_t$  and the passage  $\mathbf{d}_t$ . The prediction  $y_t$  is a pair of start and end positions over the passage denoted as  $\langle y_t^{(1)}, y_t^{(2)} \rangle$ . Following previous work (Rajpurkar et al., 2016), we use cross-entropy loss to update the model which is:

$$\mathcal{L}_t = -\left(\log(p_t^{(1)}) + \log(p_t^{(2)})\right)/2 \quad (1)$$

where  $p_t^{(1)}$  and  $p_t^{(2)}$  are the probabilities of the predicted start  $y_t^{(1)}$  and end  $y_t^{(2)}$  positions respectively.

**Motivation.** It is costly to learn an ensemble of  $K$  sources, since it has at least  $K$  times the training and inference costs, and even  $K$  times the parameters of a single source model (Guo et al., 2018; Ahmed et al., 2021). In order to adapt efficiently, we cast the problem as a stochastic decision-making process, where we aim to determine the best adapted model that can perform well in the target domain through user interaction.

<sup>2</sup>Note that our method can be easily applied in the offline scenario.

**Frameworks.** We first formulate the problem as multi-armed bandit learning (Kuleshov and Precup, 2014) and show how to solve it with Upper Confidence Bound (UCB) (Agrawal, 1995; Auer et al., 2002) (§4). We further discuss multi-armed dueling bandits to address the drawback of bandit learning, and propose a novel method Co-UCB (§5).

#### 3.2 Background

**Multi-Armed Bandit Learning.** The learning of multi-armed bandits (MAB) is a stochastic and iterative problem (Sui et al., 2018), which repeatedly selects a model from  $K$  sources. Each selected model receives a reward from the user. After  $T$  iterations, the goal of MAB is to minimize the cumulative regret compared to the best model:

$$\mathcal{R}^{\text{MAB}}(T) = \sum_{t=1}^T \left[ \mu^* - \mu(a_t) \right] \quad (2)$$

where  $a_t$  is the action at time  $t$  and  $\mu(a)$  is the expected reward of the action  $a$ .  $\mu^*$  is the expected reward of the best model.

**Multi-Armed Dueling Bandits.** In the multi-armed dueling bandits (MADB) problem, two distinct models are sampled among the  $K$  models (Yue et al., 2009). Also, the user needs to indicate a preference over the two selected models. In each comparison, a model  $a_i$  is preferred over  $a_j$  with the probability  $P(a_i > a_j)$ , which is equal to  $\epsilon(a_i, a_j) + 1/2$  where  $\epsilon(a_i, a_j) \in (-1/2, 1/2)$ . Suppose two models  $a_t^{(i)}$  and  $a_t^{(j)}$  are sampled at time  $t$ , and  $a^*$  is the overall best model. We define

---

**Algorithm 1** UCB for  $K$ -armed bandit learning

---

**Require:**  $K$  source models.

- 1:  $\bar{\mu} \leftarrow \mathbf{0}^K$ ,  $\mathbf{n} \leftarrow \mathbf{0}^K$ ,  $N \leftarrow 0$ ;
  - 2: **for**  $\mathcal{B}_t \in X$  **do**
  - 3:    $k \leftarrow \arg \max_j \bar{\mu}_j + \sqrt{2 \ln(N)/n_j}$ ;
  - 4:   Obtain the reward  $\mathbf{r}_t$  for model  $k$ ;
  - 5:    $\bar{\mu}_k \leftarrow (\bar{\mu}_k n_k + \mathbf{r}_t^\top \mathbf{r}_t) / (n_k + |\mathcal{B}_t|)$ ;
  - 6:    $n_k \leftarrow n_k + |\mathcal{B}_t|$ ;  $N \leftarrow N + |\mathcal{B}_t|$ ;
  - 7:   Update model  $k$  with loss  $\mathbf{r}_t^\top \mathcal{L}_t / |\mathcal{B}_t|$ ; //  $\mathcal{L}_t$  is as Eq. 1 shows.
  - 8: **end for**
  - 9: **Return:**  $k \leftarrow \arg \max_j \bar{\mu}_j + \sqrt{2 \ln(N)/n_j}$ .
- 

the cumulative regret at time  $T$  as:

$$\mathcal{R}^{\text{MADB}}(T) = \sum_{t=1}^T \left[ \epsilon(a^*, a_t^{(i)}) + \epsilon(a^*, a_t^{(j)}) \right] \quad (3)$$

which is a strong version discussed in Yue et al. (2009). It is the proportion of users who prefer the best model over the selected ones each time.

## 4 UCB for Bandit Learning

As illustrated by Figure 2, we apply UCB for multi-armed bandit learning, whose pseudo-code is shown in Algorithm 1.

**Action.** At each time  $t$ , the source model  $k$  is selected from  $K$  source models which maximizes  $\bar{\mu}_k + \sqrt{\frac{2 \ln(N)}{n_k}}$ , where  $\bar{\mu}_k$  represents the average reward obtained for the model  $k$  by attempting it for  $n_k$  times, and  $N$  is the number of all test data instances received so far.  $\sqrt{\frac{2 \ln(N)}{n_k}}$  represents the confidence interval to the action  $k$ , and a larger one means more uncertainty about the action, intending to explore the action more. As training proceeds, the policy becomes more confident about each action.

**Simulated Binary Feedback (👍 or 👎).** For each input, the selected model will first predict its answer, then the user leaves the feedback to the prediction. Here, we use binary feedback since it is simple enough for the user to provide and has often been used in previous work (Kratzwald et al., 2020; Gao et al., 2022). At each time  $t$ , a batch input  $\mathcal{B}_t \sim X$  is obtained for training, which is passed to the model  $k$  to obtain the predictions.

**Reward.** With a batch of predicted answers, the model  $k$  will obtain a vector of simulated reward  $\mathbf{r}_t \in \{0, 1\}^{|\mathcal{B}_t|}$  decided by the user. For each data instance in the batch, we follow Gao et al. (2022) to calculate the simulated reward by comparing

the predicted answer to the annotated span, where an index-wise exact match is used. If both the predicted start and end positions exactly match the annotated positions, the reward is 1; otherwise, 0.

**Model Update.** After obtaining the reward, the model  $k$  will be updated with a reward-enhanced loss, where the task-specific cross-entropy loss  $\mathcal{L}_t$  (in Eq. 1) will be multiplied by the reward  $\mathbf{r}_t$ .

**Inference.** After enough iterations, the best adapted model can be found to perform well in the target domain as line 9 of Algorithm 1 shows.

## 5 Collaborative UCB for Dueling Bandits

### 5.1 Co-UCB

**Motivation.** Since only one model is accessed each time in bandit learning, unlike ensemble learning (Guo et al., 2018), it cannot make use of the collaboration among sources during adaptation. To address such a drawback and not incur much extra training cost, we exploit the pairwise collaboration among  $K$  sources, where each time two distinct models will be sampled for joint learning. After adaptation, we also keep the best model for inference, to have the same cost as bandit learning.

Sampling pairs of models can be formulated as multi-armed dueling bandits (MADB) as discussed above. However, previous work on MADB only aims to determine the best model (Yue et al., 2009; Zoghi et al., 2014; Sui et al., 2017), so we further propose a novel method which is Collaborative UCB (Co-UCB) to let a pair of models collaborate, whose pseudo-code is presented in Algorithm 2, and illustrated by Figure 2.

**Action.** At each time  $t$ , with  $K$  source models, we construct  $C_2^K$  combinations for selection, where each combination is denoted by a pair of model indices  $\langle i, j \rangle$  ( $i < j$ ). The combination  $\langle i, j \rangle$  selected at time  $t$  should maximize  $(\bar{\mu}_i + \bar{\mu}_j) / 2 + \sqrt{2 \ln(N) / n_{i,j}}$ , where  $\bar{\mu}_i$  and  $\bar{\mu}_j$  are the average reward obtained up to time  $t$  of model  $i$  and  $j$  respectively, and  $n_{i,j}$  is the number of combinations  $\langle i, j \rangle$  explored so far.  $N$  is the total number of test data instances received until now.

Take model  $i$  for example. The reward of exploiting model  $i$  represents how well model  $i$  can beat the other models during dueling. The average reward  $\bar{\mu}_i$  is calculated as follows:

$$\bar{\mu}_i = \frac{\sum_{k=1, k \neq i}^K r_{i,k}}{\sum_{k=1, k \neq i}^K n_{i,k}} \quad (4)$$

where  $r_{i,k}$  denotes the overall reward that the model



---

**Algorithm 2** Co-UCB for  $K$ -armed dueling bandits

---

**Require:**  $K$  source models.

- 1:  $\bar{\boldsymbol{\mu}} \leftarrow \mathbf{0}^K, \mathbf{n} \leftarrow \mathbf{0}^{K \times K}, N \leftarrow 0;$
  - 2: **for**  $\mathcal{B}_t \in X$  **do**
  - 3:    $\langle i, j \rangle \leftarrow \arg \max_{i, j: i < j} (\bar{\mu}_i + \bar{\mu}_j)/2 + \sqrt{2 \ln(N)/n_{i,j}};$
  - 4:   Obtain the rewards  $\mathbf{r}_i$  and  $\mathbf{r}_j$  for model  $i$  and  $j$  respectively as in Eq. 5;
  - 5:    $\bar{\mu}_i \leftarrow (\bar{\mu}_i \sum_k n_{i,k} + \mathbf{r}_i^\top \mathbf{r}_i) / (\sum_k n_{i,k} + |\mathcal{B}_t|);$
  - 6:    $n_{i,j} \leftarrow n_{i,j} + |\mathcal{B}_t|;$
  - 7:    $\bar{\mu}_j \leftarrow (\bar{\mu}_j \sum_k n_{j,k} + \mathbf{r}_j^\top \mathbf{r}_j) / (\sum_k n_{j,k} + |\mathcal{B}_t|);$
  - 8:    $n_{j,i} \leftarrow n_{j,i} + |\mathcal{B}_t|;$
  - 9:    $N \leftarrow |\mathcal{B}_t|;$
  - 10:   Update the models  $i, j$  by loss  $\mathcal{L}_t$  as in Eq. 7;
  - 11: **end for**
  - 12: **Return:**  $k \leftarrow \arg \max_j \bar{\mu}_j + \sqrt{2 \ln(2N) / \sum_k n_{j,k}}.$
- 

$i$  received by dueling with model  $k$  and  $n_{i,k}$  denotes the number of times model  $i$  duels with model  $k$ .

In each selection, to calculate the average reward  $(\bar{\mu}_i + \bar{\mu}_j)/2$  for the combination  $\langle i, j \rangle$ , we expect  $\langle i, j \rangle$  to be the most worthy action (exploration-and-exploitation trade-off), where  $i$  and  $j$  can mostly beat the rest of the models, which means they are the two strongest models among the  $K$  sources so that they can better collaborate to improve them.

**Simulated Preference Feedback. (e.g., ②/③)** Since for each input, the user will receive two answer candidates instead of one, the binary feedback used in bandit learning is not directly applicable. Rather than creating new forms of user interaction, we apply preference feedback (Christiano et al., 2017; Gao et al., 2018; Ouyang et al., 2022) when faced with multiple candidates. Since there are only two candidates, leaving preference feedback will be as simple as binary feedback.

For the chosen models  $i$  and  $j$  at time  $t$ , the batch of input  $\mathcal{B}_t \sim X$  will be given to them independently, to obtain the predicted answer spans. Then the users need to compare the two predictions to indicate a preference, where the more accurate answer should be picked out.

**Reward.** For each data instance in the batch, the reward  $r \in \{0, 1\}$ .  $r = 1$  means the user prefers one answer from the two candidates;  $r = 0$  means the user has no preference – either the two answers are both good or none is good. This is a strict measurement for preference since the answers without

preference are discarded.

To simulate the preference, we calculate the quality score of the predicted answers against the annotated spans, where the answer with a higher score would be preferred or no preference is made if the scores are the same. We use the index-wise F1 value as the quality score, which calculates the F1 score over the predicted indices and the annotated indices, so the score is continuous from  $[0, 1]$ .

For the batch of input  $\mathcal{B}_t$ , the quality score for the model  $i$  and  $j$  is denoted as a vector  $\mathbf{s}_i$  and  $\mathbf{s}_j$  respectively. The rewards  $\mathbf{r}_i$  and  $\mathbf{r}_j$  for the model  $i$  and  $j$  respectively are obtained by:

$$\mathbf{r}_i = \mathbf{s}_i > \mathbf{s}_j; \quad \mathbf{r}_j = \mathbf{s}_j > \mathbf{s}_i \quad (5)$$

where  $\mathbf{r}_i$  and  $\mathbf{r}_j$  are one-hot vectors.

**Collaborative Model Update.** After obtaining the rewards, we perform collaborative model updates. If there is one preferred model, then it will be regarded as the teacher, and its prediction will be used to jointly update the two models. With the predictions from model  $i$  and  $j$ , we obtain the better one  $\langle \mathbf{y}_t^{*(1)}, \mathbf{y}_t^{*(2)} \rangle$ , each as a vector, as:

$$\langle \mathbf{y}_t^{*(1)}, \mathbf{y}_t^{*(2)} \rangle = \langle \mathbf{r}_i \mathbf{y}_t^{(i)(1)} + \mathbf{r}_j \mathbf{y}_t^{(j)(1)}, \mathbf{r}_i \mathbf{y}_t^{(i)(2)} + \mathbf{r}_j \mathbf{y}_t^{(j)(2)} \rangle \quad (6)$$

where  $\mathbf{y}_t^{*(1)}$  ( $\mathbf{y}_t^{*(2)}$ ) is a vector of the predicted start (end) positions from the preferred model for the batch of input  $\mathcal{B}_t$ .

Then we jointly update the two models by the loss:

$$\mathcal{L}_t = (\mathbf{r}_i + \mathbf{r}_j)^\top \mathcal{L}(\mathbf{y}_t^{*(1)}, \mathbf{y}_t^{*(2)}) / |\mathcal{B}_t| \quad (7)$$

Models updated in this way can better make use of the benefits of different source models during training, that is, when one model from a selected pair cannot predict a correct answer, the other one may make up for it by sharing its prediction.

**Inference.** After adaptation, there is no need to access a pair of models for inference anymore, so we just keep the best performing model by the method in line 12 of Algorithm 2.

## 5.2 Noise Simulation

Implicit preference feedback is naturally noisy since the preferred answer only needs to be better than the other and is not necessarily fully correct. However, the users may wrongly provide a preference in practice. Thus, we provide a pilot study to investigate the effect of such noise on adaptation performance.

There are three options that the user may pro-

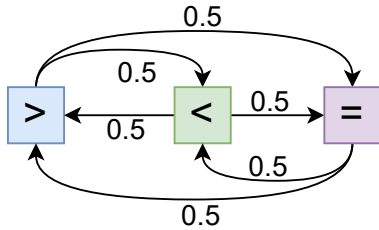


Figure 3: Transition probability for noise simulation, e.g., the correct option ‘>’ is corrupted into ‘<’ or ‘=’ equally by the same probability.

vide over two candidates, which are ‘>’, ‘<’ (with preference), and ‘=’ (no preference). For each data instance, we have a *noise rate* to randomly decide whether its feedback should be corrupted or not. If the feedback should be corrupted, then the correct option is changed to one of the remaining two options with a probability. In this work, we use the transition probabilities shown in Figure 3. We leave more complex transition probabilities to future work.

## 6 Experiments

### 6.1 Simulation Setup

**Dataset.** We conduct our experiments on MRQA (Fisch et al., 2019), which is a standard benchmark for domain generalization in extractive QA. We study six datasets (domains), which are SQuAD (Rajpurkar et al., 2016), HotpotQA (Yang et al., 2018), Natural Questions (NQ) (Kwiatkowski et al., 2019), NewsQA (Trischler et al., 2017), TriviaQA (Joshi et al., 2017), and SearchQA (Dunn et al., 2017), where each dataset forms a distinct domain. The training and development sets are used in our study.

**Setting of 5-source TTA.** To establish multi-source domain adaptation from the six domains, we set each dataset as the target domain and the remaining five datasets as the source domains. For each adaptation, the training set of the target domain is used as the unlabeled test data by discarding the labels, and the development set of the target domain is held out to evaluate the adaptation performance.

**Evaluation Metric.** We use F1 score to evaluate the performance on the held-out development set.

**Training Details.** We use the training set of each domain to train each source model, which follows the training details of Hu et al. (2020). We utilize XLMR (Conneau et al., 2020) and SpanBERT (Joshi et al., 2020) as the pre-trained lan-

guage model. In each multi-source domain adaptation, we set the batch size as 16 and use a constant learning rate of  $5e-7$ . The number of unlabeled test data instances is limited to 100K. The embedding layer is frozen to save computation. Experiments were conducted on one NVIDIA A100 GPU.

**Baselines.** We first present the results of the best source model without adaptation (**Best source**). Since our work is the first to study multi-source TTA, there are no existing baselines that address the same problem, so for the multi-source scenario, we mainly compare the two frameworks discussed above. **UCB** addresses the problem as bandit learning from binary feedback. **Co-UCB** is for dueling bandits from simulated preference feedback<sup>3</sup>.

We further compare to single-source TTA which has been studied in Gao et al. (2022). We first find the best source model before adaptation by evaluating each model on the held-out development set, then adapt the best source model from simulated binary feedback following the method of Gao et al. (2022). This baseline is denoted as **Bandit**.

### 6.2 Main Results

We first show the results of 5-source TTA in Figure 4. First, consistent with the findings of Gao et al. (2022), successful adaption is hard to see on TriviaQA and SearchQA just as the baseline of Bandit (Gao et al., 2022) indicates, so the following observations are based on the results of the remaining four target domains.

**Bandit and dueling bandits learning are effective in determining useful sources.** We find both UCB and Co-UCB can effectively improve the adaptation results compared to the best source without adaptation, which indicates that useful sources are found for adaptation during the training process.

**Leveraging multiple sources by Co-UCB performs the best.** Even without learning a  $K$ -ensemble model, Co-UCB still improves over the baselines by a large margin. Co-UCB can effectively utilize the benefits of different sources to outperform the Bandit baseline that adapts only one source model. On the contrary, UCB is not effective in making use of multiple sources, since it only achieves results similar to Bandit.

**Results during adaptation.** Figure 5 and Figure 11 plot the F1 scores vs. logging steps, where

<sup>3</sup>Due to the limitation of computing resources, we are not able to train an ensemble model of 5 source models, so we do not show the results of the ensemble model.

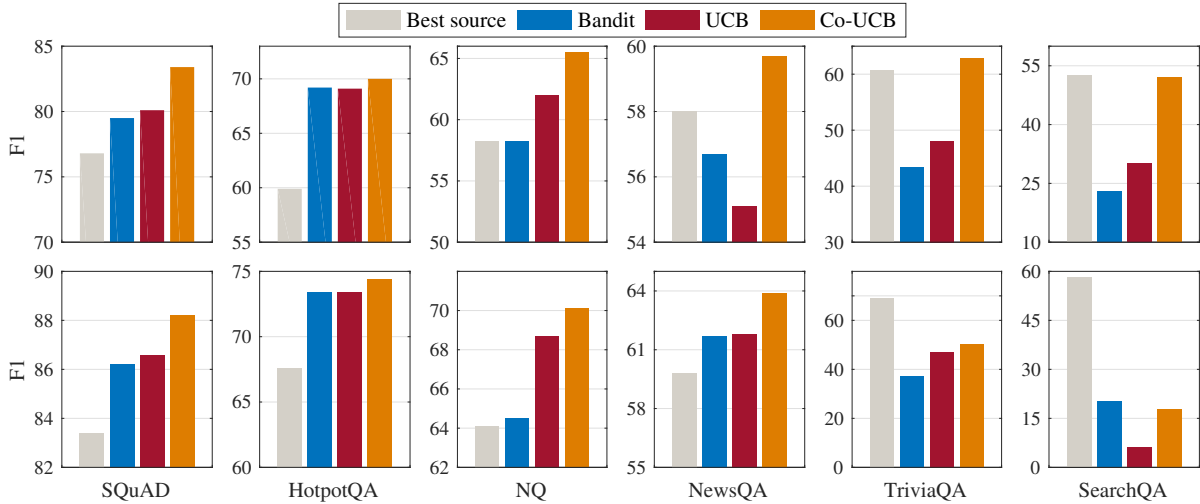


Figure 4: Results of 5-source test-time adaptation. Experiments for the adaptation methods are run three times with random seeds, and the average results are reported. The results of the first row are based on XLMR, and those of the second row are based on SpanBERT.

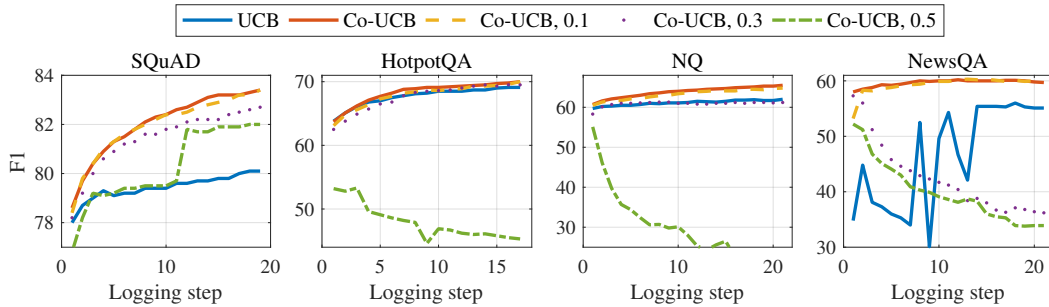


Figure 5: F1 performance on XLMR w.r.t. logging steps. 0.1, 0.3, and 0.5 are the noise rates. Results of TriviaQA and SearchQA are shown in Figure 11.

we find that UCB shows a large variance during adaptation on NewsQA, TriviaQA, and SearchQA, i.e., it is slow to find the best source model on these target domains. Co-UCB exhibits better performance and lower variance than UCB during adaptation.

### 6.3 Noise Simulation

We use the transition probabilities in Figure 3 to simulate the noise, e.g., flip the feedback from “>” to “=” with probability of 0.5. Results are presented in Figure 6. As the noise rate increases and more test data’s feedback is corrupted, the performance decreases. The result on XLMR drops dramatically with a noise rate larger than 0.5, but on SpanBERT, the rate of decline is not as fast. SpanBERT is a stronger LM than XLMR for extractive QA, so it is more robust to noisy feedback. As shown in Figure 5, a large noise rate (e.g., 0.5) will make the adaptation fail quickly on XLMR.

|          | Targets<br>Baselines | SQuAD<br>F1         | HotpotQA<br>F1      | NQ<br>F1            | NewsQA<br>F1        |
|----------|----------------------|---------------------|---------------------|---------------------|---------------------|
| XLMR     | UCB                  | 80.1 <sub>0.1</sub> | 69.1 <sub>0.1</sub> | 62.0 <sub>0.9</sub> | 55.1 <sub>0.7</sub> |
|          | Co-UCB               | 83.4 <sub>0.1</sub> | 70.0 <sub>0.1</sub> | 65.5 <sub>0.1</sub> | 59.7 <sub>0.1</sub> |
|          | w/o co.              | 79.3 <sub>0.1</sub> | 69.1 <sub>0.4</sub> | 61.4 <sub>0.2</sub> | 51.9 <sub>0.4</sub> |
| SpanBert | UCB                  | 86.6 <sub>0.1</sub> | 73.4 <sub>0.1</sub> | 68.7 <sub>0.3</sub> | 61.8 <sub>0.3</sub> |
|          | Co-UCB               | 88.2 <sub>0.0</sub> | 74.4 <sub>0.1</sub> | 70.1 <sub>0.3</sub> | 63.9 <sub>0.2</sub> |
|          | w/o co.              | 85.2 <sub>0.1</sub> | 73.0 <sub>0.1</sub> | 68.2 <sub>0.2</sub> | 61.1 <sub>0.4</sub> |

Table 1: Results of the ablation study. ‘w/o co.’ means removing the collaborative update in Co-UCB.

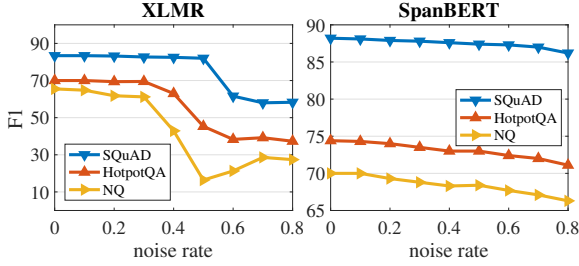


Figure 6: Adaptation results of Co-UCB w.r.t. noise rates using the noise transition in Figure 3.

|        | SQuAD       | HotpotQA    | NQ          | NewsQA      | TriviaQA    | SearchQA    |
|--------|-------------|-------------|-------------|-------------|-------------|-------------|
| UCB    | 4.77        | 3.99        | 3.55        | 2.36        | 2.03        | 2.16        |
| Co-UCB | <b>7.25</b> | <b>5.53</b> | <b>8.60</b> | <b>8.91</b> | <b>8.11</b> | <b>8.17</b> |

Table 2: Overall rewards ( $\times 10^4$ ) obtained during adaptation based on SpanBERT.

## 6.4 Further Analysis

**Ablation study.** Firstly, as Table 1 shows, without collaborative update, Co-UCB clearly degrades and it cannot compete with UCB. Considering that preference feedback is naturally noisy, Co-UCB without collaboration does not perform better than UCB.

**Overall rewards.** We calculate the overall rewards that UCB and Co-UCB obtain during adaptation in Table 2. The overall rewards are the sum of the cumulated rewards from each model. We observe that Co-UCB has a higher reward than UCB. In Co-UCB, for a certain input, when one model could not obtain the reward 1, the other model may make up for it by sharing, so that this model can also be updated and improved. The results support why Co-UCB performs better than UCB: higher rewards can better instruct the model to adapt.

**Case study.** We show the average reward and chosen count for each source model during adaptation in Figure 7. For Co-UCB, the models 0 and 2 are the best and second best model respectively, so its combination  $\langle 0, 2 \rangle$  is chosen most of the time.  $\langle 0, 1 \rangle$  and  $\langle 0, 3 \rangle$  are also often accessed since their payoff is close to the best one. However, UCB would mostly only focus on updating the best model which is model 0. As shown in Figure 8, UCB is able to quickly find the best source model (model 0), and the other models would be discarded without updating. For Co-UCB, since the models are dueling with each other, the changing of rewards behaves differently. The reward of model 0, 1, and 2 decreases, while the reward of model 3 and 4 increases, since dueling bandits learning

|  | 0     | 1     | 2     | 3     | 4     |
|--|-------|-------|-------|-------|-------|
|  | 0.494 | 0.226 | 0.263 | 0.136 | 0.055 |

(a.1) Average reward of each source model

|  | 0     | 1   | 2   | 3   | 4   |
|--|-------|-----|-----|-----|-----|
|  | 80028 | 336 | 464 | 176 | 128 |

(b.1) Chosen counts of each source model

|  | 0     | 1     | 2     | 3     | 4     |
|--|-------|-------|-------|-------|-------|
|  | 0.220 | 0.163 | 0.176 | 0.156 | 0.130 |

(a.2) Average reward of each source model

|   | 0     | 1     | 2     | 3     | 4     |
|---|-------|-------|-------|-------|-------|
| 0 |       | 0.191 | 0.198 | 0.188 | 0.175 |
| 1 | 14476 |       | 0.169 | 0.159 | 0.146 |
| 2 | 21264 | 6064  |       | 0.166 | 0.153 |
| 3 | 12256 | 4432  | 5440  |       | 0.143 |
| 4 | 7280  | 3184  | 3792  | 2944  |       |

(b.2) Chosen counts and average reward of each combination

Figure 7: Average rewards and chosen counts of UCB (a.1, b.1) and Co-UCB (a.2, b.2) obtained during adaptation on the target of HotpotQA.

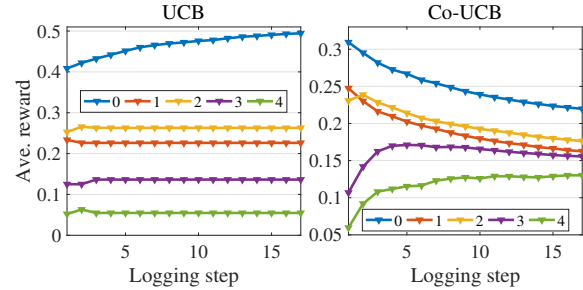


Figure 8: Average reward w.r.t. logging step for each model during adaptation on HotpotQA.

is a zero-sum game in general (Yue et al., 2009), where one model winning in dueling means another model loses. However, reward sharing happens in Co-UCB during training.

**Effects of the number of source domains.** From the results in Figure 9, we can see that the adaptation results have a very slight drop when the number of sources increases. No matter how the number of source models changes, Co-UCB still consistently performs better than UCB.

We also discuss the effects of preference feedback on UCB in Appendix A.

## 7 Conclusion

We present the first work on multi-source test-time adaptation from human feedback, where we cast it as an iterative decision-making problem. We first formulate it as multi-armed bandit learning. More importantly, to utilize pairwise collaboration, we further regard it as dueling bandits. Co-UCB is a novel method proposed in this work. Though we study online adaptation from the online data stream, our work can also be applied in offline



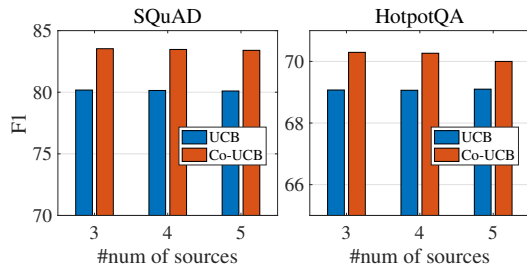


Figure 9: Effects of the number of source models on adaptation performance based on XLMR.

model refinement. For the offline setting, we do not update the model online but only update the policy for model selection when receiving user feedback each time. After collecting enough user feedback, we fine-tune the found best model offline with user feedback.

### Limitations

Learning an ensemble of multiple source models is expensive, especially for large language models. Hence, to adapt to the new target domain, we cast the problem as an iterative-decision making process. While our work reduces the model access frequency to 1 or 2 at each training step, continually updating the language model from a stream of test data is still costly. Future work can explore better methods for efficient optimization for a single LM. Besides, in some cases, the distribution of test data may change dynamically over the stream, but our work considers only the situation where the test data is from one specific distribution. More complex cases of test distribution can be studied in future work.

### Acknowledgements

This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG-RP-2018-007 and AISG2-PhD-2021-08-016[T]). The computational work for this article was partially performed on resources of the National Supercomputing Centre, Singapore (<https://www.nsc.sg>).

### References

Rajeev Agrawal. 1995. Sample mean based index policies by  $O(\log n)$  regret for the multi-armed bandit problem. *Advances in Applied Probability*, 27(4):1054–1078.

Sk Miraj Ahmed, Dripta S. Raychaudhuri, Sujoy Paul, Samet Oymak, and Amit K. Roy-Chowdhury. 2021. [Unsupervised multi-source domain adaptation without access to source data](#). In *CVPR*.

Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. [Finite-time analysis of the multiarmed bandit problem](#). *Mach. Learn.*

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martić, Shane Legg, and Dario Amodei. 2017. [Deep reinforcement learning from human preferences](#). In *NeurIPS*.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *ACL*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Matthew Dunn, Levent Sagun, Mike Higgins, V. Ugur Güney, Volkan Cirik, and Kyunghyun Cho. 2017. [SearchQA: A new q&a dataset augmented with context from a search engine](#). *ArXiv preprint*, abs/1704.05179.

Ahmed Elgohary, Christopher Meek, Matthew Richardson, Adam Fourney, Gonzalo Ramos, and Ahmed Hassan Awadallah. 2021. [NL-EDIT: Correcting semantic parse errors through natural language interaction](#). In *NAACL*.

Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. MRQA 2019 shared task: Evaluating generalization in reading comprehension. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*.

- Ge Gao, Eunsol Choi, and Yoav Artzi. 2022. [Simulating bandit learning from user feedback for extractive question answering](#). In *ACL*.
- Yang Gao, Christian M. Meyer, and Iryna Gurevych. 2018. [APRIL: interactively learning to summarise by combining active preference learning and reinforcement learning](#). In *EMNLP*.
- Han Guo, Ramakanth Pasunuru, and Mohit Bansal. 2020. [Multi-source domain adaptation for text classification via distancenet-bandits](#). In *AAAI*.
- Jiang Guo, Darsh J. Shah, and Regina Barzilay. 2018. [Multi-source domain adaptation with mixture of experts](#). In *EMNLP*.
- Timothy M. Hospedales, Antreas Antoniou, Paul Miccaelli, and Amos J. Storkey. 2022. [Meta-learning in neural networks: A survey](#). *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(9):5149–5169.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation](#). In *ICML*.
- Yusuke Iwasawa and Yutaka Matsuo. 2021. [Test-time classifier adjustment module for model-agnostic domain generalization](#). In *NeurIPS*.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: improving pre-training by representing and predicting spans](#). *Trans. Assoc. Comput. Linguistics*, 8:64–77.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *ACL*.
- Bernhard Kratzwald, Stefan Feuerriegel, and Huan Sun. 2020. [Learning a cost-effective annotation policy for question answering](#). In *EMNLP*.
- Julia Kreutzer and Stefan Riezler. 2019. [Self-regulated interactive sequence-to-sequence learning](#). In *ACL*.
- Volodymyr Kuleshov and Doina Precup. 2014. [Algorithms for multi-armed bandit problems](#). *arXiv preprint arXiv*, abs/1402.6028.
- Jogendra Nath Kundu, Naveen Venkat, Rahul M. V., and R. Venkatesh Babu. 2020. [Universal source-free domain adaptation](#). In *CVPR*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural Questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*.
- Carolin Lawrence and Stefan Riezler. 2018. [Improving a neural semantic parser by counterfactual learning from human bandit feedback](#). In *ACL*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *arXiv preprint arXiv*, abs/2107.13586.
- Vânia Mendonça, Ricardo Rei, Luísa Coheur, Alberto Sardinha, and Ana Lúcia Santos. 2021. [Online learning meets machine translation evaluation: Finding the best systems with the least human effort](#). In *ACL/IJCNLP*.
- Khanh Nguyen, Hal Daumé III, and Jordan L. Boyd-Graber. 2017. [Reinforcement learning for bandit neural machine translation with simulated human feedback](#). In *EMNLP*.
- OpenAI. 2023. [Gpt-4 technical report](#). *ArXiv*, abs/2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *NeurIPS*.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. [Adapterfusion: Non-destructive task composition for transfer learning](#). In *EACL*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *EMNLP*.
- Alan Ramponi and Barbara Plank. 2020. [Neural unsupervised domain adaptation in NLP—A survey](#). In *COLING*.
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, Morteza Behrooz, William Ngan, Spencer Poff, Naman Goyal, Arthur Szlam, Y-Lan Boureau, Melanie Kambadur, and Jason Weston. 2022. [Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage](#). *CoRR*, abs/2208.03188.

- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F. Christiano. 2020. [Learning to summarize with human feedback](#). In *NeurIPS*.
- Yanan Sui, Vincent Zhuang, Joel W. Burdick, and Yisong Yue. 2017. [Multi-dueling bandits with dependent arms](#). In *UAI*.
- Yanan Sui, Masrour Zoghi, Katja Hofmann, and Yisong Yue. 2018. [Advancements in dueling bandits](#). In *IJCAI*.
- Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei A. Efros, and Moritz Hardt. 2020. [Test-time training with self-supervision for generalization under distribution shifts](#). In *ICML*.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. [NewsQA: A machine comprehension dataset](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP*.
- Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno A. Olshausen, and Trevor Darrell. 2021a. [Tent: Fully test-time adaptation by entropy minimization](#). In *ICLR*.
- Xinyi Wang, Yulia Tsvetkov, Sebastian Ruder, and Graham Neubig. 2021b. [Efficient test time adapter ensembling for low-resource language varieties](#). In *Findings of the ACL: EMNLP*.
- Xuezhi Wang, Haohan Wang, and Diyi Yang. 2021c. [Measure and improve robustness in NLP models: A survey](#). *ArXiv preprint*, abs/2112.08313.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *CoRR*, abs/1910.03771.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *EMNLP*.
- Ziyu Yao, Yiqi Tang, Wen-tau Yih, Huan Sun, and Yu Su. 2020. [An imitation game for learning semantic parsers from user interaction](#). In *EMNLP*.
- Hai Ye, Yuyang Ding, Juntao Li, and Hwee Tou Ng. 2022. [Robust question answering against distribution shifts with test-time adaptation: An empirical study](#). In *Findings of the ACL: EMNLP 2022*.
- Hai Ye, Qingyu Tan, Ruidan He, Juntao Li, Hwee Tou Ng, and Lidong Bing. 2020. [Feature adaptation of pre-trained language models across languages and domains with robust self-training](#). In *EMNLP*.
- Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. 2009. [The k-armed dueling bandits problem](#). In *COLT*.
- Masrour Zoghi, Shimon Whiteson, Rémi Munos, and Maarten de Rijke. 2014. [Relative upper confidence bound for the k-armed dueling bandit problem](#). In *ICML*.

## A Effects of Preference Feedback on UCB

In the main content of this paper, we use preference feedback for Co-UCB, since each time the user has a pair of predictions to provide feedback. For UCB, the user only has one candidate to leave feedback, so we use binary feedback.

Here, we further study how preference feedback would affect the performance of UCB. To enable preference feedback, for each input data instance, the model first generates its top two predictions (to be comparable to Co-UCB), then the user needs to provide preference feedback to the two candidates. We follow the same procedure of Co-UCB to simulate preference feedback for UCB.

The results are presented in Figure 10. As we can see, UCB with preference feedback improves over UCB with binary feedback in some cases (not a consistent improvement), since top two predictions give the user more choices to select a good label. However, UCB with preference feedback cannot compete with Co-UCB. Co-UCB aims to leverage the benefits of different source models instead of the model’s top several predictions, which is different from UCB with preference feedback. Similar to UCB with binary feedback, UCB with preference also lacks collaboration among source models, since the top two predictions, though expanding the options to select a good label, are just from one model. This finding further demonstrates the effectiveness and importance of leveraging multiple source models during test-time adaptation.

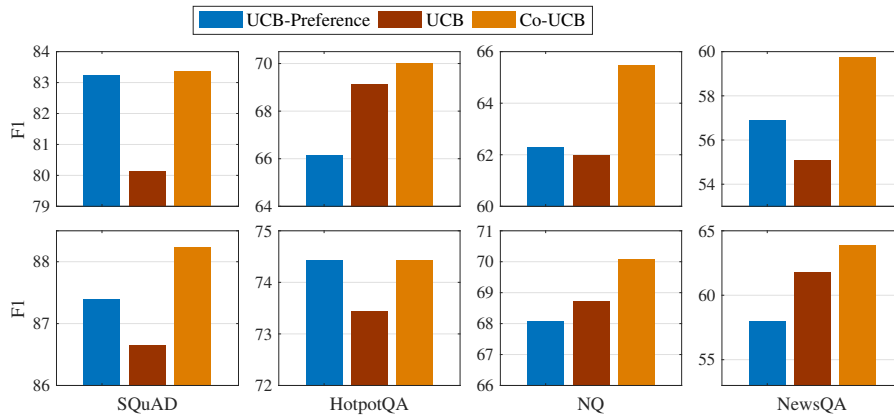


Figure 10: The results of UCB with preference feedback (i.e., UCB-Preference) with comparison to UCB and Co-UCB. The first row is based on XLMR, and the second row is based on SpanBERT. Experiments are run three times with random seeds, and the average results are reported.

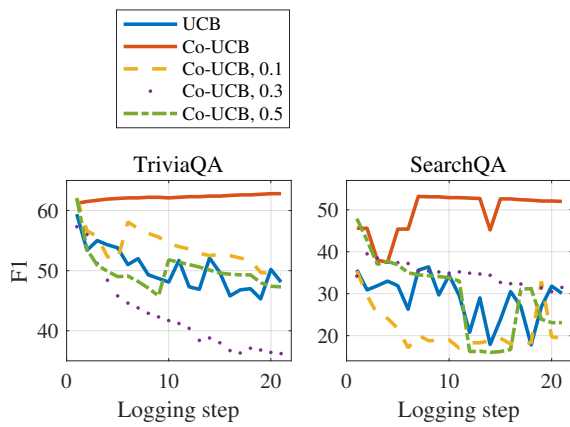


Figure 11: F1 performance on XLMR w.r.t. logging steps. 0.1, 0.3, and 0.5 are the noise rates.



## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

- A1. Did you describe the limitations of your work?  
*see the section of Limitations in the paper*
- A2. Did you discuss any potential risks of your work?  
*Our paper is about multi-source test-time adaptation, it doesn't deal with any harmful inputs and won't generate any harmful outputs, either.*
- A3. Do the abstract and introduction summarize the paper's main claims?  
*In the last of the section 1*
- A4. Have you used AI writing assistants when working on this paper?  
*We used Grammarly to check the grammar of the whole paper.*

### B Did you use or create scientific artifacts?

*Left blank.*

- B1. Did you cite the creators of artifacts you used?  
*No response.*
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?  
*No response.*
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?  
*No response.*
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?  
*No response.*
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?  
*No response.*
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.  
*No response.*

### C Did you run computational experiments?

*Section 5.*

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?  
*Section 5.1*

---

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.*

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

*Section 5.1*

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

*Appendix B*

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

*Section 5.1*

**D  Did you use human annotators (e.g., crowdworkers) or research with human participants?**

*Left blank.*

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

*No response.*

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

*No response.*

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

*No response.*

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

*No response.*

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

*No response.*