

Continual Machine Reading Comprehension via Uncertainty-aware Fixed Memory and Adversarial Domain Adaptation

Zhijing Wu^{1,2}, Hua Xu^{1,2*}, Jingliang Fang^{1,3} and Kai Gao³

¹State Key Laboratory of Intelligent Technology and Systems, Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China

²Beijing National Research Center for Information Science and Technology(BNRist)

³School of Information Science and Engineering, Hebei University of Science and Technology

wuzj18@mails.tsinghua.edu.cn

xuhua@tsinghua.edu.cn

Abstract

Continual Machine Reading Comprehension aims to incrementally learn from a continuous data stream across time without access the previous seen data, which is crucial for the development of real-world MRC systems. However, it is a great challenge to learn a new domain incrementally without catastrophically forgetting previous knowledge. In this paper, MA-MRC, a continual MRC model with uncertainty-aware fixed Memory and Adversarial domain adaptation, is proposed. In MA-MRC, a fixed size memory stores a small number of samples in previous domain data along with an uncertainty-aware updating strategy when new domain data arrives. For incremental learning, MA-MRC not only keeps a stable understanding by learning both memory and new domain data, but also makes full use of the domain adaptation relationship between them by adversarial learning strategy. The experimental results show that MA-MRC is superior to strong baselines and has a substantial incremental learning ability without catastrophically forgetting under two different continual MRC settings.

1 Introduction

Recently, Machine Reading Comprehension (MRC) has attracted wide attention and achieved remarkable success when solving specific tasks in stationary environments, such as answering factual questions with wikipedia articles or answering narrative questions with web search logs (Seo et al., 2017; Seonwoo et al., 2020; Zhang et al., 2021; Wu and Xu, 2020). However, the answering scenario changes over time in real-world applications. For example, the dialog system should continuously adapt to new user requirements (Abujabal et al., 2018; Madotto et al., 2021). In this paper, we focus on one of the most typical scenario changes for MRC tasks: the domain data shift. Existing




Figure 1: Illustration of a continual MRC task.

stationary-trained MRC systems are usually trained with in-domain data but are applied to new domain data (Fisch et al., 2019). Therefore, it is necessary to build a non-stationary MRC that continually learns with incremental domain data.

We formulate such a challenging problem as *Continual MRC* task, which is required to incrementally learn over sequential domain data and perform well on all seen domain data. Figure 1 illustrates the incremental learning and testing processing. In this scenario, the MRC system can only trained on the latest coming domain data without access the previous seen data. To tackle this issue, if we directly fine-tune the MRC model on each new incoming domain, the performance on earlier domains will significantly drop (Su et al., 2020). Another naive approach is to retrain the whole MRC model from scratch, but it is costly and time-consuming. Hence it is a great challenge for incrementally learning without largely forgetting previously acquired knowledge.

Existing studies for continual MRC can mainly be divided into three categories. The first class is model expansion techniques that design domain-individual classifier for each in coming domain (Su et al., 2020). However, it is expensive and unpractical in real-world. The second class borrows the regularization idea, which utilizes an ad-

*Hua Xu is the corresponding author.

ditional loss term to aid knowledge consolidation when learning new domains. For example, [Su et al. \(2020\)](#) added a penalty that restricts the change of important parameters to prevent forgetting previous knowledge. The third class is episodic memory based methods. For example, [de Masson d’Autume et al. \(2019\)](#) introduced a key-value memory module that stored previously seen examples for sparse experience replay and gradient-based local adaptation. [Abujabal et al. \(2018\)](#) proposed template-based Never-Ending KB-QA that learned new templates by capturing new syntactic structures with a semantic similarity function between questions and user feedback.

However, the above methods still have limitations, mainly including two aspects. On the one hand, to prevent catastrophic forgetting, these methods only design consistent constraints of model output or gradient for previous and new domains, while ignoring the domain adaptation relationship between them. However, transfer learning can help the MRC model generalize to other domains. On the other hand, the memory update strategy for continual MRC is limited. Some previous work stores fixed examples for each incoming domain. It greatly grows the number of samples kept in memory and leads to expensive costs. Other methods that limit the maximum number of memory for all seen domains usually update the memory by random sampling, ignoring the forgotten degree of different samples. In fact, the continual model should pay more attention to samples that are more likely to be forgotten.

To handle the above limitations, this paper proposes MA-MRC, an incremental model that solves continual MRC task via Uncertainty-aware fixed Memory and Adversarial Domain Adaptation. Concretely, MA-MRC 1) introduces a fixed-size memory to store a small number of samples in previous domain, which are later periodically replayed when learning new domain; and 2) updates the memory with an uncertainty-aware strategy that takes the forgotten degree of previous data into account; 3) leverages an adversarial learning strategy to make full use of the domain adaptation relationship between different domains with a domain discriminator, so as to help generalize and avoid overfitting very small memorized examples. The intuition behind this is to mimic the human learning process that replays the memory and adapts to new domains.

The key contributions of this work are: (1) This paper proposes a continual MRC model, MA-MRC, which learns new domain data incrementally. (2) Applying uncertainty-aware Memory and Adversarial learning and to MRC model contributes to strong incremental learning ability. (3) The experimental results on two different continual MRC settings indicate that MA-MRC obtains good incremental learning ability without largely forgetting and significantly outperforms strong baselines.

2 Related Work

2.1 Continual Learning

Continual Learning (CL) mainly aims to overcome the catastrophic forgetting problem when learning on sequential new tasks incrementally ([French, 1999](#)). Existing work follows three directions: architectural, regularization, and memory-based approaches. The architectural methods change the network’s architecture and add task-specific parameters, e.g., Dynamically Expandable Network ([Yoon et al., 2018](#)) and Reinforced Continual Learning ([Xu and Zhu, 2018](#)). The regularization-based techniques aid knowledge consolidation when learning new tasks. For instance, EWC and Online EWC ([Kirkpatrick et al., 2017](#); [Schwarz et al., 2018](#)) slow down the learning of parameters important for previous tasks. The third class is to save previous samples and learn a new task with a forgetting loss defined on them, such as GEM ([Lopez-Paz and Ranzato, 2017](#)), A-GEM ([Chaudhry et al., 2019](#)), DER ([Buzzega et al., 2020](#)), and MER ([Riemer et al., 2019](#)).

CL in MRC. Few previous studies apply continual learning to MRC. [Su et al. \(2020\)](#) adapted EWC method and enlarged the MRC architecture when a new domain arrives. [Su et al. \(2020\)](#) added a penalty regularization that restricts the change of important parameters to prevent forgetting. [de Masson d’Autume et al. \(2019\)](#) and [Abujabal et al., 2018](#)) designed episodic memory based methods that store training samples from previously seen data, which are later rehearsed to learn new domains. In this paper, we solve the continual MRC problem of incrementally learning over sequential domains, and build our continual model based on the above memory-based and penalty regularization paradigm.

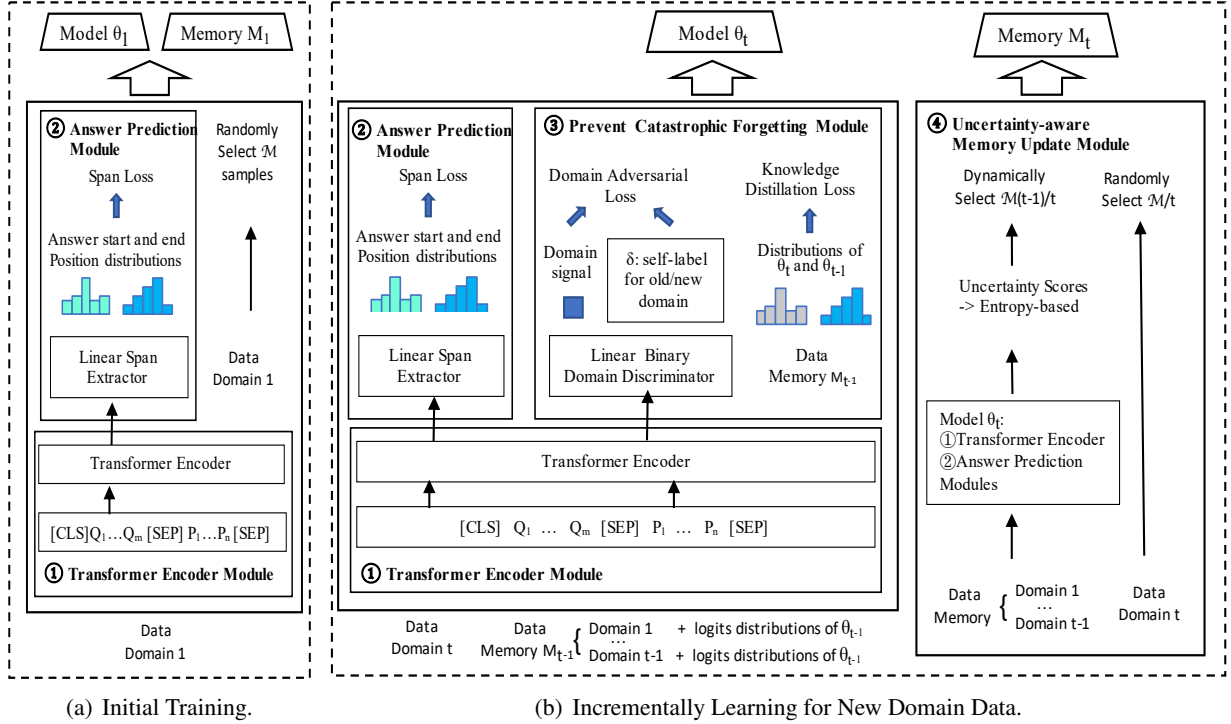


Figure 2: The overview of the proposed MA-MRC framework with sequential domain data.

2.2 Domain Adaptation

Domain adaptation aims to learn discriminative feature features and generalize to other new domains and is usually achieved by learning domain-invariant features (Ben-David et al., 2010; Ganin et al., 2016). For MRC under domain shift, previous methods (Wang et al., 2019; Cao et al., 2020; Lee et al., 2019) reduce domain discrepancy by a discriminator network that is trained to distinguish features of the target from source domains. However, recent work usually transfers the model from source to target. In this paper, we explore the domain adaptation under continual setting.

3 Proposed Method

3.1 Problem Statement

For the task of continual MRC, we assume that the learning framework has access to streams of MRC data from T different domains, denoted by $Data = \{D_1, D_2, \dots, D_T\}$. Each domain data, e.g. $D_t = \{P_n, Q_n, A_n\}_{n=1}^{|D_t|}$, consists of a series of <passage, question, answer> triples, where $|D_t|$ is the sample number of D_t . The MRC model is required to continually learn over each incoming domain data. More specifically, at each step t , the MRC model only observes D_t domain data, and is required to perform well on both the t -th domain data and

previous $t - 1$ domain data. Hence, after training on D_t , the model will be evaluated on all seen t domains. To make the MRC model perform well on previous tasks, during the continual learning process, a memory M is to set to store a small number of samples in previous domain data in a fixed-size memory $\mathcal{M} \ll |D_{old}|$.

This paper focuses on the span extractive MRC. Inspired by Su et al. (2020), we perform two different continual domain settings: CDA-C and CDA-Q, which define different domains according to paragraph type and question type, respectively.

3.2 Method overview

Figure 2 shows the overview of MA-MRC. As shown in Figure 2(a), we first train a backbone MRC model (Transformer Encoder and Answer Prediction modules) on the initial domain data and randomly select \mathcal{M} training samples as the initial memory M . Figure 2(b) shows the process of learning new domain data. When the t -th domain data is arising, MA-MRC synchronously 1) fine-tunes the backbone MRC model with both the t -th domain and memory data; 2) adversarially learns domain-invariant and transfer representations via a domain discriminator that distinguishes memory from current data, so as to generalize well on new domain and avoid overfitting very small memory samples;

Algorithm 1 Incremental Training for MA-MRC

Input:

T Training domain data $Data = \{D_1, D_2, \dots, D_T\}$

A memory M of size \mathcal{M}

Output:

The Continual MRC model θ_T

// *Initial Training*

- 1: Train the backbone MRC model θ_1 with D_1
 - 2: $M \leftarrow RandomSampling(D_1, \mathcal{M})$
// *Incrementally Learning for New Domain*
 - 3: **for** $t = 2, \dots, T$ **do**
 - 4: **Define** the domain discriminator θ_D
 - 5: **Incrementally learning** $\theta_t \leftarrow$ update θ_{t-1}
and θ_D with $\nabla \mathcal{L}_{final}$ on D_t, M_{t-1}
 - 6: Calculate need importance based on uncertainty for each sample in M_{t-1}
 - 7: $M_t \leftarrow WeightSampling(M_{t-1}, \frac{\mathcal{M}(t-1)}{t}) + RandomSampling(D_t, \frac{\mathcal{M}}{t})$
 - 8: **end for**
 - 9: **return** the final MA-MRC model θ_T ;
-

3) utilizes knowledge distillation to encourage the model to mimic its original responses for past domain samples. After finishing training at t -th step, we update the memory with an uncertainty-aware sampling strategy that focuses on remembering what the model most needs. More details about the training process are shown in Algorithm 1.

3.3 Initial Training

Note that there is only one domain data at the beginning, so the initial training with D_1 can be seen as a special case of formal extractive MRC tasks. Therefore, we build a normal backbone MRC model θ_1 (a standard BERT-MRC model with Transformer Encoder and Answer Prediction Modules) on the first domain data D_1 . We initialize a fixed-size memory M_1 to keep previous training samples that are periodically replayed while learning the new domain.

3.3.1 Prepare Backbone model

Transformer Encoder Module: First, a pre-trained L Transformer encoder blocks is used to convert the input sequence $S = [\langle CLS \rangle, Q, \langle SEQ \rangle, P, \langle SEQ \rangle]$ into contextual representations. Then, the last block output $H_L = \text{BERT}(S) \in \mathbb{R}^{(l) \times h}$ is taken as the contextual representation, where h is the hidden dimension of BERT, l is the sequence length.

Answer Prediction Module: A linear layer is applied on the contextual representation H_L to calculate the probability distribution of start and end positions of candidate answer:

$$p^{start} = \text{Softmax}(H_L W^s), \quad (1)$$

$$p^{end} = \text{Softmax}(H_L W^e), \quad (2)$$

where $W^s, W^e \in \mathbb{R}^h$ are learnable parameters.

Objective Function: The loss function of the backbone MRC is the cross-entropy:

$$\min_{\theta_{T,AP}} \mathcal{L}_{span} = -\log p_{y^s}^{start} - \log p_{y^e}^{end}, \quad (3)$$

where y^s and y^e are the ground-truth start and end indices of the corresponding sample, respectively.

3.3.2 Memory Initialization

For the first domain, to preserve the data distribution of the current domain as much as possible, we randomly select \mathcal{M} training samples from D_1 as the initial memory M_1 .

3.4 Learning for New Domain

When the t -th domain data is coming, we can access the current data D_t and the memory M_{t-1} . If we just finetune the backbone model, the model is hard to have good transfer ability and may overfit on the few memorized samples. Hence the MA-MRC utilizes the adversarial learning strategy to fully make use of the domain adaptation relationship between the previous and the current domain. After the t -th training step, we dynamically update memory M with an uncertainty-aware strategy to store the training samples that the model most needs to replay.

3.4.1 Memory-based Adversarial Training

To fully use the domain adaptation relationship, we adversarially learn the domain-invariant and transfer representations of the memory and current domain data. Inspired by domain adaptation theory, MA-MRC introduces a domain discriminator and build a two-player min-max game. The first player, a domain discriminator D , distinguishes memory data from the current new domain data. Here D is a basic binary discriminator that has a three linear layer followed by a sigmoid activation function. The second player, the Transformer Encoder T , aims to learn features that confuse D . We utilize empirical Maximum Mean Discrepancy as distance

measure to reduce the difference between marginal representation distributions and make them similar:

$$d = \text{MMD}(T(M_{t-1}), T(D_t)). \quad (4)$$

Finally, this learning procedure can be described by the following *minimax* game:

$$\min_{\theta_T} \max_{\theta_D} \mathcal{L}_{adv} = -\frac{1}{\mathcal{M}} \sum_{i=1}^{\mathcal{M}} \log D(T(P_i, Q_i)) - \frac{1}{|D_t|} \sum_{j=1}^{|D_t|} \log(1 - D(T(P_j, Q_j))) + d, \quad (5)$$

where \mathcal{M} and $|D_t|$ are the number of training samples from memory and the current domain.

3.4.2 Knowledge Distillation

It is obvious that a good continual model can learn a new domain data well while approximating the behavior observed in the old ones. Then, we leverage knowledge distillation constraints to encourage MA-MRC to mimic its original responses for past samples. Specifically, we seek to minimize the distance between the corresponding pre-softmax responses to preserve the knowledge about previous memory data:

$$\mathcal{L}_{KL} = KL(\text{logits}_{\theta_{t-1}}^{start} || \text{logits}_{\theta_t}^{start}) + KL(\text{logits}_{\theta_{t-1}}^{end} || \text{logits}_{\theta_t}^{end}), \quad (6)$$

where *logits* are the probability distribution before softmax and *KL* is the KL divergence.

3.4.3 Uncertainty-aware Memory Updating

Unlike other approaches that store fixed examples for each domain, we use a fixed memory for all domains to avoid memory growing. Therefore, when a domain arises, it is necessary to remove some old samples in memory to store new data. We design a dynamic sampling strategy that focuses on samples what the model most needs to replay. Based on existing trained MRC model θ_t , we first compute the uncertainty for each sample in memory by a unsupervised *Entropy-based* strategy:

$$u^{entropy} = \log p_{y^s}^{start} + \log p_{y^e}^{end}. \quad (7)$$

Then we calculate the gap between the above uncertainty and 1) the respective previous *best* uncertainty 2) the *average* uncertainty of all memory data, and normalize these metric differentials to create a probability distribution. For ease of exposition, we define these two normalization calculations as $norm^1$ and $norm^2$. Note that we will

Domains	#train	#test	p	q	a
CDA-Q setting					
what	37593	4749	118	9.9	3.5
which	4159	454	123.7	10.3	2.7
where	3291	433	120.4	8.2	3.1
when	5459	696	123.3	8.6	2.3
how	8124	1090	120.6	9.9	3
why	1201	151	123.9	9.6	8
other	19622	1938	118	11.6	2.7
who	8150	1059	126.1	9	2.8
CDA-C setting					
SQuAD	10000	10570	119.8	10.1	3.2
NaturalQA	10000	12836	152.4	9.2	4.3
HotpotQA	10000	5901	154.1	19.5	2.3
NewsQA	10000	4212	495.1	6.6	4
TriviaQA	10000	7785	674.2	13.2	1.6

Table 1: Dataset statistic of CDA-Q and CDA-C.

sample $\frac{\mathcal{M}}{t}$ data for each previously seen domain in memory with the above distribution separately. Finally, we get $\frac{\mathcal{M}(t-1)}{t}$ data from memory and randomly sample $\frac{\mathcal{M}}{t}$ data from current domain. In this way, the more forgotten a memory sample is, the more it will be retained.

3.4.4 Objective Function

When incrementally learning for a new domain, the span loss in answer prediction module \mathcal{L}'_{span} considers all current domain data D_t and the memory data M , while the KL loss only takes memory data M into account. Finally, the overall loss function at the t -th step is formulated as:

$$\min_{\theta_{T,AP}} \max_{\theta_D} \mathcal{L}_{final} = \mathcal{L}'_{span} + \mathcal{L}_{adv} + \mathcal{L}_{KL}. \quad (8)$$

4 Experiment

4.1 Continual MRC Datasets

Inspired by Su et al. (2020), this paper deals with two continual MRC under domain adaptation: CDA-C and CDA-Q. For CDA-C setting, we regard MRC datasets with different passage corpora (e.g., Wikipedia, News, and Web snippets) as different domains and choose five datasets: SQuAD 1.1 (Rajpurkar et al., 2016), HotpotQA (Yang et al., 2018), Natural Questions (Trischler et al., 2017), NewsQA (Trischler et al., 2017), and TriviaQA (Joshi et al., 2017). Due to computational limits, we use the curated version provided by Fisch et al. (2019). For each dataset, we randomly sample 10,000 <question, context, answer> triples from the original training datasets for continual training, and the original dev sets for testing. For CDA-Q setting, we make use of the original SQuAD

Methods	SQuAD	NaturalQA	HotpotQA	NewsQA	TriviaQA	$F1_{avg}$	$F1_{all}$
UpperBound	78.23	67.41	68	58.63	57.59	65.97	67.52
LowerBound	44.67	44.46	54.24	38.29	58.35	48	47.9
EWC	59.28	51.25	57.6	41.56	59.54	53.85	54.79
OnlineEWC	67.64	55.81	59.61	46.68	59.28	57.8	59.1
AGEM*	64.25	55.34	58.72	45.05	58.92	56.45	57.73
DER*	41.53	63.9	63.21	56.48	58.75	56.77	56.35
DER++*	46.96	63.05	63.5	56.62	57.94	57.61	53.37
MA-MRC($norm^1$)*	68.13	60.33	61.25	52.2	60.03	60.39	61.57
MA-MRC($norm^2$)*	70.11	59.35	61.92	52.18	59.51	60.61	61.77

Table 2: The overall results under CDA-C setting. “*” indicates that the \mathcal{M} memory data is used.

Methods	what	which	where	when	how	why	other	who	$F1_{avg}$	$F1_{all}$
UpperBound	76.91	80.64	72.02	87.51	73.09	54.49	77.18	84.58	75.8	77.67
LowerBound	63.8	72.42	55.98	63.62	65.79	23.7	69.27	82.89	62.18	66.39
EWC	63.27	74.4	53.17	60.7	63.77	34.62	66.92	82.19	62.38	65.37
OnlineEWC	65.72	72.3	59.48	60.75	60.91	34.21	69.09	83.74	63.28	66.89
AGEM*	66.88	74.88	66.95	76.52	66.72	49.78	70.36	85.28	69.67	70.08
DER*	55.95	75.91	71.47	85.58	72.94	54.45	76.69	84.17	72.15	67.75
DER++*	56.06	78.25	70.42	85.97	72.36	54.69	77.29	84.66	72.46	67.99
MA-MRC($norm^1$)*	72.31	75.815	69.2	85.59	72	56.76	75.65	85.86	74.15	74.92
MA-MRC($norm^2$)*	72.23	77.39	68.74	85.12	70.6	53.1	76.28	85.11	73.57	74.75

Table 3: The overall results under CDA-Q setting. “*” indicates that the \mathcal{M} memory data is used.

1.1 (Rajpurkar et al., 2016) and split it into eight domains according to the question type such as what, why, and how. The detailed statistics of these datasets are shown in Table 1.

4.2 Methods Compared

- **Bounds.** We design a standard BERT-MRC (the same as the backbone MRC model in Sec 3.3) as the basic model and then define two bounds on it: 1) **Lower Bound** continually fine-tunes the BERT-MRC model for each new domain without memorizing any historical examples. 2) **Upper Bound** remembers all examples in history and continually re-train the BERT-MRC model with all data.

- **EWC** (Kirkpatrick et al., 2017) restricts the change of model parameters for previous domains via elastic weight consolidation and a special L_2 regularization. Hence it can slow down the learning of parameters important for all previous domains.

- **Online EWC** (Schwarz et al., 2018), the extension of EWC, which only consider the restriction for the latest model parameters.

- **DER** (Buzzega et al., 2020), memory-based approaches, leverages knowledge distillation for retaining past experience.

- **DER++** (Buzzega et al., 2020), the extension of DER, uses an additional term on memory.

- **AGEM** (Chaudhry et al., 2019), memory-based approaches, uses a constraint that enables the projected gradient to decrease the average loss on previous seen domains.

4.3 Evaluation Metrics

Exact Match (EM) and word-level F1 score are used to evaluate the performance of MRC model in a single domain data. As for the continual domain adaptation setting, two common evaluation settings in continual learning theory are adopted: the average and the whole performance:

$$F1_{avg} = \frac{1}{T} \sum_{i=1}^T F1(D_{test}^i), \quad (9)$$

$$F1_{all} = F1(D_{test}^{1:T}). \quad (10)$$

The former is the average F1 score on test sets of all seen domain, and the latter is the whole F1 score on the test sets.

4.4 Implementation Details

We initialize the transformer encoder layer with the pre-trained BERT_{BASE} model officially released by Google¹. The maximum sequence length is 384, and the batch size is 30. We set memory size \mathcal{M} 400 default that means the memory stores up to 400 training samples for previous seen domain. When incrementally learn the new incoming domain data at t -th step, we first reinitialize the parameters of the domain discriminator θ_D . Then use Adam optimizer (Kingma and Ba, 2015) with a learning rate of 3e-5 and training MA-MRC model for 3 epochs.


¹<https://github.com/google-research/bert>

Methods	SQuAD	NaturalQA	HotpotQA	NewsQA	TriviaQA	$F1_{avg}$	$F1_{all}$
Full Model	68.13	60.33	61.25	52.2	60.03	60.39	61.57
w/o Adv	65.68	56.99	59.81	52.47	59.93	58.97	59.71
w/o KL	62.65	57.57	60.34	51.48	60.11	58.43	59.12
w/o Adv+KL	53.10	49.28	55.12	47.06	58.53	52.62	52.61


Table 4: The ablation study results under CDA-C setting.

Methods	what	which	where	when	how	why	other	who	$F1_{avg}$	$F1_{all}$
Full Model	72.31	75.815	69.2	85.59	72	56.76	75.65	85.86	74.15	74.92
w/o Adv	66.87	75.55	67.09	81.08	68.76	47.98	72.41	84.02	70.47	70.85
w/o KL	70.93	78.83	70.16	84.58	72.09	55.23	75.92	85.57	74.16	74.42
w/o Adv+KL	67.60	73.10	60.10	80.62	65.13	44.75	70.65	84.25	68.28	70.03

Table 5: The ablation study results under CDA-Q setting.



(a) Under CDA-C setting.



(b) Under CDA-Q setting.

Figure 3: The average F1 changes with increasing domains through the continual learning process.

4.5 Results and Discussion

4.5.1 Main Results

For the task of continual MRC, the overall results under CDA-C and CDA-Q setting and are shown in Table 2 and Table 3. Taking $F1_{avg}$ and $F1_{all}$ as two overall performance evaluations, we have the following findings: (1) The proposed MA-MRC methods outperform other baselines by a significant margin on both two continual settings. (2) The performance trend of different methods are not consistent across two continual settings. For exam-

Data	D_1	D_2	D_3	D_4	D_5
MA-MRC:					
step 1:	80.13	-	-	-	-
step 2:	74.92	69.16	-	-	-
step 3:	64.94	59.53	69.25	-	-
step 4:	74.13	62.87	65.12	60.36	-
step 5:	70.11	59.35	61.92	52.18	59.51
UpperBound:					
step 1:	80.13	-	-	-	-
step 2:	79.51	68.66	-	-	-
step 3:	78.42	67.64	70.03	-	-
step 4:	76.91	66.60	67.51	59.36	-
step 5:	78.24	67.41	68.00	58.63	57.59
MA-MRC w/o Adv:					
step 1:	80.13	-	-	-	-
step 2:	76.58	68.82	-	-	-
step 3:	68.83	61.14	69.74	-	-
step 4:	73.67	60.37	66.5	60.13	-
step 5:	65.68	56.99	59.81	52.47	59.93

Table 6: Catastrophic forgetting phenomenon of proposed MA-MRC(w/o Adv) and UpperBound methods under CDA-C setting.

ple, OnlineEWC performs well under CDA-C but performs poorly CDA-Q, and all memory-based methods outperform consolidation-based methods under CDA-Q but not CDA-C. The reason is that the characteristics of the domains and the continual learning difficulty are different. Concretely, it is obvious that the domain data in CDA-C is more different and more difficult than that in CDA-Q. (3) There is a big gap between all the models and the upper bound. We further demonstrate the evaluation results of the proposed MA-MRC and UpperBound methods on each domain at each continual learning step t in Table 6. The results indicate that MA-MRC indeed forgets knowledge learned on previously seen domain data, and there remain lots of things to be explored for continual MRC.

Besides, we plot the average F1 performance of models during the whole continual learning process in Figure 3 to investigate how performance

Sampling	SQuAD	NaturalQA	HotpotQA	NewsQA	TriviaQA	$F1_{avg}$	$F1_{all}$
$u^{entropy}, norm^1$	68.13	60.33	61.25	52.2	60.03	60.39	61.57
$u^{entropy}, norm^2$	70.11	59.35	61.92	52.18	59.51	60.61	61.77
$u^{prob}, norm^1$	68.71	59.30	62.00	52.79	59.94	60.55	61.55
$u^{prob}, norm^2$	63.78	60.42	59.88	52.72	59.54	59.27	60.25
Random	65.68	56.99	59.81	52.47	59.93	58.97	59.71

Table 7: The results of different sampling straties under CDA-C setting.

\mathcal{M}	methods	$F1_{avg}$	$F1_{all}$
200	MA-MRC($norm^1$)	53.76	53.91
	MA-MRC($norm^2$)	56.13	56.66
300	MA-MRC($norm^1$)	58.39	59.08
	MA-MRC($norm^2$)	58.14	58.88
400	MA-MRC($norm^1$)	60.39	61.57
	MA-MRC($norm^2$)	60.62	61.77

Table 8: Results with different \mathcal{M} under CDA-C.

		$F1_{avg}$	$F1_{all}$
order 1	DER++	70.21	64.35
	AGEM	57.57	30.18
	MA-MRC($norm^1$)	71.83	71.41
	MA-MRC($norm^2$)	72.77	72.48
order 2	DER++	69.9	76.2
	AGEM	74.97	77.02
	MA-MRC($norm^1$)	75.41	77.5
	MA-MRC($norm^2$)	75.27	78.01
order 3	AGEM	66.82	68.82
	MA-MRC($norm^1$)	72.64	74.27
	MA-MRC($norm^2$)	72.21	73.73

Table 9: The results of different domain order under CDA-Q setting.

changes. We observe that the performance of all models decreases in some degree with increasing numbers of domains under both CDA-C and CDA-Q settings. However, the proposed MA-MRC methods outperform other baselines and achieve better performance on the whole domain data.

4.5.2 Ablation Study

To better understand our proposed model, we conduct ablation studies to see the effectiveness of each model component. The results in Table 4 and Table 5 demonstrate that both knowledge distillation and adversarial training contribute to avoiding largely forgetting. For incremental domain learning, knowledge distillation is a naive way to enforce the model to remember its original responses for previous domains. What is more, the adversarial domain adaptation, can make MRC model learn domain-invariant and transfer representations better. Table 6 shows the evaluation results of w/o Adv at each continual step. Though w/o Adv obtains a slightly higher performance on the last two domains than MA-MRC model, it has a more harm-

	Para	Training Time / 1 epoch	
		CDA-C	CDA-Q
AGEM	$\theta_{T,AV}$	77 m	49 m
DER	$\theta_{T,AV}$	57 m	38 m
DER++	$\theta_{T,AV}$	56 m	38 m
MA-MRC($norm^1$)	$\theta_{T,AV,D}$	61 m	40 m
MA-MRC($norm^2$)	$\theta_{T,AV,D}$	60 m	39 m

Table 10: Parameters and speed comparison.

ful forgetfulness on the very previous seen domain. The results prove that the adversarial learning indeed helps for remembering previous knowledge.

4.5.3 Effect of Memory

Memory Size \mathcal{M} . Table 8 shows the performance with three different memory size \mathcal{M} : 200, 300, and 400. In low memory scenario, i.e., $\mathcal{M} = 200$, the proposed method performs poorly. The reason is that the adversarial domain adaption has difficulty transferring well with a too small memory. Therefore, as the number of memory samples increases, it will be more conducive to transfer to a new domain, and the overall performance will be better. We believe that an appropriate memory size could lead to better performance.

Uncertainty-aware Sampling. We replace the uncertainty-aware memory updating strategy with another two strategies. First, we use another uncertainty measurement that takes heuristic max softmax probability of spans as the uncertainty: $u^{prob} = \max_{i,j}(p_i^{start} + p_j^{end})$. The second strategy is random sampling. The experimental results in Table 7 indicate that the uncertainty-aware sampling (both *Entropy-based* and *Probability-based* uncertainty) is better than random sampling.

4.5.4 Effect of Domain Order

Table 9 shows the results of different domain orders. Order 1 is a descending order based on the number of training samples in each domain and order 2 is an ascending order, and order 3 is a random order. The performance of AGEM and DER++ degrades severely in order 1. However, the proposed MA-MRC methods is superior to baselines and are stable and robust under different orders.

4.5.5 Efficiency Analysis

We compare the parameters and training speed of methods with the same size of the memory in Table 10. The MA-MRC have more parameters for domain discriminator. Nevertheless, considering the larger number of parameters of the Transformer, we conclude all methods have almost the same number of parameters. As for training time per epoch, MA-MRC is slower than DER++ (4%/5% under CDA-C/CDA-Q). It can be accepted because of the significant improvement of MA-MRC.

5 Conclusion

In this paper, an incremental learning MRC model with uncertainty-aware fixed memory and adversarial domain adaptation, MA-MRC, is proposed for continual MRC and alleviating catastrophically forgetting. Inspired by the human learning process, There are two main ideas of MA-MRC: a memory that stores a small number of samples in previous seen domain data and always focuses on what the model most needs to replay; and adversarial learning the domain adaptation in a two-player game to learn better transfer representations between previous and current domain data. Experimental results show that the proposed MA-MRC can achieve a good continuous learning performance without catastrophically forgetting under CDA-C and CDA-Q settings.

In the future, we would like to explore a more effective sampling strategy, domain adaptation strategy, and balance training strategy for multiple objectives to enhance the continual MRC model.

Acknowledgment

This paper is funded by National Natural Science Foundation of China (Grant No. 62173195). This work is partly supported by seed fund of Tsinghua University (Department of Computer Science and Technology) -Siemens Ltd., China Joint Research Center for Industrial Intelligence and Internet of Things.

References

Abdalghani Abujabal, Rishiraj Saha Roy, Mohamed Yahya, and Gerhard Weikum. 2018. [Never-ending learning for open-domain question answering over knowledge bases](#). In *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France, April 23-27, 2018*, pages 1053–1062.

Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. 2010. [A theory of learning from different domains](#). *Mach. Learn.*, 79(1-2):151–175.

Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. 2020. [Dark experience for general continual learning: a strong, simple baseline](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

Yu Cao, Meng Fang, Baosheng Yu, and Joey Tianyi Zhou. 2020. [Unsupervised domain adaptation on reading comprehension](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7480–7487.

Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. 2019. [Efficient lifelong learning with A-GEM](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.

Cyprien de Masson d’Autume, Sebastian Ruder, Lingpeng Kong, and Dani Yogatama. 2019. [Episodic memory in lifelong language learning](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13122–13131.

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, MRQA@EMNLP 2019, Hong Kong, China, November 4, 2019*, pages 1–13.

Robert M. French. 1999. [Catastrophic forgetting in connectionist networks](#). *Trends in Cognitive Sciences*, 3(4):128–135.

Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky. 2016. [Domain-adversarial training of neural networks](#). *J. Mach. Learn. Res.*, 17:59:1–59:35.

Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. [Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1601–1611.

Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015*.

- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526.
- Seanie Lee, Donggyu Kim, and Jangwon Park. 2019. Domain-agnostic question-answering with adversarial training. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering, MRQA@EMNLP 2019, Hong Kong, China, November 4, 2019*, pages 196–202.
- David Lopez-Paz and Marc’Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6467–6476.
- Andrea Madotto, Zhaoyang Lin, Zhenpeng Zhou, Seungwhan Moon, Paul A. Crook, Bing Liu, Zhou Yu, Eunjoon Cho, Pascale Fung, and Zhiguang Wang. 2021. Continual learning in task-oriented dialogue systems. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 7452–7467.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100, 000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2383–2392.
- Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. 2019. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- Jonathan Schwarz, Wojciech Czarnecki, Jolena Luketina, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. 2018. Progress & compress: A scalable framework for continual learning. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 4535–4544.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017*.
- Yeon Seonwoo, Ji-Hoon Kim, Jung-Woo Ha, and Alice Oh. 2020. Context-aware answer extraction in question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2418–2428.
- Lixin Su, Jiafeng Guo, Ruqing Zhang, Yixing Fan, Yanyan Lan, and Xueqi Cheng. 2020. Continual domain adaptation for machine reading comprehension. In *CIKM ’20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*, pages 1395–1404.
- 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, Rep4NLP@ACL 2017, Vancouver, Canada, August 3, 2017*, pages 191–200.
- Huazheng Wang, Zhe Gan, Xiaodong Liu, Jingjing Liu, Jianfeng Gao, and Hongning Wang. 2019. Adversarial domain adaptation for machine reading comprehension. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 2510–2520.
- Zhijing Wu and Hua Xu. 2020. Improving the robustness of machine reading comprehension model with hierarchical knowledge and auxiliary unanswerability prediction. *Knowl. Based Syst.*, 203:106075.
- Ju Xu and Zhanxing Zhu. 2018. Reinforced continual learning. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 907–916.
- 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 *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2369–2380.
- Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. 2018. Lifelong learning with dynamically expandable networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.
- Zhuosheng Zhang, Junjie Yang, and Hai Zhao. 2021. Retrospective reader for machine reading comprehension. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 14506–14514.