

Continual Learning for Natural Language Generation in Task-oriented Dialog Systems

Fei Mi^{1*}, Liangwei Chen¹, Mengjie Zhao², Minlie Huang^{3*} and Boi Faltings¹

¹LIA, EPFL, Lausanne, Switzerland

²CIS, LMU, Munich, Germany

³CoAI, DCST, Tsinghua University, Beijing, China

{fei.mi, liangwei.chen, boi.faltings}@epfl.ch
mzhao@cis.lmu.de, aihuang@tsinghua.edu.cn

Abstract

Natural language generation (NLG) is an essential component of task-oriented dialog systems. Despite the recent success of neural approaches for NLG, they are typically developed in an offline manner for particular domains. To better fit real-life applications where new data come in a stream, we study NLG in a “continual learning” setting to expand its knowledge to new domains or functionalities incrementally. The major challenge towards this goal is catastrophic forgetting, meaning that a continually trained model tends to forget the knowledge it has learned before. To this end, we propose a method called *ARPER* (Adaptively Regularized Prioritized Exemplar Replay) by replaying prioritized historical exemplars, together with an adaptive regularization technique based on Elastic Weight Consolidation. Extensive experiments to continually learn new domains and intents are conducted on MultiWoZ-2.0 to benchmark *ARPER* with a wide range of techniques. Empirical results demonstrate that *ARPER* significantly outperforms other methods by effectively mitigating the detrimental catastrophic forgetting issue.

1 Introduction

As an essential part of task-oriented dialog systems (Wen et al., 2015b; Bordes et al., 2016), the task of Natural Language Generation (NLG) is to produce a natural language utterance containing the desired information given a *semantic representation* (so-called dialog act). Existing NLG models (Wen et al., 2015c; Tran and Nguyen, 2017; Tseng et al., 2018) are typically trained offline using annotated data from a single or a fixed set of domains. However, a desirable dialog system in real-life applications often needs to expand its knowledge to new domains and functionalities. Therefore, it is crucial to develop an NLG approach with the capability

of continual learning after a dialog system is deployed. Specifically, an NLG model should be able to continually learn new utterance patterns without forgetting the old ones it has already learned.

The major challenge of continual learning lies in catastrophic forgetting (McCloskey and Cohen, 1989; French, 1999). Namely, a neural model trained on new data tends to forget the knowledge it has acquired on previous data. We diagnose in Section 4.4 that neural NLG models suffer such detrimental catastrophic forgetting issues when continually trained on new domains. A naive solution is to retrain the NLG model using all historical data every time. However, it is not scalable due to severe computation and storage overhead.

To this end, we propose storing a small set of representative utterances from previous data, namely *exemplars*, and replay them to the NLG model each time it needs to be trained on new data. Methods using exemplars have shown great success in different continual learning (Rebuffi et al., 2017; Castro et al., 2018; Chaudhry et al., 2019) and reinforcement learning (Schaul et al., 2016; Andrychowicz et al., 2017) tasks. In this paper, we propose a *prioritized* exemplar selection scheme to choose representative and diverse exemplar utterances for NLG. We empirically demonstrate that the prioritized exemplar replay helps to alleviate catastrophic forgetting by a large degree.

In practice, the number of exemplars should be reasonably small to maintain a manageable memory footprint. Therefore, the constraint of not forgetting old utterance patterns is not strong enough. To enforce a stronger constraint, we propose a regularization method based on the well-known technique, Elastic Weight Consolidation (EWC (Kirkpatrick et al., 2017)). The idea is to use a quadratic term to elastically regularize the parameters that are important for previous data. Besides the wide application in computer vision, EWC has been recently

* Correspondence Author

applied to the domain adaptation task for Neural Machine Translation (Thompson et al., 2019; Saunders et al., 2019). In this paper, we combine EWC with exemplar replay by approximating the Fisher Information Matrix w.r.t. the carefully chosen exemplars so that not all historical data need to be stored. Furthermore, we propose to adaptively adjust the regularization weight to consider the difference between new and old data to flexibly deal with different new data distributions.

To summarize our contribution, (1) to the best of our knowledge, this is the first attempt to study the practical continual learning configuration for NLG in task-oriented dialog systems; (2) we propose a method called Adaptively Regularized Prioritized Exemplar Replay (*ARPER*) for this task, and benchmark it with a wide range of state-of-the-art continual learning techniques; (3) extensive experiments are conducted on the MultiWoZ-2.0 (Budzianowski et al., 2018) dataset to continually learn new tasks, including domains and intents using two base NLG models. Empirical results demonstrate the superior performance of *ARPER* and its ability to mitigate catastrophic forgetting. Our code is available at <https://github.com/MiFei/Continual-Learning-for-NLG>

2 Related Work

Continual Learning. The major challenge for continual learning is catastrophic forgetting (McCloskey and Cohen, 1989; French, 1999), where optimization over new data leads to performance degradation on data learned before. Methods designed to mitigate catastrophic forgetting fall into three categories: *regularization*, *exemplar replay*, and *dynamic architectures*. Methods using dynamic architectures (Rusu et al., 2016; Maltoni and Lomonaco, 2019) increase model parameters throughout the continual learning process, which leads to an unfair comparison with other methods. In this work, we focus on the first two categories.

Regularization methods add specific regularization terms to consolidate knowledge learned before. Li and Hoiem (2017) introduced the knowledge distillation (Hinton et al., 2015) to penalize model logit change, and it has been widely employed in Rebuffi et al. (2017); Castro et al. (2018); Wu et al. (2019); Hou et al. (2019); Zhao et al. (2019). Another direction is to regularize parameters crucial to old knowledge according to various importance measures (Kirkpatrick et al., 2017;

Zenke et al., 2017; Aljundi et al., 2018).

Exemplar replay methods store past samples, a.k.a *exemplars*, and replay them periodically. Instead of selecting exemplars at random, Rebuffi et al. (2017) incorporated the *Herding* technique (Welling, 2009) to choose exemplars that best approximate the mean feature vector of a class, and it is widely used in Castro et al. (2018); Wu et al. (2019); Hou et al. (2019); Zhao et al. (2019); Mi et al. (2020a,b). Ramalho and Garnelo (2019) proposed to store samples that the model is least confident. Chaudhry et al. (2019) demonstrated the effectiveness of exemplars for various continual learning tasks in computer vision.

Catastrophic Forgetting in NLP. The catastrophic forgetting issue in NLP tasks has raised increasing attention recently (Mou et al., 2016; Chronopoulou et al., 2019). Yogatama et al. (2019); Arora et al. (2019) identified the detrimental catastrophic forgetting issue while fine-tuning ELMo (Peters et al., 2018) and BERT (Devlin et al., 2019). To deal with this issue, He et al. (2019) proposed to replay pre-train data during fine-tuning heavily, and Chen et al. (2020) proposed an improved Adam optimizer to recall knowledge captured during pre-training. The catastrophic forgetting issue is also noticed in domain adaptation setups for neural machine translation (Saunders et al., 2019; Thompson et al., 2019; Varis and Bojar, 2019) and the reading comprehension task (Xu et al., 2019).

Lee (2017) firstly studied the continual learning setting for dialog state tracking in task-oriented dialog systems. However, their setting is still a one-time adaptation process, and the adopted dataset is small. Shen et al. (2019) recently applied progressive network (Rusu et al., 2016) for the semantic slot filling task from a continual learning perspective similar to ours. However, their method is based on a dynamic architecture that is beyond the scope of this paper. Liu et al. (2019) proposed a Boolean operation of “conceptor” matrices for continually learning sentence representations using linear encoders. Li et al. (2020) combined continual learning and language systematic compositionality for sequence-to-sequence learning tasks.

Natural Language Generation (NLG). In this paper, we focus on NLG for task-oriented dialog systems. A series of neural methods have been proposed to generate accurate, natural, and diverse utterances, including HLSTM (Wen et al., 2015a),

SCLSTM (Wen et al., 2015c), Enc-Dec (Wen et al., 2015b), RALSTM (Tran and Nguyen, 2017), SCVAE (Tseng et al., 2018).

Recent works have considered the domain adaptation setting. Tseng et al. (2018); Tran and Nguyen (2018b) proposed to learn domain-invariant representations using VAE (Kingma and Welling, 2013). They later designed two domain adaptation critics (Tran and Nguyen, 2018a). Recently, Mi et al. (2019); Qian and Yu (2019); Peng et al. (2020) studied learning new domains with limited training data. However, existing methods only consider a one-time adaptation process. The continual learning setting and the corresponding catastrophic forgetting issue remain to be explored.

3 Model

In this section, we first introduce the background of neural NLG models in Section 3.1, and the continual learning formulation in Section 3.2. In Section 3.3, we introduce the proposed method *ARPER*.

3.1 Background on Neural NLG Models

The NLG component of task-oriented dialog systems is to produce natural language utterances conditioned on a *semantic representation* called dialog act (DA). Specifically, the dialog act \mathbf{d} is defined as the combination of *intent* \mathbf{I} and a set of slot-value pairs $S(\mathbf{d}) = \{(s_i, v_i)\}_{i=1}^p$:

$$\mathbf{d} = \left[\underbrace{\mathbf{I}}_{\text{Intent}}, \underbrace{(s_1, v_1), \dots, (s_p, v_p)}_{\text{Slot-value pairs}} \right], \quad (1)$$

where p is the number of slot-value pairs. Intent \mathbf{I} controls the utterance functionality, while slot-value pairs contain information to express. For example, “There is a restaurant called [La Margherita] that serves [Italian] food.” is an utterance corresponding to a DA “[Inform, (name=La Margherita, food=Italian)]”

Neural models have recently shown promising results for NLG tasks. Conditioned on a DA, a neural NLG model generates an utterance containing the desired information word by word. For a DA \mathbf{d} with the corresponding ground truth utterance $\mathbf{Y} = (y_1, y_1, \dots, y_K)$, the probability of generating \mathbf{Y} is factorized as below:

$$f_\theta(\mathbf{Y}, \mathbf{d}) = \prod_{k=1}^K p_{y_k} = \prod_{k=1}^K p(y_k | y_{<k}, \mathbf{d}, \theta), \quad (2)$$

where f_θ is the NLG model parameterized by θ , and p_{y_k} is the output probability (i.e. softmax of




Figure 1: An example for a NLG model to continually learn new domains. The model needs to perform well on all domains it has seen before. For example f_{θ_3} needs to deal with all three previous domains (*Attraction*, *Restaurant*, *Hotel*).

logits) of the ground truth token y_k at position k . The typical objective function for an utterance \mathbf{Y} with DA \mathbf{d} is the average cross-entropy loss w.r.t. all tokens in the utterance (Wen et al., 2015c,b; Tran and Nguyen, 2017; Peng et al., 2020):

$$L_{CE}(\mathbf{Y}, \mathbf{d}, \theta) = -\frac{1}{K} \sum_{k=1}^K \log(p_{y_k}) \quad (3)$$

3.2 Continual Learning of NLG

In practice, an NLG model needs to continually learn new domains or functionalities. Without loss of generality, we assume that new data arrive *task by task* (Rebuffi et al., 2017; Kirkpatrick et al., 2017). In a new task t , new data \mathbf{D}_t are used to train the NLG model $f_{\theta_{t-1}}$ obtained till the last task. The updated model f_{θ_t} needs to perform well on *all* tasks so far. An example setting of continually learning new domains is illustrated in Figure 1. A task can be defined with different modalities to reflect diverse real-life applications. In subsequent experiments, we consider continually learning new domains and intents in Eq. (1).

We emphasize that the setting of continual learning is different from that of domain adaptation. The latter is a one-time adaptation process, and the focus is to optimize performance on a target domain transferred from source domains but without considering potential performance drop on source domains (Mi et al., 2019; Qian and Yu, 2019; Peng et al., 2020). In contrast, continual learning requires a NLG model to continually learn new tasks in multiple transfers, and the goal is to make the model perform well on all tasks learned so far.

3.3 Adaptively Regularized Prioritized Exemplar Replay (*ARPER*)

We introduce the proposed method (*ARPER*) with prioritized exemplar replay and an adaptive regu-

larization technique to further alleviate the catastrophic forgetting issue.

3.3.1 Prioritized Exemplar Replay

To prevent the NLG model catastrophically forgetting utterance patterns in earlier tasks, a small subset of a task’s utterances are selected as *exemplars*, and exemplars in previous tasks are *replayed* to the later tasks. During training the NLG model f_{θ_t} for task t , the set of exemplars in previous tasks, denoted as $\mathbf{E}_{1:t-1} = \{\mathbf{E}_1, \dots, \mathbf{E}_{t-1}\}$, is replayed by joining with the data \mathbf{D}_t of the current task. Therefore, the training objective with exemplar replay can be written as:

$$L_{ER}(\theta_t) = \sum_{\{\mathbf{Y}, \mathbf{d}\} \in \mathbf{D}_t \cup \mathbf{E}_{1:t-1}} L_{CE}(\mathbf{Y}, \mathbf{d}, \theta_t). \quad (4)$$

The set of exemplars of task t , referred to as \mathbf{E}_t , is selected after f_{θ_t} has been trained and will be replayed to later tasks.

The quality of exemplars is crucial to preserve the performance on previous tasks. We propose a *prioritized* exemplar selection method to select representative and diverse utterances as follows.

Representative utterances. The first criterion is that exemplars \mathbf{E}_t of a task t should be representative of \mathbf{D}_t . We propose to select \mathbf{E}_t as a *priority list* from \mathbf{D}_t that *minimize* a priority score:

$$U(\mathbf{Y}, \mathbf{d}) = L_{CE}(\mathbf{Y}, \mathbf{d}, \theta_t) \cdot |S(\mathbf{d})|^\beta, \quad (5)$$

where $S(\mathbf{d})$ is the set of slots in \mathbf{Y} , and β is a hyper-parameter. This formula correlates the representativeness of an utterance to its L_{CE} . Intuitively, the NLG model f_{θ_t} trained on \mathbf{D}_t should be confident with representative utterances of \mathbf{D}_t , i.e., low L_{CE} . However, L_{CE} is agnostic to the number of slots. We found that an utterance with many common slots in a task could also have very low L_{CE} , yet using such utterances as exemplars may lead to overfitting and thus forgetting of previous general knowledge. The second term $|S(\mathbf{d})|^\beta$ controls the importance of the number of slots in an utterance to be prioritized as exemplars. We empirically found in Appendix A.1 that the best β is greater than 0.

Diverse utterances. The second criterion is that exemplars should contain diverse slots of the task, rather than being similar or repetitive. A drawback of the above priority score is that similar or duplicated utterances containing the same set of frequent slots could be prioritized over utterances w.r.t. a

Algorithm 1 *select_exemplars*: Prioritized exemplar selection procedure of *ARPER* for task t

```

1: procedure select_exemplars( $\mathbf{D}_t, f_{\theta_t}, m$ )
2:    $\mathbf{E}_t \leftarrow$  new Priority_list()
3:    $\mathbf{D}_t \leftarrow$  sort( $\mathbf{D}_t, key = U, order = asc$ )
4:   while  $|\mathbf{E}_t| < m$  do
5:      $\mathbf{S}_{seen} \leftarrow$  new Set()
6:     for  $\{\mathbf{Y}, \mathbf{d}\} \in \mathbf{D}_t$  do
7:       if  $S(\mathbf{d}) \in \mathbf{S}_{seen}$  then continue
8:       else
9:          $\mathbf{D}_t.remove(\{\mathbf{Y}, \mathbf{d}\})$ 
10:         $\mathbf{E}_t.insert(\{\mathbf{Y}, \mathbf{d}\})$ 
11:         $\mathbf{S}_{seen.insert}(S(\mathbf{d}))$ 
12:        if  $|\mathbf{E}_t| == m$  then
13:          return  $\mathbf{E}_t$ 

```

diverse set of slots. To encourage diversity of selected exemplars, we propose an iterative approach to add data from \mathbf{D}_t to the priority list \mathbf{E}_t based on the above priority score. At each iteration, if the set of slots of the current utterance is already covered by utterances in \mathbf{E}_t , we skip it and move on to the data with the next best priority score.

Algorithm 1 shows the procedure to select m exemplars as a priority list \mathbf{E}_t from \mathbf{D}_t . The outer loop allows multiple passes through \mathbf{D}_t to select various utterances for the same set of slots $S(\mathbf{d})$.

3.3.2 Reducing Exemplars in Previous Tasks

Algorithm 1 requires the number of exemplars to be given. A straightforward choice is to store the same and fixed number of exemplars for each task as in Castro et al. (2018); Wu et al. (2019); Hou et al. (2019). However, there are two drawbacks in this method: (1). the memory usage increases linearly with the number of tasks; (2) it does not discriminate tasks with different difficulty levels.

To this end, we propose to store a *fixed* number of exemplars throughout the entire continual learning process to maintain a bounded memory footprint as in Rebuffi et al. (2017). As more tasks are continually learned, exemplars in previous tasks are gradually reduced by only keeping the ones in the *front* of the priority list¹, and the exemplar size of a task is set to be proportional to the training data size of the task to differentiate the task’s difficulty. To be specific, suppose M exemplars are kept in

¹the priority list implementation allows reducing exemplars in constant time for each task

total. The number of exemplars for a task is:

$$|\mathbf{E}_i| = M \cdot \frac{|\mathbf{D}_i|}{\sum_{j=1}^t |\mathbf{D}_j|}, \forall i \in 1, \dots, t, \quad (6)$$

where we choose 250/500 for M in experiments.

3.3.3 Constraint with Adaptive Elastic Weight Consolidation

Although exemplars of previous tasks are stored and replayed, the size of exemplars should be reasonably small ($M \ll |\mathbf{D}_{1:t}|$) to reduce memory overhead. As a consequence, the constraint we have made to prevent the NLG model catastrophically forgetting previous utterance patterns is not strong enough. To enforce a stronger constraint, we propose a regularization method based on the well-known Elastic Weight Consolidation (EWC, Kirkpatrick et al., 2017) technique.

Elastic Weight Consolidation (EWC). EWC utilizes a quadratic term to elastically regularize parameters important for previous tasks. The loss function of using the EWC regularization together with exemplar replay for task t can be written as:

$$L_{ER_EWC}(\theta_t) = L_{ER}(\theta_t) + \lambda \sum_i^N F_i(\theta_{t,i} - \theta_{t-1,i})^2 \quad (7)$$

where N is the number of model parameters; $\theta_{t-1,i}$ is the i -th converged parameter of the model trained till the previous task; $F_i = \nabla^2 L_{CE}^{\mathbf{E}_{1:t-1}}(\theta_{t-1,i})$ is the i -th diagonal element of the Fisher Information Matrix approximated w.r.t. the set of previous exemplars $\mathbf{E}_{1:t-1}$. F_i measures the importance of $\theta_{t-1,i}$ to previous tasks represented by $\mathbf{E}_{1:t-1}$. Typical usages of EWC compute F_i w.r.t. a uniformly sampled subset from historical data. Yet, we propose to compute F_i w.r.t. the carefully chosen $\mathbf{E}_{1:t-1}$ so that not all historical data need to be stored. The scalar λ controls the contribution of the quadratic regularization term. The idea is to elastically penalize changes on parameters important (with large F_i) to previous tasks, and more plasticity is assigned to parameters with small F_i .

Adaptive regularization. In practice, new tasks have different difficulties and similarities compared to previous tasks. Therefore, the degree of need to preserve the previous knowledge varies. To this end, we propose an adaptive weight (λ) for the EWC regularization term as follows:

$$\lambda = \lambda_{base} \sqrt{V_{1:t-1}/V_t}, \quad (8)$$

Algorithm 2 *learn_task*: Procedure of *ARPER* to continually learn task t

```

1: procedure learn_task( $\mathbf{D}_t, \mathbf{E}_{1:t-1}, f_{\theta_{t-1}}, M$ )
2:    $\theta_t \leftarrow \theta_{t-1}$ 
3:   while  $\theta_t$  not converged do
4:      $\theta_t \leftarrow \text{update}(L_{ER\_EWC}(\theta_t))$ 
5:      $m \leftarrow M \cdot \frac{|\mathbf{D}_t|}{\sum_{j=1}^t |\mathbf{D}_j|}$ 
6:      $\mathbf{E}_t \leftarrow \text{select\_exemplars}(\mathbf{D}_t, f_{\theta_t}, m)$ 
7:     for  $j = 1$  to  $t - 1$  do
8:        $\mathbf{E}_j \leftarrow \mathbf{E}_j.\text{top}(M \cdot \frac{|\mathbf{D}_j|}{\sum_{j=1}^t |\mathbf{D}_j|})$ 
9:   return  $f_{\theta_t}, \mathbf{E}_t$ 

```

where $V_{1:t-1}$ is the *old* word vocabulary size in previous tasks, and V_t is the *new* word vocabulary size in the current task t ; λ_{base} is a hyper-parameter. In general, λ increases when the ratio of the size of old word vocabularies to that of new ones increases. In other words, the regularization term becomes more important when the new task contains fewer new vocabularies to learn.

Algorithm 2 summarizes the continual learning procedure of *ARPER* for task t . θ_t is initialized with θ_{t-1} , and it is trained with prioritized exemplar replay and adaptive EWC in Eq. (7). After training θ_t , exemplars \mathbf{E}_t of task t are computed by Algorithm 1, and exemplars in previous tasks are reduced by keeping the most prioritized ones to preserve the total exemplar size.

4 Experiments

4.1 Dataset

We use the MultiWoZ-2.0 dataset² (Budzianowski et al., 2018) containing six domains (*Attraction, Hotel, Restaurant, Booking, Taxi and Train*) and seven DA intents (“*Inform, Request, Select, Recommend, Book, Offer-Booked, No-Offer*”). The original train/validation/test splits are used. For methods using exemplars, both training and validation set are continually expanded with exemplars extracted from previous tasks.

To support experiments on continual learning new domains, we pre-processed the original dataset by segmenting multi-domain utterances into single-domain ones. For instance, an utterance “*The ADC Theatre is located on Park Street. Before I find your train, could you tell me where you would like to go?*” is split into two utterances with domain

²extracted for NLG at <https://github.com/andy194673/nlg-sclstm-multiwoz>




Figure 2: Venn diagram visualizing intents in different domains. The number of utterances of each domain (bold) and intents (italic) is indicated in parentheses.

“Attraction” and “Train” separately. If multiple sentences of the same domain in the original utterance exist, they are still kept in one utterance after pre-processing. In each continual learning task, all training data of one domain are used to train the NLG model, as illustrated in Figure 1. Similar pre-processing is done at the granularity of DA intents for experiments in Section 4.6. The statistics of the pre-processed MultiWoZ-2.0 dataset is illustrated in Figure 2. The resulting datasets and the pre-processing scripts are open-sourced.

4.2 Evaluation Metrics

Following previous studies, we use the slot error rate (SER) and the BLEU-4 score (Papineni et al., 2002) as evaluation metrics. SER is the ratio of the number of *missing* and *redundant* slots in a generated utterance to the total number of ground truth slots in the DA.

To better evaluate the continual learning ability, we use two additional commonly used metrics (Kemker et al., 2018) for both SER and BLEU-4:

$$\Omega_{all} = \frac{1}{T} \sum_{i=1}^T \Omega_{all,i}, \quad \Omega_{first} = \frac{1}{T} \sum_{i=1}^T \Omega_{first,i}$$

where T is the total number of continual learning tasks; $\Omega_{all,i}$ is the test performance on *all* the tasks after the i^{th} task has been learned; $\Omega_{first,i}$ is that on the *first* task after the i^{th} task has been learned. Since Ω can be either SER or BLEU-4, both Ω_{all} and Ω_{first} have two versions. Ω_{all} evaluates the

overall performance, while Ω_{first} evaluates the ability to alleviate catastrophic forgetting.

4.3 Baselines

Two methods without exemplars are as below:

- **Finetune**: At each task, the NLG model is initialized with the model obtained till the last task, and then fine-tuned with the data from the current task.
- **Full**: At each task, the NLG model is trained with the data from the current and *all* historical tasks. This is the “upper bound” performance for continual learning w.r.t. Ω_{all} .

Several exemplar replay (*ER*) methods trained with Eq. (4) using different exemplar selection schemes are compared:

- **$ER_{herding}$** (Welling, 2009; Rebuffi et al., 2017): This scheme chooses exemplars that best approximate the mean DA vector over all training examples of this task.
- **ER_{random}** : This scheme selects exemplars at random. Despite its simplicity, the distribution of the selected exemplars is the same as the distribution of the current task in expectation.
- **ER_{prio}** : The proposed prioritized scheme (c.f. Algorithm 1) to select representative and diverse exemplars.

Based on ER_{prio} , four regularization methods (including ours) to further alleviate catastrophic forgetting are compared:

- **L2**: A static L2 regularization by setting $F_i = 1$ in Eq. (7). It regularizes all parameters equally.
- **KD** (Rebuffi et al., 2017; Wu et al., 2019; Hou et al., 2019): The widely-used knowledge distillation (KD) loss (Hinton et al., 2015) is adopted by distilling the prediction logit of current model w.r.t. the prediction logit of the model trained till the last task. More implementation details are included in Appendix A.1.
- **Dropout** (Mirzadeh et al., 2020): Dropout Hinton et al. (2012) is recently shown by (Mirzadeh et al., 2020) that it effectively alleviates catastrophic forgetting. We tuned different dropout rates assigned to the non-recurrent connections.
- **ARPER** (c.f. Algorithm 2): The proposed method using adaptive EWC with ER_{prio} .

We utilized the well-recognized semantically-conditioned LSTM (SCLSTM Wen et al., 2015c) as

	250 exemplars in total				500 exemplars in total			
	Ω_{all}		Ω_{first}		Ω_{all}		Ω_{first}	
	SER%	BLEU-4	SER%	BLEU-4	SER%	BLEU-4	SER%	BLEU-4
<i>Finetune</i>	64.46	0.361	107.27	0.253	64.46	0.361	107.27	0.253
<i>ER_{herding}</i>	16.89	0.535	9.89	0.532	12.25	0.555	4.53	0.568
<i>ER_{random}</i>	10.93	0.552	6.96	0.553	8.36	0.569	4.41	0.572
<i>ER_{prio}</i>	9.67**	0.578	5.28**	0.578	7.48**	0.597	3.59*	0.620
<i>ER_{prio}+L2</i>	14.94	0.579	5.31**	0.587	10.51	0.596	4.28**	0.605
<i>ER_{prio}+KD</i>	8.65**	0.586	6.87	0.601	7.37**	0.596	4.89	0.617
<i>ER_{prio}+Dropout</i>	7.15**	0.588	5.53**	0.594	6.09*	0.595	4.51**	0.616
<i>ARPER</i>	5.22	0.590	2.99	0.624	5.12	0.598	2.81	0.627
<i>Full</i>	4.26	0.599	3.60	0.616	4.26	0.599	3.60	0.616

Table 1: Average performance of continually learning 6 domains using 250/500 exemplars. Best Performance excluding “Full” are in bold in each column. In each column, * indicates $p < 0.05$ and ** indicates $p < 0.01$ for a one-tailed t-test comparing *ARPER* to the three top-performing competitors except *Full*.

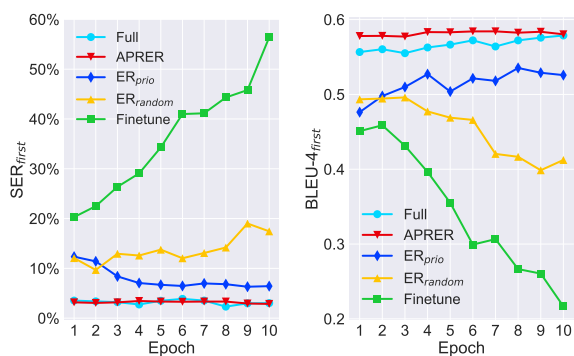


Figure 3: Diagnose the catastrophic forgetting issue in NLG. SER (Left) and BLEU-4 (Right) on the test data of “Attraction” at different epochs when a model pre-trained on the “Attraction” domain is continually trained on another “Train” domain.

the base NLG model f_{θ} ³ with one hidden layer of size 128. Dropout is not used by default, which is evaluated as a separate regularization technique (c.f. *ER_{prio}+Dropout*). For all the above methods, the learning rate of Adam is set to $5e-3$, batch size is set to 128, and the maximum number of epochs used to train each task is set to 100. Early stop to avoid over-fitting is adopted when the validation loss does not decrease for 10 consecutive epochs. To fairly compare different methods, they are trained with the identical configuration on the first task to have a consistent starting point. Hyper-parameters of different methods are included in Appendix A.1.

4.4 Diagnose Forgetting in NLG

Before proceeding to our main results, we first diagnose whether the catastrophic forgetting issue exists when training an NLG model continually. As

³Comparisons based on other base NLG models are included in Section 4.9.

an example, a model pre-trained on the “Attraction” domain is continually trained on the “Train” domain. We present test performance on “Attraction” at different epochs in Figure 3 with 250 exemplars.

We can observe: (1) catastrophic forgetting indeed exists as indicated by the sharp performance drop of *Finetune*; (2) replaying carefully chosen exemplars helps to alleviate catastrophic forgetting by a large degree, and *ER_{prio}* does a better job than *ER_{random}*; and (3) *ARPER* greatly mitigates catastrophic forgetting by achieving similar or even better performance compared to *Full*.

4.5 Continual Learning New Domains

In this experiment, the data from six domains are presented sequentially. We test 6 runs with different domain order permutations. Each domain is selected as the first task for one time, and the remaining five domains are randomly ordered⁴. Results averaged over 6 runs using 250 and 500 total exemplars are presented in Table 1. Several interesting observations can be noted:

- All methods except *Finetune* perform worse on all seen tasks (Ω_{all}) than on the first task (Ω_{first}). This is due to the diverse knowledge among different tasks, which increases the difficulty of handling all the tasks. *Finetune* performs poorly in both metrics because of the detrimental catastrophic forgetting issue.
- Replaying exemplars helps to alleviate the catastrophic forgetting issue. Three *ER* methods substantially outperform *Finetune*. Moreover, the proposed prioritized exemplar selection scheme is effective, indicated by the superior perfor-

⁴Exact domain orders are provided in Appendix A.2




Figure 4: SER on all seen domains (**solid**) and on the first domain (**dashed**) when more domains are continually learned using 250 exemplars.

mance of ER_{prio} over $ER_{herding}$ and ER_{random} .

- $ARPER$ significantly outperforms three ER methods and other regularization-based baselines. Compared to the three closest competitors, $ARPER$ is significantly better with p -value < 0.05 w.r.t SER.
- The improvement margin of $ARPER$ is significant w.r.t SER that is critical for measuring an output’s fidelity to a given dialog act. Different methods demonstrate similar performance w.r.t BLEU-4, where several of them approach $Full$, thus are very close to the upper bound performance.
- $ARPER$ achieves comparable performance w.r.t to the upper bound ($Full$) on all seen tasks (Ω_{all}) even with a very limited number of exemplars. Moreover, it outperforms $Full$ on the first task (Ω_{first}), indicating that $ARPER$ better mitigates forgetting the first task than $Full$, and the latter is still interfered by data in later domains.

Dynamic Results in Continual Learning In Figure 4, several representative methods are compared as more domains are continually learned. With more tasks continually learned, $ARPER$ performs consistently better than other methods on all seen tasks (solid lines), and it is comparable to $Full$. On the first task (dashed lines), $ARPER$ outperforms all the methods, including $Full$, at every continual learning step. These results illustrate the advantage of $ARPER$ through the entire continual learning process.

4.6 Continual Learning New DA Intents

It is also essential for a task-oriented dialog system to continually learn new functionalities, namely, supporting new DA intents. To test this ability,

	Ω_{all}		Ω_{first}	
	SER%	BLEU-4	SER%	BLEU-4
<i>Finetune</i>	49.94	0.382	44.00	0.375
$ER_{herding}$	13.96	0.542	8.50	0.545
ER_{random}	8.58	0.626	5.53	0.618
ER_{prio}	8.21	0.684	5.20	0.669
$ER_{prio}+L2$	6.87	0.693	4.92	0.661
$ER_{prio}+KD$	10.59	0.664	10.87	0.649
$ER_{prio}+Dropout$	6.32	0.689	5.55	0.658
$ARPER$	3.63	0.701	3.52	0.685
<i>Full</i>	3.08	0.694	2.98	0.672

Table 2: Performance of continually learning 7 DA intents using 250 exemplars. Best Performance excluding “ $Full$ ” are in bold.

	Ω_{all}		Ω_{first}	
	SER%	BLEU-4	SER%	BLEU-4
$ARPER$	4.82	0.592	3.88	0.569
w/o ER	6.41	0.584	5.85	0.559
w/o PE	5.53	0.587	5.85	0.562
w/o AR	5.57	0.587	4.57	0.563

Table 3: Ablation study for $ARPER$. ER / PE / AR stands for the Exemplar Replay loss / Prioritized Exemplars / Adaptive Regularization, respectively.

the data of seven DA intents are presented sequentially in the order of decreasing data size, i.e., “*Inform, Request, Book, Recommend, Offer-Booked, No-Offer, Select*”. Results using 250 exemplars are presented in Table 2. We can observe that $ARPER$ still largely outperforms other methods, and similar observations for $ARPER$ can be made as before. Therefore, we conclude that $ARPER$ is able to learn new functionalities continually.

Compared to previous experiments, the performance of $ER_{prio}+KD$ degrades, while the performance of $ER_{prio}+L2$ improves due to the very large data size in the first task (“*Inform*”), which means that they are sensitive to task orders.

4.7 Ablation Study

In Table 3, we compare several simplified versions of $ARPER$ to understand the effects of different components. Comparisons are based on continually learning 6 domains starting with “*Attraction*”. We can observe that: (1). L_{ER} is beneficial because dropping it (“w/o ER”) degrades the performance of $ARPER$. (2). Using prioritized exemplars is advantageous because using random exemplars (“w/o PE”) for $ARPER$ impairs its performance. (3). Adaptive regularization is also effective, indicated by the superior performance of $ARPER$ compared to using fixed regularization weights (“w/o AR”).

Recommend (Addr=regent street, Fee=free, Name=Downing College)	
Reference	<i>[Downing College] is my favorite. It's located on [regent street] and it's [free] to get in.</i>
$ER_{prio}+Dropout$	<i>[Downing College] is located in the city and it's located in the [regent street]. it's located at located at! it's located in the [Slot-Hotel-Area]. (missing: Fee=free)</i>
ARPER	<i>I would recommend [Downing College]. It is located at [regent street] and has a entrance fee of [free]. (correct)</i>
Recommend (Area=centre of town, Name=saints church, Type=architecture destination)	
Reference	<i>There is a [saints church] that is an [architecture destination] in the [centre of town], would you like that?</i>
$ER_{prio}+Dropout$	<i>I recommend [saints church] in the [centre of town]. it is a nice. it is a guest house in a in a [Slot-Restaurant-Food]. (missing: Type=architecture destination)</i>
ARPER	<i>[saints church] is a [architecture destination] in the [centre of town]. (correct)</i>

Table 4: Sample utterances generated for the first domain (“Attraction”) after the NLG is continually trained on all 6 domains using 250 exemplars. Redundant and missing slots are colored in orange and blue respectively. Obvious grammar mistakes (redundant repetitions) are colored in purple.

	SCVAE		GPT-2	
	Ω_{all}	Ω_{first}	Ω_{all}	Ω_{first}
<i>Finetune</i>	60.83	98.86	28.69	31.76
$ER_{herding}$	17.95	11.48	11.95	10.48
ER_{random}	9.31	7.52	9.87	8.85
ER_{prio}	8.92	6.16	8.72	8.20
$ER_{prio}+L2$	12.47	6.67	10.51	9.20
$ER_{prio}+KD$	6.32	6.09	8.41	8.09
$ER_{prio}+Dropout$	8.01	8.77	7.60	7.72
ARPER	4.45	4.04	5.32	5.05
<i>Full</i>	3.99	4.03	4.75	4.53

Table 5: SER in % of using SCVAE and GPT-2 as f_{θ} . Best Performance excluding “Full” are in bold.

4.8 Case Study

Table 4 shows two examples generated by ARPER and the closest competitor ($ER_{prio}+Dropout$) on the first domain (“Attraction”) after the NLG model is continually trained on all 6 domains starting with “Attraction”. In both examples, $ER_{prio}+Dropout$ fails to generate slot “Fee” or “Type”, instead, it mistakenly generates slots belonging to later domains (“Hotel” or “Restaurant”) with several obvious redundant repetitions colored in purple. It means that the NLG model is *interfered* by utterance patterns in later domains, and it forgets some old patterns it has learned before. In contrast, ARPER succeeds in both cases without forgetting previously learned patterns.

4.9 Results using Other NLG Models

In this experiment, we changed the base NLG model From SCLSTM to SCVAE (Tseng et al., 2018) and GPT-2 (Radford et al., 2019). For GPT-

2, we used the pre-trained model with 12 layers and 117M parameters. As in Peng et al. (2020), exact slot values are not replaced by special placeholders during training as in SCLSTM and SCVAE. The dialog act is concatenated with the corresponding utterance before feeding into GPT-2. More details are included in Appendix A.1.

Results of using 250 exemplars to continually learn 6 domains starting with “Attraction” are presented in Table 5. Thanks to the large-scale pre-trained language model, GPT-2 suffers less from the catastrophic forgetting issue because of the better performance of *Finetune*. In general, the relative performance patterns of different methods are similar to that we observed in Section 4.5 and 4.6. Therefore, we can claim that the superior performance of ARPER can generalize to different base NLG models.

5 Conclusion

In this paper, we study the practical continual learning setting of language generation in task-oriented dialog systems. To alleviate catastrophic forgetting, we present ARPER which replays representative and diverse exemplars selected in a prioritized manner, and employs an adaptive regularization term based on EWC (Elastic Weight Consolidation). Extensive experiments on MultiWoZ-2.0 in different continual learning scenarios reveal the superior performance of ARPER. The realistic continual learning setting and the proposed technique may inspire further studies towards building more scalable task-oriented dialog systems.

References

- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. 2018. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154.
- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. 2017. Hindsight experience replay. In *Advances in neural information processing systems*, pages 5048–5058.
- Gaurav Arora, Afshin Rahimi, and Timothy Baldwin. 2019. Does an lstm forget more than a cnn? an empirical study of catastrophic forgetting in nlp. In *Proceedings of the The 17th Annual Workshop of the Australasian Language Technology Association*, pages 77–86.
- Antoine Bordes, Y-Lan Boureau, and Jason Weston. 2016. Learning end-to-end goal-oriented dialog. *arXiv preprint arXiv:1605.07683*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026.
- Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. 2018. End-to-end incremental learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 233–248.
- Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. 2019. Continual learning with tiny episodic memories. *arXiv preprint arXiv:1902.10486*.
- Sanyuan Chen, Yutai Hou, Yiming Cui, Wanxiang Che, Ting Liu, and Xiangzhan Yu. 2020. Recall and learn: Fine-tuning deep pretrained language models with less forgetting. *arXiv preprint arXiv:2004.12651*.
- Alexandra Chronopoulou, Christos Baziotis, and Alexandros Potamianos. 2019. An embarrassingly simple approach for transfer learning from pre-trained language models. 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 2089–2095.
- 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.
- Robert M French. 1999. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, pages 128–135.
- Tianxing He, Jun Liu, Kyunghyun Cho, Myle Ott, Bing Liu, James Glass, and Fuchun Peng. 2019. Mix-review: Alleviate forgetting in the pretrain-finetune framework for neural language generation models. *arXiv preprint arXiv:1910.07117*.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. 2012. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*.
- Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. 2019. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 831–839.
- Ronald Kemker, Marc McClure, Angelina Abitino, Tyler L Hayes, and Christopher Kanan. 2018. Measuring catastrophic forgetting in neural networks. In *Thirty-second AAAI conference on artificial intelligence*.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- 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.
- Sungjin Lee. 2017. Toward continual learning for conversational agents. *arXiv preprint arXiv:1712.09943*.
- Yuanpeng Li, Liang Zhao, Kenneth Church, and Mohamed Elhoseiny. 2020. Compositional continual language learning.
- Zhizhong Li and Derek Hoiem. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947.
- Tianlin Liu, Lyle Ungar, and João Sedoc. 2019. Continual learning for sentence representations using conceptors. 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 3274–3279.

- Davide Maltoni and Vincenzo Lomonaco. 2019. Continuous learning in single-incremental-task scenarios. *Neural Networks*, 116:56–73.
- Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- Fei Mi, Minlie Huang, Jiyong Zhang, and Boi Faltings. 2019. Meta-learning for low-resource natural language generation in task-oriented dialogue systems. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 3151–3157. AAAI Press.
- Fei Mi, Lingjing Kong, Tao Lin, Kaicheng Yu, and Boi Faltings. 2020a. Generalized class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 240–241.
- Fei Mi, Xiaoyu Lin, and Boi Faltings. 2020b. Ader: Adaptively distilled exemplar replay towards continual learning for session-based recommendation. In *Fourteenth ACM Conference on Recommender Systems*, pages 408–413.
- Seyed-Iman Mirzadeh, Mehrdad Farajtabar, and Hassan Ghasemzadeh. 2020. Dropout as an implicit gating mechanism for continual learning. *arXiv preprint arXiv:2004.11545*.
- Lili Mou, Zhao Meng, Rui Yan, Ge Li, Yan Xu, Lu Zhang, and Zhi Jin. 2016. How transferable are neural networks in nlp applications? In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 479–489.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pages 311–318.
- Baolin Peng, Chenguang Zhu, Chunyuan Li, Xijun Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. 2020. Few-shot natural language generation for task-oriented dialog. *arXiv preprint arXiv:2002.12328*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237.
- Kun Qian and Zhou Yu. 2019. Domain adaptive dialog generation via meta learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2639–2649.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Tiago Ramalho and Marta Garnelo. 2019. Adaptive posterior learning: few-shot learning with a surprise-based memory module. In *International Conference on Learning Representations (ICLR)*.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2001–2010.
- Matthew Riemer, Tim Klinger, Djallel Bouneffouf, and Michele Franceschini. 2019. Scalable recollections for continual lifelong learning. In *AAAI*, volume 33, pages 1352–1359.
- Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. 2016. Progressive neural networks. *arXiv preprint arXiv:1606.04671*.
- Danielle Saunders, Felix Stahlberg, Adrià de Gispert, and Bill Byrne. 2019. Domain adaptive inference for neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 222–228.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. 2016. Prioritized experience replay.
- Yilin Shen, Xiangyu Zeng, and Hongxia Jin. 2019. A progressive model to enable continual learning for semantic slot filling. 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)*, pages 1279–1284.
- Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. 2017. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, pages 2990–2999.
- Brian Thompson, Jeremy Gwinnup, Huda Khayrallah, Kevin Duh, and Philipp Koehn. 2019. Overcoming catastrophic forgetting during domain adaptation of neural machine translation. 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 2062–2068.
- Van-Khanh Tran and Le-Minh Nguyen. 2017. Natural language generation for spoken dialogue system using rnn encoder-decoder networks. In *Proceedings of the 21st Conference on Computational Natural Language Learning*, pages 442–451.

- Van-Khanh Tran and Le-Minh Nguyen. 2018a. Adversarial domain adaptation for variational neural language generation in dialogue systems. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1205–1217.
- Van-Khanh Tran and Le-Minh Nguyen. 2018b. Dual latent variable model for low-resource natural language generation in dialogue systems. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 21–30.
- Bo-Hsiang Tseng, Florian Kreyszig, Paweł Budzianowski, Iñigo Casanueva, Yen-chen Wu, Stefan Ultes, and Milica Gasic. 2018. Variational cross-domain natural language generation for spoken dialogue systems. In *19th Annual SIG-Dial Meeting on Discourse and Dialogue*, pages 338–343.
- Dusan Varis and Ondřej Bojar. 2019. Unsupervised pretraining for neural machine translation using elastic weight consolidation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 130–135.
- Max Welling. 2009. Herding dynamical weights to learn. In *Proceedings of the 26th International Conference on Machine Learning*, pages 1121–1128. ACM.
- Tsung-Hsien Wen, Milica Gašić, Dongho Kim, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015a. Stochastic language generation in dialogue using recurrent neural networks with convolutional sentence reranking. In *16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, page 275.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Lina M Rojas-Barahona, Pei-Hao Su, David Vandyke, and Steve Young. 2015b. Toward multi-domain language generation using recurrent neural networks. In *NIPS Workshop on Machine Learning for Spoken Language Understanding and Interaction*.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015c. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1711–1721.
- Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. 2019. Large scale incremental learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 831–839.
- Ying Xu, Xu Zhong, Antonio Jose Jimeno Yepes, and Jey Han Lau. 2019. Forget me not: Reducing catastrophic forgetting for domain adaptation in reading comprehension. *arXiv preprint arXiv:1911.00202*.
- Dani Yogatama, Cyprien de Masson d’Autume, Jerome Connor, Tomas Kocisky, Mike Chrzanowski, Lingpeng Kong, Angeliki Lazaridou, Wang Ling, Lei Yu, Chris Dyer, et al. 2019. Learning and evaluating general linguistic intelligence. *arXiv preprint arXiv:1901.11373*.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual learning through synaptic intelligence. In *Proceedings of the 34th International Conference on Machine Learning*, pages 3987–3995. JMLR.org.
- Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shutao Xia. 2019. Maintaining discrimination and fairness in class incremental learning. *arXiv preprint arXiv:1911.07053*.

Appendix

A Reproducibility Checklist

A.1 Model Details and Hyper-parameters

We first elaborate implementation details of the knowledge distillation (KD) baseline compared in our paper. We used the below loss term:

$$L_{KD}(\mathbf{Y}, \mathbf{d}, f_{\theta_{t-1}}, f_{\theta_t}) = - \sum_{k=1}^K \sum_{i=1}^{|L|} \hat{p}_{k,i} \cdot \log(p_{k,i})$$

where L is the vocabulary that appears in previous tasks but not in task t . At each position k of \mathbf{Y} , $[\hat{p}_{k,1}, \dots, \hat{p}_{k,|L|}]$ is the predicted distribution⁵ over L given by $f_{\theta_{t-1}}$, and $[p_{k,1}, \dots, p_{k,|L|}]$ is the distribution given by f_{θ_t} . L_{KD} penalizes prediction changes on the vocabularies specific to earlier tasks. For all $\{\mathbf{Y}, \mathbf{d}\} \in \mathbf{D}_t \cup \mathbf{E}_{1:t-1}$, L_{KD} is linearly interpolated with L_{ER} by $L_{ER} + \eta \cdot L_{KD}$, with the η tuned as a hyper-parameter.

Hyper-parameters of SCVAE reported in Section 4.9 are set by default according to <https://github.com/andy194673/nlg-scvae>, except that the learning rate is set to $2e-3$. For GPT-2, we used the implementation pipeline from <https://github.com/pengbaolin/SC-GPT>. We pre-processed the dialog act \mathbf{d} into the format of: $\mathbf{d}' = [\mathbf{I} (s_1 = v_1, \dots, s_p = v_p)]$, and the corresponding utterance \mathbf{Y} is appended to be \mathbf{Y}' with a special start token [BOS] and an end token [EOS]. \mathbf{d}' and \mathbf{Y}' are concatenated before feeding into GPT-2. The learning rate of Adam optimizer is set to $5e-5$ without weight decay. As GPT-2 converges faster, we train maximum 10 epochs for each task with early stop applied to 3 consecutive epochs.

Hyper-parameters of different methods are tuned to maximize SER_{all} using grid search, and the optimal settings of different methods in various experiments are summarized in Table 6.

A.2 Domain Order Permutations

In Table 7, we provide the exact domain order permutations of the 6 runs used in the experiments in Table 1 and Figure 4.

A.3 Computation Resource

All experiments are conducted using a single GPU (GeForce GTX TITAN X). In Table 8, we compared the average training time of one epoch using

⁵The temperature in (Hinton et al., 2015) is set to 1 due to its minimum impact based on our experiments.

	Domains	DA Intents
$ER_{prio} (\beta)$	0.5/0.5/0.5/0.5	0.5
$L2$ (weight on $L2$)	1e-3/1e-3/1e-3/5e-4	1e-2
KD (weight on L_{KD})	2.0/3.0/2.0/0.5	5.0
$Dropout$ (rate)	0.25/0.25/0.25/0.1	0.25
$ARPER (\lambda_{base})$	300k/350k/200k/30k	100k

Table 6: Optimal hyper-parameters of methods experimented in this paper. Four different values in the column “Domains” correspond to using 250 exemplars in both Table 1 and Table 2 / 500 exemplars in Table 1 / using SCVAE / GPT-2 as $f(\theta)$ in Table 5, respectively.

Run 1	0	5	2	1	3	4
Run 2	1	4	0	5	3	2
Run 3	2	0	3	1	4	5
Run 4	3	2	4	0	1	5
Run 5	4	2	1	5	0	3
Run 6	5	3	2	0	1	4

Table 7: Each row corresponds to a domain order permutation. The mapping from domain to id is: {“Attraction”: 0, “Booking”: 1, “Hotel”: 2, “Restaurant”: 3, “Taxi”: 4, “Train”: 5.}

Finetune	ER_{prio}	$L2$	KD	$Dropout$	$ARPER$	Full
17.5s	18.5s	19.5s	24.6s	15.5s	39.5s	242.5s

Table 8: Average training time of one epoch at the last task when continually learning 6 domains starting with “Attraction” using 250 exemplars. Methods other than *Finetune* and *Full* are applied on top of ER_{prio} .

different methods. *Full* spends more than 200s of extra computation overhead per epoch than other methods using bounded exemplars. *ARPER* takes a slightly longer time to train than the methods except for *Full*. Nevertheless, considering its superior performance, we contend that *ARPER* achieves desirable resource-performance trade-off. In addition, 250 exemplars are less than 1% of historical data at the last task, and the memory usage to store a small number of exemplars is trivial.

B Supplementary Empirical Results

B.1 Comparison to Pseudo Exemplar Replay

Instead of storing raw samples as exemplars, Shin et al. (2017); Riemer et al. (2019) generate “pseudo” samples akin to past data. The NLG model itself can generate pseudo exemplars. In this experiment, we replace the 500 raw exemplars of ER_{random} , ER_{prio} , and *ARPER* by pseudo samples generated by the continually trained NLG model using the dialog acts of the same raw exemplars




Figure 5: An visualization of the change of SCLSTM’s hidden layer weights obtained from two consecutive tasks of *ARPER* (**Top**) and *ER_{prio}+Dropout* (**Bottom**). Two sample task transitions (“from *Attraction*” to “*Train*”, and then from “*Train*” to “*Hotel*”) are shown. High temperature areas of *ARPER* is highlighted by red bounding boxes for better visualization.

	Ω_{all}		Ω_{first}	
	SER%	BLEU-4	SER%	BLEU-4
<i>ER_{random}</i>	9.82	0.495	8.64	0.405
<i>Pseudo-ER_{random}</i>	9.26	0.551	6.88	0.519
<i>ER_{prio}</i>	7.84	0.573	6.20	0.523
<i>Pseudo-ER_{prio}</i>	8.87	0.557	6.37	0.521
<i>ARPER</i>	4.43	0.597	3.40	0.574
<i>Pseudo-ARPER</i>	5.07	0.590	3.51	0.570

Table 9: Comparison with Pseudo Exemplar Replay.

as input. Result comparing using pseudo or raw exemplars to continually learn 6 domains starting with “*Attraction*” are illustrated in Table 9. We can see that using pseudo exemplars performs better for *ER_{random}*, but worse for *ER_{prio}* and *ARPER*. It means that pseudo exemplars are better when exemplars are chosen randomly, while carefully chosen exemplars (c.f. algorithm 1) is better than pseudo exemplars. Explorations on utilizing pseudo exemplars for NLG is orthogonal to our work, and it is left as future work.

B.2 Flow of Parameters Update

To further understand the superior performance of *ARPER*, we investigated the update of parameters throughout the continual learning process. Specifically, we compared SCLSTM’s hidden layer weights obtained from consecutive tasks, and the pairwise L_1 difference of two sample transitions is shown in Figure 5.

We can observe that *ER_{prio}+Dropout* tends to update almost all parameters, while *ARPER* only updates a small fraction of them. Furthermore, *ARPER* has different sets of important parameters for distinct tasks, indicated by different high-temperature areas in distinct weight updating heat maps. In comparison, parameters of *ER_{prio}+Dropout* seem to be updated uniformly in different task transitions. The above observations verify that *ARPER* indeed elastically allocates different network parameters to distinct NLG tasks to mitigate catastrophic forgetting.