

Overcoming Catastrophic Forgetting by Exemplar Selection in Task-oriented Dialogue System

Chen Chen¹, Ruizhe Li², Yuchen Hu¹, Yuanyuan Chen¹, Chengwei Qin¹, Qiang Zhang³

¹Nanyang Technological University, Singapore ²University of Aberdeen, UK

³Zhejiang University, China

Abstract

Intelligent task-oriented dialogue systems (ToDs) are expected to continuously acquire new knowledge, also known as Continual Learning (CL), which is crucial to fit ever-changing user needs. However, catastrophic forgetting dramatically degrades the model performance in face of a long streamed curriculum. In this paper, we aim to overcome the forgetting problem in ToDs and propose a method (HESIT) with hyper-gradient-based exemplar strategy, which samples influential exemplars for periodic retraining. Instead of unilaterally observing data or models, HESIT adopts a profound exemplar selection strategy that considers the general performance of the trained model when selecting exemplars for each task domain. Specifically, HESIT analyzes the training data influence by tracing their hyper-gradient in the optimization process. Furthermore, HESIT avoids estimating Hessian to make it compatible for ToDs with a large pre-trained model. Experimental results show that HESIT effectively alleviates catastrophic forgetting by exemplar selection, and achieves state-of-the-art performance on the largest CL benchmark of ToDs in terms of all metrics.

1 Introduction

Serving as a core technique of smart assistants, task-oriented dialogue systems (ToDs) are expected to continuously acquire new knowledge through time regarding user needs (Madotto et al., 2020), e.g., adding fresh slot-value pairs or handling dissimilar tasks. This ability is also known as Continual Learning (CL) (Mundt et al., 2020), which has recently attracted a surge of interest in NLP community (Biesialska et al., 2020), as well as other machine learning techniques (Qu et al., 2021).

In this setting, the central problem of CL is *catastrophic forgetting* (Kirkpatrick et al., 2017), as the data is streamed and the neural model inevitably forgets previously learned knowledge when fitting

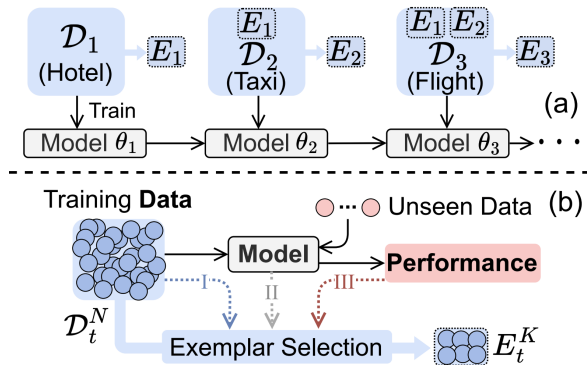


Figure 1: (a) Rehearsal-based CL in Task-oriented dialogue system. Exemplars E_t are sampled from t -domain training data for episodic rehearsal. (b) Exemplar selection in terms of influence chain “Data (I) – Model (II) – Performance (III)”. Our method penetrates into the performance perspective.

new training data in the sequential order. Such a phenomenon is particularly conspicuous in ToDs tasks due to the obvious distributional shift between task scenarios, e.g., setting an alarm clock and booking a flight are completely irrelevant tasks. As a result, it is challenging for a single model that is learned from a unidirectional curriculum to simultaneously handle multi-domain tasks.

To mitigate the problem of forgetting, existing efforts (De Lange et al., 2021) follow the three lines: (i) adding regularization terms to consolidate learned knowledge, (ii) developing dynamic architectures for the task-specific domain, and (iii) applying historical data rehearsal. Previous works (Madotto et al., 2020) have demonstrated that regularization methods gradually lose effectiveness if faced with a long curriculum, and we confirm this belief with a relative experiment. Dynamic architecture methods (e.g. adapter tuning) are usually considered as an **unfair baseline** (Rusu et al., 2016; Maltoni and Lomonaco, 2019; Mi et al., 2020) due to multiple sets of parameters, and it requires a further step during inference to determine

the correct parameters resulting in high latency for ToDs. Therefore, this work focuses on rehearsal methods (Fig. 1 (a)) (Rolnick et al., 2019) that store a small number of past samples (a.k.a. exemplars) in episodic memory and replay them periodically, which is simple but effective. Then we raise our research question: *which utterances are suitable to be selected as rehearsal exemplars for ToDs?*

Although statistic methods (e.g., reservoir algorithm (Chaudhry et al., 2019)) conduce to select better exemplars than random sampling, according to the accredited influence chain “*Data-Model-Performance*” (Sun et al., 2022), an exemplar selection strategy should further consider the general performance of models during testing. In other words, we should select the exemplars that positively influence the *performance* of the trained network on unseen data, as shown in Fig. 1 (b).

To explore the influence of training data from the performance perspective, one pioneering work uses the Influence Functions (IFs) based on the derivation chain rule (Koh and Liang, 2017) to quantify the contribution of an individual training sample, by observing its impact on the test loss when removing it. Though IFs have shown effectiveness in image classification tasks, this method requires estimating the inverse Hessian matrix, leading to high computing costs (Guo et al., 2020) and unstable results for large pre-trained models that are widely used in ToDs. Furthermore, IFs only measure the contribution around the model parameters at the final training epoch, which fails to trace the data contribution in the dynamic optimization process (Chen et al., 2021).

In this paper, we propose a rehearsal-based CL method, named HESIT (Hyper-gradient-based Exemplar Selection by Influence Tracing) to effectively overcome the catastrophic forgetting in ToDs. Specifically, instead of only focusing on the model parameters of the final epoch, HESIT traces and analyzes the hyper-gradient of training examples in the complicated optimization process. More importantly, its exemplar selection strategy roots in model *performance* on unseen data to measure the influence of these traced samples. In this way, influential and representative exemplars can be determined and stored, and then utilized to remind the model of learned knowledge. Moreover, HESIT is Hessian-free to avoid the instability when estimating Hessian matrix of a large pre-trained models, which has been widely used in ToDs. Experimental

results demonstrate that HESIT achieves state-of-the-art performance on the largest CL benchmark of ToDs (Madotto et al., 2020) including 37 domain tasks in terms of all metrics. Furthermore, comparative experiments show that the exemplar selection in HESIT surpasses all other selection strategies from other perspectives.

The main contributions of this paper can be summarized as follows:

- We propose HESIT – a rehearsal-based CL method with a novel exemplar selection strategy that can effectively overcome catastrophic forgetting in ToDs.
- HESIT dynamically traces the hyper-gradient of candidate data in the training process and selects exemplars in terms of data influence to model performance.
- Compared with other influence function-based methods, HESIT avoids estimating the Hessian matrix without degrading effect, which is compatible for ToDs with a large-scale pre-trained model.
- Extensive experiments show that HESIT can achieve state-of-the-art performance on the largest CL benchmark of ToDs (37 domain tasks).

2 Related Work

Continual learning. In the past few years, CL has achieved remarkable progress in mitigating catastrophic forgetting (McCloskey and Cohen, 1989), which can be classified into three categories: (1) *regularization-based* methods (Kirkpatrick et al., 2017; Serra et al., 2018; Huang et al., 2021; Qin and Joty, 2022a; Qin et al., 2024), which focus on regularizing the parameters corresponding to be inherited from the old tasks and penalizing the feature drift. (2) *architecture-based* methods (Xu and Zhu, 2018; Li et al., 2019; Ermis et al., 2022; Qin et al., 2023), which develop dynamic parameter isolation or expansion during CL training, where each task domain learns a specific architecture. (3) *rehearsal-based* methods (Cui et al., 2019; Chaudhry et al., 2019; Verwimp et al., 2021; Qin and Joty, 2022b), which utilize a small replay buffer to retain a fraction of learned training data and utilize them to retain the task knowledge. One shortcoming of architecture-based and rehearsal-based methods

is that they require extra memory space for storing parameters and old data samples (De Lange et al., 2021), while regularization-based methods might be overwhelmed when handling many task domains (Mai et al., 2022).

Continual learning in ToDs. Early work on CL for ToDs was introduced by Lee (2017), where the elastic weight consolidation (EWC) method was utilized to alleviate catastrophic forgetting on 3 sequential domains. Wu et al. (2019) studied CL for DST sub-task on the MultiWOZ dataset, where several baselines have been compared. Mi et al. (2020) combined data rehearsal and EWC on NLG sub-task and expanded the learned knowledge to 13 task domains. Furthermore, Madotto et al. (2020) developed a large benchmark for all INTENT, DST, and NLG tasks, where mainstream CL methods (e.g., L2, AGEM, and Adapter) were compared on a total of 37 domains. However, it shows that regularization-based methods lose effectiveness when ToDs encounter dozens of tasks (Madotto et al., 2020). Meanwhile, architecture-based methods require a further step to determine which set of parameters to use (Wortsman et al., 2020), thus blocking the real-time response from ToDs. In this work, we focus on rehearsal-based CL methods as it is simple and effective, with an acceptable cost of extra memory.

Exemplar selection. To sample representative or informative exemplars from a large dataset, reservoir sampling has been introduced in (Isele and Cosgun, 2018; Chaudhry et al., 2019; Hu et al., 2024) so that the data distribution in the buffer follows the data distribution that has already been seen. Rebuffi et al. (2017) proposed a herding-based strategy to maintain an online coreset. Similarly, gradient-based methods have also been proposed in (Broderick et al., 2013; Aljundi et al., 2019) to maximize the variance in the buffer. For ToDs with clear task boundaries, Mi et al. (2020) defined representative exemplars as a small set of utterances that fulfill a loss-based criterion on the current domain. The above-mentioned methods leverage models to choose high-quality data points, but fail to provide an insightful perspective in terms of model performance on unseen data.

3 Methodology

3.1 Background: End-to-end ToDs

We formulate task-oriented dialogue systems as a sequence-to-sequence generation problem that gen-

erates both API-calls and system responses. The API-call consists of the user intent and the current dialogue state, which can be empty or the system speech-act, to generate the system response. Thanks to the recent advance in ToDs, all ground-truth information at each turn is provided by existing annotated dialogue datasets. In this setting, the data format of API-call \mathcal{C} is shown as follows:

$$\mathcal{C}(\mathbf{H}) = \underbrace{\mathbf{I}}_{\text{Intent}} \underbrace{(s_1 = v_1, \dots, s_p = v_p)}_{\text{Slot-value pairs}}, \quad (1)$$

where \mathbf{H} is the dialogue history, \mathbf{I} denotes the user intent, and (s_i, v_i) stands for a slot-value pair from dialogue state tracking.

ToDs can be typically decomposed into different modules, including intent recognition (INTENT), history state tracking (DST), and natural language generation (NLG). In this paper, besides the sub-modules, we define the end-to-end (E2E) manner that directly generates the system response \mathbf{R} :

$$\mathbf{H} + \mathcal{C}(\mathbf{H}) \rightarrow \mathbf{R} \quad (2)$$

where $\mathcal{C}(\mathbf{H})$ is often empty, and thus the model can directly generate a response according to dialogue history ($\mathbf{H} \rightarrow \mathbf{R}$). In addition, our method HESIT can work for both E2E and modular ToDs. More details are attached in Appendix B.

3.2 Rehearsal-based CL in ToDs

Given the training data \mathcal{D} consisting of T ToDs task domains $\mathcal{D} = \{\mathcal{D}_1 \cdots \mathcal{D}_T\}$, continual learning aims to train a neural model f_θ on \mathcal{D} in a sequence of tasks. In each task t , new data \mathcal{D}_t is used to update model $f_{\theta_{t-1}}$, while the updated model f_{θ_t} needs to perform well on all tasks so far. In general, the crux of learning f is to overcome the catastrophic forgetting problem where learned knowledge is continually overwritten by streamed data.

To suppress it, the rehearsal-based CL methods construct a small size of memory buffer \mathcal{M}_t from each task domain \mathcal{D}_t . When training task t , the data in $\mathcal{M}_{1:t-1}$ will be retrained together with \mathcal{D}_t . Accordingly, a specific training step of task t is written as:

$$\hat{\theta} = \arg \min_{\theta} \frac{1}{N} \sum_{(x_i, y_i) \in \mathcal{D}_t^{\text{tm}}} \ell((x_i, y_i), \theta) \quad (3)$$

$$\text{s.t. } (x_i, y_i) \in \mathcal{M}_{1:t-1} \cup \mathcal{D}_t^{\text{tm}}$$

where N is the total number of training data points from T tasks, ℓ denotes the empirical loss and

(x_i, y_i) is a training data point of task t . Considering the storage burden of \mathcal{M} , the retained samples, called exemplars, should be informative or influential for $\mathcal{D}_t^{\text{trn}}$.

3.3 HESIT: Exemplar Selection

In this subsection, we present the details of the exemplar selection strategy to find influential utterances in \mathcal{M}_t , which also serves as the core technique in HESIT.

Intuitively, given a training sample $z_i := (x_i, y_i) \in \mathcal{D}_t^{\text{trn}}$, we want to quantify its influence on the validation loss $\mathcal{L}(\mathcal{D}_t^{\text{val}}, \hat{\theta})$ via the resulting model $\hat{\theta}$. One possible way is to find out the loss difference $\mathcal{I}(z_i, \mathcal{D}_t^{\text{val}}) := \mathcal{L}(\mathcal{D}_t^{\text{val}}, \hat{\theta}_{\bar{z}_i}) - \mathcal{L}(\mathcal{D}_t^{\text{val}}, \hat{\theta})$ where $\hat{\theta}_{\bar{z}_i}$ would be the resulting model if z_i were not present in $\mathcal{D}_t^{\text{trn}}$. For this purpose, we introduce a weight variable ϵ_i and consider the new full-set training loss:

$$\mathcal{L}_{\text{tm}}(\theta) = \frac{1}{N} \sum_{(x_i, y_i) \in \mathcal{D}_t^{\text{trn}}} (\ell((x_i, y_i), \theta) + \epsilon_i \ell(z_i, \theta)). \quad (4)$$

Now $\hat{\theta}$ and $\mathcal{I}(z_i, \mathcal{D}_t^{\text{val}})$ are taken as functions of ϵ_i and $\mathcal{I}(z_i, \mathcal{D}_t^{\text{val}})$ can be approximated by the first order Taylor expansion:

$$\begin{aligned} \mathcal{I}(z_i, \mathcal{D}_t^{\text{val}}) &\approx -\frac{1}{N} \left. \frac{d\mathcal{L}(\mathcal{D}_t^{\text{val}}, \hat{\theta})}{d\epsilon_i} \right|_{\epsilon_i=0} \\ &= -\frac{1}{N} \left. \frac{\partial \mathcal{L}(\mathcal{D}_t^{\text{val}}, \hat{\theta})}{\partial \hat{\theta}} \frac{d\hat{\theta}}{d\epsilon_i} \right|_{\epsilon_i=0}. \end{aligned} \quad (5)$$

While $\frac{\partial \mathcal{L}(\mathcal{D}_t^{\text{val}}, \hat{\theta})}{\partial \hat{\theta}}$ can be exactly calculated using gradient back-propagation on the validation data, computing $\nabla_i := \frac{d\hat{\theta}}{d\epsilon_i}$, which is also named hyper-gradient (Bengio, 2000), is not straightforward because the parameter ϵ_i is involved in the entire optimization process.

In gradient descent based optimization, for each iteration step r , we have $\theta_r = \theta_{r-1} - \gamma g_{r-1}$ where γ is the learning rate and g_{r-1} is the batch gradient. Then ∇_i can be computed recursively as:

$$\nabla_{i,r} = \nabla_{i,r-1} - \gamma (H_{r-1} \nabla_{i,r-1} + \frac{\partial g_{r-1}}{\partial \epsilon_i}) \quad (6)$$

where H_{r-1} denotes the Hessian of the batch loss with respect to θ_{r-1} . Most related works impose some assumptions on the model to ensure that H is invertible and attempt to approach the inverse Hessian vector product by numerical methods (Koh and Liang, 2017). However, the assumptions are too

Algorithm 1 Influential Example Tracing

Input: Traced example set $\mathcal{Z}_t = \{z_i\}_{i=1}^I$ of task t , training dataset $\mathcal{D}_t^{\text{trn}}$, validation dataset $\mathcal{D}_t^{\text{val}}$, batch size B , training dataset size N , total iteration steps R , learning rate $\gamma(t)$

Output: Model $\hat{\theta}$, influence values $\mathcal{I}(\mathcal{Z}, \mathcal{D}_t^{\text{val}})$

- 1: Setup reproducible training environment
 - 2: Train model with $\mathcal{D}_t^{\text{trn}}$ and obtain $\hat{\theta}$
 - 3: $\mathbf{v} \leftarrow \frac{\partial \mathcal{L}(\mathcal{D}_t^{\text{val}}, \hat{\theta})}{\partial \hat{\theta}}$
 - 4: Reset training environment and retrain model with $\mathcal{D}_t^{\text{trn}}$ for second identical training
 - 5: **for** $i = 1$ to I **do**
 - 6: $\mathbf{v} \nabla_{i,0} \leftarrow 0$
 - 7: **end for**
 - 8: **for** $r = 1$ to R **do**
 - 9: **if** current batch contains z_i **then**
 - 10: $\mathbf{v} \nabla_{i,r} \leftarrow \gamma(t) \mathbf{v} \nabla_{i,r-1} - \frac{N}{B} \mathbf{v} g_{i,r-1}$
 - 11: **else**
 - 12: $\mathbf{v} \nabla_{i,r} \leftarrow \gamma(t) \mathbf{v} \nabla_{i,r-1}$
 - 13: **end if**
 - 14: **end for**
 - 15: **for** $i = 1$ to I **do**
 - 16: $\mathcal{I}(z_i, \mathcal{D}_t^{\text{val}}) \leftarrow -\frac{1}{N} \mathbf{v} \nabla_{i,R}$
 - 17: **end for**
-

strong to be feasible in practice, and with increasing model and dataset size, the numerical methods are time-consuming and may lead to diverging results. Chen et al. (2021) propose a faster scheme by discarding the Hessian computations, which have a bounded error. Therefore, HESIT can recurrently update the $\nabla_{i,r}$ through the optimization trajectory with an acceptable time cost. We further design a validation experiment in Section 5.3 to demonstrate that such a manner can successfully estimate data influence without the estimation of Hessian.

As illustrated in Algorithm 1, the influence scores of training examples are computed via reproducible retraining to reduce the space complexity. In this way, the time complexity to trace hyper-gradient is $\mathcal{O}(I \cdot R \cdot \omega)$, where ω is the time spent in computing parameter gradient, but the space complexity has been reduced to $\mathcal{O}(I)$. Finally, a set of exemplars E_t^K are picked out according to the influence scores, with top- K strategy.

3.4 HESIT: Training Schedule

We present the training schedule of HESIT in Algorithm 2, which integrates the exemplar selection strategy and rehearsal-based CL into ToDs. Specifi-

Algorithm 2 HESIT Training

Input: Sequential T -domain training data $\mathcal{D} = \{(\mathcal{D}_1^{\text{trn}}, \mathcal{D}_1^{\text{val}}), \dots, (\mathcal{D}_T^{\text{trn}}, \mathcal{D}_T^{\text{val}})\}$

Output: Parameter θ_T that handles T -domain tasks

```
1: Initialize the model parameter  $\theta_0$  and replay
   buffer  $\mathcal{M}_0$ 
2: for  $t = 1$  to  $T$  do
3:   Sample  $\mathcal{Z}_t \sim \mathcal{D}_t^{\text{trn}}$ 
4:   Calculate  $\mathcal{I}(\mathcal{Z}_t, \mathcal{D}_t^{\text{val}})$  and  $\theta_{t-1}$  by Algo.1
5:   Select  $E_t^K \subset \mathcal{Z}_t$  by top  $\|\mathcal{I}(\mathcal{Z}_t, \mathcal{D}_t^{\text{val}})\|$ 
6:   while not converge do
7:     Update  $\theta_{t-1}$  using  $\{\mathcal{D}_t^{\text{trn}} \cup \mathcal{M}_{t-1}\}$ 
8:   end while
9:    $\mathcal{M}_t \leftarrow E_t^K \cup \mathcal{M}_{t-1}$ 
10: end for
```

cally, for each domain t , we select E_t^K from \mathcal{D}_t using a hyper-gradient based strategy and feed them into the replay buffer \mathcal{M} . Then \mathcal{M} will be involved in the subsequent training process for retraining the neural model.

To further accelerate training, we adopt two simplifications. Firstly, instead of tracing all training examples, we sample and trace a subset $\mathcal{Z}_t \gg E_t^K$ from each $\mathcal{D}_t^{\text{trn}}$ to reduce the cost. Secondly, tracing is performed only in the first R iterations rather than the entire training process, as the main benefits of optimization have been obtained in these iterations. In addition, the subsequent training (line 6-8) is depend on the size of domain t .

4 Experiment

4.1 Dataset

To evaluate the performance of our method, we employ the largest CL benchmark of the ToDs task developed in (Madotto et al., 2020). Specifically, four ToD datasets are merged: TaskMaster 2019 (TM19) (Byrne et al., 2019), TaskMaster 2020 (TM20) (Byrne et al., 2019), Schema Guided Dialogue (SGD) (Lin et al., 2021), and MultiWoZ (Budzianowski et al., 2018). Consequently, a curriculum of 37 domains is to be learned continuously in ToDs, where three modules (INTENT, DST, and NLG) are well annotated after pre-processing. We consider two settings which are **Modularized setting** that learns three modules separately, and **E2E setting** that learns these modules in a unified manner, as illustrated in 3.1.

In addition, we summarize the main statistics

and detailed sample numbers for each domain in Appendix D. It is noted that the amount of examples is highly imbalanced across different domains, which ranges from a few hundred to more than 36k. This distribution is more proximate to reality, as some task domains lack training data in practice.

4.2 Evaluation Metrics

Based on three modules, we employ the following well-defined metrics to evaluate system performance:

- INTENT recognition is directly measured in terms of accuracy between the predicted intent and ground-truth intent.
- DST is evaluated by Joint Goal Accuracy (JGA) (Wu et al., 2019) over the ground-truth dialogue.
- NLG is evaluated by two metrics: 1) BLEU score (Papineni et al., 2002) which is calculated by the distance between generated response and reference sentence, and 2) slot error rate (EER) (Wen et al., 2015) which is computed as the ratio between the total number of slots and the values not appearing in the response.

For all metrics except EER, a higher value denotes better performance. In addition, same as (Kale and Rastogi, 2020), we do not calculate EER count for the SGD dataset, since some slots of the SGD dataset have only binary values, *e.g.*, yes or no, which is unfair to calculate average with others.

4.3 Baselines

In order to compare the effects of different CL methods, we investigate both regularization-based and rehearsal-based baselines in the ToDs benchmark. Architecture-based methods are excluded as they require a further decision step to determine the agnostic domain of the test set, resulting in a slow system response.

Regularization baseline. We consider two kinds of regularization terms to constrain the model parameter update, which is identity function (**L2**) and Fisher information matrix (**EWC**) proposed in Kirkpatrick et al. (2017).

Rehearsal baseline. We first employ 2 kinds of mainstream rehearsal-based CL methods, which are **A-GEM** (Chaudhry et al., 2018), and **LAMOL** (Sun et al., 2019). Since our contribution focuses on exemplar selection, we further

reproduce 4 coreset selection strategies in the ToDs benchmark, where 2 baselines are from data-perspective: (1) **Random** denotes that exemplars are randomly sampled from the train set, and (2) **UNIFORM** denotes that we uniformly select exemplars to ensure all user intents would be contained in the buffer. Meanwhile, the other two baselines are from the model perspective: (3) **GSS** (Aljundi et al., 2019) selects the examples that maximize the gradient variance in the replay buffer. (4) **ARPER** (Mi et al., 2020) selects exemplars from the training set that obtains minimum loss-based criterion. (1)~(4) share the same training schedule with the proposed **HESIT** but they adopt different exemplar selection strategies respectively.

Additionally, we attach a lower-bound baseline, called **VANILLA**, which is trained on each task continuously without any anti-forgetting mechanism. We also provide the **MULTI** baseline which trains models with all data in an integrated curriculum simultaneously, which is widely viewed as the upper bound of CL methods.

4.4 Experimental Setup

For all experiments, we leverage the pre-trained GPT-2 (Radford et al., 2019) as the ToDs backbone. Each domain is trained for 10 epochs with early stopping over the validation set. The learning rate is set as 0.001 with a warm-up schedule. In L2 and EWC, the regularization weight is set as 0.001. Considering the buffer memory, all data replay-based methods sample 50 exemplars for each task domain. In HESIT, the Z_t is set as 1000 for each domain, and the interested training examples are traced for the first 5 epochs. Task order in the curriculum usually has a slight impact on the final performance. To avoid contingency, all experiments are repeated three times and report the average results.

5 Result and Analysis

5.1 Main Comparison Results

E2E setting. We report the main E2E results on the full test set in Table 1, where all metrics are evaluated at the end of the curriculum (37 task domains). We observe that regularization baselines (L2 and EWC) can not alleviate catastrophic forgetting well, and even achieve worse performance than VANILLA in terms of most metrics. The learning curve of EWC is visualized in Fig. 4. The reason is that the regularization item gradually constrains

Method	Mem.	INTENT	DST	NLG	
		Accuracy \uparrow	JGA \uparrow	EER \downarrow	BLEU \uparrow
VANILLA	\emptyset	2.65	9.33	50.91	4.49
Regularization-based methods					
L2	\emptyset	2.33	6.85	56.42	5.08
EWC	\emptyset	2.46	8.98	52.61	4.70
Rehearsal-based methods					
A-GEM	$t \mathcal{M} $	31.02	11.23	60.98	4.53
LAMOL	\emptyset	2.68	9.42	66.31	3.82
Data-perspective					
RANDOM	$t \mathcal{M} $	78.22	29.47	19.36	16.92
UNIFORM	$t \mathcal{M} $	80.67	28.94	19.63	17.84
Model-perspective					
GSS	$t \mathcal{M} $	81.42	30.33	17.48	17.97
ARPER	$t \mathcal{M} $	77.60	27.82	20.52	16.44
Performance-perspective (ours)					
HESIT	$t \mathcal{M} $	83.46	31.22	16.78	18.25
MULTI	-	95.45	48.90	12.56	23.61

Table 1: **E2E** results on the test set in terms of INTENT accuracy, JGA, EER, and BLEU. ‘‘Mem.’’ denotes the memory size of the buffer, where $t = 37$ and $|\mathcal{M}| = 50$.

Method	Mem.	INTENT	DST	NLG	
		Accuracy \uparrow	JGA \uparrow	EER \downarrow	BLEU \uparrow
VANILLA	\emptyset	3.06	10.28	18.09	10.42
Regularization-based methods					
L2	\emptyset	3.59	9.94	18.16	11.13
EWC	\emptyset	3.72	10.06	18.12	11.70
Rehearsal-based methods					
A-GEM	$t \mathcal{M} $	10.57	9.86	36.22	6.40
LAMOL	\emptyset	2.72	9.44	35.83	4.43
Data-perspective					
RANDOM	$t \mathcal{M} $	79.92	39.54	5.82	21.33
UNIFORM	$t \mathcal{M} $	81.10	39.92	5.42	21.17
Model-perspective					
GSS	$t \mathcal{M} $	81.36	39.80	5.17	21.62
ARPER	$t \mathcal{M} $	79.63	39.21	6.08	21.12
Performance-perspective (ours)					
HESIT	$t \mathcal{M} $	82.71	40.04	5.11	21.48
MULTI	-	87.50	50.03	3.42	26.15

Table 2: **Modularized** results on the test set in terms of INTENT accuracy, JGA, EER, and BLEU.

the optimization space of the neural model, making it difficult to handle subsequent tasks in a long curriculum. Similarly, due to clear task boundaries of ToDs, A-GEM and LAMOL baselines also lose effectiveness as their learning schedules depend on correlations across different domains. In this case, regardless of the selection strategy, storing some raw data and retraining the model periodically can obtain remarkable performance gain. As a performance-based selection strategy, our proposed HESIT performs significantly better than the RANDOM baseline, which respectively achieves 6.70%, 5.94%, 13.33%, and 7.86% relative improvement in terms of INTENT accuracy, JGA, EER, and BLEU. In addition, compared with the model perspective GSS baseline, HESIT respec-




Figure 2: Learning curve for INTENT accuracy in E2E setting. Each test point is evaluated on the already learned task in the curriculum.

tively achieves 2.04%, 2.93%, 4.0%, and 1.48% relative improvement.

Modularized setting. We summarize modularized results in Table 2 that INTENT, DST, and NLG modules are evaluated separately. Compared with the E2E setting, we observe some result fluctuations in different evaluation metrics. On the one hand, each neural model focuses on a single module leading to performance gains. Especially in the NLG sub-task, the modularized model surpasses the E2E model by a large margin, where the EER of MULTI reduces nearly fourfold (12.56%→3.42%). On the other hand, the modularized model ignores the relevance of each sub-task, resulting in performance degradation according to INTENT accuracy (95.45%→87.50%). Regardless of E2E or modularized settings, HESIT shows superiority in alleviating catastrophic forgetting in most metrics.

Forgetting curve. We visualize the learning curve of INTENT accuracy in the 37-domain curriculum to observe how HESIT overcomes catastrophic forgetting. To avoid the overlapping of the chart, partial baselines are selected for comparison and shown in Fig. 2. It is worth noting that we fix the task order in the curriculum, and each test point on the x-axis evaluates the model performance on the test set of learned tasks.

From the forgetting curve, we observe that 1) Catastrophic forgetting happened in the VANILLA baseline, and the model performance obviously degrades over time. 2) EWC, A-GEM, and LAMOL can alleviate catastrophic forgetting in the first few tasks, but gradually fail to handle learned tasks with an increase in task amounts. In addition, LAMOL achieves outstanding accuracy at the beginning of the curriculum and performs better than HESIT. 3)

Method	# Er.	INTENT Accuracy \uparrow	DST JGA \uparrow	NLG EER \downarrow	NLG BLEU \uparrow
VANILLA	-	2.65	9.33	50.91	4.49
Data-perspective					
UNIFORM	20	65.15	24.08	23.31	16.47
	30	72.22	25.98	22.01	17.19
	40	77.56	27.80	19.36	17.44
	50	80.67	28.94	19.63	17.84
Model-perspective					
GSS	20	66.78	23.85	18.99	17.04
	30	74.32	26.68	21.02	17.55
	40	78.09	29.76	18.61	17.70
	50	81.42	30.33	17.48	17.97
Performance-perspective					
HESIT	20	70.32	24.97	19.94	16.81
	30	72.85	27.85	18.40	17.45
	40	80.01	30.29	18.52	17.93
	50	83.46	31.22	16.78	18.25

Table 3: E2E results using different exemplar selection strategies. “# Er.” denotes the amount of exemplars for each task domain.

For HESIT, there are some fluctuations in performance, but then stabilize at a high performance in terms of accuracy.

5.2 Analysis for Exemplar Selection

Amount of exemplar. Using different selection strategies, we vary the size of the memory buffer containing [20,30,40,50] exemplars for each domain. The results are shown in Table 3. We observe that 20 exemplars can significantly alleviate cartographic forgetting, benefiting all metrics compared with the VANILLA baseline. As the buffer size increases, all systems achieve better performance including each ToDs module. Specifically, GSS achieves the best NLG performance when the exemplar amount is 20. In other cases, HESIT performs the best regardless of the exemplar amount, demonstrating its superiority over the performance perspective methods.

Case study. As each training sample is assessed using the Algo. 1’s Influence Functions to evaluate the impact of unseen data, we conduct a case study on several detrimental examples with low influence scores to examine how HESIT eliminates these negative examples. To this end, we simultaneously trace 1k training points of each domain and randomly sample some low-score examples in Table 4. The first two examples adhere to the format of API-call, and the third example is a system response. In the first two cases, the HESIT successfully identifies the incorrect train_inform() intent and the missing value-slot pairs with extremely negative

ID	Domain	Module	Type	Utterance	Score
1	Train	INTENT	$H \rightarrow C(H)$	USER: "Yes, I would like a taxi to the Town Centre. From the train station." API: "taxi_inform() train_inform() [eos]"	-0.64
2	Taxi	INTENT & DST	$H \rightarrow C(H)$	USER: "I need the phone number and location of the nearest Red Lobster in the downtown Cambridge area?" API: "taxi_request(phone=?)" [eos]"	-0.77
3	Hotel	NLG	$H \rightarrow R$	USER: "Yes, I am attending a physician's conference and need to locate a room for tonight. API-OUT:" SYSTEM: "Do you have an area or price preference? [eos]"	-0.30

Table 4: Detrimental examples with low score influences selected by HESIT from the Multi-WoZ dataset. Scores of traced examples are re-normalized to $[-1, 1]$.

scores. The third case is likewise eliminated by the HESIT with a negative score because the model is confused by the empty API-OUT from the history.

5.3 Effect of HESIT

In this part, we conduct experiments to demonstrate that HESIT can successfully measure the training data influence without estimating Hessian. To this end, we construct two Hessian-involved IFs baselines **Conjugate gradients (CG)** and **Stochastic estimation (LISSA)** (Agarwal et al., 2017) and one Hessian-free **Hydra** (Chen et al., 2021) for comparison. As the Hessian of a large pre-trained model is likely irreversible, we further employ a CNN backbone and CIFAR-10 dataset (Krizhevsky et al., 2009) to evaluate both computation time and performance. The details of the dataset, model and baselines are attached in Appendix E.

We vary the amount of training and validation data and report the time spent in influence analysis in Table 5. "Hes." indicates whether the method needs to estimate the Hessian of the model. "Max_Iter." denotes CG iteration in (Martens et al., 2010), and is set according to T . In LISSA, "depth" and "repeat" are two hyper-parameters that influence the computation time, and generally, their products should be equal to T . By comparing the baselines and HESIT on (T, V) training-validation dataset pairs of varying sizes, we find that HESIT, Hydra and LISSA require comparable calculation time, whereas baseline CG is significantly slower.

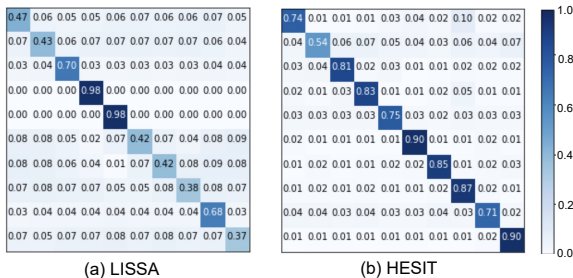


Figure 3: Inter-class and Intra-class (diagonal elements) contributions for CIFAR-10 dataset, which are measured by (a) LISSA and (b) HESIT.

Method	Hes.	# (T, V)	Setting	Time (Sec.)
CG	✓	$(10^2, 10)$	$Max_Iter = T$	23.1
		$(10^3, 10^2)$		665.8
		$(10^4, 10^3)$		-
LISSA	✓	$(10^2, 10)$	$Depth = T / 10$ $Repeat = 10$	3.6
		$(10^3, 10^2)$		13.2
		$(10^4, 10^3)$		136.7
Hydra	✗	$(10^2, 10)$	$Trace_ID = T$	2.4
		$(10^3, 10^2)$		6.9
		$(10^4, 10^3)$		161.1
HESIT	✗	$(10^2, 10)$	$Trace_ID = T$	2.2
		$(10^3, 10^2)$		6.3
		$(10^4, 10^3)$		132.9

Table 5: Computation time for data analysis methods. # (T, V) denotes that measure the influence of T training data to the performance on V validation data. The computing devices are AMD-EPYC 7763 (CPU) and NVIDIA A100 (GPU).

To further demonstrate that HESIT can still properly qualify data contribution even without estimating inverse Hessian, we then visualize the inter-class and intra-class contributions in Fig 3, where diagonal elements denote the contribution of the intra-class training data to the test class. It is observed that without estimating inverse Hessian, HESIT can still successfully measure the intra-class data influence by hyper-gradient tracing. In addition, it is observed that compared with LISSA, HESIT provides a higher estimation of intra-class contribution due to iteratively tracing.

6 Conclusion

In this paper, we propose HESIT to address the catastrophic forgetting in ToDs. It selects exemplar for rehearsal-based CL methods that examines training examples in each domain from a performance perspective. Furthermore, HESIT tracks the hyper-gradient of training examples in an optimization method that is Hessian-free and compatible with large pre-trained models. Experiments show that HESIT effectively overcomes catastrophic forgetting and delivers state-of-the-art performance on the largest ToDs benchmark for CL.

Limitations

When selecting exemplars for the current domain, HESIT only evaluates the performance on the self-domain validation set. However, we suggest that cross-domain influences should be taken into account, which means that the validation set of the later domain in the curriculum might be utilised to examine the picked example of the earlier domain. We would leave it as future work. In addition, we can only compare our model to the state-of-the-art work on ToDs from 2020, as there has been little work on ToDs employing CL in recent years compared to the extensive explorations of CL in computer vision, especially classification task.

Ethical Consideration

In developing and evaluating HESIT, our hyper-gradient-based exemplar strategy for Intelligent task-oriented dialogue systems (ToDs), we considered a number of significant ethical implications to ensure that our research conforms to generally accepted standards for the development of responsible AI.

Privacy and Data Security: While HESIT requires the selection of exemplars from the training data, we are mindful of user privacy and data security. The data used for training and testing HESIT are either publicly accessible or anonymized and used with the proper permissions, ensuring that no personally identifiable user information is exploited.

Bias and Fairness: Despite the optimistic results demonstrated by HESIT, it is essential to recognise the inherent possibility of bias in any AI system. We have attempted to mitigate this risk by assuring a diverse and representative training dataset, thereby preventing the model from favouring certain task domains disproportionately over others. Future work will include continued evaluations of potential biases with the objective of developing additional bias mitigation strategies.

Transparency and Accountability: The performance of HESIT is primarily dependent on the exemplars chosen from each task domain. This strategy is effective in terms of model performance, but if not managed carefully, it could contribute to a lack of transparency. We intend to maintain detailed documentation regarding the exemplar selection procedure and to make our algorithms, methods, and results accessible and interpretable to both practitioners and users.

Impact on Employment: The development of ToDs with the capacity for continuous learning may raise concerns regarding potential employment displacement. While our research contributes to the advancement of AI, it is not intended to supplant humans but rather to enhance their capabilities and productivity. To ensure a responsible transition to this technology, we encourage ongoing discussions regarding this issue.

Potential Misuse: As with all AI technologies, there is the potential for misuse with our method. If HESIT is used with nefarious intent or if the model is fed inappropriate data, there is a risk of abuse. We strongly advise establishing appropriate precautions and usage policies.

Long-term effects: We recognise that the long-term effects of HESIT and analogous technologies are indeterminate and require constant monitoring. As researchers, we are committed to continuous monitoring of our technology, its applications, and their societal implications.

In conclusion, we believe that the advantages of developing Continual Learning systems such as HESIT must be weighed against these ethical considerations. As a result, we do not perceive this work as a destination, but rather as a stepping stone on the path towards responsible and beneficial AI.

7 Acknowledgement

This work is funded by the Zhejiang Provincial “Jianbing” “Lingyan” Research and Development Program of China (2024C01135), National Natural Science Foundation of China (62302433, U23A20496), Zhejiang Provincial Natural Science Foundation of China (LQ24F020007) and CCF-Tencent Rhino-Bird Fund (RAGR20230122).

References

- Naman Agarwal, Brian Bullins, and Elad Hazan. 2017. Second-order stochastic optimization for machine learning in linear time. *The Journal of Machine Learning Research*, 18(1):4148–4187.
- Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. 2019. Gradient based sample selection for online continual learning. *Advances in neural information processing systems*, 32.
- Sanghwan Bae, Donghyun Kwak, Sungdong Kim, Donghoon Ham, Soyoun Kang, Sang-Woo Lee, and Woomyoung Park. 2022. Building a role specified open-domain dialogue system leveraging large-scale language models. *arXiv preprint arXiv:2205.00176*.

- Yoshua Bengio. 2000. Continuous optimization of hyper-parameters. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, volume 1, pages 305–310. IEEE.
- Magdalena Biesialska, Katarzyna Biesialska, and Marta R Costa-Jussa. 2020. Continual lifelong learning in natural language processing: A survey. *arXiv preprint arXiv:2012.09823*.
- Tamara Broderick, Nicholas Boyd, Andre Wibisono, Ashia C Wilson, and Michael I Jordan. 2013. Streaming variational bayes. *Advances in neural information processing systems*, 26.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Inigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. Multiwoz—a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. *arXiv preprint arXiv:1810.00278*.
- Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Daniel Duckworth, Semih Yavuz, Ben Goodrich, Amit Dubey, Andy Cedilnik, and Kyu-Young Kim. 2019. Taskmaster-1: Toward a realistic and diverse dialog dataset. *arXiv preprint arXiv:1909.05358*.
- Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. 2018. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*.
- Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and M Ranzato. 2019. Continual learning with tiny episodic memories. *arXiv preprint arXiv:1902.10486*.
- Yuanyuan Chen, Boyang Li, Han Yu, Pengcheng Wu, and Chunyan Miao. 2021. Hydra: Hypergradient data relevance analysis for interpreting deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7081–7089.
- Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. 2019. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277.
- Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. 2021. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385.
- Beyza Ermis, Giovanni Zappella, Martin Wistuba, Aditya Rawal, and Cédric Archambeau. 2022. Continual learning with transformers for image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3774–3781.
- Han Guo, Nazneen Fatema Rajani, Peter Hase, Mohit Bansal, and Caiming Xiong. 2020. Fastif: Scalable influence functions for efficient model interpretation and debugging. *arXiv preprint arXiv:2012.15781*.
- Yuchen Hu, Chen Chen, Chao-Han Huck Yang, Chengwei Qin, Pin-Yu Chen, Eng Siong Chng, and Chao Zhang. 2024. Self-taught recognizer: Toward unsupervised adaptation for speech foundation models. *arXiv preprint arXiv:2405.14161*.
- Yufan Huang, Yanzhe Zhang, Jiaao Chen, Xuezhi Wang, and Diyi Yang. 2021. Continual learning for text classification with information disentanglement based regularization. *arXiv preprint arXiv:2104.05489*.
- David Isele and Akansel Cosgun. 2018. Selective experience replay for lifelong learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Mihir Kale and Abhinav Rastogi. 2020. Few-shot natural language generation by rewriting templates. *arXiv preprint arXiv:2004.15006*.
- 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.
- Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In *International conference on machine learning*, pages 1885–1894. PMLR.
- Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images.
- Sungjin Lee. 2017. Toward continual learning for conversational agents. *arXiv preprint arXiv:1712.09943*.
- Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. 2019. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *International Conference on Machine Learning*, pages 3925–3934. PMLR.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, Peng Xu, Feijun Jiang, Yuxiang Hu, Chen Shi, and Pascale Fung. 2021. Bitod: A bilingual multi-domain dataset for task-oriented dialogue modeling. *arXiv preprint arXiv:2106.02787*.
- Andrea Madotto, Zhaojiang Lin, Zhenpeng Zhou, Seungwhan Moon, Paul Crook, Bing Liu, Zhou Yu, Eunjoon Cho, and Zhiguang Wang. 2020. Continual learning in task-oriented dialogue systems. *arXiv preprint arXiv:2012.15504*.
- Zheda Mai, Ruiwen Li, Jihwan Jeong, David Quispe, Hyunwoo Kim, and Scott Sanner. 2022. Online continual learning in image classification: An empirical survey. *Neurocomputing*, 469:28–51.

- Davide Maltoni and Vincenzo Lomonaco. 2019. Continuous learning in single-incremental-task scenarios. *Neural Networks*, 116:56–73.
- James Martens et al. 2010. Deep learning via hessian-free optimization. In *ICML*, volume 27, pages 735–742.
- 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, Liangwei Chen, Mengjie Zhao, Minlie Huang, and Boi Faltings. 2020. Continual learning for natural language generation in task-oriented dialog systems. *arXiv preprint arXiv:2010.00910*.
- Martin Mundt, Yong Won Hong, Iuliia Plushch, and Visvanathan Ramesh. 2020. A wholistic view of continual learning with deep neural networks: Forgotten lessons and the bridge to active and open world learning. *arXiv preprint arXiv:2009.01797*.
- 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 of the Association for Computational Linguistics*, pages 311–318.
- Garima Pruthi, Frederick Liu, Satyen Kale, and Mukund Sundararajan. 2020. Estimating training data influence by tracing gradient descent. *Advances in Neural Information Processing Systems*, 33:19920–19930.
- Chengwei Qin, Chen Chen, and Shafiq Joty. 2023. [Lifelong sequence generation with dynamic module expansion and adaptation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6701–6714, Singapore. Association for Computational Linguistics.
- Chengwei Qin, Ruirui Chen, Ruochen Zhao, Wenhan Xia, and Shafiq Joty. 2024. [Lifelong event detection with embedding space separation and compaction](#). In *2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics*.
- Chengwei Qin and Shafiq Joty. 2022a. [Continual few-shot relation learning via embedding space regularization and data augmentation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2776–2789, Dublin, Ireland. Association for Computational Linguistics.
- Chengwei Qin and Shafiq Joty. 2022b. [LFPT5: A unified framework for lifelong few-shot language learning based on prompt tuning of t5](#). In *International Conference on Learning Representations*.
- Haoxuan Qu, Hossein Rahmani, Li Xu, Bryan Williams, and Jun Liu. 2021. Recent advances of continual learning in computer vision: An overview. *arXiv preprint arXiv:2109.11369*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- 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.
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. 2019. Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32.
- 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*.
- Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. 2018. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pages 4548–4557. PMLR.
- Fan-Keng Sun, Cheng-Hao Ho, and Hung-Yi Lee. 2019. Lamol: Language modeling for lifelong language learning. *arXiv preprint arXiv:1909.03329*.
- Qing Sun, Fan Lyu, Fanhua Shang, Wei Feng, and Liang Wan. 2022. Exploring example influence in continual learning. *arXiv preprint arXiv:2209.12241*.
- Eli Verwimp, Matthias De Lange, and Tinne Tuytelaars. 2021. Rehearsal revealed: The limits and merits of revisiting samples in continual learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9385–9394.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned lstm-based natural language generation for spoken dialogue systems. *arXiv preprint arXiv:1508.01745*.
- Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. 2020. Supermasks in superposition. *Advances in Neural Information Processing Systems*, 33:15173–15184.
- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019. Transferable multi-domain state generator for task-oriented dialogue systems. *arXiv preprint arXiv:1905.08743*.
- Ju Xu and Zhanxing Zhu. 2018. Reinforced continual learning. *Advances in Neural Information Processing Systems*, 31.

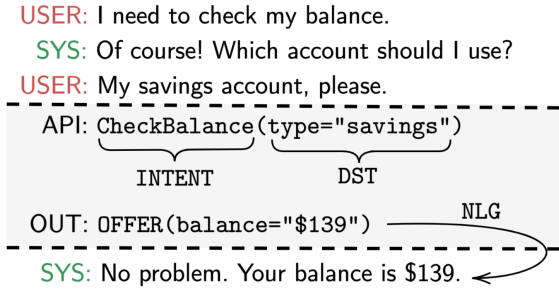


Figure 4: Example of input-out pairs, for the four settings, INTENT, DST, NLG and end-to-end (E2E).

A Large Language Model and ToDs

Large Language Models (LLMs) have emerged as an epistemic beacon in the field of natural language processing (NLP), endowing text-based tasks with substantial performance gains including ToDs (Bae et al., 2022). Furthermore, a foundation model with multiple adapters seems to be potential solution for overcoming catastrophic forgetting. However, this combination can not fit continual learning in ToDs, especially with increasing task domains. Firstly, the task domain of inference data is agnostic that requires a extra procedure to determine which of adapter should be activated for inference. Secondly, no matter what procedure is utilized, all of 37 combinations need to be traversed once, resulting in unacceptable delay for ToDs. Therefore, instead of using LLMs with adapter, we utilize single model to adapt to all domains.

B Examples of E2E ToDs

We first give an example to elaborate INTENT, DST, NLG, and E2E. As shown in Fig 4, the grey box denotes the API-call $\mathcal{C}(\mathbf{H})$ process, where the user intents (INTENT) and slot-value pairs (DST) are both predicted according to dialogue history \mathbf{H} . Then the system will generate a response (NLG) based on both \mathbf{H} and $\mathcal{C}(\mathbf{H})$: $\mathbf{H} + \mathcal{C}(\mathbf{H}) \rightarrow \mathbf{R}$. In addition, when $\mathcal{C}(\mathbf{H})$ is empty, the system can directly generate responses based on history: $\mathbf{H} \rightarrow \mathbf{R}$.

C Methodology Supplement

For each task domain, Let \mathcal{D}_t consists of training set $\mathcal{D}_t^{\text{trn}}$, validation set $\mathcal{D}_t^{\text{val}}$, and test set $\mathcal{D}_t^{\text{tst}}$. Assume $\mathcal{D}_t^{\text{trn}}$ contains N training-points which is defined as $z_i = (x_i, y_i)$. Our goal is to pick out those influential z_i s that benefit model performance on unseen test-point $z' \in \mathcal{D}_t^{\text{val}}$, as $\mathcal{D}_t^{\text{tst}}$ is unavailable during training. Leveraging derivation chain rule,

the IFs (Koh and Liang, 2017) define the influence of z on z' as:

$$\mathcal{I}(z, z') = -\nabla_{\theta} \ell(z', \hat{\theta})^T H_{\hat{\theta}}^{-1} \nabla_{\theta} \ell(z, \hat{\theta}) \quad (7)$$

where $H_{\hat{\theta}} = \nabla_{\theta}^2 \ell(z, \hat{\theta})$ is the Hessian and is positive definite (PD) by assumption. In practice, $H_{\hat{\theta}}$ can not be materialized in memory, let alone be inverted by standard linear algebra.

Mainstream methods employ LISSA (Agarwal et al., 2017), which only samples a single point per iteration as an estimator of inverse Hessian. Compared with the standard transformation of matrix inversion (Martens et al., 2010), LISSA shows significant speedups in the following iteration r :

$$H_r^{-1} v = v + (1 - H) H_{r-1}^{-1} v \quad (8)$$

where H is approximated on random batch and $v = \nabla_{\theta} \ell(z', \hat{\theta})$ is a gradient. However, this method faces two drawbacks. Firstly, only the final parameters are utilized to calculate the gradient, while the data influence is involved in a dynamic optimization process. Secondly, Hessian might not be positive definite for a large pre-trained model, and the estimation is not accurate enough. In practice, we found that using LISSA to estimate the inverse Hessian of GPT-2 is unstable and time-consuming.

To address the above defects, *TracIn* (Pruthi et al., 2020) traces the changes in the loss across all gradient steps, and avoids the estimation of inverse Hessian by the following definition:

$$\mathcal{I}_{\text{TracIn}}(z, z') = \frac{1}{C} \sum_{i=1}^C \gamma_i \nabla_{\theta_i} \ell(z, \hat{\theta}_i) \cdot \nabla_{\theta_i} \ell(z', \hat{\theta}_i) \quad (9)$$

where C is the number of the checkpoints during training. Despite considering the optimization process, the computing complexity is C times slower than that of using exact gradient similarity and, as discussed in (Pruthi et al., 2020), care needs to be taken in selecting checkpoints. Another practical obstacle of *TracIn* in CL settings is that only several epochs are trained for each domain, resulting in insufficient C for influence analysis.

D Dataset Statistics

The main dataset statistics are shown in Table 6. It contains 37 domains, 280 types of intents, and more than 31K training dialogues. The detailed example amounts of modules as well as training, validation, and test set are illustrated in Table 7.

Name	Train	Valid	Test	Dom.	Intents	Turns
TM19	4,403	551	553	6	112	19.97
TM20	13,839	1,731	1,734	7	128	16.92
MWoZ	7,906	1,000	1,000	5	15	13.93
SGD	5,278	761	1,531	19	43	14.71
Total	31,426	4,043	4,818	37	280	16.23

Table 6: Main datasets statistics.

Due to a serious imbalance of task domain, we observe that some CL methods (e.g., EWC) are vulnerable to the task order in the curriculum. If the model first learns on a small domain, then the regularization item would influence the subsequent learning on large domains. Meanwhile, we found the data replay methods are more insensitive to task order, as the data amount of each domain stored for retraining is equal. In addition, we have also attempted to sample different amounts of exemplars according to the size of the domain. However, it achieves similar results compared with sampling an equal amount of exemplars for each domain.

E Experimental Details

E.1 CIFAR-10 Dataset

The CIFAR-10 dataset (Krizhevsky et al., 2009) consists of 60000 32x32 color images in 10 classes, which can be split into 50000 training images and 10000 test images. There are 10 different classes in the CIFAR-10 dataset, including cars, birds, cats, deer, dogs, frogs, horses, ships, and trucks. Each class contains 6000 images.

E.2 CNN backbone

We employ a simple convolutional neural network, including 2×2 D CNN layers with pooling layers, then a ReLU activation function and $2 \times$ linear layer are added to predict the labels.

In order to obtain Fig. 3, the CNN network is trained by an SGD optimizer with a momentum of 0.9. The learning rate is $1 \times e^{-3}$ and 50 training epochs are repeated. Then we calculate the contribution of training data to test data regarding different categories.

Domains	DST & INTENT			NLG			End-to-End		
	<i>Train</i>	<i>Dev</i>	<i>Test</i>	<i>Train</i>	<i>Dev</i>	<i>Test</i>	<i>Train</i>	<i>Dev</i>	<i>Test</i>
TM19 movie	4733	584	500	3010	366	341	12766	1632	1481
TM19 auto	3897	448	522	2128	223	283	10918	1248	1443
TM19 restaurant	4434	568	561	2582	330	333	12862	1669	1630
TM19 pizza	2883	381	359	1326	171	171	8720	1145	1083
TM19 uber	4378	535	525	2418	290	278	11331	1362	1361
TM19 coffee	2591	302	335	1381	151	184	7429	894	936
TM20 flight	15868	1974	1940	10148	1272	1245	36778	4579	4569
TM20 food-ordering	3404	411	431	2394	277	287	7838	941	986
TM20 hotel	15029	1908	1960	6590	842	869	35022	4400	4532
TM20 music	5917	764	769	4196	537	523	13723	1773	1787
TM20 restaurant	13738	1761	1691	8356	1063	994	34560	4398	4297
TM20 sport	13072	1668	1654	12044	1553	1542	29391	3765	3723
TM20 movie	13221	1703	1567	9406	1203	1093	32423	4158	3881
MWOZ taxi	1239	234	194	402	71	56	2478	468	388
MWOZ train	1452	158	160	563	63	59	2905	316	320
MWOZ restaurant	5227	243	281	3333	141	177	10461	486	563
MWOZ hotel	2798	289	385	1924	194	258	5602	579	771
MWOZ attraction	484	43	42	295	27	26	975	86	85
sgd restaurants	2686	278	616	1720	166	386	5756	606	1354
sgd media	1411	230	458	988	167	324	3114	502	1005
sgd events	4881	598	989	3241	389	590	10555	1317	2197
sgd music	1892	275	556	1506	224	464	4040	597	1215
sgd movies	1665	181	52	996	114	44	3760	420	126
sgd flights	4766	1041	1756	2571	627	982	10429	2244	3833
sgd ridesharing	652	85	187	377	48	107	1448	188	418
sgd rentalcars	1510	250	469	865	153	280	3277	538	1009
sgd buses	1862	331	653	1102	218	412	4050	709	1393
sgd hotels	3237	394	948	1997	243	597	6983	858	2053
sgd services	3328	360	926	2225	230	611	7262	803	2016
sgd homes	2098	170	533	1312	96	338	4519	394	1158
sgd banks	1188	139	293	723	84	181	2599	319	667
sgd calendar	592	115	236	397	65	133	1313	246	501
sgd alarm	212	34	91	221	30	74	580	82	198
sgd weather	196	32	80	123	23	59	433	70	169
sgd travel	186	23	48	121	14	30	420	53	106
sgd payment	227	21	51	143	14	32	497	44	113
sgd trains	300	73	128	149	43	66	668	158	274
Total	147254	18604	22946	93273	11722	14429	347885	44047	53641

Table 7: Detained data amounts for each domian and module