

Total Recall: a Customized Continual Learning Method for Neural Semantic Parsers

Zhuang Li, Lizhen Qu,* Gholamreza Haffari

Faculty of Information Technology

Monash University

firstname.lastname@monash.edu

Abstract

This paper investigates continual learning for semantic parsing. In this setting, a neural semantic parser learns tasks sequentially without accessing full training data from previous tasks. Direct application of the SOTA continual learning algorithms to this problem fails to achieve comparable performance with re-training models with all seen tasks, because they have not considered the special properties of structured outputs, yielded by semantic parsers. Therefore, we propose TOTAL RECALL, a continual learning method designed for neural semantic parsers from two aspects: i) a sampling method for memory replay that diversifies logical form templates and balances distributions of parse actions in a memory; ii) a two-stage training method that significantly improves generalization capability of the parsers across tasks. We conduct extensive experiments to study the research problems involved in continual semantic parsing, and demonstrate that a neural semantic parser trained with TOTAL RECALL achieves superior performance than the one trained directly with the SOTA continual learning algorithms, and achieve a 3-6 times speedup compared to retraining from scratch. Code and datasets are available at: https://github.com/zhuang-li/cl_nsp.

1 Introduction

In the recent market research report published by MarketsandMarkets (INC., 2020), it is estimated that the smart speaker market is expected to grow from USD 7.1 billion in 2020 to USD 15.6 billion by 2025. Commercial smart speakers, such as Alexa and Google assistant, often need to translate users' commands and questions into actions. Therefore, semantic parsers are widely adopted in dialogue systems to map natural language (NL) utterances to executable programs or logical forms

(LFs) (Damonte et al., 2019; Rongali et al., 2020). Due to the increasing popularity of such speakers, software developers have implemented a large volume of skills for them and the number of new skills grow quickly every year. For example, as of 2020, the number of Alexa skills exceeds 100,000 and 24 new skills are introduced per day in 2020 (KINSELLA, 2020). Although machine learning-based semantic parsers achieve the state-of-the-art performance, they face the following challenges due to the fast growing number of tasks.

Given new tasks, one common practice is to re-train the parser from scratch on the training data of all seen tasks. However, it is both economically and computationally expensive to re-train semantic parsers because of a fast-growing number of new tasks (Lialin et al., 2020). To achieve its optimal performance, training a deep model on all 8 tasks of NLMap (Lawrence and Riezler, 2018) takes approximately 14 times longer than training the same model on one of those tasks. In practice, the cost of repeated re-training for a commercial smart speaker is much higher, e.g. Alexa needs to cope with the number of tasks which is over 10,000 times more than the one in NLMap¹. In contrast, *continual learning* provides an alternative cost-effective training paradigm, which learns tasks sequentially without accessing full training data from the previous tasks, such that the computational resources are utilized only for new tasks.

Privacy leakage has gradually become a major concern in many Artificial Intelligence (AI) applications. As most computing environments are not 100% safe, it is not desirable to always keep a copy of the training data including identifiable personal information. Thus, it is almost not feasible to assume that complete training data of all known tasks is always available for re-training a semantic

¹A rough estimation: re-training of our semantic parser for 100,000 tasks will take more than 138 days (2 mins of training time per NLMap task × 100,000) on a single-GPU machine.

*corresponding author

parser (Irfan et al., 2021). For the semantic parser of a privacy-sensitive AI system, e.g. personalized social robot, *continual learning* provides a solution to maintain the knowledge of all learned tasks when the complete training data of those data is not available anymore due to security reasons.

A major challenge of continual learning lies in *catastrophic forgetting* that the (deep) models easily *forget* the knowledge learned in the previous tasks when they learn new tasks (French, 1991; Mi et al., 2020). Another challenge is to learn what kind of knowledge the tasks share in common and support fast adaptation of models for new tasks. Methods are developed to mitigate catastrophic forgetting (Lopez-Paz and Ranzato, 2017; Han et al., 2020) and facilitate forward knowledge transfer (Li and Hoiem, 2017). Instead of directly measuring speedup of training, those methods assume that there is a small fixed-size memory available for storing training examples or parameters from the previous tasks. The memory limits the size of training data thus proportionally reduces training time. However, we empirically found that direct application of those methods to neural semantic parsers leads to a significant drop of test performance on benchmark datasets, in comparison to re-training them with all available tasks each time.

In this work, we investigated the applicability of existing continual learning methods to semantic parsing in-depth, and we have found that most methods have not considered the special properties of structured outputs, which distinct semantic parsing from the multi-class classification problem. Therefore, we propose TOTAL RECALL (TR), a continual learning method that is especially designed to address the semantic parsing specific problems from two perspectives. First, we customize the sampling algorithm for memory replay, which stores a small sample of examples from each previous task when continually learning new tasks. The corresponding sampling algorithm, called **D**iversified **L**ogical **F**orm **S**election (DLFS), diversifies LF templates and maximizes the entropy of the parse action distribution in a memory. Second, motivated by findings from cognitive neuroscience (Goyal and Bengio, 2020), we facilitate knowledge transfer between tasks by proposing a two-stage training procedure, called **F**ast **S**low **C**ontinual **L**earning (FSCL). It updates only unseen action embeddings in the fast-learning stage and updates all model parameters in the follow-up

stage. As a result, it significantly improves generalization capability of parsing models.

Our key contributions are as follows:

- We conduct the *first* in-depth empirical study of the problems encountered by neural semantic parsers to learn a *sequence* of tasks continually in various settings. The most related work (Lialin et al., 2020) only investigated incremental learning between two semantic parsing tasks.
- We propose DLFS, a sampling algorithm for memory replay that is customized for semantic parsing. As a result, it improves the best sampling methods of memory replay by 2-11% on Overnight (Wang et al., 2015a).
- We propose a two-stage training algorithm, coined FSCL, that improves the test performance of parsers across tasks by 5-13% in comparison with using only Adam (Kingma and Ba, 2014).
- In our extensive experiments, we investigate applicability of the SOTA continual learning methods to semantic parsing with *three* different task definitions, and show that TR outperforms the competitive baselines by 4-9% and achieves a speedup by 3-6 times compared to training from scratch.

2 Related Work

Semantic Parsing The recent surveys (Kamath and Das, 2018; Zhu et al., 2019; Li et al., 2020) cover an ample of work in semantic parsing. Most current work employ a sequence-to-sequence architecture (Sutskever et al., 2014) to map an utterance into a structured meaning representations, such as LFs, SQL, and abstract meaning representation (Banarescu et al., 2013). The output sequences are either linearized LFs (Dong and Lapata, 2016, 2018; Cao et al., 2019) or sequences of parse actions (Chen et al., 2018; Cheng et al., 2019; Lin et al., 2019; Zhang et al., 2019; Yin and Neubig, 2018; Chen et al., 2018; Guo et al., 2019; Wang et al., 2020a; Li et al., 2021). There are also work (Guo et al., 2019; Wang et al., 2020a; Li et al., 2021) exploring semantic parsing with unseen database schemas or actions. Feedback semantic parsing interactively collects data from user feedbacks as continuous data streams but does not address the problem of catastrophic forgetting or

improve forward transfer (Iyer et al., 2017; Yao et al., 2019; Labutov et al., 2018).

Continual Learning The continual learning methods can be coarsely categorized into i) regularization-based methods (Kirkpatrick et al., 2017; Zenke et al., 2017; Ritter et al., 2018; Li and Hoiem, 2017; Zhao et al., 2020; Schwarz et al., 2018) which either applies knowledge distillation (Hinton et al., 2015) to penalize the loss updates or regularizes parameters which are crucial to the old tasks; ii) dynamic architecture methods (Mallya and Lazebnik, 2018; Serra et al., 2018; Maltoni and Lomonaco, 2019; Houlsby et al., 2019; Wang et al., 2020b; Pfeiffer et al., 2021; Rusu et al., 2016) which dynamically alter the structures of models to reduce the catastrophic forgetting; iii) memory-based methods (Lopez-Paz and Ranzato, 2017; Wang et al., 2019; Han et al., 2020; Aljundi et al., 2019; Chrysakis and Moens, 2020; Kim et al., 2020) which stores historical instances and continually learn them along with instances in new tasks. There are also hybrid methods (Mi et al., 2020; Liu et al., 2020; Rebuffi et al., 2017) which integrate more than one type of such methods.

In natural language processing (NLP), continual learning is applied to tasks such as relation extraction (Wang et al., 2019; Han et al., 2020), natural language generation (Mi et al., 2020), language modelling (Sun et al., 2019) and the pre-trained language models adapting to multiple NLP tasks (Wang et al., 2020b; Pfeiffer et al., 2021). To the best of our knowledge, (Lialin et al., 2020) is the only work studying catastrophic forgetting for semantic parsing. However, they consider learning between *only* two tasks, have not proposed new methods, and also have not evaluated recently proposed continual learning methods. In contrast, we propose two novel continual learning methods customized for semantic parsing and compare them with the strong and recently proposed continual learning methods that are not applied to semantic parsing before.

3 Base Parser

A semantic parser learns a mapping $\pi_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ to convert an natural language (NL) utterance $x \in \mathcal{X}$ into its corresponding logical form (LF) $y \in \mathcal{Y}$. Most SOTA neural semantic parsers formulate this task as translating a word sequence into an output sequence, whereby an output sequence is either a sequence of LF tokens or a sequence of parse ac-

tions that construct an LF. For a fair comparison between different continual learning algorithms, we adopt the same base model for them, as commonly used in prior works (Lopez-Paz and Ranzato, 2017; Wang et al., 2019; Han et al., 2020).

Similar to (Shin et al., 2019; Iyer et al., 2019), the base parser converts the utterance x into a sequence of actions $\mathbf{a} = \{a_1, \dots, a_t\}$. As an LF can be equivalently parsed into an abstract syntax tree (AST), the actions \mathbf{a} sequentially construct an AST deterministically in the depth-first order, wherein each action a_t at time step t either i) expands an intermediate node according to the production rules from a grammar, or ii) generates a leaf node. As in (Shin et al., 2019), the idioms (frequently occurred AST fragments) are collapsed into single units. The AST is further mapped back to the target LF.

The parser employs the attention-based sequence-to-sequence (SEQ2SEQ) architecture (Luong et al., 2015) for estimating action probabilities.

$$P(\mathbf{a}|x) = \prod_{t=1}^{|\mathbf{a}|} P(a_t|\mathbf{a}_{<t}, x) \quad (1)$$

Encoder. The encoder in SEQ2SEQ is a standard bidirectional Long Short-term Memory (LSTM) network (Hochreiter and Schmidhuber, 1997), which encodes an utterance x into a sequence of contextual word representations.

Decoder. The decoder applies an LSTM to generate action sequences. At time t , the decoder produces an action representation \mathbf{s}_t , which is yielded by concatenating the hidden representation \mathbf{h}_t produced by the LSTM and the context vector \mathbf{o}_t produced by the soft attention (Luong et al., 2015).

We maintain an embedding for each action in the embedding table. The probability of an action a_t is estimated by:

$$P(a_t|\mathbf{a}_{<t}, x) = \frac{\exp(\mathbf{c}_{a_t}^\top \mathbf{s}_t)}{\sum_{a' \in \mathcal{A}_t} \exp(\mathbf{c}_{a'}^\top \mathbf{s}_t)} \quad (2)$$

where \mathcal{A}_t is the set of applicable actions at time t , and \mathbf{c}_a is the embedding of the action a_t , which is referred to as *action embedding* in the following.

4 Continual Semantic Parsing

Problem Formulation. We consider a widely adopted continual learning setting that a parser π_θ is trained continually on a sequence of K distinct tasks $\{\mathcal{T}^{(1)}, \mathcal{T}^{(2)}, \dots, \mathcal{T}^{(K)}\}$ (Lopez-Paz and Ranzato, 2017; Han et al., 2020). In both training and testing, we know which task an example belongs

to. As the definition of tasks is application specific and parallel data of semantic parsing is often created by domain experts, it is easy to identify the task of an example in practice. We further assume that there is a fixed-size memory \mathcal{M}_k associated with each task $\mathcal{T}^{(k)}$ for e.g. storing a small amount of replay instances, as adopted in (Rebuffi et al., 2017; Wang et al., 2019). This setting is practical for personalized conversational agents because it is difficult for them to re-collect past information except reusing the ones in the memories.

4.1 Challenges

We demonstrate catastrophic forgetting in continual semantic parsing by training the base parser sequentially with each task from the OVERNIGHT corpus (Wang et al., 2015a) and report the test accuracy of exactly matched LFs of all seen tasks combined (More evaluation details are in Sec. 5).

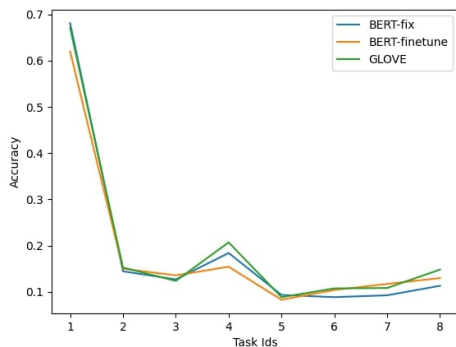


Figure 1: The accuracy on OVERNIGHT after the parser being trained on each task. The parser uses GLOVE (Pennington et al., 2014), BERT (Devlin et al., 2019) with parameters updated (BERT-finetune) and fixed (BERT-fix) as the input embedding layer.

Fig. 1 shows the performance of continually training the base parser with BERT (Devlin et al., 2019) (with and without fine-tuning BERT parameters) and GLOVE respectively by using the standard cross entropy loss. The accuracy on the combined test set drops dramatically after learning the *second* task. The training on the initial task appears to be crucial on mitigating catastrophic forgetting. The BERT-based parser with/without fine-tuning obtains no improvement over the one using GLOVE. The forgetting with BERT is even more serious compared with using GLOVE. The same phenomenon is also observed in (Arora et al., 2019) that the models with pre-trained language models obtain inferior performance than LSTM or CNN, when fine-tuning incrementally on each task. They

conjecture that it is more difficult for models with large capacity to mitigate catastrophic forgetting.




Figure 2: The average conditional probabilities $P(a_t | \mathbf{a}_{<t}, \mathbf{x})$ of the representative cross-task (solid) and task-specific (dash) actions till the seen tasks on OVERNIGHT after learning on each task sequentially. The boxes at i th task indicate the actions from the initial task also exist in the i th task.

We further investigate which parse actions are easy to forget. To measure the degree of forgetfulness of an action, after training the parser in the first task, we average the probabilities $P(a_t | \mathbf{a}_{<t}, \mathbf{x})$ produced by the parser on the training set of the first task. We recompute the same quantity after learning each task sequentially and plot the measures. Fig. 2 depicts the top two and the bottom two actions are easiest to forget on average. Both top two actions appear only in the first task, thus it is difficult for the parser to remember them after learning new tasks. In contrast, cross-task actions, such as $GEN(string <)$, may even obtain improved performance after learning on the last task. Thus, it indicates the importance of differentiating between task-specific actions and cross-task actions when designing novel continual learning algorithms.

4.2 TOTAL RECALL

To save training time for each new task, we cannot use all training data from previous tasks, thus we introduce a designated sampling method in the sequel to fill memories with the examples most likely mitigating catastrophic forgetting. We also present the two-stage training algorithm FSCl to facilitate knowledge transfer between tasks.

Sampling Method. DLFS improves Episodic Memory Replay (EMR) (Wang et al., 2019; Chaudhry et al., 2019) by proposing a designated sampling method for continual semantic parsing. EMR utilizes a memory module $\mathcal{M}_k = \{(\mathbf{x}_1^{(k)}, \mathbf{y}_1^{(k)}), \dots, (\mathbf{x}_M^{(k)}, \mathbf{y}_M^{(k)})\}$ to store a few exam-

ples sampled from the training set of task $\mathcal{T}^{(k)}$, where $(\mathbf{x}_m^{(k)}, \mathbf{y}_m^{(k)}) \in \mathcal{D}_{train}^{(k)}$ and M is the size of the memory. The training loss of EMR takes the following form:

$$\mathcal{L}_{EMR} = \mathcal{L}_{\mathcal{D}_{train}^{(k)}} + \sum_{i=1}^{k-1} \mathcal{L}_{\mathcal{M}_i} \quad (3)$$

where $\mathcal{L}_{\mathcal{D}_{train}^{(k)}}$ and $\mathcal{L}_{\mathcal{M}_i}$ denotes the loss on the training data of current task $\mathcal{T}^{(k)}$ and the memory data from task $\{\mathcal{T}\}_{i=1}^{(k-1)}$, respectively. The training methods for memory replay often adopt a subroutine called *replay training* to train models on instances in the memory. Furthermore, prior works (Aljundi et al., 2019; Wang et al., 2019; Han et al., 2020; Mi et al., 2020; Chrysakis and Moens, 2020; Kim et al., 2020) discovered that storing a small amount of *diversified* and *long-tailed* examples is helpful in tackling catastrophic forgetting for memory-based methods. Semantic parsing is a structured prediction problem. We observe that semantic parsing datasets are highly imbalanced w.r.t. LF structures. Some instances with similar LF structures occupy a large fraction of the training set. Therefore, we presume storing the diversified instances in terms of the corresponding LF structures would alleviate the problem of catastrophic forgetting in continual semantic parsing.

To sample instances with the diversified LF structures, our method DLFS partitions the LFs in \mathcal{D}_{train} into M clusters, followed by selecting representative instances from each cluster to maximize entropy of actions in a memory. To characterize differences in structures, we first compute similarities between LFs by $sim(\mathbf{y}_i, \mathbf{y}_j) = (Smatch(\mathbf{y}_i, \mathbf{y}_j) + Smatch(\mathbf{y}_j, \mathbf{y}_i))/2$, where *Smatch* (Cai and Knight, 2013) is a asymmetrical similarity score between two LFs yielded by calculating the overlapping percentage of their triples. Then we run a flat clustering algorithm using the distance function $1 - sim(\mathbf{y}_i, \mathbf{y}_j)$ and the number of clusters is the same as the size of a memory. We choose *K-medoids* (Park and Jun, 2009) in this work for easy interpretation of clustering results.

We formulate the problem of balancing action distribution and diversifying LF structures as the following constrained optimization problem. In particular, it i) aims to balance the actions of stored instances in the memory module \mathcal{M} by increasing the entropy of the action distribution, and ii) requires that each instance m in \mathcal{M} belongs to a different cluster c_j . Let the function $c(m)$ return

the cluster id of an instance in a memory \mathcal{M} and m_i denote its i th entry, we have

$$\begin{aligned} \max_{\mathcal{M}} & - \sum_{a_i \in \mathcal{A}} p_{\mathcal{M}}(a_i) \log p_{\mathcal{M}}(a_i) \\ \text{s.t.} & \forall m_i, m_j \in \mathcal{M}, c(m_i) \neq c(m_j) \end{aligned} \quad (4)$$

where $p_{\mathcal{M}}(a_i) = \frac{n_i}{\sum_{a_j \in \mathcal{A}} n_j}$, with n_i being the frequency of action a_i in \mathcal{M} and \mathcal{A} being the action set included in the training set \mathcal{D}_{train} . In some occasions, the action set \mathcal{A} is extremely large (e.g. 1000+ actions per task), so it may be infeasible to include all actions in the limited memory \mathcal{M} . We thus sample a subset of h actions, $\mathcal{A}' \subseteq \mathcal{A}$, given the distribution of $P_D(\mathcal{A})$ in \mathcal{D}_{train} where $p_D(a_i) = \frac{n_i}{\sum_{a_j \in \mathcal{A}} n_j}$, with n_i being the frequency of a_i in \mathcal{D}_{train} . In that case, our method addresses the optimization problem over the actions in \mathcal{A}' . We solve the above problem by using an iterative updating algorithm, whose details can be found in Appendix B. The closest works (Chrysakis and Moens, 2020; Kim et al., 2020) maintain only the balanced label distribution in the memory while our work maintains the balanced memory w.r.t. both the LF and action distributions.

Fast-Slow Continual Learning. Continual learning methods are expected to learn what the tasks have in common and in what the tasks differ. If there are some shared structures between tasks, it is possible to transfer knowledge from one task to another. Inspired by findings from cognitive neuroscience, the learning should be divided into slow learning of stationary aspects between tasks and fast learning of task-specific aspects (Goyal and Bengio, 2020). This is an inductive bias that can be leveraged to obtain cross-task generalization in the space of all functions.

We implement this inductive bias by introducing a two-stages training algorithm. In the base model, action embeddings \mathbf{c}_a (Eq. (7)) are task-specific, while the remaining parts of the model, which builds representations of utterances and action histories, are shared to capture common knowledge between tasks. Thus, in the *fast-learning* stage, we update only the embeddings of unseen actions $\mathbf{c}_a^{(i)}$ with the cross-entropy loss, in the *slow-learning* stage, we update all model parameters. Fast-learning helps parsers generalize to new tasks by giving the unseen actions good initialized embeddings and reduces the risk of forgetting prior knowledge by focusing on minimal changes of model parameters for task-specific patterns.

Algorithm 1: Fast-Slow Training for the k -th task

Input : Training set $\mathcal{D}_{train}^{(k)}$ of k -th task $\mathcal{T}^{(k)}$, memory data $\mathcal{M} = \{\mathcal{M}_1, \dots, \mathcal{M}_{k-1}\}$, the known action set $\mathcal{A}^{(1:k-1)}$ before learning $\mathcal{T}^{(k)}$

Extract action set $\mathcal{A}^{(k)}$ from $\mathcal{D}_{train}^{(k)}$

Obtain unseen actions $\mathcal{A}_u^{(k)}$ by excluding $\mathcal{A}^{(1:k-1)}$ from $\mathcal{A}^{(k)}$

fast-learning on unseen action embeddings

for $i \leftarrow 1$ **to** $epoch_1$ **do**

 | Update $\mathbf{c}_a^{(k)}$ with $\nabla_{\mathbf{c}_a^{(k)}} \mathcal{L}(\theta_g, \theta_s^{(k)})$ on $\mathcal{D}_{train}^{(k)}$

end

slow-learning stage

for $i \leftarrow 1$ **to** $epoch_2$ **do**

fine-tune all cross-task parameters and task-specific parameters of the current task

 Update $(\theta_g, \theta_s^{(k)})$ with $\nabla_{(\theta_g, \theta_s^{(k)})} \mathcal{L}(\theta_g, \theta_s^{(k)})$ on $\mathcal{D}_{train}^{(k)}$

replay training with task-specific parameters of the previous tasks frozen

for $\mathcal{M}_i \in \mathcal{M}$ **do**

 | Update θ_g with $\nabla_{\theta_g} \mathcal{L}(\theta_g, \theta_s^{(i)})$ on $\mathcal{M}^{(i)}$

end

end

In the fast-learning stage, the unseen actions $\mathcal{A}_u^{(k)}$ of the k -th task are obtained by excluding all historical actions from the action set of current task $\mathcal{T}^{(k)}$, namely $\mathcal{A}_u^{(k)} = \mathcal{A}^{(k)} \setminus \mathcal{A}^{(1:k-1)}$, where $\mathcal{A}^{(k)}$ denotes the action set of the k -th task. All actions are unseen in the first task, thus we update all action embeddings by having $\mathcal{A}_u^{(0)} = \mathcal{A}^{(0)}$. In the slow-learning stage, we differ updating parameters w.r.t. current task from updating parameters w.r.t. memories of previous tasks. For the former, the parameters θ_g shared across tasks are trained w.r.t. all the data while the task-specific parameters $\theta_s^{(i)}$ are trained only w.r.t. the data from task $\mathcal{T}^{(i)}$. For the latter, the task-specific parameters learned from the previous tasks are frozen to ensure they do not forget what is learned from previous tasks. More details can be found in Algo. 1. This training algorithm is closely related to Invariant Risk Minimization (Arjovsky et al., 2019), which learns invariant structures across different training environments. However, in their work, they assume the same label space across environments and have access to all training environments at the same time.

Loss During training, we augment the EMR loss with the Elastic Weight Consolidation (EWC) regularizer (Kirkpatrick et al., 2017) to obtain the training loss $L_{CL} = L_{EMR} + \Omega_{EWC}$, where $\Omega_{EWC} = \lambda \sum_{j=1}^N F_j (\theta_{k,j} - \theta_{k-1,j})^2$, N is the number of

model parameters, $\theta_{k-1,j}$ is the model parameters learned until $\mathcal{T}^{(k-1)}$ and $F_j = \nabla^2 \mathcal{L}(\theta_{k-1,j})$ w.r.t. the instances stored in \mathcal{M} . EWC slows down the updates of parameters which are crucial to previous tasks according to the importance measure F_j .

5 Experiments

Datasets and Task Definitions. In this work, we consider three different scenarios: i) different tasks are in different domains and there are task-specific predicates and entities in LFs; ii) there are task-specific predicates in LF templates; iii) there are a significant number of task-specific entities in LFs. All tasks in the latter two are in the same domain. We select Overnight (Wang et al., 2015b) and NLMapV2 (Lawrence and Riezler, 2018) to simulate the proposed three continual learning scenarios, coined OVERNIGHT, NLMAP(QT) and NLMAP(CITY), respectively.

Overnight includes around 18,000 queries involving eight domains. The data in each domain includes 80% training instances and 20% test instances. Each domain is defined as a task.

NLMapV2 includes 28,609 queries involving 79 cities and categorizes each query into one of 4 different question types and their sub-types. In the NLMAP(QT) setting, we split NLMapV2 into 4 tasks with queries in different types. In the setting of NLMAP(CITY), NLMapV2 is split into 8 tasks with queries of 10 or 9 distinct cities in each task. Each city includes a unique set of point of interest regions. In both NLMAP(CITY) and NLMAP(QT), each task is divided into 70%/10%/20% of training/validation/test sets, respectively.

We attribute different distribution discrepancies between tasks to different definitions of tasks. Overall, distribution discrepancy between tasks on OVERNIGHT is the largest while the tasks in other two settings share relatively smaller distribution discrepancies because tasks of NLMAP(QT) and NLMAP(CITY) are all in the same domain.

Baselines. Our proposed method is compared with 8 continual learning baselines. **FINE-TUNE** fine-tunes the model on the new tasks based on previous models. **EWC** (Kirkpatrick et al., 2017) adds an L2 regularization to slow down the update of model parameters important to the historical tasks. **HAT** (Serra et al., 2018) activates a different portion of parameters for each task with task-specific mask functions. **GEM** (Lopez-Paz and Ranzato, 2017) stores a small number of instances from pre-

vious tasks and uses the gradients of previous instances as the constraints on directions of gradients w.r.t. current instances. **EMR** (Chaudhry et al., 2019; Wang et al., 2019) trains the model on the data from the current task along with mini-batches of memory instances. **EMAR** (Han et al., 2020) is an extension of EMR using memory instances to construct prototypes of relation labels to prohibit the model from overfitting on the memory instances. **ARPER** (Mi et al., 2020) adds an adaptive EWC regularization on the EMR loss, where the memory instances are sampled with a unique sampling method called **PRIOR**. **ProtoParser** (Li et al., 2021) utilizes prototypical networks (Snell et al., 2017) to improve the generalization ability of semantic parsers on the unseen actions in the new task. We customize it by training the **PROTOPARSER** on the instances on current task as well as the memory instances. The **ORACLE** (All Tasks) setting trains the model on the data of all tasks combined, considered as an upper bound of continual learning.

Evaluation. To evaluate the performance of continual semantic parsing, we report accuracy of exactly matched LFs as in (Dong and Lapata, 2018). We further adopt two common evaluation settings in continual learning. One setting measures the performance by averaging the accuracies of the parser on test sets of all seen tasks $\{\mathcal{D}_{test}^{(1)}, \dots, \mathcal{D}_{test}^{(k)}\}$ after training the model on task $\mathcal{T}^{(k)}$, i.e. $ACC_{avg} = \frac{1}{k} \sum_{i=1}^k acc_{i,k}$ (Lopez-Paz and Ranzato, 2017; Wang et al., 2019). The other one evaluates the test sets of all seen tasks combined after finishing model training on task $\mathcal{T}^{(k)}$, $ACC_{whole} = acc_{\mathcal{D}_{test}^{(1:k)}}$ (Wang et al., 2019; Han et al., 2020). For reproducibility, we include the detailed **implementation details** in Appendix A.

5.1 Results and Discussion

As shown in Table 1, the base parser trained with our best setting, TR (+EWC), significantly outperforms all the other baselines ($p < 0.005$) in terms of both ACC_{avg} and ACC_{whole} . The performance of TR (+EWC) is, on average, only 3% lower than the ORACLE setting. Without EWC, TR still performs significantly better than all baselines except it is marginally better than ARPER and PROTOPARSER in the setting of NLMAP(QT). From Fig. 4 we can see that our approaches are more stable than the other methods, and demonstrates less and slower forgetting than the baselines.

Methods	OVERNIGHT		NLMAP(QT)		NLMAP(CITY)	
	W	A	W	A	W	A
Fine-tune	14.40	14.37	60.22	55.53	49.29	48.11
EWC	38.57	40.45	65.44	62.25	59.28	57.56
HAT	15.45	15.71	64.69	60.88	53.30	52.41
GEM	41.33	42.13	63.28	59.38	55.14	54.37
EMR	45.29	46.01	59.75	55.59	58.36	56.95
EMAR	46.68	48.57	14.25	12.89	51.79	50.93
ARPER	48.28	49.90	67.79	64.73	62.62	60.61
PROTOPARSER	48.15	49.71	67.10	63.73	62.58	60.87
TR	54.40	54.73	70.06	67.77	62.96	61.59
TR (+EWC)	59.02	58.04	72.36	70.66	67.15	64.89
ORACLE (All Tasks)	63.14	62.76	73.20	71.65	69.48	67.18

Table 1: LF Exact Match Accuracy (%) on two datasets with three settings after model learning on all tasks. “W” stands for the Whole performance ACC_{whole} , and “A” stands for the Average performance ACC_{avg} . All the results are statistically significant ($p < 0.005$) compared with TR (+EWC) according to the Wilcoxon signed-rank test (Woolson, 2007). All experiments are run 10 times with different sequence orders and seeds.

The dynamic architecture method, HAT, performs worst on OVERNIGHT while achieves much better performance on NLMAP(QT) and NLMAP(CITY). Though the performance of the regularization method, EWC, is steady across different settings, it ranks higher among other baselines on NLMAP(CITY) and NLMAP(QT) than on OVERNIGHT. In contrast, the memory-based methods, GEM, and EMR, rank better on OVERNIGHT than on NLMAP(QT) and NLMAP(CITY).

We conjecture that the overall performance of continual learning approaches varies significantly in different settings due to different distribution discrepancies as introduced in *Datasets*. The general memory-based methods are better at handling catastrophic forgetting than the regularization-based and dynamic architecture methods, when the distribution discrepancies are large. However, those memory-based methods are less effective when the distribution discrepancies across tasks are small. Another weakness of memory-based methods is demonstrated by EMAR, which achieves only 14.25% of ACC_{whole} on NLMAP(QT), despite it is the SOTA method on continual relation extraction. A close inspection shows that the instances in the memory are usually insufficient to include all actions when the number of actions is extremely large (i.e., more than 1000 actions per task in NLMAP(QT)) while EMAR relies on instances in memory to construct prototypes for each label. Furthermore, large training epochs for memory-based methods usually lead to severe catastrophic forgetting on the previous tasks, while the regularization method could largely alleviate this effect.

ARPER and PROTOPARSER are the two best

baselines. Similar to TR, ARPER is a hybrid method combining EMR and EWC, thus the joint benefits lead to consistent superior performance over the other baselines except PROTOPARSER in all three settings. The generalization capability to unseen actions in new tasks also seems critical in continual semantic parsing. Merely combining PROTOPARSER and EMR yields a new baseline, which performs surprisingly better than most existing continual learning baselines. From that perspective, the parser with FSCL performs well in continual learning also because of its strength in generalizing to unseen actions.

Methods	OVERNIGHT			NLMAP(CITY)		
	10	25	50	10	25	50
RANDOM	37.63	45.18	50.63	58.82	59.96	60.64
FSS	38.64	47.08	52.63	58.94	59.67	60.89
GSS	34.01	39.45	43.87	57.97	59.40	60.36
PRIOR	37.60	44.84	50.14	54.39	54.09	53.43
BALANCE	38.21	45.33	48.64	58.39	59.89	61.19
LFS	38.24	44.66	53.06	58.88	59.81	60.76
DLFS	39.24	48.21	54.73	59.31	60.58	61.59

Table 2: ACC_{avg} (%) of TR with different sampling strategies and memory sizes 10, 25 and 50.

Influence of Sampling Strategies. Table 2 reports the evaluation results of TR with different sampling strategies and sizes on OVERNIGHT and NLMAP(CITY). The results on NLMAP(QT) can be found in Appendix C. **RANDOM** randomly samples instances from the train set of each task. **FSS** (Aljundi et al., 2019; Wang et al., 2019; Mi et al., 2020), **GSS** (Aljundi et al., 2019) and **LFS** partition the instances into clusters w.r.t. the spaces of utterance encoding features, instance gradients, and LFs, respectively, and then select the instances which are closest to the centroids. **PRIOR** (Mi et al., 2020) selects instances that are most confident to the models and diversifies the entities in the utterances of memory. **BALANCE** (Chrysakis and Moens, 2020; Kim et al., 2020) balances the action distribution in a memory.

Overall, our sampling method consistently outperforms all other baselines on both OVERNIGHT and NLMAP(CITY). On OVERNIGHT with memory size 50, the gap between DLFS and GSS is even up to 11% and 2% between DLFS and FSS, the best baseline. However, on NLMAP(CITY), the performance differences across various sampling methods are smaller than those on OVERNIGHT. Similar observation applies to the influence of different sample sizes. We conclude that the smaller distribution discrepancy reduces the differences of

sampling methods as well as the sample sizes in the memory-based methods.

RANDOM performs steadily across different settings though it is usually in mediocre performance. FSS, GSS, and PRIOR are model-dependent sampling methods. The gradients and model confidence scores are not stable features for the sample selection algorithms. We inspect that the instances selected with GSS are significantly different even when model parameters are slightly disturbed. For the PRIOR, the semantic parsing model is usually confident to instances with similar LF templates. Diversifying entities do not necessarily lead to diversities of LF templates since the LFs with different entities may share similar templates. Therefore, GSS and PRIOR can only perform well in one setting. In contrast, the utterance encoding features are much more reliable. FSS can achieve the second-best performance among all methods. Either balancing action distribution (BALANCE) or selecting centroid LFs from LF clusters (LFS) alone performs no better than DLFS, proving it is advantageous to select a instance in a cluster which balances the memory action distribution over directly using the centroid.

Methods	OVERNIGHT		NLMAP(QT)		NLMAP(CITY)	
	W	A	W	A	W	A
TR (+EWC)	59.02	58.04	72.36	70.66	67.15	64.89
- fast	56.22	54.77	69.12	65.97	64.88	62.42
-/+ fast/lwf	55.80	54.54	69.45	66.37	65.14	62.70
-/+ fast/emar	56.93	56.05	69.43	66.42	64.89	62.51
TR	54.40	54.73	70.06	67.77	62.96	61.59
- fast	49.28	49.48	60.22	54.84	57.53	55.75
-/+ fast/lwf	47.47	47.77	58.96	53.54	55.24	56.89
-/+ fast/emar	49.63	48.74	64.98	61.09	56.89	55.24

Table 3: The ablation study results of FSCL.

Ablation Study of FSCL Training. Table 3 shows the ablation study of FSCL training by removing (-) or replacing (-/+) the corresponding component/step.

The fast-learning with action embeddings is the most critical step in FSCL training. Removing it causes up to 13% performance drop. To study this step in depth, we also replace our fast-learning with fine-tuning all task-specific parameters except in the first task, as done in LwF (Li and Hoiem, 2017), or fine tuning all parameters, as done in EMAR (Han et al., 2020), in the fast-learning stage. The corresponding performance is no better than removing it in most cases. We also plot the training errors and test errors with or without this step in Fig. 3. This step clearly leads to dramatically improvement of both generalization and optimization.




Figure 3: The training and test errors of the base parser with/without fast-learning on OVERNIGHT.




Figure 4: ACC_{whole} till the seen tasks on OVERNIGHT after learning on each task sequentially.

Another benefit of this fast-learning step is in the first task. We observe that a good optimization on the first task is crucial to the model learning on the following tasks. Our preliminary study shows that by applying the fast-learning only to the first task, the model can still keep the close-to-optimal performance. As shown in Fig. 4, our method with this fast-learning step is better optimized and generalized on the initial tasks than all the other baselines and largely alleviate the forgetting problem caused by learning on the second task.

Influence of Pre-trained Language Models.

We study the impact of pre-trained language models for semantic parsing in supervised learning and continual learning, respectively. In both settings, we evaluate the base parsers using BERT (Devlin et al., 2019) as its embedding layer in two configurations: fine-tuning the parameters of BERT (BERT-finetune) and freezing BERT’s parameters (BERT-fix). As in Tab. 4, BERT slightly improves the overall performance of the base parsers in supervised training (the ORACLE setting) on OVERNIGHT. In contrast, in the continual learning setting, base parsers with the BERT embedding perform worse than the ones with the GLOVE em-

bedding. On NLMAP(QT), the accuracy of FINE-TUNE with GLOVE embedding is 30% and 20% higher than that with BERT’s embedding updated and fixed, respectively. We conjecture the deeper neural models suffer more from the catastrophic forgetting. However, the average training speeds of parsers with BERT-fix and BERT-finetune are 5-10 times and 20-40 times respectively slower than those with GLOVE on each task. Overall, our method outperforms other SOTA continual learning methods except that EWC with BERT-fix performs comparably with ours on NLMAP(CITY). In contrast, the performance of PROTOPARSER, the best baseline with GLOVE, is highly unstable on NLMAP with BERT.

Methods	OVERNIGHT		NLMAP(QT)		NLMAP(CITY)	
	W	A	W	A	W	A
BERT-finetune						
Fine-tune	16.19	14.89	30.33	29.22	39.47	36.07
EWC	13.97	13.39	45.97	44.19	35.74	34.17
EMR	51.74	51.40	37.39	34.61	54.89	52.61
PROTOPARSER	47.74	46.99	15.10	13.03	34.56	32.27
TR	52.58	52.86	64.07	61.35	58.38	56.35
ORACLE (All Tasks)	65.40	64.39	72.70	71.60	67.51	68.80
BERT-fix						
Fine-tune	14.10	12.09	40.45	36.06	49.15	46.48
EWC	17.69	18.56	54.20	51.32	64.10	61.67
EMR	39.10	39.03	42.46	39.50	57.55	55.42
PROTOPARSER	41.69	41.79	51.39	48.90	45.66	43.92
TR	47.14	48.12	59.41	55.84	64.48	62.66
ORACLE (All Tasks)	64.70	63.90	73.20	71.91	69.80	67.57

Table 4: ACC_{whole} and ACC_{avg} (%) of parsers using BERT by fine-tuning (Up) and fixing (Bottom) BERT’s parameters.

6 Conclusion

We conducted the first in-depth empirical study to investigate continual learning for semantic parsing. To cope with the catastrophic forgetting and facilitate knowledge transfer between tasks, we propose TOTAL RECALL, consisting of a sampling method specifically designed for semantic parsing and a two-stage training method implementing an inductive bias for continual learning. The resulted parser achieves superior performance over the existing baselines on three benchmark settings. The ablation studies also demonstrate why it is effective.

Acknowledgements

This material is based on research sponsored by the ARC Future Fellowship FT190100039. The computational resources of this work are supported by the Multi-modal Australian Science Imaging and Visualisation Environment (MASSIVE). We appreciate the reviewers for their useful comments.

References

- Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. 2019. Gradient based sample selection for online continual learning. In *Advances in Neural Information Processing Systems*, pages 11816–11825.
- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*.
- 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.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186.
- Shu Cai and Kevin Knight. 2013. Smatch: an evaluation metric for semantic feature structures. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 748–752.
- Ruisheng Cao, Su Zhu, Chen Liu, Jieyu Li, and Kai Yu. 2019. Semantic parsing with dual learning. *arXiv preprint arXiv:1907.05343*.
- Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. 2019. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*.
- Bo Chen, Le Sun, and Xianpei Han. 2018. Sequence-to-action: End-to-end semantic graph generation for semantic parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 766–777.
- Jianpeng Cheng, Siva Reddy, Vijay Saraswat, and Mirella Lapata. 2019. Learning an executable neural semantic parser. *Computational Linguistics*, 45(1):59–94.
- Aristotelis Chrysakis and Marie-Francine Moens. 2020. Online continual learning from imbalanced data. In *International Conference on Machine Learning*, pages 1952–1961. PMLR.
- Marco Damonte, Rahul Goel, and Tagyoung Chung. 2019. Practical semantic parsing for spoken language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Industry Papers)*, pages 16–23.
- 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.
- Li Dong and Mirella Lapata. 2016. Language to logical form with neural attention. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 33–43.
- Li Dong and Mirella Lapata. 2018. Coarse-to-fine decoding for neural semantic parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 731–742.
- Robert M French. 1991. Using semi-distributed representations to overcome catastrophic forgetting in connectionist networks. In *Proceedings of the 13th annual cognitive science society conference*, volume 1, pages 173–178.
- Anirudh Goyal and Yoshua Bengio. 2020. Inductive biases for deep learning of higher-level cognition. *arXiv preprint arXiv:2011.15091*.
- Jiaqi Guo, Zecheng Zhan, Yan Gao, Yan Xiao, Jian-Guang Lou, Ting Liu, and Dongmei Zhang. 2019. Towards complex text-to-sql in cross-domain database with intermediate representation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4524–4535.
- Xu Han, Yi Dai, Tianyu Gao, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2020. Continual relation learning via episodic memory activation and reconsolidation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6429–6440.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR.
- MarketsandMarkets INC. 2020. [Smart speaker market worth \\$15.6 billion by 2025](#).

- Bahar Irfan, Aditi Ramachandran, Samuel Spaulding, Sinan Kalkan, German I Parisi, and Hatice Gunes. 2021. Lifelong learning and personalization in long-term human-robot interaction (leap-hri). In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, pages 724–727.
- Srinivasan Iyer, Alvin Cheung, and Luke Zettlemoyer. 2019. Learning programmatic idioms for scalable semantic parsing. 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 5429–5438.
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, Jayant Krishnamurthy, and Luke Zettlemoyer. 2017. Learning a neural semantic parser from user feedback. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 963–973.
- Aishwarya Kamath and Rajarshi Das. 2018. A survey on semantic parsing. In *Automated Knowledge Base Construction (AKBC)*.
- Chris Dongjoo Kim, Jinseo Jeong, and Gunhee Kim. 2020. Imbalanced continual learning with partitioning reservoir sampling. In *European Conference on Computer Vision*, pages 411–428. Springer.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- BRET KINSELLA. 2020. [Amazon alexa skill growth has slowed further in 2020](#).
- 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.
- Igor Labutov, Bishan Yang, and Tom Mitchell. 2018. Learning to learn semantic parsers from natural language supervision. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1676–1690.
- Carolin Lawrence and Stefan Riezler. 2018. Improving a neural semantic parser by counterfactual learning from human bandit feedback. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1820–1830.
- Zhizhong Li and Derek Hoiem. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947.
- Zhuang Li, Lizhen Qu, and Gholamreza Haffari. 2020. Context dependent semantic parsing: A survey. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2509–2521.
- Zhuang Li, Lizhen Qu, Shuo Huang, and Gholamreza Haffari. 2021. Few-shot semantic parsing for new predicates. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1281–1291.
- Vladislav Lialin, Rahul Goel, Andrey Simanovsky, Anna Rumshisky, and Rushin Shah. 2020. Update frequently, update fast: Retraining semantic parsing systems in a fraction of time. *arXiv preprint arXiv:2010.07865*.
- Kevin Lin, Ben Bogin, Mark Neumann, Jonathan Berant, and Matt Gardner. 2019. Grammar-based neural text-to-sql generation. *arXiv preprint arXiv:1905.13326*.
- Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. 2020. Mnemonics training: Multi-class incremental learning without forgetting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12245–12254.
- David Lopez-Paz and Marc’Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. In *Advances in neural information processing systems*, pages 6467–6476.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421.
- Arun Mallya and Svetlana Lazebnik. 2018. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7765–7773.
- Davide Maltoni and Vincenzo Lomonaco. 2019. Continuous learning in single-incremental-task scenarios. *Neural Networks*, 116:56–73.
- Fei Mi, Liangwei Chen, Mengjie Zhao, Minlie Huang, and Boi Faltings. 2020. Continual learning for natural language generation in task-oriented dialog systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 3461–3474.
- Hae-Sang Park and Chi-Hyuck Jun. 2009. A simple and fast algorithm for k-medoids clustering. *Expert systems with applications*, 36(2):3336–3341.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. Adapterfusion: Non-destructive task composition for transfer learning. In *Proceedings of the 16th*

- Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 487–503.
- 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.
- Hippolyt Ritter, Aleksandar Botev, and David Barber. 2018. Online structured laplace approximations for overcoming catastrophic forgetting. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 3742–3752.
- Subendhu Rongali, Luca Soldaini, Emilio Monti, and Wael Hamza. 2020. Don’t parse, generate! a sequence to sequence architecture for task-oriented semantic parsing. In *Proceedings of The Web Conference 2020*, pages 2962–2968.
- 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*.
- Jonathan Schwarz, Wojciech Czarnecki, Jelena Luketina, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. 2018. Progress & compress: A scalable framework for continual learning. In *International Conference on Machine Learning*, pages 4528–4537. PMLR.
- 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.
- Eui Chul Shin, Miltiadis Allamanis, Marc Brockschmidt, and Alex Polozov. 2019. Program synthesis and semantic parsing with learned code idioms. In *Advances in Neural Information Processing Systems*, pages 10824–10834.
- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087.
- Fan-Keng Sun, Cheng-Hao Ho, and Hung-Yi Lee. 2019. Lamol: Language modeling for lifelong language learning. In *International Conference on Learning Representations*.
- I Sutskever, O Vinyals, and QV Le. 2014. Sequence to sequence learning with neural networks. *Advances in NIPS*.
- Bailin Wang, Richard Shin, Xiaodong Liu, Oleksandr Polozov, and Matthew Richardson. 2020a. Rat-sql: Relation-aware schema encoding and linking for text-to-sql parsers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7567–7578.
- Hong Wang, Wenhan Xiong, Mo Yu, Xiaoxiao Guo, Shiyu Chang, and William Yang Wang. 2019. Sentence embedding alignment for lifelong relation extraction. 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 796–806.
- Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Cuihong Cao, Daxin Jiang, Ming Zhou, et al. 2020b. K-adapter: Infusing knowledge into pre-trained models with adapters. *arXiv preprint arXiv:2002.01808*.
- Yushi Wang, Jonathan Berant, and Percy Liang. 2015a. Building a semantic parser overnight. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1332–1342.
- Yushi Wang, Jonathan Berant, and Percy Liang. 2015b. Building a semantic parser overnight. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1332–1342.
- RF Woolson. 2007. Wilcoxon signed-rank test. *Wiley encyclopedia of clinical trials*, pages 1–3.
- Ziyu Yao, Yu Su, Huan Sun, and Wen-tau Yih. 2019. Model-based interactive semantic parsing: A unified framework and a text-to-sql case study. 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 5450–5461.
- Pengcheng Yin and Graham Neubig. 2018. Tranx: A transition-based neural abstract syntax parser for semantic parsing and code generation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (Demo Track)*.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR.
- Sheng Zhang, Xutai Ma, Kevin Duh, and Benjamin Van Durme. 2019. Broad-coverage semantic parsing as transduction. *arXiv preprint arXiv:1909.02607*.
- Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. 2020. Maintaining discrimination and fairness in class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13208–13217.
- Q. Zhu, X. Ma, and X. Li. 2019. [Statistical learning for semantic parsing: A survey](#). *Big Data Mining and Analytics*, 2(4):217–239.

A Reproducibility Checklist

The hyper-parameters are cross-validated on the training set of OVERNIGHT and validated on the validation set of NLMAP(QT) and NLMAP(CITY). We train the semantic parser on each task with learning rate 0.0025, batch size 64 and for 10 epochs. The fast-learning training epochs is 5. We use the 200-dimensional GLOVE embeddings (Pennington et al., 2014) to initialize the word embeddings for utterances. As different task orders influence the performance of the continual semantic parsing, all experiments are run on 10 different task orders with a different seed for each run. We report the average ACC_{avg} and ACC_{whole} of 10 runs. In addition, we use one GPU of Nvidia V100 to run all our experiments. The sizes of hidden states for LSTM and the action embeddings are 256 and 128, respectively. The default optimizer is Adam (Kingma and Ba, 2014). For our DLFS method, we sample the subsets of the action sets with size 300 and 500 on NLMAP(CITY) and NLMAP(QT), respectively. The number of our model parameters is around 1.8 million. We grid-search the training epochs from $\{10,50\}$, the learning rate from $\{0.001,0.0025\}$. The coefficient for EWC is selected from $\{50000, 200000\}$. Here we also provide the experiment results on the validation sets as in Table 5.

Methods	OVERNIGHT		NLMAP(QT)		NLMAP(CITY)	
	W	A	W	A	W	A
TR	52.38	53.30	68.78	66.83	68.51	66.23
TR (+EWC)	56.84	54.85	70.74	69.21	74.84	71.19

Table 5: Results on the validation sets.

B DLFS Algorithm

We provide the detailed DLFS as in Algo. 2.

C Results of Different Sampling Methods

Table 6 shows the performance of TR with different sampling strategies and different memory sizes on NLMAP(QT).

D Dynamic Action Representation

To differentiate the learning of cross-task and task-specific aspects, we innovatively integrate a designated dynamic architecture into the base parser along with DLFS and FSCL for continual semantic parsing, coined **Dynamic Action Representation**

Algorithm 2: DLFS

Input : Training set \mathcal{D} , memory size M
Output : The memory \mathcal{M}
 Partition \mathcal{D} into M clusters, denoted as C , with the similarity/distance metric
 Randomly sample the memory \mathcal{M} of size M from \mathcal{D}
 $H_{old} \leftarrow -\inf$
 Compute the entropy H_{new} of the action distribution $P_{\mathcal{M}}(\mathcal{A})$ in \mathcal{M} as in Eq. 4
while $H_{new} > H_{old}$ **do**
 for i -th cluster $c_i \in C$ **do**
 Compute the entropy H_c of $P_{\mathcal{M}}(\mathcal{A})$
 for instance $n \in c_i$ **do**
 Replicate \mathcal{M} with \mathcal{M}'
 Replace the i -th instance m_i in \mathcal{M}' with n
 Compute the entropy H'_c of \mathcal{M}'
 if $H'_c > H_c$ **then**
 $\mathcal{M} \leftarrow \mathcal{M}'$
 $E_c \leftarrow H'_c$
 end
 end
 $H_{old} \leftarrow H_{new}$
 Compute the entropy H_{new} of $P_{\mathcal{M}}(\mathcal{A})$
end
end

(DAR). This method could also significantly mitigate the catastrophic forgetting and improve the forward transfer in the continual semantic parsing. Due to the limited space, we did not put it into the main paper. The details and analysis of this method are listed below.

Decoder of Base Parser. The decoder of the base parser applies an LSTM to generate action sequences. At time t , the LSTM produces a hidden state $\mathbf{h}_t = \text{LSTM}(\mathbf{c}_{a_{t-1}}, \mathbf{h}_{t-1})$, where $\mathbf{c}_{a_{t-1}}$ is the embedding of the previous action a_{t-1} . We maintain an embedding for each action in the embedding table. As defined in Luong et al. (2015), we concatenate \mathbf{h}_t with a context vector \mathbf{o}_t to yield \mathbf{s}_t ,

$$\mathbf{s}_t = \tanh(\mathbf{W}_c[\mathbf{h}_t; \mathbf{o}_t]) \quad (5)$$

where \mathbf{W}_c is a weight matrix and the context vector \mathbf{c}_t is generated by the soft attention (Luong et al., 2015),

$$\mathbf{o}_t = \sum_{i=1}^n \text{softmax}(\mathbf{h}_t^T \mathbf{E}) \mathbf{e}_i \quad (6)$$

The probability of an action a_t is estimated by:

$$P(a_t | \mathbf{a}_{<t}, \mathbf{x}) = \frac{\exp(\mathbf{c}_{a_t}^T \mathbf{s}_t)}{\sum_{a' \in \mathcal{A}_t} \exp(\mathbf{c}_{a'}^T \mathbf{s}_t)} \quad (7)$$

where \mathcal{A}_t is the set of applicable actions at time t . In the following, the dense vectors \mathbf{c}_a are referred to as *action embeddings*.

Decoder of DAR. The key idea of dynamic architectures is to add new parameters for new tasks and retain the previous parameters for old

Methods	NLMAP(QT)		
	10	25	50
RANDOM	65.63	66.20	67.21
FSS	66.51	67.15	67.67
GSS	65.89	66.60	67.07
PRIOR	66.08	67.47	67.42
BALANCE	65.37	66.21	66.99
LFS	66.08	67.20	67.53
DLFS	65.21	66.60	67.77

Table 6: ACC_{whole} (%) of TR with different sampling strategies and memory sizes 10, 25 and 50 on NLMAP(QT).

tasks (Houlsby et al., 2019; Wang et al., 2020b; Pfeiffer et al., 2021). As a result, those methods can adapt to new tasks by using new parameters and still memorize the knowledge of previous tasks by keeping existing parameters. In semantic parsing, we differentiate task-specific actions $\mathcal{A}_s^{(k)}$, which generate task-specific predicates or entities, from cross-task actions, which are the remaining actions \mathcal{A}_g associated with predicates appearing in more than one tasks. We model different actions using different action embeddings (Eq. (7)). But the key challenge lies in switching between task-specific and cross-task hidden representations.

To address the problem, given an output hidden state of LSTM, $\mathbf{h}_t = \text{LSTM}(\mathbf{c}_{a_{t-1}}, \mathbf{h}_{t-1})$, we apply a task-specific adapter modules to transform the hidden state $\mathbf{h}_t \in \mathbb{R}^d$.

$$\hat{\mathbf{h}}_t = g_i(\mathbf{h}_t)\phi_i(\mathbf{h}_t) + (1 - g_i(\mathbf{h}_t))\mathbf{h}_t \quad (8)$$

where $\phi_i(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is an adapter network and $g_i(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a gating function for task $\mathcal{T}^{(i)}$. Here, we adopt the following modules for the adaptor network and the gating function,

$$\phi_i(\mathbf{h}_t) = \tanh(\mathbf{W}_\phi^i \mathbf{h}_t) \quad (9)$$

$$g_i(\mathbf{h}_t) = \text{sigmoid}(\mathbf{W}_g^i [\phi_i(\mathbf{h}_t); \mathbf{h}_t]) \quad (10)$$

where parameters $\mathbf{W}_\phi^i \in \mathbb{R}^{d \times d}$ and $\mathbf{W}_g^i \in \mathbb{R}^{2d \times d}$ are task-specific. The number of parameters introduced per task is merely $O(3d^2)$, which is parameter-efficient. Therefore, the context vector of attention and the state to infer action probability

in Eq. 5 and 6 become:

$$\hat{\mathbf{o}}_t = \sum_{i=1}^n \text{softmax}(\hat{\mathbf{h}}_t^T \mathbf{E}) \mathbf{e}_i \quad (11)$$

$$\hat{\mathbf{s}}_t = \tanh(\mathbf{W}_c [\hat{\mathbf{h}}_t; \hat{\mathbf{o}}_t]) \quad (12)$$

Methods	OVERNIGHT		NLMAP(QT)		NLMAP(CITY)	
	W	A	W	A	W	A
TR (+EWC)	59.02	58.04	72.36	70.66	67.15	64.89
- cross	57.90	56.65	71.52	69.74	66.06	63.92
- specific	55.32	54.64	71.47	69.81	65.87	63.73
TR	54.40	54.73	70.06	67.77	62.96	61.59
- cross	52.77	53.07	68.87	66.49	62.18	60.91
- specific	51.44	52.84	69.18	66.88	61.26	59.93

Table 7: The ablation study results of DAR.

Ablation Study of DAR As shown in Tab. 7, removing the task-specific representations (-specific) generally degrades the model performance by 1.5-3.5% except on NLMAP(QT). Our further inspection shows that the proportion of task-specific actions in NLMAP(QT) is only 1/20 while the ratios are 1/4 and 2/5 in OVERNIGHT and NLMAP(CITY), respectively. Using either task-specific representations (-specific) or cross-task representations (-cross) alone cannot achieve the optimal performance.

E Accuracy Curve

Methods	OVERNIGHT	NLMAP(QT)	NLMAP(CITY)
Fine-tune	54.19	360.09	110.12
EWC	71.66	541.20	223.02
HAT	53.85	212.63	128.67
GEM	94.67	389.04	259.80
EMR	102.97	399.05	214.95
EMAR	139.44	402.64	240.64
ARPER	160.05	901.33	654.51
PROTOPARSER	148.04	490.95	304.45
TR	117.63	549.22	275.09
TR (+EWC)	124.49	540.68	282.87
ORACLE (All Tasks)	712.43	1876.60	1531.05

Table 8: The average training time (seconds) of each continual learning method on one task. The training time of the ORACLE setting is reported with training on all tasks. All the methods are running on a server with one Nvidia V100 and four cores of Intel i5 5400.

Figs. 5 depicts the performance curve of semantic parsers till the seen tasks on NLMAP(CITY) and NLMAP(QT) after learning on each task sequentially.

The base parsers are the same for all training methods in comparison. However, the training methods are not exactly the same. For example, PROTOPARSER and EMAR use meta-learning




Figure 5: ACC_{whole} till the seen tasks on NLMAP(CITY) (Up) and NLMAP(QT) (Bottom) after learning on each task sequentially.

methods to train the parser. HAT manipulates parameter gradients during training and uses adapter layers to modify the weights of model parameters on different tasks. ARPER and EWC use regularization during continual training. Different training methods cause the baselines to obtain different results on the initial and subsequent tasks.

In the first task, Fast-Slow Continual Learning (FSCL) differs from the traditional supervised training by updating all action embeddings first, followed by updating all model parameters. From Fig. 4 and Fig. 5, we can tell FSCL leads to a significant performance gain over the baselines in the first task. In this way, our parser trained with FSCL lays a better foundation than the baselines for learning future tasks in terms of both the forward and backward transfer. For the new tasks, the fast-learning step of FSCL leads to minimal changes of model parameters for task-specific patterns. In contrast, the baselines modify the majority of model parameters for each new task, hence easily lead to catastrophic forgetting. As a result, our model with FSCL could




Figure 6: The training and test error points of semantic parsing models with/without fast-learning on NLMAP(CITY) (Up) and NLMAP(QT) (Bottom).

achieve better performance than all baselines on both all tasks and only the initial task as in Fig. 4 and Fig. 5.

F Training Time Analysis

The average training times of different continual learning models on each task of OVERNIGHT, NLMAP(CITY), and NLMAP(QT) are depicted in Tab. 8. On average, the training time of Fine-tune is 13, 5, and 14 times faster than training the parser from scratch on the tasks of OVERNIGHT, NLMAP(CITY), and NLMAP(QT), respectively. In general, the training times of memory-based methods are longer than regularization and dynamic architecture methods due to the replay training. Since our method, TOTAL RECALL, is a memory-based method, its training time is comparable to the other memory-based methods such as GEM, EMR and EMAR. In addition, EWC slows the convergence speed of the parser on NLMAP(CITY), and NLMAP(QT), thus increases the training time of parsers on each task to achieve their optimal per-

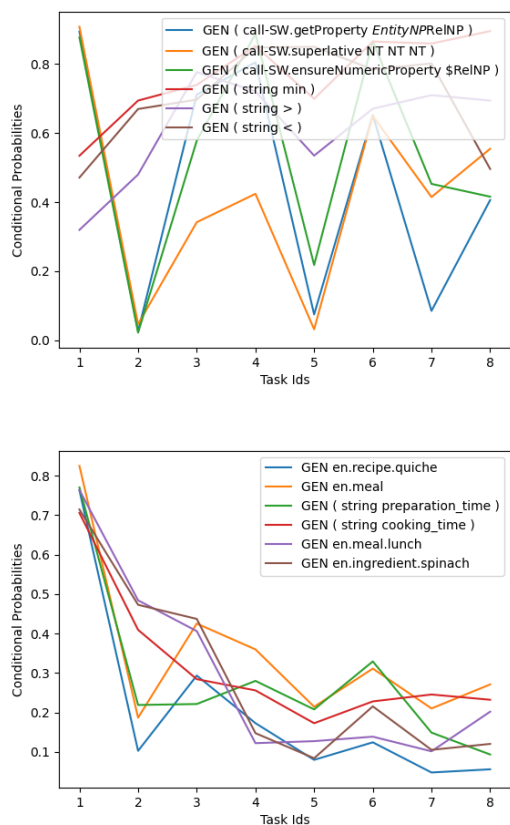


Figure 7: The conditional probabilities $P(a_t|\mathbf{a}_{<t}, \mathbf{x})$ of the representative cross-task actions (Up) and task-specific actions (Bottom) from evaluation on the initial task after parser being trained on each task on OVERNIGHT sequentially.

formance. Therefore, the hybrid method, ARPER, that utilizes both EMR and EWC takes the longest training time among all continual learning methods. However, our FSCL could speed up the convergence of the base parser even with EWC; thus, the training time of TOTAL RECALL (+EWC) is much less than the one of ARPER.

G Training and Test Error Plots

Fig. 6 provides the training and test error points of semantic parsers on NLMAP(CITY) and NLMAP(QT), respectively. As we can see, same as on OVERNIGHT, the base parser with this fast-learning step is better optimized than without this step on NLMAP(CITY) and NLMAP(QT).

H Forgetting Analysis on Actions

Following 4.1, Fig. 7 depicts the conditional probabilities, $P(a_t|\mathbf{a}_{<t}, \mathbf{x})$, of cross-task and task-specific actions, respectively, predicted by the base parser fine-tuned sequentially on each task. Overall, task-specific actions are more likely to be forgotten than cross-task actions while learning parsers on the new tasks. Due to the rehearsal training of the cross-task actions in the future tasks, the prediction performance over cross-task actions fluctuates on different tasks.