

Memory Consolidation for Contextual Spoken Language Understanding with Dialogue Logistic Inference

He Bai^{1,2}, Yu Zhou^{1,2}, Jiajun Zhang^{1,2} and Chengqing Zong^{1,2,3}

¹ National Laboratory of Pattern Recognition, Institute of Automation, CAS, Beijing, China

² University of Chinese Academy of Sciences, Beijing, China

³ CAS Center for Excellence in Brain Science and Intelligence Technology
{he.bai, yzhou, jjzhang, cqzong}@nlpr.ia.ac.cn

Abstract

Dialogue contexts are proven helpful in the spoken language understanding (SLU) system and they are typically encoded with explicit memory representations. However, most of the previous models learn the context memory with only one objective to maximizing the SLU performance, leaving the context memory under-exploited. In this paper, we propose a new dialogue logistic inference (DLI) task to consolidate the context memory jointly with SLU in the multi-task framework. DLI is defined as sorting a shuffled dialogue session into its original logical order and shares the same memory encoder and retrieval mechanism as the SLU model. Our experimental results show that various popular contextual SLU models can benefit from our approach, and improvements are quite impressive, especially in slot filling.

1 Introduction

Spoken language understanding (SLU) is a key technique in today’s conversational systems such as Apple Siri, Amazon Alexa, and Microsoft Cortana. A typical pipeline of SLU includes domain classification, intent detection, and slot filling (Tur and De Mori, 2011), to parse user utterances into semantic frames. Example semantic frames (Chen et al., 2018) are shown in Figure 1 for a restaurant reservation.

Traditionally, domain classification and intent detection are treated as classification tasks with popular classifiers such as support vector machine and deep neural network (Haffner et al., 2003; Sarikaya et al., 2011). They can also be combined into one task if there are not many intents of each domain (Bai et al., 2018). Slot filling task is usually treated as a sequence labeling task. Popular approaches for slot filling include conditional random fields (CRF) and recurrent neural network (RNN) (Raymond and Riccardi, 2007;

D	restaurants
I	reserve_restaurant
u_1	book a table for 10 people tonight
S_1	O O O O B-people O B-date

s	Which restaurant would you like to book a table for?
u_1	Cascal for 6
S_2	B-rest O B-time B-people

Figure 1: Example semantic frames of utterances u_1 and u_2 with domain (D), intent (I) and semantic slots in IOB format (S_1, S_2).

Yao et al., 2014). Considering that pipeline approaches usually suffer from error propagation, the joint model for slot filling and intent detection has been proposed to improve sentence-level semantics via mutual enhancement between two tasks (Xu and Sarikaya, 2013; Hakkani-Tür et al., 2016; Zhang and Wang, 2016; Goo et al., 2018), which is a direction we follow. To create a more effective SLU system, the contextual information has been shown useful (Bhargava et al., 2013; Xu and Sarikaya, 2014), as natural language utterances are often ambiguous. For example, the number 6 of utterance u_2 in Figure 1 may refer to either B-time or B-people without considering the context. Popular contextual SLU models (Chen et al., 2016; Bapna et al., 2017) exploit the dialogue history with the memory network (Weston et al., 2014), which covers all three main stages of memory process: encoding (write), storage (save) and retrieval (read) (Baddeley, 1976). With such a memory mechanism, SLU model can retrieve context knowledge to reduce the ambiguity of the current utterance, contributing to a stronger SLU model. However, the memory consolidation, a well-recognized operation for maintaining and updating memory in cognitive psy-




Figure 2: Architecture of our proposed contextual SLU with memory consolidation.

chology (Sternberg and Sternberg, 2016), is underestimated in previous models. They update memory with only one objective to maximizing the SLU performance, leaving the context memory under-exploited.

In this paper, we propose a multi-task learning approach for multi-turn SLU by consolidating context memory with an additional task: dialogue logistic inference (DLI), defined as sorting a shuffled dialogue session into its original logical order. DLI can be trained with contextual SLU jointly if utterances are sorted one by one: selecting the right utterance from remaining candidates based on previously sorted context. In other words, given a response and its context, the DLI task requires our model to infer whether the response is the right one that matches the dialogue context, similar to the next sentence prediction task (Logeswaran and Lee, 2018). We conduct our experiments on the public multi-turn dialogue dataset KVRET (Eric and Manning, 2017), with two popular memory based contextual SLU models. According to our experimental results, noticeable improvements are observed, especially on slot filling.

2 Model Architecture

This section first explains the memory mechanism for contextual SLU, including memory encoding and memory retrieval. Then we introduce the SLU tagger with context knowledge, the definition of DLI and how to optimize the SLU and DLI jointly. The overall model architecture is illustrated in Figure 2.

Memory Encoding

To represent and store dialogue history $\{x_1, x_2, \dots, x_k\}$, we first encode them into

memory embedding $M = \{m_1, m_2, \dots, m_k\}$ with a BiGRU (Chung et al., 2014) layer and then encode the current utterance x_{k+1} into sentence embedding c with another BiGRU:

$$m_i = \text{BiGRU}_m(x_i) \quad c = \text{BiGRU}_c(x_{k+1}) \quad (1)$$

Memory Retrieval

Memory retrieval refers to formulating contextual knowledge of the user’s current utterance x_{k+1} by recalling dialogue history. There are two popular memory retrieval methods:

The **attention based** (Chen et al., 2016) method first calculates the attention distribution of c over memories M by taking the inner product followed by a softmax function. Then the context can be represented with a weighted sum over M by the attention distribution:

$$p_i = \text{softmax}(c^T m_i) \quad m_{ws} = \sum_i p_i m_i \quad (2)$$

where p_i is the attention weight of m_i . In Chen et al., they sum m_{ws} with utterance embedding c , then multiplied with a weight matrix W_o to generate an output knowledge encoding vector h :

$$h = W_o(c + m_{ws}) \quad (3)$$

The **sequential encoder based** (Bapna et al., 2017) method shows another way to calculate h :

$$g_i = \text{sigmoid}(\text{FF}([c; m_i])) \quad (4)$$

$$h = \text{BiGRU}_g([g_1, g_2, \dots, g_k]) \quad (5)$$

where the function $\text{FF}()$ is a fully connected forward layer.

Contextual SLU

Following [Bapna et al.](#), our SLU model is a stacked BiRNN: a BiGRU layer followed by a BiLSTM layer. However, [Bapna et al.](#) only initializes the BiLSTM layer’s hidden state with h , resulting in the low participation of context knowledge. In this work, we feed h to the second layer in every time step:

$$O_1 = \text{BiGRU}_1(x_{k+1}) \quad (6)$$

$$O_2 = \text{BiLSTM}_2([O_1; h]) \quad (7)$$

where $O_1 = \{o_1^1, \dots, o_1^m\}$ is the first layer’s output and m is the length of x_{k+1} . The second layer encodes $\{[o_1^1; h], \dots, [o_1^m; h]\}$ into the final state $s_2 = [\overleftarrow{s}_2^1; \overleftarrow{s}_2^m]$ and outputs $O_2 = \{o_2^1, \dots, o_2^m\}$, which can be used in the following intent detection layer and slot tagger layer respectively.

$$P^i = \text{softmax}(U s_2) \quad P_t^s = \text{softmax}(V o_2^t) \quad (8)$$

where U and V are weight matrices of output layers and t is the index of each word in utterance x_{k+1} .

Dialogue Logistic Inference

As described above, the memory mechanism holds the key to contextual SLU. However, context memory learned only with SLU objective is under-exploited. Thus, we design a dialogue logistic inference (DLI) task that can consolidate the context memory by sharing encoding and retrieval components with SLU. DLI is introduced below:

Given a dialogue session $X = \{x_1, x_2, \dots, x_n\}$, where x_i is the i th sentence in this conversation, we can shuffle X into a random order set X' . It is not hard for human to restore X' to X by determining which is the first sentence then the second and so on. This is the basic idea of DLI: choosing the right response given a context and all candidates. For each integer j in range $k + 1$ to n , training data of DLI can be labelled automatically by:

$$P(x_j | x_1, \dots, x_k) = \begin{cases} 1 & j = k + 1 \\ 0 & j \neq k + 1 \end{cases} \quad (9)$$

where $k + 1$ is the index of the current utterance. In this work, we calculate the above probability with a 2-dimension softmax layer:

$$P(x_j | x_1, \dots, x_k) = \text{softmax}(W_d h) \quad (10)$$

where W_d is a weight matrix for dimension transformation.

Datasets	Train	Dev	Test	Avg.turns
KVRET	2425	302	304	5.25
KVRET*	1830	224	226	6.88

Table 1: Detailed information of KVRET and KVRET* datasets, including train/dev/test size and average turns per conversation.

Joint Optimization

As we depict in Figure 2, we train DLI and SLU jointly in order to benefit the memory encoder and memory retrieval components. Loss functions of SLU and DLI are as follows.

$$L_{SLU} = \log(p(y^I | x_1, \dots, x_{k+1})) + \sum_t \log(p(y_t^S | x_1, \dots, x_{k+1})) \quad (11)$$

$$L_{DLI} = \sum_{x_j} \log(p(y^D | x_j, x_1, \dots, x_k)) \quad (12)$$

where x_j is a candidate of the current response, y^I , y_t^S and y^D are training targets of intent, slot and DLI respectively. Finally, the overall multi-task loss function is formulated as

$$L = (1 - \lambda)L_{SLU} + \lambda L_{DLI} \quad (13)$$

where λ is a hyper parameter.

3 Experiments

In this section, we first introduce datasets we used, then present our experimental setup and results on these datasets.

3.1 Datasets

KVRET ([Eric and Manning, 2017](#)) is a multi-turn task-oriented dialogue dataset for an in-car assistant. This dataset was collected with the Wizard-of-Oz scheme ([Wen et al., 2017](#)) and consists of 3,031 multi-turn dialogues in three distinct domains, and each domain has only one intent, including calendar scheduling, weather information retrieval, and point-of-interest navigation.

However, all dialogue sessions of KVRET are single domain. Following [Bapna et al.](#), we further construct a multi-domain dataset KVRET* by randomly selecting two dialogue sessions with different domain from KVRET and recombining them into one conversation. The recombining probability is set to 50%. Detailed information about these two datasets is shown in Table 1.

Models	DLI	KVRET				KVRET*			
		Slot			Intent	Slot			Intent
		P	R	F1	Acc.	P	R	F1	Acc.
NoMem	No	54.8	80.0	56.7	93.4	48.9	81.0	54.7	93.8
MemNet	No	75.8	81.1	75.8	93.9	73.1	81.8	74.5	92.8
	Yes	76.0	82.3	77.4(+1.6)	93.9(+0)	75.8	81.3	76.3(+1.8)	93.8(+1.0)
SDEN	No	70.5	80.9	70.1	93.6	56.9	81.3	59.4	93.0
	Yes	64.9	80.9	70.8(+0.7)	93.8(+0.2)	56.5	81.4	60.2(+0.8)	93.5(+0.5)
SDEN [†]	No	71.9	82.2	74.0	93.7	72.7	80.8	74.9	93.2
	Yes	75.2	81.4	76.6(+2.6)	94.3(+0.6)	78.0	81.4	78.3(+3.4)	93.2(+0)

Table 2: SLU results on original KVRET and multi-domain KVRET*, including accuracy of intent detection and average precision, recall and F1 score of slot filling.

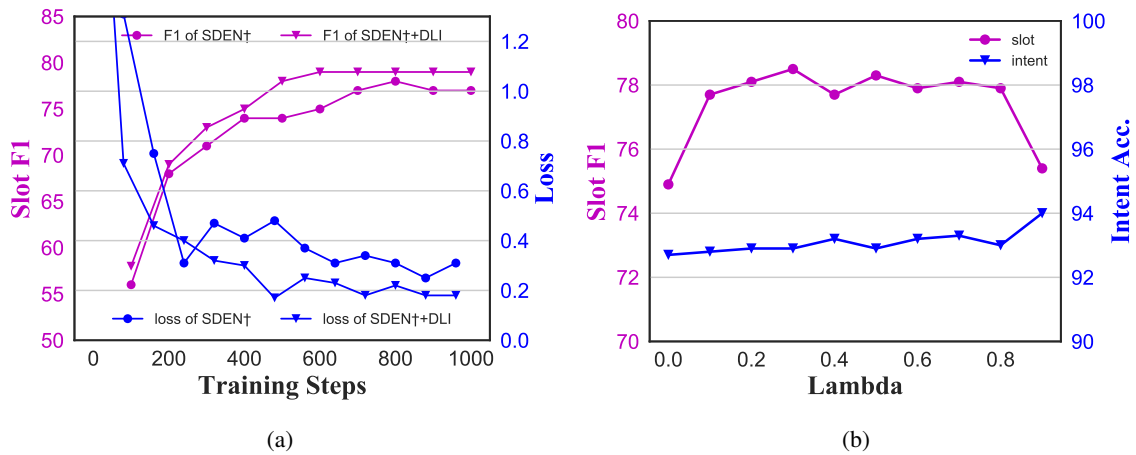


Figure 3: (a) Validation loss and slot F1 score of SDEN[†] during training. (b) Slot F1 score and intent accuracy of SDEN[†] with different lambda.

3.2 Experimental Setup

We conduct extensive experiments on intent detection and slot filling with datasets described above. The domain classification is skipped because intents and domains are the same for KVRET.

For training model, our training batch size is 64, and we train all models with Adam optimizer with default parameters (Kingma and Ba, 2014). For each model, we conduct training up to 30 epochs with five epochs’ early stop on validation loss. The word embedding size is 100, and the hidden size of all RNN layer is 64. The λ is set to be 0.3. The dropout rate is set to be 0.3 to avoid over-fitting.

3.3 Results

The following methods are investigated and their results are shown in Table 2:

NoMem: A single-turn SLU model without memory mechanism.

MemNet: The model described in Chen et al. ,

with attention based memory retrieval.

SDEN: The model described in Bapna et al. , with sequential encoder based memory retrieval.

SDEN[†]: Similar with SDEN, but the usage of h is modified with Eq.6.

As we can see from Table 2, all contextual SLU models with memory mechanism can benefit from our dialogue logistic dependent multi-task framework, especially on the slot filling task. We also note that the improvement on intent detection is trivial, as single turn information has already trained satisfying intent classifiers according to results of NoMem in Table 2. Thus, we mainly analyze DLI’s impact on slot filling task and the prime metric is the F1 score.

In Table 2, the poorest contextual model is the SDEN, as its usage of the vector h is too weak: simply initializes the BiLSTM tagger’s hidden state with h , while other models concatenate h with BiLSTM’s input during each time step. The more the contextual model is dependent on

h , the more obvious the improvement of the DLI task is. Comparing the performance of MemNet with SDEN[†] on these two datasets, we can find that our SDEN[†] is stronger than MemNet after the dialogue length increased. Finally, we can see that improvements on KVRET* are higher than KVRET. This is because retrieving context knowledge from long-distance memory is challenging and our proposed DLI can help to consolidate the context memory and improve memory retrieval ability significantly in such a situation.

We further analyze the training process of SDEN[†] on KVRET* to figure out what happens to our model with DLI training, which is shown in Figure 3(a). We can see that the validation loss of SDEN[†] + DLI falls quickly and its slot F1 score is relatively higher than the model without DLI training, indicating the potential of our proposed method.

To present the influence of hyper-parameter λ , we show SLU results with λ ranging from 0.1 to 0.9 in Figure 3(b). In this figure, we find that the improvements of our proposed method are relatively steady when λ is less than 0.8, and 0.3 is the best one. When λ is higher than 0.8, our model tends to pay much attention to the DLI task, overlook detail information within sentences, leading the SLU model to perform better on the intent detection but failing in slot filling.

4 Conclusions

In this work, we propose a novel dialogue logistic inference task for contextual SLU, with which memory encoding and retrieval components can be consolidated and further enhances the SLU model through multi-task learning. This DLI task needs no extra labeled data and consumes no extra inference time. Experiments on two datasets show that various contextual SLU model can benefit from our proposed method and improvements are quite impressive, especially on the slot filling task. Also, DLI is robust to different loss weight during multi-task training. In future work, we would like to explore more memory consolidation approaches for SLU and other memory related tasks.

Acknowledgements

The research work described in this paper has been supported by the National Key Research and Development Program of China under Grant No.

2017YFB1002103 and the Natural Science Foundation of China under Grant No. U1836221.

References

- Alan D Baddeley. 1976. The psychology of memory. *New York*.
- He Bai, Yu Zhou, Jiajun Zhang, Liang Zhao, Mei-Yuh Hwang, and Chengqing Zong. 2018. Source critical reinforcement learning for transferring spoken language understanding to a new language. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3597–3607.
- Ankur Bapna, Gokhan Tur, Dilek Hakkani-Tur, and Larry Heck. 2017. Sequential Dialogue Context Modeling for Spoken Language Understanding. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)*.
- A. Bhargava, A. Celikyilmaz, D. Hakkani-Tur, and R. Sarikaya. 2013. Easy contextual intent prediction and slot detection. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8337–8341. IEEE.
- Yun-Nung Chen, Asli Celikyilmaz, and Dilek Hakkani-Tur. 2018. Deep learning for dialogue systems. In *Proceedings of the 27th International Conference on Computational Linguistics: Tutorial Abstracts (COLING)*, pages 25–31. Association for Computational Linguistics.
- Yun-Nung Chen, Dilek Hakkani-Tr, Gokhan Tur, Jianfeng Gao, and Li Deng. 2016. End-to-End Memory Networks with Knowledge Carryover for Multi-Turn Spoken Language Understanding. In *Inter-speech*, pages 3245–3249.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS 2014 Deep Learning and Representation Learning Workshop*.
- Mihail Eric and Christopher D. Manning. 2017. Key-Value Retrieval Networks for Task-Oriented Dialogue. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)*.
- Chih-Wen Goo, Guang Gao, Yun-Kai Hsu, Chih-Li Huo, Tsung-Chieh Chen, Keng-Wei Hsu, and Yun-Nung Chen. 2018. Slot-Gated Modeling for Joint Slot Filling and Intent Prediction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 753–757, New Orleans, Louisiana.
- Patrick Haffner, Gokhan Tur, and Jerry H Wright. 2003. Optimizing svms for complex call classification. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 1, pages I–I. IEEE.

- Dilek Hakkani-Tür, Gökhan Tür, Asli Celikyilmaz, Yun-Nung Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang. 2016. Multi-domain joint semantic frame parsing using bi-directional rnn-lstm. In *Inter-speech*, pages 715–719.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *Computer Science*.
- Lajanugen Logeswaran and Honglak Lee. 2018. An efficient framework for learning sentence representations. *arXiv preprint arXiv:1803.02893*.
- Christian Raymond and Giuseppe Riccardi. 2007. Generative and discriminative algorithms for spoken language understanding. In *Eighth Annual Conference of the International Speech Communication Association*.
- Ruhi Sarikaya, Geoffrey E Hinton, and Bhuvana Ramabhadran. 2011. Deep belief nets for natural language call-routing. In *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, pages 5680–5683. IEEE.
- Robert J Sternberg and Karin Sternberg. 2016. *Cognitive psychology*. Nelson Education.
- Gokhan Tur and Renato De Mori. 2011. *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. A Network-based End-to-End Trainable Task-oriented Dialogue System. In *EACL*, pages 1–12.
- Jason Weston, Sumit Chopra, and Antoine Bordes. 2014. Memory networks. *arXiv preprint arXiv:1410.3916*.
- Puyang Xu and Ruhi Sarikaya. 2013. Convolutional neural network based triangular crf for joint intent detection and slot filling. In *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*, pages 78–83. IEEE.
- Puyang Xu and Ruhi Sarikaya. 2014. Contextual domain classification in spoken language understanding systems using recurrent neural network. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 136–140, Florence, Italy. IEEE.
- Kaisheng Yao, Baolin Peng, Yu Zhang, Dong Yu, Geoffrey Zweig, and Yangyang Shi. 2014. Spoken language understanding using long short-term memory neural networks. In *Spoken Language Technology Workshop (SLT), 2014 IEEE*, pages 189–194. IEEE.
- Xiaodong Zhang and Houfeng Wang. 2016. A joint model of intent determination and slot filling for spoken language understanding. In *IJCAI*, pages 2993–2999.