

# A Working Memory Model for Task-oriented Dialog Response Generation

Xiuyi Chen<sup>1,2,3</sup>, Jiaming Xu<sup>1,2\*</sup> and Bo Xu<sup>1,2,3,4</sup>

<sup>1</sup>Institute of Automation, Chinese Academy of Sciences (CASIA), Beijing, China

<sup>2</sup>Research Center for Brain-inspired Intelligence, CASIA

<sup>3</sup>University of Chinese Academy of Sciences

<sup>4</sup>Center for Excellence in Brain Science and Intelligence Technology, CAS, China  
{chenxiuyi2017, jiaming.xu, xubo}@ia.ac.cn

## Abstract

Recently, to incorporate external Knowledge Base (KB) information, one form of world knowledge, several end-to-end task-oriented dialog systems have been proposed. These models, however, tend to confound the dialog history with KB tuples and simply store them into one memory. Inspired by the psychological studies on working memory, we propose a working memory model (WMM2Seq) for dialog response generation. Our WMM2Seq adopts a working memory to interact with two separated long-term memories, which are the episodic memory for memorizing dialog history and the semantic memory for storing KB tuples. The working memory consists of a central executive to attend to the aforementioned memories, and a short-term storage system to store the “activated” contents from the long-term memories. Furthermore, we introduce a context-sensitive perceptual process for the token representations of the dialog history, and then feed them into the episodic memory. Extensive experiments on two task-oriented dialog datasets demonstrate that our WMM2Seq significantly outperforms the state-of-the-art results in several evaluation metrics.

## 1 Introduction

Task-oriented dialog systems, such as hotel booking or technical support service, help users to achieve specific goals with natural language. Compared with traditional pipeline solutions (Williams and Young, 2007; Young et al., 2013; Wen et al., 2017), end-to-end approaches recently gain much attention (Zhao et al., 2017; Eric and Manning, 2017a; Lei et al., 2018), because they directly map dialog history to the output responses and consequently reduce human effort for modular designs and hand-crafted state labels. To effectively incorporate KB information and perform knowledge-

based reasoning, memory augmented models have been proposed (Bordes et al., 2017; Seo et al., 2017; Eric and Manning, 2017b; Madotto et al., 2018; Raghu et al., 2018; Reddy et al., 2019; Wu et al., 2019). Bordes et al. (2017) and Seo et al. (2017) attended to retrieval models, lacking the ability of generation, while others incorporated the memory (i.e. end-to-end memory networks, abbreviated as MemNNs, Sukhbaatar et al. (2015)) and copy mechanism (Gu et al., 2016) into a sequential generative architecture. However, most models tended to confound the dialog history with KB tuples and simply stored them into one memory. A shared memory forces the memory reader to reason over the two different types of data, which makes the task harder, especially when the memory is large. To explore this problem, Reddy et al. (2019) very recently proposed to separate memories for modeling dialog context and KB results. In this paper, we adopt working memory to interact with two long-term memories. Furthermore, compared to Reddy et al. (2019), we leverage the reasoning ability of MemNNs to instantiate the external memories.

Our intuition comes from two aspects. First, psychologists tend to break down the long-term memory<sup>1</sup> into episodic memory for events (e.g. visual and textual perceptual inputs) and semantic memory for facts (world knowledge, such as KB information) as not all memory of experiences is the same (Gazzaniga and Ivry, 2013). Second, a successful task-oriented dialog system needs more intelligence, and recent works suggest that a critical component of intelligence may be working memory (Sternberg and Sternberg, 2016). Hence, leveraging the knowledge from psychological studies (Baddeley and Hitch, 1974; Baddeley, 2000; Doshier, 2003), we explore working memory for the dialog response generation. Our contributions

<sup>1</sup>Here, the long-term memory is referred to declarative memory that we have conscious access to.

\* Corresponding Author

are summarized as follows:

Firstly, inspired by the psychological studies on working memory, we propose the WMM2Seq for dialog generation which separates the storage of dialog history and KB information by using the episodic and semantic memories and then leverages the working memory to interact with them.

Secondly, we leverage two kinds of transformations (CNN and biGRU) to incorporate the context information for better token representations. This procedure can be seen as a part of perceptual processes before the episodic memory storage, and can alleviate the Out-Of-Vocabulary (OOV) problem.

Finally, our WMM2Seq outperforms the existing methods on several evaluation metrics in two task-oriented dialog datasets and shows a better reasoning ability in the OOV situation.

## 2 Model Description

Figure 1 illustrates the flow of our WMM2Seq for dialog response generation. WMM2Seq can be seen as an encoder-decoder model, where decoder is the Working Memory (WM) which could interact with two long-term memories (the episodic memory memorizing dialog history and semantic memory storing KB information). As MemNN is well-known for its multiple hop reasoning ability, we instantiate the encoder and the two memories with three different MemNNs (MemNN Encoder, E-MemNN and S-MemNN). Furthermore, we augment E-MemNN and S-MemNN with copy mechanism from where we need to copy tokens or entities. The encoder encodes the dialog history to obtain the high-level signal, a distributed intent vector. The WM consists of a Short-Term Storage system (STS) and a Central-EXE including an Attention Controller (Attn-Ctrl) and a rule-based word selection strategy. The Attn-Ctrl dynamically generates the attention control vector to query and reason over the two long memories and then stores three “activated” distributions into STS. Finally a generated token is selected from the STS under the word selection strategy at each decoder step.

The symbols are defined in Table 1, and more details can be found in the supplementary material. We omit the subscript E or S<sup>2</sup>, following Madotto et al. (2018) to define each pointer index set:

$$ptr_i = \begin{cases} \max(z) & \text{if } \exists z \text{ s.t. } y_i = xb_z \\ n_{xb} + 1 & \text{otherwise} \end{cases}, \quad (1)$$

<sup>2</sup>Note, all variables belonging to the episodic memory are with subscript E, and semantic memory are with subscript S.

Symbol	Definition
$x_i$ or $y_i$	a token in the dialog history or system response
$\$$	a special token used as a sentinel (Madotto et al., 2018)
$X$	$X = \{x_1, \dots, x_n, \$\}$ , the dialog history
$Y$	$Y = \{y_1, \dots, y_m\}$ , the expected response
$b_i$	one KB tuple, actually the corresponding entity
$B$	$B = \{b_1, \dots, b_l, \$\}$ , the KB tuples
$PTR_E$	$PTR_E = \{ptr_{E,1}, \dots, ptr_{E,m}\}$ , dialog pointer index set. supervised information for copying words in dialog history
$PTR_S$	$PTR_S = \{ptr_{S,1}, \dots, ptr_{S,m}\}$ , KB pointer index set. supervised information for copying entities in KB tuples

Table 1: Notation Table.

where  $xb_z \in X$  or  $B$  is the dialog history or KB tuples according to the subscript (E or S) and  $n_{xb} + 1$  is the sentinel position index as  $n_{xb}$  is equal to the dialog history length  $n$  or the number of KB triples  $l$ . The idea behind Eq. 1 is that we can obtain the positions of where to copy by matching the target text with the dialog history or KB information. Furthermore, we hope this provides the model with an accurate guidance of how to activate the two long-term memories.

### 2.1 MemNN Encoder

Here, on the context of our task, we give a brief description of  $K$ -hop MemNN with adjacent weight tying and more details can be found in (Sukhbaatar et al., 2015). The memory of MemNN is represented by a set of trainable embedding matrices  $C = \{C^1, \dots, C^{K+1}\}$ . Given input tokens in the dialog history  $X$ , MemNN first writes them into memories by Eq. 2 and then uses a query to iteratively read from them with multi hops to reason about the required response by Eq. 3 and Eq. 4. For each hop  $k$ , we update the query by Eq. 5 and the initial query is a learnable vector as like Yang et al. (2016). The MemNN encoder finally outputs a user intent vector  $o^K$ .

$$A_i^k = C^k(x_i) \quad (2)$$

$$p_i^k = \text{Softmax}((q^k)^T A_i^k) \quad (3)$$

$$o^k = \sum_i p_i^k A_i^{k+1} \quad (4)$$

$$q^{k+1} = q^k + o^k \quad (5)$$

To incorporate the context information, we explore two context-aware transformation  $\text{TRANS}(\cdot)$  by replacing Eq. 2 with  $A_i^k = \text{TRANS}(C^k(x_i))$ , which is defined as follows:

$$h_i = \text{TRANS}(\phi^e(x_i)) \\ = \text{CNN}([\phi^e(x_{i-2}), \dots, \phi^e(x_{i+2})]) \quad (6)$$

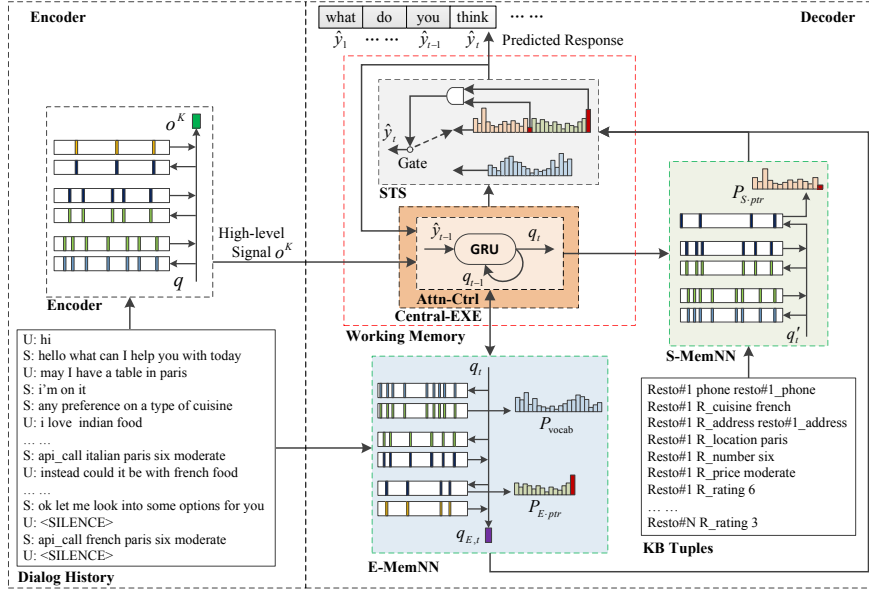


Figure 1: The Working Memory (WM) interacts with two long-term memories to generate the response.

or

$$\begin{aligned}
 h_i &= \text{TRANS}(\phi^e(x_i)) \\
 &= \begin{bmatrix} \vec{h}_i \\ \vec{h}_i \end{bmatrix} = \begin{bmatrix} \overrightarrow{\text{GRU}}(\phi^e(x_i), \vec{h}_{i-1}) \\ \overleftarrow{\text{GRU}}(\phi^e(x_i), \vec{h}_{i+1}) \end{bmatrix}, \quad (7)
 \end{aligned}$$

where  $h_i$  is the context-aware representation, and  $\phi^e$  is a trainable embedding function. We combine MemNNs with  $\text{TRANS}(\cdot)$  to alleviate the OOV problem when reasoning about memory contents.

## 2.2 Working Memory Decoder

Inspired by the studies on the working memory, we design our decoder as an attentional control system for dialog generation which consists of the working memory and two long-term memories. As shown in Figure 1, we adopt the E-MemNN to memorize the dialog history  $X$  as described in Section 2.1, and then store KB tuples into the S-MemNN without  $\text{TRANS}(\cdot)$ . We also incorporate additional temporal information and speaker information into dialog utterances as (Madotto et al., 2018) and adopt a (subject, relation, object) representation of KB information as (Eric and Manning, 2017b). More details can be found in the supplementary material.

Having written dialog history and KB tuples into E-MemNN and S-MemNN, we then use the WM to interact with them (to query and reason over them) to generate the response. At each decoder step, the Attn-Ctrl, instantiated as a GRU, dynamically generates the query vector  $q_t$  as follows:

$$q_t = \text{GRU}(C_E^1(\hat{y}_{t-1}), q_{t-1}). \quad (8)$$

Here, query  $q_t$  is used to access E-MemNN activating the final query  $q_E = o_E^K$ , vocabulary distribution  $P_{\text{vocab}}$  by Eq. 9 and copy distribution for dialog history  $P_{E\text{-ptr}}$ . When querying S-MemNN, we consider the dialog history by using query  $q'_t = q_E + q_t$  and then obtain the copy distribution for KB entities  $P_{S\text{-ptr}}$ . The two copy distributions are obtained by augmenting MemNNs with copy mechanism that is  $P_{E\text{-ptr}} = p_{E,t}^K$  and  $P_{S\text{-ptr}} = p_{S,t}^K$ .

$$P_{\text{vocab}}(\hat{y}_t) = \text{Softmax}(W_1[q_t; o_E^1]). \quad (9)$$

Now, three distributions,  $P_{\text{vocab}}$ ,  $P_{E\text{-ptr}}$  and  $P_{S\text{-ptr}}$ , are activated and moved into the STS, and then a proper word is generated from the activated distributions. We here use a rule-based word selection strategy by extending the sentinel idea in (Madotto et al., 2018), which is shown in Figure 1. If the expected word is not appearing either in the episodic memory or the semantic memory, the two copy pointers are trained to produce the sentinel token and our WMM2Seq generates the token from  $P_{\text{vocab}}$ ; otherwise, the token is generated by copying from either the dialog history or KB tuples and this is done by comparing the two copy distributions. We always select the other distribution if one of the two distributions points to the sentinel or select to copy the token corresponding to the biggest probability of the two distributions. Hence, during the training stage, all the parameters are jointly learned by minimizing the sum of three standard cross-entropy losses with the corresponding targets ( $Y$ ,  $PTR_E$  and  $PTR_S$ ).

Task	Ptr-Unk	Mem2Seq	HyP-MN	GLMP	WMM2Seq+CNN	WMM2Seq+biGRU	WMM2Seq	WMM2Seq+biGRU (H1)
T1	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)
T2	100 (100)	100 (100)	99.9 (99.8)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)
T3	85.1 (19.0)	94.7 (62.1)	94.9 (63.2)	<b>96.3 (75.6)</b>	95.03 (63.6)	95.32 (68.2)	94.94 (63.9)	95.01 (64.6)
T4	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)
T5	<b>99.4 (91.5)</b>	97.9 (69.6)	97.7 (67)	99.2 (88.5)	98.49 (76.6)	99.34 (90.3)	97.95 (71.2)	99.26 (88.8)
T1-OOV	92.5 (54.7)	94.0 (62.2)	<b>100 (100)</b>	<b>100 (100)</b>	<b>100 (100)</b>	<b>100 (100)</b>	91.28 (57.2)	<b>100 (100)</b>
T2-OOV	83.2 (0)	86.5 (12.4)	<b>100 (100)</b>	<b>100 (100)</b>	<b>100 (100)</b>	<b>100 (100)</b>	83.28 (0)	<b>100 (100)</b>
T3-OOV	82.9 (13.4)	90.3 (38.7)	<b>95.6 (63.9)</b>	95.5 (65.7)	94.87 ( <b>66.2</b> )	94.64 (61.6)	94.54 (60.5)	94.80 (62.2)
T4-OOV	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)	100 (100)
T5-OOV	73.6 (0)	84.5 (2.3)	89.3 (9.7)	92.0 (21.7)	92.32 ( <b>24.3</b> )	<b>92.56 (24.3)</b>	84.45 (3.6)	91.86 (22.0)

Table 2: Per-response and per-dialog (in the parentheses) accuracy on bAbI dialogs.

### 3 Experiments

We conduct experiments on the simulated bAbI Dialogue dataset (Bordes et al., 2017) and the Dialog State Tracking Challenge 2 (DSTC2) (Henderson et al., 2014). We actually adopt the refined version of DSTC2 from Bordes et al. (2017) and their statistics are given in the supplementary material.

Our model is trained end-to-end using Adam optimizer (Kingma and Ba, 2014), and the responses are generated using greedy search without any rescoring techniques. The shared size of embedding and hidden units is selected from [64, 512] and the default hop  $K = 3$  is used for all MemNNs. The learning rate is simply fixed to 0.001 and the dropout ratio is sampled from [0.1, 0.4]. Furthermore, we randomly mask some memory cells with the same dropout ratio to simulate the OOV situation for both episodic and semantic memories. The hyper-parameters for best models are given in the supplementary material.

#### 3.1 Results and Analysis

We use Per-response/dialog Accuracy (Bordes et al., 2017), BLEU (Papineni et al., 2002) and Entity F1 (Madotto et al., 2018) to compare the performance of different models. And the baseline models are Seq2Seq+Attn (Luong et al., 2015), Pointer to Unknown (Ptr-Unk, Gulcehre et al. (2016)), Mem2Seq (Madotto et al., 2018), Hierarchical Pointer Generator Memory Network (HyP-MN, Raghuram et al. (2018)) and Global-to-Local Memory Pointer (GLMP, Wu et al. (2019)).

**Automatic Evaluation:** The results on the bAbI dialog dataset are given in Table 2. We can see that our model does much better on the OOV situation and is on par with the best results on T5. Moreover, our model can perfectly issue API calls (task 1), update API calls (task 2) and provide extra information (task 4). As task 5 is a combination of tasks 1-4, our best performance on T5-OOV exhibits the powerful reasoning ability to the unseen

	Ent. F1	BLEU	Per-Resp.(Dial.)
Seq2Seq	69.7	55.0	46.4 (1.5)
Seq2Seq+Attn	67.1	56.6	46.0 (1.4)
Seq2Seq+Copy	71.6	55.4	47.3 (1.3)
Mem2Seq	75.3	55.3	45.0 (0.5)
HyP-MN	73.9	55.4	46.4 (1.7)
<b>WMM2Seq+CNN</b>	<b>80.73</b>	57.33	48.80 (1.61)
<b>WMM2Seq+biGRU</b>	80.23	58.39	<b>49.02</b> (1.25)
WMM2Seq	75.45	56.81	45.25 (1.25)
WMM2Seq+biGRU (H1)	78.87	<b>58.57</b>	48.81 (1.61)

Table 3: Automatic Evaluation on DSTC2.

dialog history and KB tuples. And this reasoning ability is also proved by the performance improvements on the DSTC2 dataset according to several metrics in Table 3. Especially, a significant improvement on entity F1 scores indicates that our model can choose the right entities and incorporate them into responses more naturally (with highest BLEU scores). Furthermore, there is no significant difference between the two kinds of the transformation TRANS( $\cdot$ ).

**Ablation Study:** To better understand the components used in our model, we report our ablation studies from three aspects. First, we remove the context-sensitive transformation TRANS( $\cdot$ ) and then find significant performance degradation. This suggests that perceptual processes are a necessary step before storing perceptual information (the dialog history) into the episodic memory and it is important for the performance of working memory. Second, we find that WMM2Seq outperforms Mem2Seq, which uses a unified memory to store dialog history and KB information. We can safely conclude that the separation of context memory and KB memory benefits the performance, as WMM2Seq performs well with less parameters than Mem2Seq on task 5. Finally, we additionally analysis how the multi-hop attention mechanism helps by showing the performance differences between the hop  $K = 1$  and the default hop  $K = 3$ . Though multi-hop attention strengthens the reasoning ability and improves the results, we find that the performance difference between the hops  $K = 1$  and  $K = 3$  is not so obvious as shown in



	Mem2Seq	WMM2Seq	Gold
Appropriate	4.31	4.47	4.61
Humanlike	4.37	4.48	4.80

Table 4: Human Evaluation.

(Madotto et al., 2018; Wu et al., 2019). Furthermore, our model performs well even with one hop, which we mainly attribute to the reasoning ability of working memory. The separation of memories and stacking S-MemNN on E-MemNN also help a lot, because the whole external memory, consisting of the episodic and semantic memories, can be seen as a multi-hop (two-level) structure (the first level is the episode memory and the second level is the semantic memory).

**Attention Visualization:** As an intuitive way to show the model’s dynamics, attention weight visualization is also used to understand how the Central-EXE controls the access to the two long-term memories (E-MemNN and S-MemNN). Figure 2 shows the episodic and semantic memory attention vectors at the last hop for each generated token. Firstly, our model generates a different but still correct response as the customer wants a moderately priced restaurant in the west and does not care about the type of food. Secondly, the generated response has tokens from the vocabulary (e.g. “is” and “a”), dialog history (e.g. “west” and “food”) and KB information (e.g. “saint\_johns\_chop\_house” and “british”), indicating that our model learns to interact well with the two long-term memories by two sentinels.

**Human Evaluation:** Following the methods in (Eric and Manning, 2017b; Wu et al., 2019), we report human evaluation of the generated responses in Table 4. We adopt Mem2Seq as the baseline for human evaluation considering its good performance and code release<sup>3</sup>. First we randomly select 100 samples from the DSTC2 test set, then generate the corresponding responses using WMM2Seq and Mem2Seq, and finally ask two human subjects to judge the quality of the generated responses according to the appropriateness and humanlikeness on a scale from 1 to 5. As shown in Table 4, WMM2Seq outperforms Mem2Seq in both measures, which is coherent to the automatic evaluation. More details about human evaluation are reported in the supplementary material.

<sup>3</sup>We thank the authors for releasing their code at <https://github.com/HLTCHKUST/Mem2Seq>.

GOLD: saint\_johns\_chop\_house is a nice place in the west of town and the prices are moderate  
 PRED: saint\_johns\_chop\_house is a nice restaurant in the west of town serving british food




Figure 2: Last hop semantic and episodic memory attention visualization from the DSTC2 dataset.

## 4 Conclusion

We leverage the knowledge from the psychological studies and propose our WMM2Seq for dialog response generation. First, the storage separation of the dialog history and KB information is very important and we explore two context-sensitive perceptual processes for the word-level representations of the dialog history. Second, working memory is adopted to interact with the long-term memories and then generate the responses. Finally, the improved performance on two task-oriented datasets demonstrates the contributions from the separated storage and the reasoning ability of working memory. Our future work will focus on how to transfer the long-term memory across different tasks.

## Acknowledgments

We would like to thank the anonymous reviewers for their insightful comments. This work was supported by the National Natural Science Foundation of China (61602479), the Strategic Priority Research Program of the Chinese Academy of Sciences (XDB32070000) and the Beijing Brain Science Project (Z181100001518006).

## References

- Alan Baddeley. 2000. [The episodic buffer: a new component of working memory?](#) *Trends in Cognitive Sciences*, 4(11):417–423.
- Alan D Baddeley and Graham Hitch. 1974. Working memory. In *Psychology of learning and motivation*, volume 8, pages 47–89. Elsevier.
- Antoine Bordes, Y-Lan Boureau, and Jason Weston. 2017. [Learning end-to-end goal-oriented dialog](#). In *International Conference on Learning Representations*.
- Dosher. 2003. Working memory. In *L. Nadel (Ed.), Encyclopedia of cognitive science*, volume 4, pages 569–577. London: Nature Publishing Group.
- Mihail Eric and Christopher Manning. 2017a. [A copy-augmented sequence-to-sequence architecture gives good performance on task-oriented dialogue](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, volume 2, pages 468–473.
- Mihail Eric and Christopher D Manning. 2017b. [Key-value retrieval networks for task-oriented dialogue](#). *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 37–49.
- Michael Gazzaniga and Richard B Ivry. 2013. *Cognitive Neuroscience: The Biology of the Mind: Fourth International Student Edition*. WW Norton.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. [Incorporating copying mechanism in sequence-to-sequence learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1631–1640, Berlin, Germany. Association for Computational Linguistics.
- Caglar Gulcehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. 2016. [Pointing the unknown words](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 140–149.
- Matthew Henderson, Blaise Thomson, and Jason D Williams. 2014. [The second dialog state tracking challenge](#). In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 263–272.
- Diederik P Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *arXiv preprint arXiv:1412.6980*.
- Wenqiang Lei, Xisen Jin, Min-Yen Kan, Zhaochun Ren, Xiangnan He, and Dawei Yin. 2018. [Sequicity: Simplifying task-oriented dialogue systems with single sequence-to-sequence architectures](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1437–1447. Association for Computational Linguistics.
- 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.
- Andrea Madotto, Chien-Sheng Wu, and Pascale Fung. 2018. [Mem2seq: Effectively incorporating knowledge bases into end-to-end task-oriented dialog systems](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1468–1478. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Dinesh Raghu, Nikhil Gupta, and Mausam. 2018. [Hierarchical pointer memory network for task oriented dialogue](#). *arXiv preprint arXiv:1805.01216*.
- Revanth Reddy, Danish Contractor, Dinesh Raghu, and Sachindra Joshi. 2019. [Multi-level memory for task oriented dialogs](#). *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)*.
- Minjoon Seo, Sewon Min, Ali Farhadi, and Hannaneh Hajishirzi. 2017. [Query-reduction networks for question answering](#). In *International Conference on Learning Representations*.
- Robert J Sternberg and Karin Sternberg. 2016. *Cognitive psychology*. Nelson Education.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. [End-to-end memory networks](#). In *Advances in neural information processing systems*, pages 2440–2448.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, 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 *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, volume 1, pages 438–449.
- Jason D Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.

Chien-Sheng Wu, Richard Socher, and Caiming Xiong. 2019. [Global-to-local memory pointer networks for task-oriented dialogue](#). *International Conference on Learning Representations*.

Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. [Hierarchical attention networks for document classification](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489.

Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. 2013. [Pomdp-based statistical spoken dialog systems: A review](#). *Proceedings of the IEEE*, 101(5):1160–1179.

Tiancheng Zhao, Allen Lu, Kyusong Lee, and Maxine Eskenazi. 2017. [Generative encoder-decoder models for task-oriented spoken dialog systems with chatting capability](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*.