

DYNAMIC TIME-AWARE ATTENTION TO SPEAKER ROLES AND CONTEXTS FOR SPOKEN LANGUAGE UNDERSTANDING

Po-Chun Chen* Ta-Chung Chi* Shang-Yu Su† Yun-Nung Chen*

*Department of Computer Science and Information Engineering

†Graduate Institute of Electrical Engineering

National Taiwan University, Taipei, Taiwan

{b02902019,r06922028,r05921117}@ntu.edu.tw y.v.chen@ieee.org

ABSTRACT

Spoken language understanding (SLU) is an essential component in conversational systems. Most SLU component treats each utterance independently, and then the following components aggregate the multi-turn information in the separate phases. In order to avoid error propagation and effectively utilize contexts, prior work leveraged history for contextual SLU. However, the previous model only paid attention to the content in history utterances without considering their temporal information and speaker roles. In the dialogues, the most recent utterances should be more important than the least recent ones. Furthermore, users usually pay attention to 1) self history for reasoning and 2) others utterances for listening, the speaker of the utterances may provides informative cues to help understanding. Therefore, this paper proposes an attention-based network that additionally leverages temporal information and speaker role for better SLU, where the attention to contexts and speaker roles can be automatically learned in an end-to-end manner. The experiments on the benchmark Dialogue State Tracking Challenge 4 (DSTC4) dataset show that the time-aware dynamic role attention networks significantly improve the understanding performance¹.

Index Terms— dialogue, language understanding, SLU, temporal, role, attention, deep learning.

1. INTRODUCTION

Spoken dialogue systems that can help users to solve complex tasks such as booking a movie ticket have become an emerging research topic in artificial intelligence and natural language processing area. With a well-designed dialogue system as an intelligent personal assistant, people can accomplish certain tasks more easily via natural language interactions. Today, there are several virtual intelligent assistants, such as Apple’s Siri, Google’s Home, Microsoft’s Cortana, and Amazon’s Echo. Recent advance of deep learning has

inspired many applications of neural models to dialogue systems. Prior work introduced network-based end-to-end trainable task-oriented dialogue systems [1, 2, 3, 4].

A key component of the understanding system is a spoken language understanding (SLU) module—it parses user utterances into semantic frames that capture the core meaning, where three main tasks of SLU are domain classification, intent determination, and slot filling [5]. A typical pipeline of SLU is to first decide the domain given the input utterance, and based on the domain, to predict the intent and to fill associated slots corresponding to a domain-specific semantic template. With the power of deep learning, there are emerging better approaches of SLU [6, 7, 8, 9]. However, the above work focused on single-turn interactions, where each utterance is treated independently.

The contextual information has been shown useful for SLU [10, 11, 12, 13]. For example, Figure 1 shows conversational utterances, where the intent of the highlighted tourist utterance is to ask about location information, but it is difficult to understand without contexts. Hence, it is more likely to estimate the location-related intent given the contextual utterances about location recommendation. Contextual information has been incorporated into the recurrent neural network (RNN) for improved domain classification, intent prediction, and slot filling [11, 14, 15, 16].

Most of previous dialogue systems did not take speaker role into consideration. However, different speaker roles can cause notable variance in speaking habits and later affect the system performance differently. From Figure 1, the benchmark dialogue dataset, Dialogue State Tracking Challenge 4 (DSTC4) [17]², contains two specific roles, a tourist and a guide. Under the scenario of dialogue systems and the communication patterns, we take the tourist as a user and the guide as the dialogue agent (system). During conversations, the user may focus on not only *reasoning (user history)* but also *listening (agent history)*, so different speaker roles could provide various cues for better understanding [18].

In addition, neural models incorporating attention mech-

¹The first three authors have equal contributions.

²The released code: <https://github.com/MiuLab/Time-SLU>

²<http://www.colips.org/workshop/dstc4/>

Guide: and you were saying that you wanted to come to singapore	FOL_CONFIRM; FOL_INFO
Guide: uh maybe can i have a little bit more details like uh when will you be coming	QST_INFO; QST_WHEN
Guide: and like who will you be coming with	QST_WHO
Tourist: uh yes	FOL_CONFIRM
Tourist: um i'm actually planning to visit	RES_WHEN
Tourist: uh on august	RES_WHEN

Fig. 1. The human-human conversational utterances and their associated semantic labels from DSTC4.

anisms have had great successes in machine translation [19], image captioning [20], and various tasks. Attentional models have been successful because they separate two different concerns: 1) deciding which input contexts are most relevant to the output and 2) actually predicting an output given the most relevant inputs. For example, the highlighted current utterance from the tourist, “*uh on august*”, in the conversation of Figure 1 is to respond the question about *when*, and the content-aware contexts that can help current understanding are the first two utterances from the guide “*and you were saying that you wanted to come to singapore*” and “*un maybe can i have a little bit more details like uh when will you be coming*”. Although content-aware contexts may help understanding, the most recent contexts may be more important than others. In the same example, the second utterance is more related to the *when* question, so the temporal information can provide additional cues for the attention design.

This paper focuses on investigating various attention mechanism in neural models with contextual information and speaker role modeling for language understanding. In order to comprehend what tourist is talking about and imitate how guide react to these meanings, this work proposes a role-based contextual model by modeling role-specific contexts differently for improving the system performance and further design associated time-aware and content-aware attention mechanisms.

2. PROPOSED APPROACH

The model architecture is illustrated in Figure 2. First, the previous utterances are fed into the contextual model to encode into the history summary, and then the summary vector and the current utterance are integrated for helping understanding. The contextual model leverages the attention mechanisms illustrated in the red block, which implements different attention types and attention levels. The whole model is trained in an end-to-end fashion, where the history summary vector and the attention weights are automatically learned based on the downstream SLU task. The objective of the proposed model is to optimize the conditional probability $p(\hat{y} | \mathbf{x})$, so that the difference between the predicted distribution and the target distribution, $q(y_k = z | \mathbf{x})$, can be

minimized:

$$\mathcal{L} = - \sum_{k=1}^K \sum_{z=1}^N q(y_k = z | \mathbf{x}) \log p(\hat{y}_k = z | \mathbf{x}), \quad (1)$$

where the labels \mathbf{y} are the labeled intent tags for understanding.

2.1. Contextual Language Understanding

Given the current utterance $\mathbf{x} = \{w_t\}_1^T$, the goal is to predict the user intents of \mathbf{x} , which includes the speech acts and associated attributes shown in Figure 1. We apply a bidirectional long short-term memory (BLSTM) model [21] to integrate preceding and following words to learn the probability distribution of the user intents.

$$\mathbf{v}_{\text{cur}} = \text{BLSTM}(\mathbf{x}, W_{\text{his}} \cdot \mathbf{v}_{\text{his}}), \quad (2)$$

$$\mathbf{o} = \text{sigmoid}(W_{\text{SLU}} \cdot \mathbf{v}_{\text{cur}}), \quad (3)$$

where W_{his} is a dense matrix and \mathbf{v}_{his} is the history summary vector, \mathbf{v}_{cur} is the context-aware vector of the current utterance encoded by the BLSTM, and \mathbf{o} is the intent distribution. Note that this is a multi-label and multi-class classification, so the *sigmoid* function is employed for modeling the distribution after a dense layer. The user intent labels \mathbf{y} are decided based on whether the value is higher than a threshold θ tuned by the development set.

2.2. Speaker Role Contextual Module

In order to leverage the contextual information, we utilize the prior contexts from two roles to learn history summary representations, \mathbf{v}_{his} in (2). The illustration is shown in the left-right part of Figure 1.

In a dialogue, there are at least two roles communicating with each other, each individual has his/her own goal and speaking habit. For example, the tourists have their own desired touring goals and the guides’ goal is try to provide the sufficient touring information for suggestions and assistance. Prior work usually ignored the speaker role information or only modeled a single speaker’s history for various tasks [16, 22]. The performance may be degraded due to the possibly unstable and noisy input feature space.

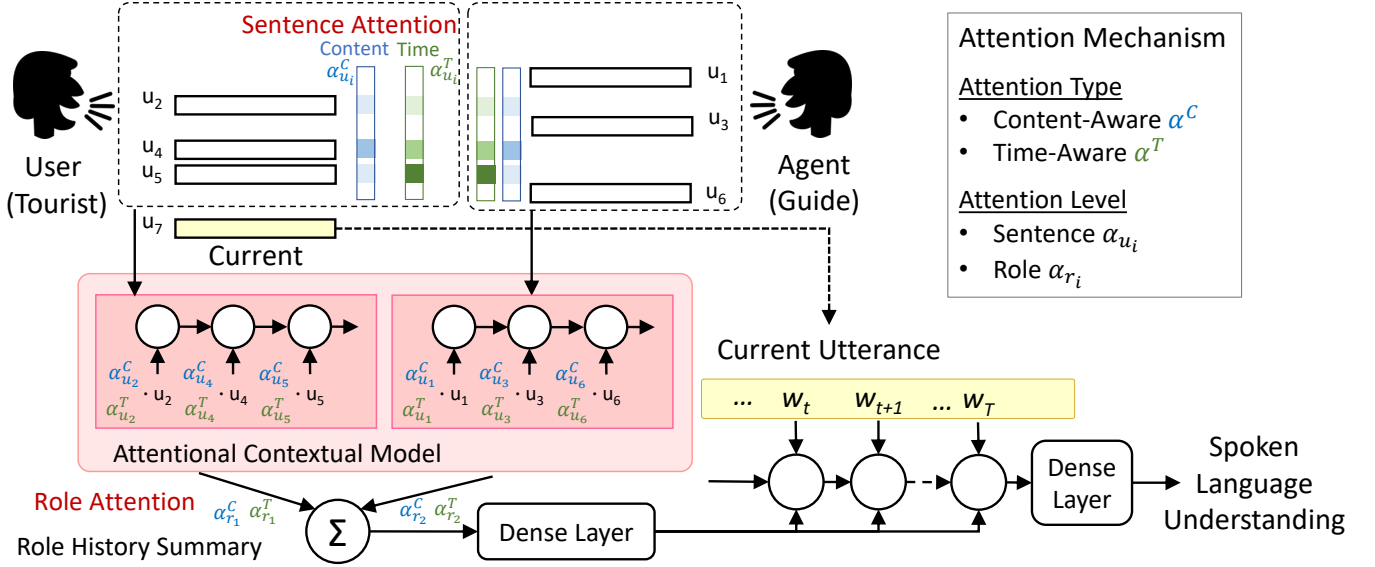


Fig. 2. Illustration of the proposed attentional contextual model.

To address this issue, this work proposes the role-based contextual model: instead of using only a single BLSTM model for the history, we utilize one individual contextual module for each speaker role. Each role-dependent recurrent unit $\text{BLSTM}_{\text{role}_i}$ receives corresponding inputs, x_{t,role_i} , which includes multiple utterances u_i ($i = [1, \dots, t-1]$) preceding to the current utterance u_t from the specific role, role_i , and have been processed by an encoder model.

$$\begin{aligned} \mathbf{v}_{\text{his}} &= \mathbf{v}_{\text{his},\text{role}_a} + \mathbf{v}_{\text{his},\text{role}_b} \\ &= \text{BLSTM}_{\text{role}_a}(x_{t,\text{role}_a}) + \text{BLSTM}_{\text{role}_b}(x_{t,\text{role}_b}), \end{aligned} \quad (4)$$

where x_{t,role_i} is vector after one-hot encoding that represents the annotated intent and the attribute features. Note that this model requires the ground truth annotations for history utterances for training and testing. Therefore, each role-based contextual module focuses on modeling role-dependent goals and speaking style, and \mathbf{v}_{cur} from (2) would contain role-based contextual information.

2.3. Neural Attention Mechanism

One of the earliest work with a memory component applied to language processing is memory networks [15, 23], which encode mentioned facts into vectors and store them in the memory for question answering (QA). The idea is to encode important knowledge and store it into memory for future usage with attention mechanisms. Attention mechanisms allow neural network models to selectively pay attention to specific parts. There are also various tasks showing the effectiveness of attention mechanisms [24, 16]. This paper focuses on two attention types: content-aware (α^C) and time-aware (α^T), and two attention levels: sentence (α_{u_i}) and role (α_{r_i}) illustrated in

Figure 2. The following sections first describe content-aware and time-aware attention mechanism using the sentence-level structure, and the role-level attention is explained afterwards.

2.3.1. Content-Aware Attention

Given the utterance contexts, the model can learn the attention to decide where to focus more for better understanding the current utterance [15, 16]. In the content-aware attention, it is intuitive to use the semantic similarity between the current utterance representations and the history context vector as the attention weight, which is used as a measurement of how much the model should focus on a specific preceding utterance:

$$\alpha_{u_i}^C = \text{softmax}(M_S(\mathbf{v}_{\text{cur}} + \mathbf{x}_i)), \quad (5)$$

where $\alpha_{u_i}^C$ is the content-aware attention vector highlighted as blue texts in Figure 2, the softmax function is used to normalize the attention values of history utterances, and M_S is an MLP for learning the attention weight given the current representation \mathbf{v}_{cur} and the \mathbf{x}_i is the vector of u_i in the history.

Then when computing the history summary vector, the BLSTM encoder additionally considers the corresponding attention weight for each history utterance in order to emphasize the content-related contexts. (4) can be rewritten into

$$\begin{aligned} \mathbf{v}_{\text{his}} &= \text{BLSTM}_{\text{role}_a}(x_{t,\text{role}_a}, \alpha_{\text{role}_a}^C) \\ &+ \text{BLSTM}_{\text{role}_b}(x_{t,\text{role}_b}, \alpha_{\text{role}_b}^C), \end{aligned} \quad (6)$$

2.3.2. Time-Aware Attention

Intuitively, most recent utterances contain more relevant information; therefore we introduce time-aware attention mech-

anism which computes attention weights by the time of utterance occurrence explicitly. We first define the time difference between the current utterance and the preceding sentence as $d(u_i)$ for each preceding utterance u_i , and then simply use its reciprocal as the attention value, $\alpha_{u_i}^T$:

$$\alpha_{u_i}^T = \frac{1}{d(u_i)}. \quad (7)$$

Here the importance of a earlier history sentence would be considerably compressed. For the sentence-level attention, before feeding into the contextual module, each history vector is weighted by its reciprocal of distance.

$$\begin{aligned} \mathbf{v}_{\text{his}} &= \text{BLSTM}_{\text{role}_a}(x_{t,\text{role}_a}, \alpha_{\text{role}_a}^C \cdot \alpha_{\text{role}_a}^T) \\ &+ \text{BLSTM}_{\text{role}_b}(x_{t,\text{role}_b}, \alpha_{\text{role}_b}^C \cdot \alpha_{\text{role}_b}^T), \end{aligned} \quad (8)$$

2.3.3. Dynamic Speaker Role Attention

Switching to role-level attention, a dialogue is disassembled from a different perspective about which speaker’s information is more important. From (4), we have the role-dependent history summary representations, $\mathbf{v}_{\text{his},\text{role}_i}$. The role-level attention is to decide how much to address on different speaker roles’ contexts in order to better understand the current utterance.

$$\alpha_{r_i} = \text{softmax}(M_R(\mathbf{v}_{\text{cur}} + \mathbf{v}_{\text{his},\text{role}_i})), \quad (9)$$

where α_{r_i} is the attention weight of the role i , and M_R is the MLP for learning the role-level content-aware attention weights, while M_S in (5) is at the sentence level.

In terms of time-aware attention, experimental results show that the importance of a speaker given the contexts can be approximated to the minimum distance among the speaker’s utterances³.

$$\alpha_{r_i}^T = \min \frac{1}{d(u_j)}, \quad (10)$$

where u_j includes all contextual utterances from the speaker r_i .

With content-aware or time-aware attention at the role level, the sentence-level history result from (6) can be rewritten into

$$\mathbf{v}_{\text{his}} = \alpha_{\text{role}_a}^{C/T} \cdot \mathbf{v}_{\text{his},\text{role}_a} + \alpha_{\text{role}_b}^{C/T} \cdot \mathbf{v}_{\text{his},\text{role}_b}, \quad (11)$$

for combining role-dependent history vectors with their attention weights.

2.4. End-to-End Training

The objective is to optimize SLU performance, predicting multiple speech acts and attributes described in 2.1. In the proposed model, all encoders, prediction models, and attention weights (except time-aware attention) can be automatically learned in an end-to-end manner.

³Different settings (min, avg,) were attempted.

3. EXPERIMENTS

To evaluate the effectiveness of the proposed model, we conduct the language understanding experiments on human-human conversational data.

3.1. Setup

The experiments are conducted on DSTC4, which consists of 35 dialogue sessions on touristic information for Singapore collected from Skype calls between 3 tour guides and 35 tourists [17]. All recorded dialogues with the total length of 21 hours have been manually transcribed and annotated with speech acts and semantic labels at each turn level. The speaker information (guide and tourist) is also provided. Unlike previous DSTC series collected human-computer dialogues, human-human dialogues contain rich and complex human behaviors and bring much difficulty to all the tasks. Given the fact that different speaker roles behave differently and longer contexts, DSTC4 is a suitable benchmark dataset for evaluation.

We choose the mini-batch *adam* as the optimizer with the batch size of 256 examples. The size of each hidden recurrent layer is 128. We use pre-trained 200-dimensional word embeddings *GloVe* [25]. We only apply 30 training epochs without any early stop approach. We focus on predicting multiple labels including intents and attributes, so the evaluation metric is average F1 score for balancing recall and precision in each utterance. The experiments are shown in Table 1, where we report the average results over five runs.

3.2. Baseline Results

The baselines (rows (a) and (b)) are two of the best participants of DSTC4 in IWSDS 2016 [17] (best results for tourist and guide understanding respectively). It is obvious that tourist intents are much more difficult than guide intents (most systems achieved higher than 60% of F1 for guide intents but lower than 50% for tourist intents), because the guides usually follow similar interaction patterns. In our experiments, we focus more on the tourist part, because SLU in a dialogue system is to understand the users, who ask for assistance.

The baseline (c) only takes the current utterance into account without any history information, and then applies a simple BLSTM for SLU. The baselines (d) and (e) leverage the semantic labels of contexts for learning history summary vectors, achieving the improvement compared to the results without contexts. The baseline (e) that uses history information and conducts our role-based contextual model to capture role-specific information separately slightly improves the SLU performance for both tourist and guide, achieving 69.1% and 74.4% for tourist and guide understanding respectively.

Model	Attention Level	Tourist	Guide	All
Baseline	(a) <i>DSTC4-Best 1</i>	52.1	61.2	57.8
	(b) <i>DSTC4-Best 2</i>	51.1	67.4	61.4
	(c) w/o context	63.2	69.2	66.6
	(d) w/ context w/o speaker role	68.3	74.4	71.6
	(e) w/ context w/ speaker role	69.1	74.4	72.1
Content-Aware Attention	(f) Sentence	68.5	73.6	71.6
	(g) Role	68.6	74.0	71.8
	(h) Both	67.9	73.5	71.5
Time-Aware Attention	(i) Sentence	70.5[†]	77.7[†]	74.6[†]
	(j) Role	70.0 [†]	76.8 [†]	74.2 [†]
	(k) Both	70.4 [†]	77.3 [†]	74.5 [†]
Content-Aware + Time-Aware Attention	(l) Sentence	69.7 [†]	76.7 [†]	74.0 [†]
	(m) Role	70.1[†]	77.1[†]	74.1[†]
	(n) Both	69.5 [†]	76.2 [†]	73.5 [†]

Table 1. Spoken language understanding performance reported on F-measure in DSTC4 (%). [†] indicates the significant improvement compared to all baseline methods.

3.3. Content-Aware Attention

For the content-aware attention, our model learns the importance of each utterance on both sentence-level and role-level in an end-to-end fashion. The table shows that the learned attention for both sentence-level and role-level does not yield improvement compared to the baseline (e). The probable reason is that the content-aware importance can be handled by the BLSTM when producing the history summary, and the learned attention does not provide additional cues for improvement.

3.4. Time-Aware Attention

With the time-aware attention, using the reciprocal of distance as attention value significantly improves the performance to about 70% (tourist) and about 77% (guide) using either sentence-level or role-level attention. The reason that the performance among different levels of time-aware mechanisms are close may be that the closest utterance is capable of capturing the importance of the history. Using the reciprocal of minimum distance among tourist/guide history as their attention weights in the role-level attention can effectively pay the correct attention to the salient information, achieving promising performance.

However, the table shows that the sentence-level attention is slightly better than role-level attention for both tourist and guide parts when considering time-aware attention. The reason may be that the intents are very diverse and hence require the more precise focus on the content-related history in order to achieve correct understanding results. Also, the results using both sentence-level and role-level attention do not differ a lot from the results using only sentence-level attention.

In addition to the current setting for time-aware attention, other trends about temporal decade can be further investi-

gated. We leave this part as the future work.

3.5. Content-Aware and Time-Aware Attention

The proposed model includes two attention types, content-aware and time-aware, and we integrate both types into a single model to analyze whether their advantages can be combined. The rows (l)-(n) are the results using both content-aware and time-aware attention weights, but the performance is similar to the time-aware results and the difference is not significant. For the results using only role-level attention, combining content-aware and time-aware attention weights obtains the improvement (row (j) v.s row (m)).

3.6. Comparison between Sentence-Level and Role-Level

Among the experiments using either content-aware attention or time-aware attention or both, we introduce both sentence-level and role-level mechanisms for analysis. From Table 1, rows (i) and (l) show that the sentence-level attention can benefit the results even though we do not leverage the content-aware attention. On the other hand, rows (j) and (m) show that the role-level attention requires content-related information in order to achieve better performance. The reason is probably that the role-level attention is not precise enough to capture the salience of utterances. Therefore, combining with content-aware cues can effectively focus on the correct part.

3.7. Dynamic Speaker Role Attention Analysis

To further analyze the dynamic role-level attention, we compute the mean of the learned attention weights. The results are shown in Table 2.

For tourist understanding, the tourist and the guide attention weights are 0.48 and 0.52 respectively. We can there-

Task	Role-Level Attention Weight	
	Tourist Context	Guide Context
Tourist Understanding	0.48	0.52
Guide Understanding	0.30	0.70

Table 2. The average role-level attention weight learned from the proposed model using the role-level attention mechanism.

fore conclude that both tourist and guide history are likewise important in order to understand the tourist utterances. The observation matches our analysis about tourist understanding in 3.6 — since tourist intents are highly diverse, SLU needs to extract as much information as possible from both speaker roles to help understanding. In addition, it is reasonable that tourist understanding focuses on guide history a little bit more than tourist history under a dialogue scenario.

For guide understanding, the tourist and the guide attention weights are 0.30 and 0.70 respectively. The remarkable difference of role-level attention can be explained by the data characteristics of DSTC4. In the tourist and guide dialogues, it is obvious that the guides usually follow the same interactive patterns, such as suggesting a famous location first and then explaining it. In other words, the tourist contexts do not matter a lot when understanding guide utterances, because the guide only needs to follow a pattern to interact with the tourist. This also explains why guide understanding always achieves higher performance compared to the tourist part in our experiments.

To further verify whether the learned attention focuses on the right part, we conduct a similar experiment as the row (d) in Table 1. Instead of using preceding utterances from both speakers as history information, when understanding tourist utterances, we use only tourist history; likewise, guide understanding only takes guide history into account. The performance of tourist understanding drops from 68.3% to 65.9%, while the performance of guide understanding remains almost the same (from 74.4% to 73.8%). The trend proves that the content-aware attention mechanism is capable of focusing on the correct part for better understanding. In the future, we would like to investigate whether the proposed model can generalize to the scenarios with more than two speaker roles.

3.8. Overall Results

From the above experimental results, the proposed time-aware attention model significantly improves the performance, and sentence-level and role-level attention has different capacities of improving understanding performance, where sentence-level attention is useful even though we do not consider content-aware information, and role-level attention is more effective when combining both content-aware and time-aware attention.

In sum, the best results are time-aware attention with sentence-level attention (row (i)) and content-aware and time-

aware attention with role-level attention (row (m)), reaching higher than 70% and 77% on F-measure for tourist understanding and guide understanding respectively. Therefore, the proposed attention mechanisms are demonstrated to be effective for improving SLU in such complex human-human conversations.

4. CONCLUSION

This paper proposes an end-to-end attentional role-based contextual model that automatically learns speaker-specific contextual encoding and investigates various content-aware and time-aware attention mechanisms on it. Experiments on a benchmark multi-domain human-human dialogue dataset show that the time-aware and role-level attention mechanisms provide additional cues to guide the model to focus on the salient contexts, and achieve better performance on spoken language understanding for both speaker roles. Moreover, the proposed time-aware and role-level attention mechanisms are easily extendable to multi-party conversations, and we leave the extension of other variants for the future work.

5. ACKNOWLEDGEMENTS

We would like to thank reviewers for their insightful comments on the paper. The authors are supported by the Ministry of Science and Technology of Taiwan and MediaTek Inc..

6. REFERENCES

- [1] Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, David Vandyke, and Steve Young, “A network-based end-to-end trainable task-oriented dialogue system,” in *Proceedings of EACL*, 2017, pp. 438–449.
- [2] Antoine Bordes, Y-Lan Boureau, and Jason Weston, “Learning end-to-end goal-oriented dialog,” in *Proceedings of ICLR*, 2017.
- [3] Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng, “Towards end-to-end reinforcement learning of dialogue agents for information access,” in *Proceedings of ACL*, 2017, pp. 484–495.
- [4] Xiujun Li, Yun-Nung Chen, Lihong Li, Jianfeng Gao, and Asli Celikyilmaz, “End-to-end task-completion neural dialogue systems,” in *Proceedings of The 8th International Joint Conference on Natural Language Processing*, 2017.
- [5] Gokhan Tur and Renato De Mori, *Spoken language understanding: Systems for extracting semantic information from speech*, John Wiley & Sons, 2011.

- [6] Dilek Hakkani-Tür, Gökhan Tür, Asli Celikyilmaz, Yun-Nung Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang, “Multi-domain joint semantic frame parsing using bi-directional rnn-lstm,” in *Proceedings of INTER-SPEECH*, 2016, pp. 715–719.
- [7] Yun-Nung Chen, Dilek Hakkani-Tur, Gokhan Tur, Asli Celikyilmaz, Jianfeng Gao, and Li Deng, “Knowledge as a teacher: Knowledge-guided structural attention networks,” *arXiv preprint arXiv:1609.03286*, 2016.
- [8] Yun-Nung Chen, Dilek Hakanni-Tür, Gokhan Tur, Asli Celikyilmaz, Jianfeng Guo, and Li Deng, “Syntax or semantics? knowledge-guided joint semantic frame parsing,” in *Proceedings of 2016 IEEE Spoken Language Technology Workshop*, 2016, pp. 348–355.
- [9] Zhangyang Wang, Yingzhen Yang, Shiyu Chang, Qing Ling, and Thomas S Huang, “Learning a deep l encoder for hashing,” in *Proceedings of IJCAI*, 2016, pp. 2174–2180.
- [10] Anshuman Bhargava, Asli Celikyilmaz, Dilek Hakkani-Tur, and Ruhi Sarikaya, “Easy contextual intent prediction and slot detection,” in *Proceedings of ICASSP*, 2013, pp. 8337–8341.
- [11] Puyang Xu and Ruhi Sarikaya, “Contextual domain classification in spoken language understanding systems using recurrent neural network,” in *Proceedings of ICASSP*, 2014, pp. 136–140.
- [12] Yun-Nung Chen, Ming Sun, Alexander I. Rudnicky, and Anatole Gershan, “Leveraging behavioral patterns of mobile applications for personalized spoken language understanding,” in *Proceedings of ICMI*, 2015, pp. 83–86.
- [13] Ming Sun, Yun-Nung Chen, and Alexander I. Rudnicky, “An intelligent assistant for high-level task understanding,” in *Proceedings of IUI*, 2016, pp. 169–174.
- [14] Yangyang Shi, Kaisheng Yao, Hu Chen, Yi-Cheng Pan, Mei-Yuh Hwang, and Baolin Peng, “Contextual spoken language understanding using recurrent neural networks,” in *Proceedings of ICASSP*, 2015, pp. 5271–5275.
- [15] Jason Weston, Sumit Chopra, and Antoine Bordes, “Memory networks,” in *Proceedings of ICLR*, 2015.
- [16] Yun-Nung Chen, Dilek Hakkani-Tür, Gökhan Tür, Jianfeng Gao, and Li Deng, “End-to-end memory networks with knowledge carryover for multi-turn spoken language understanding,” in *Proceedings of INTER-SPEECH*, 2016, pp. 3245–3249.
- [17] Seokhwan Kim, Luis Fernando DHaro, Rafael E Banchs, Jason D Williams, and Matthew Henderson, “The fourth dialog state tracking challenge,” in *Proceedings of IWSDS*, 2016.
- [18] Ta-Chung Chi, Po-Chun Chen, Shang-Yu Su, and Yun-Nung Chen, “Speaker role contextual modeling for language understanding and dialogue policy learning,” in *Proceedings of IJCNLP*, 2017.
- [19] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [20] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio, “Show, attend and tell: Neural image caption generation with visual attention,” in *Proceedings of ICML*, 2015, pp. 2048–2057.
- [21] Mike Schuster and Kuldip K Paliwal, “Bidirectional recurrent neural networks,” *IEEE Transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, 1997.
- [22] Xuesong Yang, Yun-Nung Chen, Dilek Hakkani-Tür, Paul Crook, Xiujun Li, Jianfeng Gao, and Li Deng, “End-to-end joint learning of natural language understanding and dialogue manager,” in *Proceedings of ICASSP*, 2017, pp. 5690–5694.
- [23] Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al., “End-to-end memory networks,” in *Proceedings of NIPS*, 2015, pp. 2431–2439.
- [24] Caiming Xiong, Stephen Merity, and Richard Socher, “Dynamic memory networks for visual and textual question answering,” *arXiv preprint arXiv:1603.01417*, 2016.
- [25] Jeffrey Pennington, Richard Socher, and Christopher D Manning, “Glove: Global vectors for word representation,” in *Proceedings of EMNLP*, 2014, vol. 14, pp. 1532–1543.