

Improving Dialogue State Tracking by Joint Slot Modeling

Ting-Rui Chiang and Yi-Ting Yeh

Carnegie Mellon University

{tingruic, yitingye}@andrew.cmu.edu

Abstract

Dialogue state tracking models play an important role in a task-oriented dialogue system. However, most of them model the slot types conditionally independently given the input. We discover that it may cause the model to be confused by slot types that share the same data type. To mitigate this issue, we propose TripPy-MRF and TripPy-LSTM that models the slots jointly. Our results show that they are able to alleviate the confusion mentioned above, and they push the state-of-the-art on dataset Multi-WoZ 2.1 from 58.7 to 61.3. Our implementation is available at <https://github.com/CTinRay/Trippy-Joint>.

1 Introduction

Multi-domain dialogue state tracking (DST) has been an important challenge of conversational artificial intelligence (Ram et al., 2018). Compared to single-domain DST, the belief state of a multi-domain DST contains slot types across multiple domains. It causes difficulty to the traditional DST models that assume the complete ontology is available because a complete ontology becomes hard to obtain (Wu et al., 2019b) when there are many domains,

To tackle these challenges, many of the mainstream approaches for DST formulate this as a span prediction task (Xu and Hu, 2018; Wu et al., 2019a; Kumar et al., 2020). Instead of predicting a predefined value in the ontology, the span-based DST models predict slot values by extracting a span from the conversation between the user and the system. By directly extracting spans as slot values, the DST models are able to handle unseen slot values and are potentially transferable to different domains. However, the span-based DST models can only deal with the slot values that are explicitly expressed as a sequence of tokens and fall short in dealing with coreference ("I'd like a

restaurant in the same area") and implicit choice ("Any of those is ok") The current state-of-the-art DST model TripPy (Heck et al., 2020) is designed to deal with these problems. In addition to extracting values directly from the user utterance, TripPy maintains two additional memories on the fly and uses them to deal with the coreference and implicit choice challenges. The design details of TripPy will be further illustrated in Section 2.2

Despite the success of these span-based models, one concern is that they often consider the slots conditionally independently. For example, in TripPy, the conversation is first encoded by BERT (Devlin et al., 2019), and then the extracted features are used by n modules for the n slot types independently. That implies, conditioning the features extracted by BERT, the slots are predicted independently. This may increase the difficulty of this task because these slot types are indeed not independent of each other. If the value of one slot type in a domain is present in an utterance, then it is more likely that the value of another slot type of the same domain is present.

In particular, in this work, we focus on the cases where modeling the joint probability of the slots could possibly be helpful. We notice that there are some slot types that share the same data type. The data types of both the slot type "hotel-people" and the slot type "restaurant-book people" are both integers. Therefore, when an integer is present in an utterance, a DST model needs to predict which slot(s) this integer is for. Given the fact that some slots co-occur more often than others, modeling the slot types jointly may be helpful when making this kind of prediction. Based on this intuition, we group the slot types into 4 groups by their data type: time, place name, integer. We then analyze the current state-of-the-art model TripPy (Heck et al., 2020; Mehri et al., 2020). Our analysis shows that TripPy is indeed often confused with slot types in the same group.

We then investigate whether modeling the relation between slot types explicitly can alleviate this problem. By modeling the slot types jointly, information from the slots that the model is more confident with may reduce the confusion for the other slots. Specifically, to achieve this goal, we propose to model the joint probability of slots with MRF and LSTM models. They can be plugged into any DST model that models the slots conditionally independently. In this work, we choose the current state-of-the-art DST model, TripPy pretrained with DialogGLUE (Mehri et al., 2020). We experiment on the most widely used multi-domain DST dataset MultiWoZ 2.1 (Zang et al., 2020b; Eric et al., 2019). The results show that both the MRF and the LSTM models are able to mitigate the confusion, and can improve the performance consistently. We push the state-of-the-art performance on MultiWoZ 2.1 from 58.7 to 61.3.

2 Background

2.1 Problem Formulation

The dialogue state tracking (DST) task is to predict values of slot types for each turn in task-oriented dialogue. Specifically, dialogue consists of a sequence of user and system utterances. The input of the task is a dialogue $X = \{(U_1, M_1), \dots, (U_T, M_T)\}$, where U_t and M_t are utterances by user and system respectively, and T is the total number of turns in this dialogue. The output of the task is the dialogue state at each time step. A dialogue state consists of N slot types. The goal of a DST task is to predict the slot values of those N slot types at each time step. The slot value $S_s^{(t)}$ of the s -th slot type at time step t might be a span in the previous conversation or simply *none*. We denote the state at time step t as $S^{(t)} = \{S_1^{(t)}, S_2^{(t)}, \dots, S_N^{(t)}\}$.

2.2 TripPy: Triple Copy Strategy for DST

As described in Section 1, in addition to extract values from the user utterance, TripPy maintains two memories to tackle the coreference and the implicit choice problems in the span-based DST model. Specifically, we use *system inform memory* to remember the slots which the model previously informed, and *DS memory* to store the seen slots in the dialogue history. At dialogue turn t , TripPy predicts the belief state $S^{(t)}$ with the following steps: 1) The model first uses a pretrained BERT model to encode the dialogue history until the current time

step t . 2) For each slot type s , it uses a classifier to predict the *slot class* $C_s^{(t)}$. For slot types whose possible values are "yes" or "no", possible values of $C_s^{(t)}$ are $\{none, dontcare, yes, no\}$. For the other slot types, possible values of $C_s^{(t)}$ are $\{none, dontcare, span, inform, refer\}$. These slot classes determine which source the slot value should be copied from. 3) The first two slot classes express special cases. *none* denotes that the slot does not take a value in this turn, and *dontcare* states that any value is acceptable for this slot. If $C_s^{(t)}$ is *span*, then TripPy uses a span prediction model to copy a span of text from the dialogue history as the prediction of a slot type. On the other hand, if $C_s^{(t)}$ is *inform* and *refer*, the model will instead copy the value from *system inform memory* and *DS memory* respectively. The *system inform memory* allows the model to solve the implicit choice issue and the *DS memory* helps the model solve coreference problems.

2.3 Experimental Setup

We train and test our model on MultiWoZ 2.1 (Eric et al., 2019), which is the most challenging dataset for DST and is widely used for the evaluation of multi-domain DST models. MultiWoZ 2.1 is comprised of over 10000 dialogues in 5 domains and has 30 different slots with over 45000 possible values. We compute the joint goal accuracy (JGA) on all test samples to evaluate our models. The JGA is the ratio of dialog turns in the dataset for which all slots have been filled with the correct value according to the ground truth. Our baseline is the TripPy model in (Mehri et al., 2020), which is additionally pretrained with more dialogue tasks. To reproduce their result, we follow the setting of their hyperparameters and successfully achieve JGA 58.0, which is close to the reported result of 58.7 in the paper.

3 Analysis

3.1 Accuracy per Copy Class

We analyze the accuracy of each slot by the ground truth copy strategy to use on the validation set. For each slot type, we also calculate the joint goal accuracy presuming it is predicted correctly. From the results in Table 1 we can have a few observations. 1) The class *none* takes most of the cases. Therefore, it is important to predict whether the value is *none* correctly. 2) On the other hand, improving other classes may be less effective, since the number of mistakes is much smaller.

Slot Type	none	don't care	span	inform	refer	true	false	Oracle	Δ
taxi-leaveAt	.99 (7218)	-. (0)	.88 (131)	.91 (23)	.00 (2)	-. (0)	-. (0)	.581	.001
taxi-destination	.98 (7159)	-. (0)	.98 (108)	.94 (18)	.61 (89)	-. (0)	-. (0)	.587	.007
taxi-departure	.98 (7169)	-. (0)	.89 (100)	1.0 (21)	.60 (84)	-. (0)	-. (0)	.587	.007
taxi-arriveBy	.99 (7278)	.50 (2)	.85 (55)	1.0 (10)	.76 (29)	-. (0)	-. (0)	.581	.001
restaurant-book_people	.99 (7001)	-. (0)	.99 (319)	1.0 (41)	.85 (13)	-. (0)	-. (0)	.583	.002
restaurant-book_day	1.0 (7011)	-. (0)	.99 (286)	1.0 (53)	.92 (24)	-. (0)	-. (0)	.581	.001
restaurant-book_time	.99 (6979)	-. (0)	.98 (328)	1.0 (67)	-. (0)	-. (0)	-. (0)	.583	.003
restaurant-food	.99 (6801)	.45 (20)	1.0 (422)	.98 (131)	-. (0)	-. (0)	-. (0)	.593	.013
restaurant-pricerange	.99 (6906)	.78 (9)	.98 (343)	.98 (113)	1.0 (3)	-. (0)	-. (0)	.587	.006
restaurant-name	.96 (6926)	.00 (2)	.93 (162)	1.0 (283)	.00 (1)	-. (0)	-. (0)	.598	.018
restaurant-area	.98 (6881)	.71 (17)	.97 (314)	.96 (137)	.88 (25)	-. (0)	-. (0)	.594	.014
hotel-book_people	.99 (7031)	-. (0)	.99 (300)	.97 (35)	.62 (8)	-. (0)	-. (0)	.581	.001
hotel-book_day	1.0 (7035)	-. (0)	.98 (289)	1.0 (40)	.70 (10)	-. (0)	-. (0)	.582	.002
hotel-book_stay	.99 (7015)	-. (0)	.99 (314)	1.0 (45)	-. (0)	-. (0)	-. (0)	.582	.001
hotel-name	.96 (6884)	.00 (1)	.98 (170)	1.0 (319)	-. (0)	-. (0)	-. (0)	.59	.009
hotel-area	.97 (7028)	.45 (29)	1.0 (205)	.98 (92)	.65 (20)	-. (0)	-. (0)	.595	.014
hotel-parking	.97 (7204)	.33 (6)	-. (0)	-. (0)	-. (0)	.86 (154)	.50 (10)	.587	.007
hotel-pricerange	.98 (6986)	.67 (21)	.97 (271)	1.0 (91)	.60 (5)	-. (0)	-. (0)	.59	.01
hotel-stars	.99 (7094)	.50 (6)	.98 (210)	.91 (64)	-. (0)	-. (0)	-. (0)	.584	.004
hotel-internet	.98 (7200)	1.0 (2)	-. (0)	-. (0)	-. (0)	.89 (169)	.67 (3)	.588	.007
hotel-type	.93 (7135)	.00 (5)	-. (0)	-. (0)	-. (0)	.59 (80)	.85 (154)	.612	.031
attraction-type	.98 (6835)	.38 (8)	.98 (430)	1.0 (101)	-. (0)	-. (0)	-. (0)	.595	.015
attraction-name	.98 (7006)	.00 (2)	.97 (137)	1.0 (229)	-. (0)	-. (0)	-. (0)	.59	.01
attraction-area	.98 (6946)	.39 (18)	.98 (248)	.96 (135)	.93 (27)	-. (0)	-. (0)	.593	.013
train-book_people	.98 (7040)	-. (0)	.96 (314)	1.0 (8)	.42 (12)	-. (0)	-. (0)	.591	.011
train-leaveAt	.97 (7125)	.50 (2)	.96 (228)	.68 (19)	-. (0)	-. (0)	-. (0)	.612	.032
train-destination	.98 (6693)	-. (0)	.99 (503)	1.0 (176)	.00 (2)	-. (0)	-. (0)	.586	.006
train-day	.99 (6791)	-. (0)	1.0 (481)	1.0 (86)	.94 (16)	-. (0)	-. (0)	.585	.005
train-arriveBy	.99 (7034)	.75 (4)	.98 (322)	.75 (12)	.00 (2)	-. (0)	-. (0)	.592	.012
train-departure	.98 (6749)	-. (0)	.99 (478)	1.0 (146)	.00 (1)	-. (0)	-. (0)	.587	.007
sum	3928	77	168	73	29	116	37		

Table 1: Accuracy of each slot type by ground truth class. The number in the parenthesis is the number of occurrence of the class for the slot type. The oracle column is the performance calculated when presuming the slot value is predicted correctly.

The above results encourage us to inspect the prediction of the class *none* more closely. Here we focus on the cases where the model is to decide whether it is going to give a value to a slot type. We ignore cases where the value of the slot type in the previous turn is not none. We calculate the number of incorrect value predictions when *none* is predicted incorrectly. From Table 2, we can see that many incorrect predictions are resulted from incorrect *none* prediction.

3.2 Relation Between Slots

We further examine the model’s behavior when *none* is not predicted correctly. We check the frequency of the following cases of incorrect prediction: 1) p_1 : when the ground truth class is *none*, the cases where the predicted slot value is the ground truth value of another slot type. 2) p_2 : when the predicted class is *none*, cases of the slot value that is the ground truth value of another slot type. We plot the frequency $p_1 + p_2$ in Figure 1. From the figure, we can see that the model is often confused by similar slot types. For example, the value of the slot type "taxi-departure" and the value of the slot type "taxi-arriveBy" is time. It may be the reason why the model often fails to give non-*none* value to the correct slot type.

4 Proposed Method

We conjecture some mistakes made by TripPy maybe because it models the slots conditionally independently. Based on Table 2, the model does not predict *none* very accurately, and Figure 1 indicates that it is because the model is often confused by similar slot types. In the original design, given a dialogue X , TripPy models the class prediction independently as follow:

$$P(C_1^{(t)}, C_2^{(t)}, \dots, C_N^{(t)} | X) = \prod_{s=1}^N P(C_s^{(t)} | X). \quad (1)$$

TripPy ignores the potential correlation between different slots. Thus, for all s , $P(C_s^{(t)} | X)$ needs to predict the class correctly at the same time. It may be difficult since there are numerous slots.

To alleviate this problem, we propose to directly model the joint probability $P(C_1^{(t)}, C_2^{(t)}, \dots, C_N^{(t)} | X)$. If c_1, c_2, \dots, c_N are the correct classes, modeling the joint probability does not require $P(C_s^{(t)} = c_s) > P(C_s^{(t)} = c'_s)$ for all s and all incorrect assignments $\{c'_s\}_{s=1}^N$. Instead, to predict the belief state correctly, the model is only required to ensure $P(C_1^{(t)} = c_1, C_2^{(t)} = c_2, \dots, C_N^{(t)} = c_N | X) > P(C_1^{(t)} = c'_1, C_2^{(t)} = c'_2, \dots, C_N^{(t)} = c'_N | X)$. In this way, information from the other states



Figure 3: The value of $p_1 + p_2$ (Section 3.2) predicted by the TripPy + LSTM model.

is parameter-efficient and incurs little computation overhead for inference. However, having a large number of vertices in an MRF makes training intractable. Therefore, we follow the intuition above, grouping the slots by their domains, and model the joint distribution within each group respectively. Since a large portion of the mistakes made by the original TripPy model is due to the incorrect prediction of whether a slot value is present in an utterance, we further decompose the probability of a class as

$$\begin{aligned}
 P(C_s^{(t)} = c) \\
 = P(C_s^{(t)} = c | C_s^{(t)} \neq \text{"none"}) P(C_s^{(t)} \neq \text{"none"})
 \end{aligned} \tag{2}$$

Let $C_s'^{(t)}$ be the random variable indicating whether $C_s^{(t)} = \text{"none"}$. Modeling only the $P(\{C_s'^{(t)}\}_{s=1}^N)$ part jointly makes the training and inference of the MRF model feasible: given input x , for each domain d and the assignment of slots c_1, c_2, \dots, c_N for the N slots in d , the probability of the assignment is

$$\begin{aligned}
 P(c_1, c_2, \dots, c_N | x) \\
 = \frac{1}{Z(x)} \Phi_d(\{C_s'^{(t)} = c'_s\}_{s=1}^N) \\
 \prod_{s=1}^N \Phi_s(c_s, x) P(C_s^{(t)} = c_s | C_s'^{(t)} = c'_s),
 \end{aligned} \tag{3}$$

where Φ_d map each assignment c'_1, c'_2, \dots, c'_N to a learnable scalar, and $\Phi_s(c'_s, x)$ is a scalar resulted

from applying a linear layer over the pooled BERT output. Since the number of slot types in a domain is at most 10, and all of the variables are binary, the partition function $Z_i(x)$ can be calculated by enumerating all the possible assignments of c'_1, c'_2, \dots, c'_N . Therefore, we can train this model with the plain MLE loss.

4.2 Modeling with LSTM models

Another approach we explore is to use the LSTM model (Schmidhuber and Hochreiter, 1997). Similar to language modeling, we model the joint probability by predicting the value of each slot sequentially:

$$\begin{aligned}
 P(C_1^{(t)}, C_2^{(t)}, \dots, C_N^{(t)} | X) \\
 = \prod_{s=1}^N P(C_s^{(t)} | X, C_1^t, \dots, C_{s-1}^t),
 \end{aligned} \tag{4}$$

For the implementation of this model, we first encode the dialogue history X with BERT and extract its CLS token representation h_X . The original TripPy model will feed h_X into a linear layer f_s to predict slot class for each slot.

$$P(C_s^{(t)} | X) = \text{Softmax}(f_s(h_X)) \tag{5}$$

Here, for all s , we additionally train another linear layer f_s^{lstm} to transform h_X and feed it into a LSTM layer:

$$h_s = \text{LSTM}(f_s^{lstm}(h_X), h_{s-1}) \tag{6}$$

	Other	FN	FP
taxi-leaveAt	13	7	25
taxi-destination	10	32	13
taxi-departure	14	37	13
taxi-arriveBy	1	15	19
restaurant-book_people	1	5	37
restaurant-book_day	1	3	16
restaurant-book_time	3	4	12
restaurant-food	3	13	48
restaurant-pricerange	4	8	57
restaurant-name	9	7	45
restaurant-area	3	19	57
hotel-book_people	2	6	34
hotel-book_day	1	7	8
hotel-book_stay	2	2	17
hotel-name	6	3	72
hotel-area	2	24	94
hotel-parking	11	20	88
hotel-pricerange	5	13	38
hotel-stars	4	10	24
hotel-internet	1	18	85
hotel-type	12	49	220
attraction-type	2	11	104
attraction-name	21	5	31
attraction-area	0	25	53
train-book_people	1	19	37
train-leaveAt	3	14	57
train-destination	6	4	12
train-day	1	2	17
train-arriveBy	3	9	24
train-departure	3	6	8
	148	397	1365

Table 2: Number of incorrect slot value predictions in different cases of *none* prediction. FN (false negative)/FP (false positive) means that the ground truth class is *none* is not *none*, but the model predicts the class incorrectly.

Model	Dev.	Test
TripPy (Mehri et al., 2020)		58.7
TripPy (reproduced)	61.3	58.0
+ MRF	59.3	60.1
+ LSTM	62.8	61.3

Table 3: Joint goal accuracy on MultiWoZ 2.1.

We then concatenate h_s and h_x to model the conditional probability:

$$P(C_s^{(t)}|X, C_1^{(t)}, \dots, C_{s-1}^{(t)}) = f_s([h_x; h_s]) \quad (7)$$

where $[x; y]$ denotes the concatenation of vectors x and y . We name this model TripPy+LSTM.

5 Results and Discussion

Our main results are shown in Table 3. Both MRF and LSTM modules have a consistent improvement over the baseline model on the test set of MultiWoZ 2.1, which proves the effectiveness of our proposed approaches.

As a more fine-grained analysis, we plot the frequency of incorrect prediction as described in Section 3.2. Compared to Figure 1, Figure 2 and Figure 3 show that both TripPy+MRF and TripPy+LSTM can relieve the confusion between slot types that have the same data type. This is coherent to our intuition for modeling the slot types jointly.

The results also show that TripPy+LSTM performs better than TripPy+MRF. It may be because TripPy+LSTM models all the assignment $\{C_s\}_{s=1}^N$ of the classes jointly, while TripPy+MRF only models $\{C'_s\}_{s=1}^N$ jointly. Also, TripPy+LSTM has the potential to model the joint distribution across the domains. However, this also sacrifices efficiency. Since TripPy+LSTM predicts the slots sequentially, it cannot predict the slots in parallel. TripPy+MRF could be a better choice when latency is an important issue.

6 Related Work

Williams et al. (2013) is the first benchmark for DST models. After that, Henderson et al. (2014) proposed the DSTC2 dataset which includes more linguistic phenomena. Ramadan et al. (2018a) proposed MultiWoZ, which is the first large-scale human-human multi-domain task-oriented DST dataset. Its successive versions (Eric et al., 2019;

Zang et al., 2020a) cleanse the noisy conversations in it. Since then, it has become the standard DST benchmark dataset.

Early approaches for DST rely on a list candidate slot values (Mrkšić et al., 2017; Liu and Lane, 2017; Ramadan et al., 2018b; Rastogi et al., 2017). Some recent approaches formulate it as a sequence generation task. TRADE (Wu et al., 2019a) and MA-DST (Kumar et al., 2020) generates slot values with a pointer generator with a copy mechanism. SOM-DST (Kim et al., 2020) uses a BERT model to encode the dialogue. To improve the efficiency, the state from the previous turn instead of the whole dialogue history is fed into the model. SUMBT (Lee et al., 2019) utilizes a BERT model (Devlin et al., 2019) to encode both the dialogue history and candidate slot values. In Gao et al. (2019) they formulate DST as a question answering task. DS-DST (Zhang et al., 2020) adopts a dual strategy: they predict slot values from candidates but also predict slot values by copying from the dialogue. TripPy (Heck et al., 2020) predicts slot values by copying text span from three different sources. In Mehri et al. (2020), they improve the performance of TripPy by pre-training the BERT with multiple dialogue tasks. Simple-TOD (Hosseini-Asl et al., 2020) and SOLOIST (Peng et al., 2020) utilize GPT-2 for end-to-end slot value and response generation. In Li et al. (2020), they investigate the generalization capability of DST models for unseen slot values. MinTL (Lin et al., 2020) augments data by substituting slot values in the utterances and modifying the syntax structure.

7 Conclusion

In this work, we explore the relation between slot types. Our contributions are two-fold. 1) We identify a problem with the current state-of-the-art model TripPy. We observe that it is often confused by slots that share the same type of value. 2) We propose TripPy+MRF and TripPy+LSTM that aim at mitigating this problem. One can choose one of them based on the trade-off between efficiency and accuracy. Our preliminary results show that both of them are able to alleviate the confusion. Our findings in this work lead future researchers in a promising direction to improve the performance of multi-domain DST.

Acknowledgments

We would like to thank Graham Neubig for his in-depth discussions. We are also thankful to the anonymous reviewers for their comments on the paper.

References

- 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, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mihail Eric, Rahul Goel, Shachi Paul, Adarsh Kumar, Abhishek Sethi, Peter Ku, Anuj Kumar Goyal, Sanchit Agarwal, Shuyang Gao, and Dilek Hakkani-Tur. 2019. Multiwoz 2.1: A consolidated multi-domain dialogue dataset with state corrections and state tracking baselines. *arXiv preprint arXiv:1907.01669*.
- Shuyang Gao, Abhishek Sethi, Sanchit Agarwal, Tagyoung Chung, and Dilek Hakkani-Tur. 2019. [Dialog state tracking: A neural reading comprehension approach](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 264–273, Stockholm, Sweden. Association for Computational Linguistics.
- Michael Heck, Carel van Niekerk, Nurul Lubis, Christian Geishauer, Hsien-Chin Lin, Marco Moresi, and Milica Gasic. 2020. [TripPy: A triple copy strategy for value independent neural dialog state tracking](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 35–44, 1st virtual meeting. Association for Computational Linguistics.
- 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, Philadelphia, PA, U.S.A. Association for Computational Linguistics.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. A simple language model for task-oriented dialogue. *arXiv preprint arXiv:2005.00796*.
- Sungdong Kim, Sohee Yang, Gyuwan Kim, and Sangwoo Lee. 2020. [Efficient dialogue state tracking by selectively overwriting memory](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 567–582, Online. Association for Computational Linguistics.

- Adarsh Kumar, Peter Ku, Anuj Goyal, Angeliki Metallinou, and Dilek Hakkani-Tur. 2020. [Ma-dst: Multi-attention-based scalable dialog state tracking](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8107–8114.
- Hwaran Lee, Jinsik Lee, and Tae-Yoon Kim. 2019. [SUMBT: Slot-utterance matching for universal and scalable belief tracking](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5478–5483, Florence, Italy. Association for Computational Linguistics.
- Shiyang Li, Semih Yavuz, Kazuma Hashimoto, Jia Li, Tong Niu, Nazneen Rajani, Xifeng Yan, Yingbo Zhou, and Caiming Xiong. 2020. [Coco: Controllable counterfactuals for evaluating dialogue state trackers](#). *ArXiv*, abs/2010.12850.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, and Pascale Fung. 2020. [Mintl: Minimalist transfer learning for task-oriented dialogue systems](#). *arXiv preprint arXiv:2009.12005*.
- Bing Liu and Ian Lane. 2017. [An end-to-end trainable neural network model with belief tracking for task-oriented dialog](#). In *Proc. Interspeech 2017*, pages 2506–2510.
- S. Mehri, M. Eric, and D. Hakkani-Tur. 2020. [Dialogue: A natural language understanding benchmark for task-oriented dialogue](#). *ArXiv*, abs/2009.13570.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2017. [Neural belief tracker: Data-driven dialogue state tracking](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1777–1788, Vancouver, Canada. Association for Computational Linguistics.
- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayan-deh, Lars Liden, and Jianfeng Gao. 2020. [Soloist: Few-shot task-oriented dialog with a single pre-trained auto-regressive model](#). *arXiv preprint arXiv:2005.05298*.
- Ashwin Ram, Rohit Prasad, Chandra Khatri, Anu Venkatesh, Raefer Gabriel, Qing Liu, Jeff Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, et al. 2018. [Conversational ai: The science behind the alexa prize](#). *arXiv preprint arXiv:1801.03604*.
- Osman Ramadan, Paweł Budzianowski, and Milica Gašić. 2018a. [Large-scale multi-domain belief tracking with knowledge sharing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, volume 2, pages 432–437.
- Osman Ramadan, Paweł Budzianowski, and Milica Gašić. 2018b. [Large-scale multi-domain belief tracking with knowledge sharing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 432–437, Melbourne, Australia. Association for Computational Linguistics.
- Abhinav Rastogi, Dilek Hakkani-Tür, and Larry Heck. 2017. [Scalable multi-domain dialogue state tracking](#). In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 561–568. IEEE.
- Jürgen Schmidhuber and Sepp Hochreiter. 1997. [Long short-term memory](#). *Neural Comput*, 9(8):1735–1780.
- Jason Williams, Antoine Raux, Deepak Ramachandran, and Alan Black. 2013. [The dialog state tracking challenge](#). In *Proceedings of the SIGDIAL 2013 Conference*, pages 404–413, Metz, France. Association for Computational Linguistics.
- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019a. [Transferable multi-domain state generator for task-oriented dialogue systems](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 808–819, Florence, Italy. Association for Computational Linguistics.
- Chien-Sheng Wu, Richard Socher, and Caiming Xiong. 2019b. [Global-to-local memory pointer networks for task-oriented dialogue](#). *arXiv preprint arXiv:1901.04713*.
- Puyang Xu and Qi Hu. 2018. [An end-to-end approach for handling unknown slot values in dialogue state tracking](#). In *Proceedings of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1448–1457, Melbourne, Australia. Association for Computational Linguistics.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020a. [MultiWOZ 2.2 : A dialogue dataset with additional annotation corrections and state tracking baselines](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 109–117, Online. Association for Computational Linguistics.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020b. [Multiwoz 2.2: A dialogue dataset with additional annotation corrections and state tracking baselines](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI, ACL 2020*, pages 109–117.
- Jianguo Zhang, Kazuma Hashimoto, Chien-Sheng Wu, Yao Wang, Philip Yu, Richard Socher, and Caiming Xiong. 2020. [Find or classify? dual strategy for slot-value predictions on multi-domain dialog state tracking](#). In *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, pages 154–167, Barcelona, Spain (Online). Association for Computational Linguistics.