

Overview of the Ninth Dialog System Technology Challenge: DSTC9

Chulaka Gunasekara, Seokhwan Kim, Luis Fernando D’Haro, Abhinav Rastogi, Yun-Nung Chen, Mihail Eric, Behnam Hedayatnia, Karthik Gopalakrishnan, Yang Liu, Chao-Wei Huang, Dilek Hakkani-Tür, Jinchao Li, Qi Zhu, Lingxiao Luo, Lars Liden, Kaili Huang, Shahin Shayandeh, Runze Liang, Baolin Peng, Zheng Zhang, Swadheen Shukla, Minlie Huang, Jianfeng Gao, Shikib Mehri, Yulan Feng, Carla Gordon, Seyed Hossein Alavi, David Traum, Maxine Eskenazi, Ahmad Beirami, Eunjoon (EJ) Cho, Paul A. Crook, Ankita De, Alborz Geramifard, Satwik Kottur, Seungwhan Moon, Shivani Poddar, Rajen Subba

Abstract—This paper introduces the Ninth Dialog System Technology Challenge (DSTC-9). This edition of the DSTC focuses on applying end-to-end dialog technologies for four distinct tasks in dialog systems, namely, 1. Task-oriented dialog Modeling with unstructured knowledge access, 2. Multi-domain task-oriented dialog, 3. Interactive evaluation of dialog, and 4. Situated interactive multi-modal dialog. This paper describes the task definition, provided datasets, baselines and evaluation set-up for each track. We also summarize the results of the submitted systems to highlight the overall trends of the state-of-the-art technologies for the tasks.

I. INTRODUCTION

The Dialog System Technology Challenge (DSTC) is a one of the leading series of research competitions in the space of dialog systems. Since the inception in 2013, DSTC has been accelerating the development of dialog technologies, by bringing the leading researchers and engineers together to solve important problems in dialog systems. The challenge has been evolving every year to cater the demand and the interest of the dialog community to foster the development of technology.

The first Dialog System Technology Challenge [1] used human-to-bot dialogs in the bus timetable domain. Dialog State Tracking Challenges 2 [2] and 3 [3] used restaurant reservation application which introduced more complicated and dynamic dialog states. Dialog State Tracking Challenge 4 [4] and Dialog State Tracking Challenge 5 [5] moved to tracking human-to-human dialogs in mono and cross-language settings. From the sixth challenge [6], the DSTC has rebranded itself as “Dialog System Technology Challenge” and organized multiple tracks in parallel to address a wider variety of dialog related problems. The tracks in DSTC-6 were focused on end-to-end conversation modeling and dialog breakdown detection. DSTC-7 [7] focused on developing end-to-end dialog technologies for noetic response selection [8], [9], grounded response generation [10], and audio visual scene aware dialog [11]. More recently in DSTC-8 [12] the focus has been on diverse set of four tracks including, multi-domain task completion, predicting responses, audio visual scene-aware dialog and schema-guided dialog state tracking.

For the ninth edition of the DSTC, we received nine track proposals from the leading research organizations and top universities. The proposals went through a formal peer review process focusing on each task’s potential for, (a) impact to the community, (b) novelty of the task, (c) feasibility of the proposal, and (d) potential participants. The DSTC-8 participants were also asked to provide their feedback on the presented track proposals through a survey, and the responses were also considered in the evaluation. Finally, we ended up with the four main tracks including three newly introduced tasks and one follow-up task from DSTC-8.

The track, *Beyond Domain APIs: Task-oriented Conversational Modeling with Unstructured Knowledge Access* (Track 1), aims to support frictionless task-oriented scenarios, where the flow of the conversation does not break when users have requests that are out of the scope of APIs/DB but potentially are already available in external knowledge sources. Track 2, *Multi-domain Task-oriented Dialog Challenge II*, is a continuation of last year, and focuses on end-to-end multi-domain task completion dialog and cross-lingual multi-domain dialog state tracking. The track 3 of this year, *Interactive Evaluation of Dialog*, aims to take the first step in expanding dialog research beyond datasets and challenges the participants to develop dialog systems that can converse effectively in interactive environments with real users. *SIMMC: Situated Interactive Multi-Modal Conversational AI* (track 4) is aimed at laying the foundations for the real-world assistant agents that can handle multi-modal inputs, and perform multi-modal actions.

The following sections describe the details of each track.

II. TRACK 1 - BEYOND DOMAIN APIS - TASK-ORIENTED CONVERSATIONAL MODELING WITH UNSTRUCTURED KNOWLEDGE ACCESS

A. Track Overview

Most prior work on task-oriented dialog systems has been restricted to a limited coverage of domain APIs. However, users often have domain related requests that are not covered by the APIs. This challenge track aims to expand the coverage of task-oriented dialog systems by incorporating external unstructured knowledge sources. There are three main

TABLE I: Summary of Track 1 tasks

Task #1	Knowledge-seeking Turn Detection
Goal	To decide whether to continue existing flow or trigger the knowledge access branch for a given utterance and dialog history
Input	Current user utterance, dialog context, and domain API and knowledge sources
Output	Binary class (requires knowledge access or not)
Task #2	Knowledge Selection
Goal	To select proper knowledge sources from the domain knowledge-base given dialog context at each turn with knowledge access
Input	Current user utterance, dialog context, and the entire set of knowledge candidates
Output	Ranking of top- k knowledge candidates
Task #3	Knowledge-grounded Response Generation
Goal	To generate a system response for a given triple of input utterance, dialog context, and the selected knowledge sources
Input	Current user utterance, dialog context, and selected knowledge sources
Output	Generated system response

TABLE II: Statistics of the Track 1 data sets

Source	Split	# dialogs	total # instances	# knowledge seeking turns
MultiWOZ	Train	7,190	71,348	19,184
	Valid	1,000	9,663	2,673
	Test	977	2,084	977
SF	Written	900	1,834	900
	Spoken	107	263	104

tasks in this track as introduced in [13]: knowledge-seeking turn detection, knowledge selection, and knowledge-grounded response generation (Table I).

B. Data

This challenge track uses two different data sets (Table II). The first data is an augmented version of MultiWOZ 2.1 [14] that includes newly introduced knowledge-seeking turns in the MultiWOZ conversations. The data augmentation was incrementally done by the crowdsourcing tasks described in [13]. A total of 22,834 utterance pairs were newly collected based on 2,900 knowledge candidates from the FAQ webpages about the domains and the entities in MultiWOZ databases. For the challenge track, we divided the whole data into three subsets: train, validation and test. The first two sets were released in the development phase along with the ground-truth annotations and human responses for participants to develop their models.

In the evaluation phase, we released the test split of the augmented MultiWOZ 2.1 and the other conversations collected from scratch about touristic information for San Francisco. To evaluate the generalizability of models, the new conversations cover knowledge, locale and domains that are unseen from the train and validation data sets. In addition, this test set includes not only written conversations, but also spoken dialogs to evaluate system performance across different modalities [15]. All the backend resources for this data collection were also released, which includes 9,139 knowledge snippets and 855 database entries for San Francisco.

TABLE III: Objective evaluation metrics for the Track 1 tasks

Task	Metrics
Task #1	Precision/Recall/F-measure
Task #2	MRR@5, Recall@1, Recall@5
Task #3	BLEU-1, BLEU-2, BLEU-3, BLEU-4, METEOR ROUGE-1, ROUGE-2, ROUGE-L

C. Evaluation Criteria

Each participating team submitted up to five system outputs each of which contains the results for all three tasks on the unlabeled test instances. We first evaluated each submission using the task-specific objective metrics (Table III) by comparing to the ground-truth labels and responses. Considering the dependencies between the tasks in the pipelined architecture, the final scores for knowledge selection and knowledge-grounded response generation are computed by considering the first step knowledge-seeking turn detection recall and precision performance, as follows:

$$f(x) = \begin{cases} 1 & \text{if } x \text{ is a knowledge-seeking turn,} \\ 0 & \text{otherwise} \end{cases}$$

$$\tilde{f}(x) = \begin{cases} 1 & \text{if } x \text{ is predicted as a knowledge-seeking turn,} \\ 0 & \text{otherwise} \end{cases}$$

$$S_p(X) = \frac{\sum_{x_i \in X} (s(x_i) \cdot f(x_i) \cdot \tilde{f}(x_i))}{\sum_{x_i \in X} \tilde{f}(x_i)},$$

$$S_r(X) = \frac{\sum_{x_i \in X} (s(x_i) \cdot f(x_i) \cdot \tilde{f}(x_i))}{\sum_{x_i \in X} f(x_i)},$$

$$S_f(X) = \frac{2 \cdot S_p(X) \cdot S_r(X)}{S_p(X) + S_r(X)}, \quad (1)$$

where $s(x)$ is the knowledge selection or response generation score in a target metric for a single instance $x \in X$.

Then, we aggregated a set of multiple scores across different tasks and metrics into a single overall score computed by the mean reciprocal rank, as follows:

$$S_{overall}(e) = \frac{1}{|M|} \sum_{i=1}^{|M|} \frac{1}{rank_i(e)}, \quad (2)$$

where $rank_i(e)$ is the ranking of the submitted entry e in the i -th metric against all the other submissions and M is the number of metrics we considered.

Based on the overall objective score, we selected the finalists to be manually evaluated by the following two crowd sourcing tasks:

- **Appropriateness:** This task asks crowd workers to score how well a system output is naturally connected to a given conversation on a scale of 1-5.
- **Accuracy:** This task asks crowd workers to score the accuracy of a system output based on the provided reference knowledge on a scale of 1-5.

In both tasks, we assigned each instance to three crowd workers and took their average as the final human evaluation score for the instance. Those scores were then aggregated

TABLE IV: Objective evaluation results of the best entry from each team in the featured metrics for the Track 1 tasks. Team 0 is the baseline. Bold denotes the best result in each column and * indicates the finalists.

Team	Task#1		Task#2		Task#3	
	F1	R@1	BLEU-1	METEOR	ROUGE-L	
0	0.9455	0.6201	0.3031	0.2983	0.3039	
1	0.9721	0.8255	0.3368	0.3342	0.3364	
2	0.9644	0.8584	0.3338	0.3322	0.3330	
3*	0.9911	0.9013	0.3879	0.3914	0.3885	
4	0.8998	0.6950	0.3025	0.3001	0.2990	
5	0.9428	0.7055	0.3218	0.3266	0.3174	
6	0.9838	0.8531	0.3371	0.3341	0.3362	
7*	0.9702	0.8988	0.3752	0.3854	0.3702	
8	0.9530	0.7403	0.3135	0.3097	0.3119	
9	0.9242	0.7931	0.3154	0.3159	0.3101	
10*	0.9726	0.9158	0.3684	0.3719	0.3692	
11*	0.9675	0.8702	0.3743	0.3854	0.3797	
12	0.9540	0.8011	0.3411	0.3526	0.3401	
13*	0.9819	0.8434	0.3787	0.3902	0.3619	
14	0.9252	0.6466	0.3019	0.2974	0.3003	
15*	0.9803	0.8975	0.3779	0.3931	0.3765	
16	0.9549	0.7327	0.3351	0.3334	0.3364	
17*	0.9839	0.8734	0.3699	0.3724	0.3687	
18*	0.9635	0.8994	0.3806	0.3864	0.3726	
19*	0.9886	0.9235	0.3803	0.3869	0.3738	
20*	0.9711	0.8628	0.3619	0.3637	0.3535	
21*	0.9396	0.8269	0.3551	0.3594	0.3576	
22	0.7803	0.2831	0.1792	0.1583	0.1852	
23*	0.9618	0.8959	0.3534	0.3565	0.3519	
24	0.3471	0.0017	0.0835	0.0796	0.0866	

over the entire test set following Equation 1, i.e., weighted by the knowledge-seeking turn detection performance. Finally, we used the average of the Appropriateness and Accuracy scores to determine the official ranking of the systems in the challenge track.

D. Results

We received 105 entries in total submitted from 24 participating teams. To preserve anonymity, the teams were identified by numbers from 1 to 24, while our baseline [13] was marked as team 0. Table IV shows the objective evaluation results of the best entry from each team in the featured metrics. The full scores with all the submitted entries and the other metrics are available on the track repository¹. Most entries outperformed the baseline in all three tasks. In particular, the best entry from Team 3 achieved over 99% F-measure for knowledge-seeking turn detection, and also the highest scores in the BLEU and ROUGE variants for the response generation task. On the other hand, Team 19 was the best in the knowledge selection metrics and Team 15 was better than all the other teams in METEOR for generation. We calculated the overall score (Equation 2) of each entry and selected 12 finalists, corresponding to the best entry from each of the top 12 teams.

Table V shows the final ranking of the Track 1 participating teams based on the human evaluation scores of the finalist entries. The top three teams (Team 19, 3 and 10) commonly used ensemble of large-scale pre-trained language models in

TABLE V: Human evaluation results for the Track 1 finalists

Rank	Team	Entry	Accuracy	Appropriateness	Average
		Ground-truth	4.5930	4.4513	4.5221
1	19	2	4.3917	4.3922	4.3920
2	3	1	4.3480	4.3634	4.3557
3	10	0	4.3544	4.3201	4.3373
4	15	3	4.3793	4.2755	4.3274
5	17	0	4.3360	4.3076	4.3218
6	7	4	4.3308	4.2989	4.3149
7	18	3	4.3309	4.2859	4.3084
8	13	3	4.3763	4.2360	4.3061
9	23	0	4.3082	4.2665	4.2874
10	11	3	4.2722	4.2619	4.2670
11	20	4	4.2283	4.2486	4.2384
12	21	3	4.1060	4.1560	4.1310
		Baseline	3.7155	3.9386	3.8271

their best entries. Team 19 won the challenge track with the highest scores for both Accuracy and Appropriateness, most likely because of their better performance in the knowledge selection task (as in the objective evaluation results). To compare the importance of each task towards end-to-end performance, we calculated the Spearman’s rank correlation coefficient of the ranked lists of all the entries in every pair of objective and human evaluation metrics. As a result, Recall@1 for the knowledge selection task shows a strong correlation with the averaged human evaluation ranking at 0.8601, which is significantly higher than 0.7692 and 0.6503 with F-measure for the knowledge-seeking turn detection and BLEU-1 for the response generation, respectively. This implies that the knowledge-selection is a key task to improve end-to-end performance.

III. TRACK 2 - MULTI-DOMAIN TASK-ORIENTED DIALOG CHALLENGE II

We provide two tasks in the multi-domain task-oriented dialog setting. One is the end-to-end task-oriented dialog task aiming to solve the complexity of building end-to-end dialog systems. The other is cross-lingual dialog state tracking (DST) to address the language adaption problem for the DST task.

A. End-to-end Task-oriented Dialog Task

This task is a continuation of last year at DSTC8 [16]. Participants will develop an end-to-end task-oriented dialog system that takes natural language as input and generates natural language response as output in the travel planning setting. Both the evaluation result of last year’s challenge [17] and empirical analysis of models in ConvLab [18] show that the best rule-based pipeline systems outperform systems assembled using state-of-the-art component-wise machine learning models. From the results, we’ve also observed a discrepancy between the performance of component-wise models using corpus-based evaluation and that of the entire system using the end-to-end evaluation. These findings are consistent with the landscape of dialog development technology stacks in the industry. However, interestingly, the winning team built their model based on GPT-2 [19], and achieved significant improvement over other teams with regards to success rate, understanding score, and response score at the human evaluation

¹<https://github.com/alexa/alexa-with-dstc9-track1-dataset>

phase. Meanwhile, by using similar model training paradigms, SOLOIST [20] and SimpleTOD [21] shortly achieved top performance in the MultiWOZ leaderboard by leveraging GPT-2 [22].

This year, we continue with the end-to-end task-oriented dialog task, aiming to promote the technology of building end-to-end dialog systems one step further. Like last year, participants are encouraged to explore all possible approaches, and there is no restriction on dialog system architecture.

1) *Data*: Participants are expected to build dialog systems based on MultiWOZ 2.1 [14], a multi-domain dialog dataset spanning 7 distinct domains containing over 10,000 dialogs under the travel planning setting. Compared with MultiWOZ 2.0 [23], MultiWOZ 2.1 re-annotated states to fix the noisy annotation and incorporated user dialog act annotation. Although the dialog system is evaluated under MultiWOZ 2.1, participants can leverage any public datasets, pre-trained models, or other resources to build the dialog system.

2) *Evaluation Criteria*: ConvLab-2 [24] is employed as the platform for dialog development and evaluation. As the successor of ConvLab [25], ConvLab-2 provides a user simulator and evaluator for MultiWOZ 2.1 so that the participants can effectively run offline experiments and evaluations. Specifically, we offer two evaluation approaches:

a) *Automatic Evaluation*: The dialog system is evaluated via conversing with an end-to-end user simulator. The simulator is constructed by assembling a BERT-based natural language understanding model [26], an agenda-based user simulator [27] and a rule-based natural language generation module. A dialog is successful only if all requested slots are filled with grounded values in the database, and the booking is successful. We report metrics including success rate, book rate, number of turns for each dialog, and precision/recall/F1 score for slot prediction.

b) *Human Evaluation*: In human evaluation, Amazon Mechanical Turkers communicate with the dialog systems via natural language, judge whether the dialog is successful, and provide scores based on language understanding correctness and response appropriateness on a 5 point Likert-scale. Since MTurkers do not directly access the back-end database, we also report the success rate with grounding after verifying whether the requested slot values returned by the dialog systems match the database record. We take the average value of success rate with grounding and without grounding as the final ranking.

3) *Results*: As per our submission policy, each team is allowed to submit up to 5 models. We received 34 models in total from 10 teams. Table VI lists the automatic evaluation result for the best models of each team. We filtered out low-performance models based on the automatic evaluation result while keeping the best model for each team and sent the remaining models for human evaluation. With this process, 21 models were evaluated in human evaluation, with the performance of each team’s best model listed in Table VII.

Team 1 achieves the top 1 performance in both automatic and human evaluation by constructing an end-to-end dialog system with the pre-trained dialog generation model PLATO-2 [28]. This model generates the dialog state, system action, and

TABLE VI: Automatic Evaluation Result (Best Submissions)

Team	SR	CR	BR	Inform P/R/F1	Turn S/A
1	93	95.2	94.6	84.1/96.2/88.1	12.5/12.7
2	91.4	96.9	96.2	80.2/97.3/86.0	15.3/15.7
3	90.8	94.4	96.7	81.0/95.4/85.9	13.4/13.6
4	89.8	94.6	96.3	72.4/96.0/80.1	15.1/15.8
5	83.3	88.5	89.1	81.1/90.3/83.5	13.5/13.8
6	67.7	88.5	90.8	70.4/85.6/75.2	12.8/14.2
7	57.8	87.1	85	68.7/81.6/72.6	13.7/16.4
8	52.6	66.9	66.7	57.5/80.7/64.8	13.2/22.5
9	44.4	50	26.5	57.9/64.5/58.9	12.2/14.6
10	21.4	40.7	0	55.4/60.0/54.1	11.0/25.9
Baseline	85	92.4	91.4	79.3/94.9/84.5	13.8/14.9

SR: Success Rate, CR: Complete Rate, BR: Book Rate, Inform P/R/F1: Prec./Recall/F1 score of slots prediction, Turn S/A: Turns for successful and all dialogs, respectively.

TABLE VII: Human Evaluation Result (Best Submissions)

Team	SRa	SRwg	SRog	Under.	Appr.	Turn	Rank
1	74.8	70.2	79.4	4.54	4.47	18.5	1
2	74.8	68.8	80.8	4.51	4.45	19.4	1
7	72.3	62	82.6	4.53	4.41	17.1	3
6	70.6	60.8	80.4	4.41	4.41	20.1	4
3	67.8	60	75.6	4.56	4.42	21	5
4	60.3	51.4	69.2	4.49	4.22	17.7	6
5	58.4	50.4	66.4	4.15	4.06	19.7	7
9	55.2	43.2	67.2	4.15	3.98	19.2	8
8	35	26	44	3.27	3.15	18.5	9
10	19.5	6	33	3.23	2.93	18.8	10
Baseline	69.6	56.8	82.4	4.34	4.18	18.5	N/A

SRa: average success rate, SRwg: success rate w/ grounding, SRog: success rate w/o grounding, Under.: understanding score, Appr.: appropriateness score.

system response simultaneously, given the dialog context. The dialog state is used as the constraint for database query, and the system action is then refreshed according to the queried results to re-generate the final system response. Team 2 achieves the same ranking as Team 1 in the human evaluation using a similar hybrid end-to-end neural model. It borrows idea from [20] and [29], uses GPT-2 as the backend for pre-training and fine-tuning and add various pre/post-processing modules to improve model generalization ability. An additional fault tolerance mechanism is also added to correct errors.

4) *Summary*: Compared with the challenge results at DSTC8, there is a trend of shifting from building dialogs by assembling component-wise modules to end-to-end learning. In DSTC8, out of 11 teams with valid submissions, 1 team uses GPT-2 based models, 1 team uses word DST + word policy, with the rest 9 team uses component-wise models. This year, out of 10 teams, 8 teams used the end-to-end learning mechanism by leveraging transformer-based models. The top three systems in both automatic evaluation and human evaluation are all built using transformer-based end-to-end learning, and they have achieved much better performance in human evaluation than the systems at DSTC8².

B. Cross-lingual Dialog State Tracking Task

We introduce the task of cross-lingual dialog state tracking, requiring the participants to build a dialog state tracker for

²The success rate in DSTC8 human evaluation is success rate w/o grounding.

the target language with a training set in the source resource language and a small development set in the target language. Based on newly proposed large scale multi-domain task-oriented dialog datasets, MultiWOZ 2.1 [14] and CrossWOZ [30], we offer two sub-tasks: 1) cross-lingual transfer from English to Chinese using MultiWOZ 2.1 dataset and 2) cross-lingual transfer from Chinese to English using CrossWOZ dataset.

Following a similar scheme as in DSTC-5 [5], we provided machine translations of the original dataset. We collected 500 new dialogs in the target language as the test set. The performance of each dialog state tracker is evaluated on the test set and compared with reference annotation.

1) *Data*: Compared with previous datasets [5], [31], [32] for cross-lingual transfer learning in task-oriented dialog, MultiWOZ 2.1 and CrossWOZ are much larger. MultiWOZ 2.1 contains over 10,000 dialogs, and CrossWOZ contains over 6,000 dialogs. They are also more challenging due to the multi-domain setting. For each sub-task, we prepared data in a similar way: a) collected 500 new dialogs in the source language, b) translated the ontology to the target language, and c) translated the original dialogs and the new dialogs. We released 250 new dialogs without any annotation as a public test set and reserved the other 250 dialogs as a private test set.

a) *Test Data Collection*: To collect new dialogs, we adapted the data collection website of CrossWOZ where paired workers can converse synchronously and make annotations. New user goals were generated by the goal generator from ConvLab-2. Following the Wizard-of-Oz setting, one worker acts as the user who needs to accomplish the allocated goal, and the other acts as the system that uses the database to provide information. During the conversation, both sides need to annotate the dialog acts of their utterances, and the system should also annotate the dialog states that are queries over the database.

b) *Ontology Translation*: We extracted the ontology from dialog act and dialog state annotations of both the original and test datasets. Then we used Google Translate to translate them to the target language. For some slots that may not be faithfully translated, such as “name” and “address”, we employed human translators to correct the translations. This process is vital to ensure the translation consistency of the same values in different contexts.

c) *Dialog Translation*: To make sure that the translations of values in a dialog are faithful to the ontology dictionary, we first replaced the values that appeared in the dialog with their translations in the dictionary. Then we used Google Translate to translate the resulting code-switching sentences from the original dataset and test set. In this way, translated dialogs and corresponding annotations do not conflict. 250 dialogs were sampled from the original dataset as the development set. Human translators were employed to proofread the translations of the development and test set.

2) *Evaluation Criteria*: We evaluate the performance of the dialog state tracker using the following metrics: a) Joint Goal Accuracy. This metric evaluates whether the predicted dialog state is exactly equal to the ground truth. b) Slot Accuracy. This metric evaluates whether each slot’s predicted

label is exactly equal to the ground truth, averaged over all slots. c) Slot Precision/Recall/F1. These metrics evaluate the overlap between the predicted labels and the ground truth for non-empty slots, micro-averaged over dialog turns. Each submission contains the predictions for the public test set and the model that is used to make predictions for the private test set. The results are averaged over the public and private test set. The final ranking is solely based on the joint goal accuracy.

3) *Results*: The results of MultiWOZ (en→zh) and CrossWOZ (zh→en) sub-tasks are shown in Table VIII and IX respectively. During the evaluation, we found that the CrossWOZ test data miss many “name” labels when the user accepts the attraction/hotel/restaurant recommended by the system. Therefore, we utilized the database search results and heuristic rules to correct empty “name” labels and provided an updated leaderboard for CrossWOZ in Table X. Both of the CrossWOZ leaderboards are valid, but the updated one is preferred.

TABLE VIII: MultiWOZ Leaderboard (Best Submissions).

Team	JGA	SA	Slot P/R/F1	JGA(pub/pri)	Rank
1	62.37	98.09	92.15/94.02/93.07	62.70/62.03	1
2	62.08	98.10	90.61/96.20/93.32	63.25/60.91	2
3	30.13	94.40	87.07/74.67/80.40	30.53/29.72	3
BS	55.56	97.68	92.02/91.10/91.56	55.81/55.31	N/A

JGA: joint goal accuracy, SA: slot accuracy, Slot P/R/F1: slot precision/recall/F1, pub/pri: public/private test set.

TABLE IX: CrossWOZ Leaderboard (Best Submissions).

Team	JGA	SA	Slot P/R/F1	JGA(pub/pri)	Rank
3	16.86	89.11	68.26/62.85/65.45	16.82/16.89	1
1	15.28	90.37	65.94/78.87/71.82	15.19/15.37	2
2	13.99	91.92	72.63/78.90/75.64	14.41/13.58	3
BS	7.21	85.13	55.27/46.15/50.30	7.41/7.00	N/A

TABLE X: CrossWOZ Leaderboard (Updated Evaluation, Best Submissions).

Team	JGA	SA	Slot P/R/F1	JGA(pub/pri)	Rank
2	32.30	94.35	81.39/82.25/81.82	32.70/31.89	1
1	23.96	92.94	74.96/83.41/78.96	23.45/24.47	2
3	15.31	89.70	74.78/64.06/69.01	14.25/16.37	3
BS	13.02	87.97	67.18/52.18/58.74	13.30/12.74	N/A

We adapted SUMBT [33] as the baseline model and used the translated training set of the original dataset to train for both sub-tasks. We have received 10 models for MultiWOZ (en→zh) and 8 models for CrossWOZ (zh→en) from the same 3 teams. We briefly introduce their best models here. Team 1 incorporated a four-class state operation prediction task into CHAN model [34]. Team 2 modified SOM-DST [35] and used ontology and some handcraft rules to post-process the generated values. Team 3 formulated the dialog state tracking as a sequence generation problem and used mBART to generate pairs of slot names and slot values. All of their best models were trained using the translated data in the target language.

4) *Summary*: To our surprise, all the best models are trained on monolingual machine translated data instead of

both the original data and translations. Team 2 and 3 even got negative results when training XLM/mBART on the original data and the translations simultaneously. The performance of “Translate-Train” partially depends on the machine translator, which may be why team 1 and 2 augment the data by using another translator to translate the original dataset. Team 1 and 2 modified DST models that are state-of-the-art on English MultiWOZ 2.1 dataset and got strong performance on Chinese MultiWOZ 2.1, verifying these models’ language portability.

IV. TRACK 3 - INTERACTIVE EVALUATION OF DIALOG

A. Track overview

The aim of dialog research is to create systems that can be effectively used in interactive settings by real users [36]. Despite this, the majority of research is performed on static datasets. For example, the task of response generation is typically done by producing a response for a static dialog context [37]. This track is intended to move dialog research beyond datasets and evaluate models in interactive environments with real users.

This track consists of two sub-tasks: (1) static evaluation and (2) interactive evaluation. The first subtask challenges participants to build response generation models which are evaluated in a static manner, using the Topical-Chat corpus [38]. The second subtask aims to extend dialog models beyond datasets and assess them in an interactive setting with real users, using DialPort [39]. In the first subtask, models must generate a response to a fixed dialog context. In contrast in the second subtask, they must have a back-and-forth interaction with a real user. Through the two subtasks, this track challenges participants to take strong response generation models and develop strategies of making them effective in interactive settings.

B. Data

Participants in this track were free to train on any publicly available data or use any pre-trained models. The static evaluation in the first subtask was carried out on the Topical-Chat corpus [38]. Topical-Chat is a large collection of human-human knowledge-grounded open-domain conversations that consists of 11,319 dialogs and 248,014 utterances. For each conversational turn, several relevant facts are provided. Models must leverage these facts and generate a response. This dataset was chosen because it is the largest, knowledge-grounded open-domain dataset presently available, to our knowledge. Additionally, the choice of usable facts provides a mechanism for systems to tailor responses to a specific user’s interests.

Since we continuously performed human evaluation over the duration of the challenge and used reference free evaluation metrics [40], it was not strictly necessary for models to be trained on the Topical-Chat corpus. A strong pre-trained dialog model may perform well on the first subtask, despite not training on the corpus.

The second subtask was not tied to a dataset. The interactive evaluation was carried out on DialPort³ [39] with real users recruited through Facebook Advertising.

³<http://dialog.speech.cs.cmu.edu:3000/>

C. Evaluation Criteria

The first subtask was evaluated using ongoing (1) human evaluation and (2) three automatic metrics: METEOR [41], BERTscore [42] and USR [40]. Human evaluation was carried out on Amazon Mechanical Turk with the annotation questionnaire used to collect the FED dataset [43]. Over the duration of the challenge, we carried out evaluation on the Topical-Chat frequent validation set. For human evaluation, 30 context-response pairs were sampled and each one was labeled by 3 annotators. For the final evaluation, we carried out automatic evaluation on the frequent test set and perform human evaluation on 100 randomly sampled context-response pairs. For the final evaluation, the 100 dialog contexts used for evaluation were consistent across the different systems.

The evaluation for the second subtask consists of (1) collecting dialogs through conversations with real users on DialPort and (2) post-hoc assessment of the collected dialogs. Participants submitted dialog models (via an API) to DialPort. Real users were recruited through Facebook Advertising to interact with the submitted dialog systems. After gathering a sufficient number of conversations, we performed post-hoc assessment of the dialogs with the FED metric [43] and human evaluation on Amazon Mechanical Turk with the annotation questionnaire used to collect the FED dataset [43].

Throughout the challenge, we aimed to collect at least 100 conversations for each submitted system discounting any dialogs with offensive terms (e.g., curse words, racist phrases). For each system, 100 conversations were evaluated with the FED metric and on Amazon Mechanical Turk, with 3 annotators labeling each dialog.

For the final submission, we gather dialogs for all systems over the same time period. Ultimately, given a Facebook Advertising budget of \$2500 and 11 systems (including two baselines), we obtained 4651 conversations (after removing offensive dialogs) with a total of 41,640 turns. We consider only the conversations that are at least four turns in length (total of 2960) for the final post-hoc assessment. For each system, we carry out human evaluation with 200 conversations of suitable length. Throughout the challenge, all individuals who interact with the system on DialPort *do so for free, of their own volition*, thereby avoiding common problems observed with paid users [44].

D. Results

The challenge received **33** submissions to the first subtask and **9** submissions to the second subtask.

Table XI shows the results of the static evaluation on the Topical-Chat corpus [38], for the 10 best performing systems according to the human evaluation. All of the top 10 systems used either pre-trained models or additional data, highlighting the importance of pre-training for open-domain response generation. This observation aligns with previous research, which has seen strong performance in open-domain response generation through the use of large-scale pre-training [45], [46].

In addition to performing ongoing human evaluation throughout the challenge, we assess systems in the first subtask

TABLE XI: Results for subtask 1. For brevity, we only show the top 10 submissions (out of 33) according to the human evaluation. This table only reports the overall USR metric and the overall impression of the response from the human evaluation. The full evaluation results may be found here.

System	METEOR	BERTscore	USR	Human	Rank
1	9.06	84.91	4.26	4.281	1
2	13.11	86.17	4.59	4.280	1
3	6.83	84.36	3.86	4.280	1
4	8.96	85.15	4.26	4.260	4
5	12.37	86.21	4.83	4.253	5
6	12.31	86.32	4.73	4.231	6
7	13.96	86.84	4.48	4.229	6
8	12.51	85.91	4.45	4.229	6
9	12.14	85.91	4.46	4.216	9
10	10.87	85.65	4.53	4.210	10

TABLE XII: Results for subtask 2. This table reports for each system: the overall FED metric, the overall impression of the dialogs from the human evaluation, as well as the average number of dialog turns. The full results be found here. System 6 and 11 are our DialoGPT and Transformer baselines, respectively.

System	Avg. Turns	FED	Human	Rank
1	12.44	4.97	4.15	1
2	13.47	4.79	4.14	2
3	8.89	4.61	4.08	3
4	9.36	4.68	4.03	4
5	9.82	4.53	3.93	5
6	8.75	4.72	3.87	6
7	8.51	4.41	3.85	7
8	7.67	4.30	3.85	7
9	6.53	4.64	3.83	9
10	7.35	4.80	3.69	10
11	5.80	3.69	3.60	11

using three evaluation metrics. METEOR [41] and BERTscore [42], are referenced evaluation metrics that compare a generated output to a *ground-truth response*. In contrast, USR [40] is a reference free evaluation metric that uses pre-trained models and self-supervised training objectives to estimate the quality of a response. Though none of the evaluation metrics is a perfect predictor of the final ranking, we find that USR better correlates with the system-level human performance (Spearman: 0.35, $p < 0.05$) than either METEOR (Spearman: 0.23, $p > 0.05$) or BERTscore (Spearman: 0.22, $p > 0.05$). The relatively low system-level correlation highlights the importance of performing ongoing human evaluation throughout the challenge.

The poor performance of automatic metrics, may in part be a consequence of the fact that several submissions did not fine-tune on the Topical-Chat corpus and instead relied on open-domain response generation capabilities learned through large-scale pre-training. As such, while the responses were favored by human annotators - the automatic metrics penalized them for either not having high word-overlap with the ground truth (METEOR, BERTscore) or not resembling the utterances in the Topical-Chat corpus (USR).

The results for the second subtask are shown in Table XII. System 6 is our DialoGPT baseline [45], fine-tuned on the

Topical-Chat corpus without knowledge grounding. System 11 is our Transformer baseline which was trained on the Topical-Chat corpus and uses tf-idf sentence similarity to retrieve relevant knowledge at inference time. The best performing model, **System 1**, leverages large-scale pre-training in addition to strategies for producing more diverse responses. This system achieved first place in both subtasks: **System 1** in Table XII corresponds to **System 2** in Table XI.

FED [43], which is an *unsupervised* evaluation metric for interactive dialog is shown to be a moderate predictor of the final ranking with a system-level Spearman correlation of 0.49 ($p = 0.13$), though it correctly predicts the top two systems. We also note that the average number of turns for a particular system is a strong indicator of its quality here (Spearman: 0.94, $p < 0.01$). Real users are more inclined to interact with a better system, making it an important metric for assessing systems in interactive settings [47].

While many of the submissions in the first subtask perform similarly, the scores in Table XII are much more varied. This signifies that interactive evaluation more exhaustively tests the capabilities of systems and is therefore a more indicative measure of a system’s capabilities. This observation has been shown by prior work [43], when analyzing dialogs from Meena [46].

The *Interactive Evaluation of Dialog* track demonstrates both the feasibility and the importance of evaluating dialog systems in interactive settings with real users. We show that with an advertising budget of \$2500, we collect more than 4000 dialogs on DialPort (2960 dialogs with at least 4 turns). The results of interactive evaluation are more varied (Table XII) suggesting that back-and-forth interactions with real users are challenging to dialog systems and that interactive evaluation is a better reflection of a system’s capabilities.

V. TRACK 4 - SIMMC: SITUATED INTERACTIVE MULTI-MODAL CONVERSATIONAL AI

A. Track overview

The SIMMC challenge aims to lay the foundations for the real-world assistant agents that can handle multimodal inputs, and perform multimodal actions. We thus focus on task-oriented dialogs that encompass a situated multimodal user context in the form of a co-observed image or virtual reality (VR) environment. The context is dynamically updated on each turn based on the user input and the assistant action. Moon *et al.* [48] provides more details on the datasets and the models we provide.

B. Data

SIMMC contains about 13k human-to-human dialogs (totaling about 169k utterances). We chose shopping experiences—specifically furniture and fashion—as the domain for the SIMMC datasets because of the dynamic environment created by these domains, where rich multimodal interactions happen around visually grounded items.

SIMMC offers four key advantages over previous multimodal dialog datasets:

TABLE XIII: Summary of each team’s results on Test-Std split, average of Furniture and Fashion (*Team 5 submitted results only for Fashion). Best results from each team are shown. (1) **API prediction** via accuracy, perplexity and attribute accuracy, and, (2) **Response prediction** via BLEU, recall@k (k=1,5,10), mean rank, and mean reciprocal rank (MRR). (3) **Dialog State Tracking (DST)**, via slot and intent prediction F1. \uparrow : higher is better, \downarrow : lower is better.

Teams	Subtask 1. API Prediction			Subtask 2. Response Generation						Subtask 3. DST	
	Acc \uparrow	Perp \downarrow	A.Acc \uparrow	BLEU \uparrow	r@1 \uparrow	r@5 \uparrow	r@10 \uparrow	Mean \downarrow	MRR \uparrow	Slot F1 \uparrow	Intent F1 \uparrow
Baseline	79.3	63.7	1.9	0.0061	0.145	7.2	19.8	27.3	39.2	62.4	62.1
Team 1	80.2	74.6	2.0	0.105	0.326	21.1	43.6	56.8	18.8	77.8	76.7
Team 2	82.5	69.8	1.8	0.082	0.074	2.5	8.3	13.6	47.7	-	-
Team 3	79.4	73.2	-	0.128	0.381	26.3	50.3	61.8	15.5	79.1	78.1
Team 4	81.3	73.9	3.5	0.108	0.673	52.6	87.4	95.1	3.2	78.6	77.7
Team 5*	-	-	-	-	0.39	26.7	52.1	66.0	14.8	-	-

- 1) SIMMC assumes a co-observed multimodal context between a user and an assistant and records the ground-truth item appearance logs of each item that appears. SIMMC tasks emphasize semantic processing of the input modalities, while work in this area has traditionally focused heavily on raw image processing.
- 2) Compared with the conventional task-oriented conversational datasets, the agent actions in the SIMMC datasets span across a diverse multimodal action space (e.g. “rotate”, “search”, and “add to cart”).
- 3) Agent actions can be enacted on both the object level (e.g. changing the view of a specific object within a scene) and the scene level (e.g. introducing a new scene or an image).
- 4) SIMMC emphasizes semantic processing. The proposed SIMMC annotation schema allows for a more systematic and structural approach for visual grounding of conversations, which is essential for solving challenging problems in real-world scenarios.

Datasets were collected through the SIMMC Platform [49], an extension to ParlAI [50] for multimodal conversational data collection and system evaluation that allows human annotators to each play the role of either the assistant or the user.

C. Evaluation Criteria

We present three subtasks primarily aimed at replicating human-assistant actions in order to enable rich and interactive shopping scenarios.

Subtask 1: Structural API Call Prediction focuses on predicting the assistant action as an API call given the dialog and the multimodal contexts as inputs. Since accuracy does not account for the existence of multiple valid actions, we use perplexity (defined as the exponentiation of the Shannon entropy) alongside accuracy. To also measure the correctness of the predicted action (API) arguments, we use attribute accuracy compared to the collected datasets.

Subtask 2: Response Prediction examines the relevance of the assistant response in the current turn. We evaluate in two ways; (a) as a conditional language modeling problem, where the closeness between the generated and ground-truth response is measured through using BLEU-4 score, and, (b) as a retrieval problem, where we measure the model performance when retrieving ground-truth responses from a pool of 100 candidates (randomly chosen and unique to each turn).

Subtask 3: Dialog State Tracking (DST) aims to systematically track the dialog acts and the associated slot pairs across multiple turns, as represented in the flexible ontology developed to represent the SIMMC multimodal context. We use the intent and slot prediction metrics (F1), inline with prior work in DST.

D. Results

The challenge saw a total of 11 model entries from 5 teams across the world, setting a new state-of-the-art in all three subtasks (Table XIII).

For each subtask, we listed metrics in a priority order and the entry with the most favorable performance on the highest priority metric was considered to be a candidate winner. Any entries within one standard error of this candidate’s performance were also considered as candidates. Where there were more than one candidate, as in subtask 1, we used the next metric in the priority list and repeat this process until we had a single winner.

The winner of the structural API call prediction subtask (subtask 1) was a BART [51] model (BART-Large) from Team 4 that jointly predicted the dialog state (subtask 3), API call (subtask 1) and response (subtask 2a) as single target given the dialog history, multimodal context and user utterance. This model was one of two runners up on subtask 2a, and the runner up on subtask 3.

The winner of the response retrieval subtask (subtask 2b) was a BART-based Bi-encoder [52], [53], [54], also from Team 4, whose weights were initialized from the jointly trained BART model that won subtask 1. This model achieved a mean reciprocal rank (MRR) of 0.67, a lead of 0.29 points ahead of the runner up team on this subtask.

The winner of the response generation and DST subtasks (subtask 2a and subtask 3) was an ensemble of GPT-2 [22] models from Team 3 that were of differing sizes (large and small) and used differing portions of the training and development sets. Each GPT-2 model was independently trained on the joint tasks—subtask 2a and subtask 3—using a simple language model loss that optimized over the concatenated dialog history, multimodal context, user utterance, dialog state and response. Preprocessing over dialogue states was done before training, and an ensemble beam search over each model’s prediction was used to generate the final prediction.

VI. CONCLUSIONS

This paper summarizes the four tracks of the ninth dialog system technology challenges (DSTC9). Beyond Domain APIs track expands the coverage of current task-oriented dialog systems by incorporating external unstructured knowledge sources. Multi-domain Task-oriented Dialog Challenge II, focuses on end-to-end multi-domain task completion dialog and cross-lingual multi-domain dialog state tracking. Interactive Evaluation of Dialog Track, expands dialog research beyond datasets encourages to develop dialog systems that can converse effectively in interactive environments. The Situated Interactive Multi-Modal Conversational AI track focuses on real-world assistant agents that can handle multi-modal inputs, and perform multi-modal actions. All the datasets and resources introduced for every track will be publicly available even after the challenge period to support future dialog system research.

REFERENCES

- [1] J. Williams, A. Raux, D. Ramachandran, and A. Black, "The dialog state tracking challenge," in *Proceedings of the SIGDIAL 2013 Conference*, 2013, pp. 404–413.
- [2] M. Henderson, B. Thomson, and J. D. Williams, "The second dialog state tracking challenge," in *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, 2014, pp. 263–272.
- [3] —, "The third dialog state tracking challenge," in *Spoken Language Technology Workshop (SLT), 2014 IEEE*. IEEE, 2014, pp. 324–329.
- [4] S. Kim, L. F. D'Haro, R. E. Banchs, J. D. Williams, and M. Henderson, "The fourth dialog state tracking challenge," in *Dialogues with Social Robots*. Springer, 2017, pp. 435–449.
- [5] S. Kim, L. F. D'Haro, R. E. Banchs, J. D. Williams, M. Henderson, and K. Yoshino, "The fifth dialog state tracking challenge," in *Spoken Language Technology Workshop (SLT), 2016 IEEEover*. IEEE, 2016, pp. 511–517.
- [6] C. Hori, J. Perez, R. Higashinaka, T. Hori, Y.-L. Boureau, M. Inaba, Y. Tsunomori, T. Takahashi, K. Yoshino, and S. Kim, "Overview of the sixth dialog system technology challenge: Dstc6," 2018.
- [7] K. Yoshino, C. Hori, J. Perez, L. F. D'Haro, L. Polymenakos, C. Gunasekara, W. S. Lasecki, J. K. Kummerfeld, M. Galley, C. Brockett *et al.*, "Dialog system technology challenge 7," *arXiv preprint arXiv:1901.03461*, 2019.
- [8] C. Gunasekara, J. K. Kummerfeld, L. Polymenakos, , and W. S. Lasecki, "Dstc7 task 1: Noetic end-to-end response selection," in *7th Edition of the Dialog System Technology Challenges at AAI 2019*, January 2019. [Online]. Available: http://workshop.colips.org/dstc7/papers/dstc7_task1_final_report.pdf
- [9] C. Gunasekara, J. K. Kummerfeld, L. Polymenakos, and W. Lasecki, "Dstc7 task 1: Noetic end-to-end response selection," in *Proceedings of the First Workshop on NLP for Conversational AI*, 2019, pp. 60–67.
- [10] M. Galley, C. Brockett, X. Gao, J. Gao, and B. Dolan, "Grounded response generation task at dstc7," in *AAAI Dialog System Technology Challenges Workshop*, 2019.
- [11] H. Alamri, C. Hori, T. K. Marks, D. Batr, and D. Parikh, "Audio visual scene-aware dialog (avsd) track for natural language generation in dstc7," in *DSTC7 at AAI2019 Workshop*, vol. 2, 2018.
- [12] S. Kim, M. Galley, C. Gunasekara, S. Lee, A. Atkinson, B. Peng, H. Schulz, J. Gao, J. Li, M. Adada *et al.*, "The eighth dialog system technology challenge," *arXiv preprint arXiv:1911.06394*, 2019.
- [13] S. Kim, M. Eric, K. Gopalakrishnan, B. Hedayatnia, Y. Liu, and D. Hakkani-Tur, "Beyond domain APIs: Task-oriented conversational modeling with unstructured knowledge access," in *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 1st virtual meeting: Association for Computational Linguistics, Jul. 2020, pp. 278–289. [Online]. Available: <https://www.aclweb.org/anthology/2020.sigdial-1.35>
- [14] M. Eric, R. Goel, S. Paul, A. Sethi, S. Agarwal, S. Gao, and D. Hakkani-Tur, "Multiwoz 2.1: Multi-domain dialogue state corrections and state tracking baselines," *arXiv preprint arXiv:1907.01669*, 2019.
- [15] K. Gopalakrishnan, B. Hedayatnia, L. Wang, Y. Liu, and D. Hakkani-Tur, "Are neural open-domain dialog systems robust to speech recognition errors in the dialog history? an empirical study," *INTERSPEECH*, pp. 911–915, 2020.
- [16] S. Kim, M. Galley, R. C. Gunasekara, S. Lee, A. Atkinson, B. Peng, H. Schulz, J. Gao, J. Li, M. Adada, M. Huang, L. A. Lastras, J. K. Kummerfeld, W. S. Lasecki, C. Hori, A. Cherian, T. K. Marks, A. Rastogi, X. Zang, S. Sunkara, and R. Gupta, "The eighth dialog system technology challenge," *CoRR*, vol. abs/1911.06394, 2019.
- [17] J. Li, B. Peng, S. Lee, J. Gao, R. Takanobu, Q. Zhu, M. Huang, H. Schulz, A. Atkinson, and M. Adada, "Results of the multi-domain task-completion dialog challenge," in *Proceedings of the 34th AAAI Conference on Artificial Intelligence, Eighth Dialog System Technology Challenge Workshop*, 2020.
- [18] R. Takanobu, Q. Zhu, J. Li, B. Peng, J. Gao, and M. Huang, "Is your goal-oriented dialog model performing really well? empirical analysis of system-wise evaluation," in *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, Jul. 2020, pp. 297–310.
- [19] D. Ham, J.-G. Lee, Y. Jang, and K.-E. Kim, "End-to-end neural pipeline for goal-oriented dialogue systems using gpt-2," *ACL*, 2020.
- [20] B. Peng, C. Li, J. Li, S. Shayandeh, L. Liden, and J. Gao, "Soloist: Few-shot task-oriented dialog with a single pre-trained auto-regressive model," *arXiv preprint arXiv:2005.05298*, 2020.
- [21] E. Hosseini-Asl, B. McCann, C.-S. Wu, S. Yavuz, and R. Socher, "A simple language model for task-oriented dialogue," *arXiv preprint arXiv:2005.00796*, 2020.
- [22] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [23] P. Budzianowski, T. Wen, B. Tseng, I. Casanueva, S. Ultes, O. Ramadan, and M. Gasic, "Multiwoz - A large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling," 2018, pp. 5016–5026.
- [24] Q. Zhu, Z. Zhang, Y. Fang, X. Li, R. Takanobu, J. Li, B. Peng, J. Gao, X. Zhu, and M. Huang, "Convlab-2: An open-source toolkit for building, evaluating, and diagnosing dialogue systems," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, ACL 2020, Online, July 5-10, 2020*, pp. 142–149.
- [25] S. Lee, Q. Zhu, R. Takanobu, Z. Zhang, Y. Zhang, X. Li, J. Li, B. Peng, X. Li, M. Huang, and J. Gao, "Convlab: Multi-domain end-to-end dialog system platform," in *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28 - August 2, 2019, Volume 3: System Demonstrations*, 2019, pp. 64–69.
- [26] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pp. 4171–4186.
- [27] J. Schatzmann, B. Thomson, K. Weilhammer, H. Ye, and S. Young, "Agenda-based user simulation for bootstrapping a pomdp dialogue system," in *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics: Companion Volume, Short Papers*, 2007, pp. 149–152.
- [28] S. Bao, H. He, F. Wang, H. Wu, H. Wang, W. Wu, Z. Guo, Z. Liu, and X. Xu, "Plato-2: Towards building an open-domain chatbot via curriculum learning," *arXiv preprint arXiv:2006.16779*, 2020.
- [29] S. Gururangan, A. Marasovic, S. Swayamdipta, K. Lo, I. Beltagy, D. Downey, and N. A. Smith, "Don't stop pretraining: Adapt language models to domains and tasks," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pp. 8342–8360.
- [30] Q. Zhu, K. Huang, Z. Zhang, X. Zhu, and M. Huang, "CrossWOZ: A large-scale chinese cross-domain task-oriented dialogue dataset," *Transactions of the Association for Computational Linguistics*, 2020. [Online]. Available: <https://arxiv.org/abs/2002.11893>
- [31] N. Mrksic, I. Vulic, D. Ó. Séaghdha, I. Leviant, R. Reichart, M. Gasic, A. Korhonen, and S. Young, "Semantic specialization of distributional word vector spaces using monolingual and cross-lingual constraints," *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 309–324, 2017.
- [32] S. Schuster, S. Gupta, R. Shah, and M. Lewis, "Cross-lingual transfer learning for multilingual task oriented dialog," in *NAACL-HLT*, 2019.
- [33] H. Lee, J. Lee, and T.-Y. Kim, "Sumbt: Slot-utterance matching for universal and scalable belief tracking," 2019.

- [34] Y. Shan, Z. Li, J. Zhang, F. Meng, Y. Feng, C. Niu, and J. Zhou, “A contextual hierarchical attention network with adaptive objective for dialogue state tracking,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, Jul. 2020, pp. 6322–6333. [Online]. Available: <https://www.aclweb.org/anthology/2020.acl-main.563>
- [35] S. Kim, S. Yang, G. Kim, and S.-W. Lee, “Efficient dialogue state tracking by selectively overwriting memory,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, Jul. 2020, pp. 567–582. [Online]. Available: <https://www.aclweb.org/anthology/2020.acl-main.53>
- [36] M. Eskenazi, S. Mehri, E. Razumovskaia, and T. Zhao, “Beyond turing: Intelligent agents centered on the user,” *arXiv preprint arXiv:1901.06613*, 2019.
- [37] O. Vinyals and Q. Le, “A neural conversational model,” *arXiv preprint arXiv:1506.05869*, 2015.
- [38] K. Gopalakrishnan, B. Hedayatnia, Q. Chen, A. Gottardi, S. Kwatra, A. Venkatesh, R. Gabriel, and D. Hakkani-Tür, “Topical-chat: Towards knowledge-grounded open-domain conversations,” in *INTERSPEECH*, 2019, pp. 1891–1895.
- [39] T. Zhao, K. Lee, and M. Eskenazi, “Dialport: Connecting the spoken dialog research community to real user data,” in *2016 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2016, pp. 83–90.
- [40] S. Mehri and M. Eskenazi, “Usr: An unsupervised and reference free evaluation metric for dialog generation,” *arXiv preprint arXiv:2005.00456*, 2020.
- [41] S. Banerjee and A. Lavie, “Meteor: An automatic metric for mt evaluation with improved correlation with human judgments,” in *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005, pp. 65–72.
- [42] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi, “Bertscore: Evaluating text generation with bert,” *arXiv preprint arXiv:1904.09675*, 2019.
- [43] S. Mehri and M. Eskenazi, “Unsupervised evaluation of interactive dialog with dialogpt,” *arXiv preprint arXiv:2006.12719*, 2020.
- [44] H. Ai, A. Raux, D. Bohus, M. Eskenazi, and D. Litman, “Comparing spoken dialog corpora collected with recruited subjects versus real users,” in *Proceedings of the 8th SIGdial Workshop on Discourse and Dialogue*, 2007, pp. 124–131.
- [45] Y. Zhang, S. Sun, M. Galley, Y.-C. Chen, C. Brockett, X. Gao, J. Gao, J. Liu, and B. Dolan, “Dialogpt: Large-scale generative pre-training for conversational response generation,” *arXiv preprint arXiv:1911.00536*, 2019.
- [46] D. Adiwardana, M.-T. Luong, D. R. So, J. Hall, N. Fiedel, R. Thoppilan, Z. Yang, A. Kulshreshtha, G. Nemade, Y. Lu *et al.*, “Towards a human-like open-domain chatbot,” *arXiv preprint arXiv:2001.09977*, 2020.
- [47] A. Ram, R. Prasad, C. Khatri, A. Venkatesh, R. Gabriel, Q. Liu, J. Nunn, B. Hedayatnia, M. Cheng, A. Nagar *et al.*, “Conversational ai: The science behind the alexa prize,” *arXiv preprint arXiv:1801.03604*, 2018.
- [48] S. Moon, S. Kottur, P. A. Crook, A. De, S. Poddar, T. Levin, D. Whitney, D. Difrancio, A. Beirami, E. Cho, R. Subba, and A. Geramifard, “Situating and interactive multimodal conversations,” *The 28th International Conference on Computational Linguistics (COLING)*, 2020.
- [49] P. A. Crook, S. Poddar, A. De, S. Shafi, D. Whitney, A. Geramifard, and R. Subba, “Simmc: Situated interactive multi-modal conversational data collection and evaluation platform,” *arXiv preprint arXiv:1911.02690*, 2019.
- [50] A. H. Miller, W. Feng, A. Fisch, J. Lu, D. Batra, A. Bordes, D. Parikh, and J. Weston, “Parlai: A dialog research software platform,” *arXiv preprint arXiv:1705.06476*, 2017.
- [51] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, “BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, Jul. 2020, pp. 7871–7880. [Online]. Available: <https://www.aclweb.org/anthology/2020.acl-main.703>
- [52] S. Humeau, K. Shuster, M. Lachaux, and J. Weston, “Poly-encoders: Architectures and pre-training strategies for fast and accurate multi-sentence scoring,” in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. [Online]. Available: <https://openreview.net/forum?id=SkxgnnNFvH>
- [53] P.-E. Mazaré, S. Humeau, M. Raison, and A. Bordes, “Training millions of personalized dialogue agents,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium: Association for Computational Linguistics, Oct.-Nov. 2018, pp. 2775–2779. [Online]. Available: <https://www.aclweb.org/anthology/D18-1298>
- [54] E. Dinan, S. Roller, K. Shuster, A. Fan, M. Auli, and J. Weston, “Wizard of Wikipedia: Knowledge-powered conversational agents,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.