

RADDLE: An Evaluation Benchmark and Analysis Platform for Robust Task-oriented Dialog Systems

Baolin Peng¹, Chunyuan Li¹,
Zhu Zhang^{1,2†}, Chenguang Zhu¹, Jinchao Li¹, Jianfeng Gao¹

¹Microsoft Research, Redmond, WA

²Iowa State University / Ames, IA

{bapeng, chunyl, chezhu, jincli, jfgao}@microsoft.com
zhuzhang@iastate.edu

Abstract

For task-oriented dialog systems to be maximally useful, it must be able to process conversations in a way that is (1) generalizable with a small number of training examples for new task domains, and (2) robust to user input in various styles, modalities, or domains. In pursuit of these goals, we introduce the RADDLE¹ benchmark², a collection of corpora and tools for evaluating the performance of models across a diverse set of domains. By including tasks with limited training data, RADDLE is designed to favor and encourage models with a strong generalization ability. RADDLE also includes a diagnostic checklist that facilitates detailed robustness analysis in aspects such as language variations, speech errors, unseen entities, and out-of-domain utterances. We evaluate recent state-of-the-art systems based on pre-training and fine-tuning, and find that grounded pre-training on heterogeneous dialog corpora performs better than training a separate model per domain. Adversarial training is also proposed to improve model robustness against noisy inputs. Overall, existing models are less than satisfactory in robustness evaluation, which suggests opportunities for future improvement.

1 Introduction

Dialogs constitute a crucial communication channel in completing a broad range of tasks, such as weather query, flight and restaurant booking, movie booking, IT help desk, *etc.* Comparing to chat systems that are usually modeled with single-turn context-response pairs, task-oriented dialog systems involve retrieving information from knowledge bases and reasoning over multiple dialog turns. This makes it especially important for a system to

be able to produce response that are grounded on tasks goals and user intents. In a bid to support human-computer interactions, task-oriented dialog systems have been built to allow users to converse with a computer system using natural language, such as Siri, Google Assistant, Amazon Alexa, Microsoft XiaoIce (Zhou et al., 2020). Traditionally, a task-oriented dialog system uses a modularized pipeline with four modules that execute sequentially (Gao et al., 2019). A natural language understanding (NLU) module identifies user intents and extracts associated information such as slots and corresponding values from user input. A dialog state tracker (DST) infers the belief state (or user goal) from dialog history. The belief state is often used to query a task-specific database (DB) to obtain the DB state, such as the number of entities that match the user goal. The dialog state and DB state are then passed to a dialog policy (POL) module to select the next system action. A natural language generation (NLG) module converts the action to a natural language response.

The human ability to converse is general, flexible, and robust. In contrast, most popular tools for dialog system development adopting the above modular systems are designed for specific tasks and struggle with out-of-scope data. If we aspire to develop models beyond extensively hand-crafted rules and annotated data for each single domain/task, it is critical to develop a more *unified, efficient* and *robust* model that can more quickly learn to execute a range of tasks in different domains.

To fuel research in this direction, we present the RADDLE benchmark. It includes a collection of task-oriented dialog tasks in diverse domains (e.g. end-to-end modeling, dialog state tracking). The benchmark also has a companion online platform for model evaluation, comparison, and robustness analysis. Importantly, RADDLE exhibits two

[†] Work was done when Zhu Zhang was visiting MSR

¹ Robust tAsk-orientedD DiaLog systems Evaluation

² Benchmark link: <http://aka.ms/raddle>

unique advantages that pave the way for building more pragmatic dialog systems: (i) **Limited data setting** is the major focus of RADDLE, to evaluate the generalization ability of models. It aims at simulating the real-world application scenarios where only very limited amount of labelled data is available for new domains. Given this focus, RADDLE is therefore a favorable benchmark to evaluate recent models in the pre-training and fine-tuning paradigm, which learn to represent linguistic knowledge in a way that facilitates sample-efficient learning and effective knowledge transfer. (ii) **Robustness analysis** is introduced to study model performance in various challenging scenarios, where models are evaluated with anomalous user input such as language variations, speech errors, unseen entities and out-of-domain utterances. Failing to handle these inputs often produce inappropriate responses leading to frustrating user experience. These scenarios are common for deployed systems in the real world, but are largely ignored in existing dialog benchmarks. To the best of our knowledge, RADDLE presents the first work to fill this gap.

To better understand the challenges posed by RADDLE, we conduct experiments with simple baselines and state-of-the-art task-oriented dialog models. We find that grounded pre-trained models with a unified multi-task learning objective outperform models separately trained on each domain. Moreover, even the best performing model (SOLOIST (Peng et al., 2020a)) in our evaluation achieves a fairly low score in robustness analysis. This suggests that our baseline models can handle common inputs with strong regularities, but struggle with anomalous inputs that require deeper reasoning.

In summary, our key contributions are: (i) A novel dialog benchmark with an emphasis on limited data and multiple domains/tasks, which formally creates a scenario to evaluate the grounding and generalization ability of pre-trained models. (ii) A crowd-sourced diagnostic evaluation dataset to cover a broad range of real-world sophistication to study model robustness. (iii) An online evaluation platform and leaderboard to track research progress, with human evaluation services to be granted to top-ranked submissions on a bi-monthly basis. (iv) Baseline results for major existing approaches to task-oriented dialogs are reported. An adversarially robust model is proposed to improve the generalization ability in noisy environments.

Starter codes, pre-trained models, and scripts to reproduce the results will be provided together with the benchmark.

2 Related Work

2.1 Dialog Benchmarks

To drive the progress of building dialog systems using data-driven approaches, a number of conversational corpora have been released. They are roughly grouped into two categories: (i) Corpora with structured semantic labels (Wen et al., 2017; Shah et al., 2018). These datasets are often specifically annotated, and used to study an individual module in the dialog pipeline. For example, DialoGLUE (Mehri et al., 2020) is a recently proposed benchmark with a focus on NLU and DST tasks. (ii) Corpora with an implicit user goal (Lowe et al., 2015). These datasets are often without semantic labels but can be used in end-to-end (E2E) dialog modeling (Li et al., 2016; Zhu, 2020; Wu et al., 2019; Zhu et al., 2019a; Lee et al., 2019; Zhu et al., 2020).

MultiWOZ (Budzianowski et al., 2018) is the most related work to RADDLE. It is a large-scale multi-turn conversational corpus across several domains. It can be used to develop individual dialog modules as separate tasks for existing modular-based methods, or serves as a benchmark for E2E dialog modeling methods. RADDLE inherits the advantages of MultiWOZ in its flexibility for separate/joint task modeling and its comprehensiveness in multi-domain data coverage, but differs significantly in two aspects: an emphasis on limited data settings and a unique robustness checklist. Both are essential qualities in building task bots at scale.

Further, RADDLE provides an online platform for model evaluation and fair comparison based on privately-held test data, inspired by GLUE (Wang et al., 2018). To the best of our knowledge, RADDLE is the first online platform for DST and E2E tasks in the dialog community. This can reduce the inconsistency caused by different researchers/teams using varying processing/evaluation scripts to dilute where the gain comes from.

2.2 Evaluation of Pre-Trained Models

Pre-trained language models (PLMs) have substantially advanced the state of the art across a variety of language understanding and generation tasks (Peters et al., 2018; Devlin et al., 2019; Yang et al., 2019; Liu et al., 2019; Radford et al., 2019;

Domain	Standard				Language Variations / Speech Errors				Unseen	OOD
	Attraction	Train	Hotel	Restaurant	Attraction	Train	Hotel	Restaurant	Reminder	Attraction
#Train	50	50	50	50	-	-	-	-	50	50
#Test	100	200	200	200	100	200	200	200	400	800
Task	Dialog State Tracking / End-to-End Modeling							DST / IC	DST / OOD	
Metrics	Joint Goal Accuracy / Combined Score							JGA / Acc.	JGA / F1	

Table 1: Dataset descriptions and statistics. DST is short for Dialog State Tracking, E2E denotes End-to-End modeling, and IC stands for Intent Classification. Joint Goal Accuracy (JGA) is used for DST and Combined score is used for E2E.

Keskar et al., 2019; Dong et al., 2019; Peng et al., 2020b,c; Li et al., 2020a). PLMs are often trained to predict words based on their context on massive text data, and the learned models can be fine-tuned to quickly adapt to various downstream tasks, exhibiting strong generalization capacity even with just a few in-domain training examples. Building task bots at scale requires the model to deal with the limited data problem for each domain, which can be used as a testbed to evaluate the generalization ability of PLMs. To this end, we limit the number of task-specific training examples in RADDLE to evaluate the sample-efficiency of models.

Meanwhile, task-oriented dialogs pose a unique set of challenges for PLMs (Gao et al., 2020): a dialog is intrinsically goal-driven, multi-turn and often informal/noisy. Indeed, dialog-specific PLMs are proposed (Wu et al., 2020a; Peng et al., 2020a). However, the robustness of PLMs to linguistic perturbations often occurring in dialog settings (See Section 4 for details) is largely unexplored. Note that our notion of robustness emphasizes *natural* language variations, which is different from *adversarial* examples/training that aim to fool a trained model (Nie et al., 2019). From this perspective, RADDLE provides an unique benchmark for assessing PLMs with a robustness orientation.

3 Tasks

RADDLE is centered on five English dialog scenarios in daily life, which cover a broad range of data collection schemes, task types and complexities. As our first goal of RADDLE is to spur development of generalizable dialog systems, we design the benchmark such that a good performance requires a model to leverage substantial knowledge (e.g., pre-trained parameters) learned from its previous life cycle, while still maintaining some task-specific components (Coope et al., 2020; Henderson et al., 2020; Peng et al., 2020a; Wu et al., 2020b). Specifi-

cally, we deliberately keep a small number of training examples for each scenario. This is consistent with the common practice that only limited labelled data is provided when deploying a dialog system to new domains. Table 1 shows the data statistics. Four domains in the standard-setting are sampled from MultiWOZ 2.0 (Budzianowski et al., 2018). *Reminder* is intentionally only utilized for unseen entity tracking. Because it is a human-machine corpus with a relatively smaller action space meaning that the impact of policy learning on models is largely alleviated. Therefore, the performance of models on this corpus will mostly reflect its capability of unseen entity tracking. Note that the number of training examples is limited to 50, an accepted scale that users can provide. Though it is possible to train a single model for each task from scratch without outside sources of knowledge, we expect that our focus on data-scarce settings will render this approach uncompetitive.

Furthermore, a typical task-oriented dialog system uses a modularized pipeline that has four modules and executes sequentially. Recent research has shown promising results on parameterizing the modularized pipeline using a single neural autoregressive model, and training it in an end-to-end manner (Peng et al., 2020a; Ham et al., 2020; Hosseini-Asl et al., 2020). In fact, a single autoregressive model can significantly ease the workflow of training and deploying dialog systems for new tasks, compared to existing modularized tools and methods. Therefore, we design the benchmark to allow evaluations on end-to-end dialog modeling, in addition to the modularized evaluation on dialog state tracking. To reveal the gap between the complexity of dialogs in lab environments and that in real scenarios, we construct a suite of tasks to study the robustness of models. We describe these tasks below and in Table 1.

On the evaluation front, we concentrate on

simulation-based methodologies, in order to facilitate automation. Though we only offer human evaluations (Gao et al., 2019) to top-ranked submissions at this point, we emphasize realistic scenarios in pursuit of system robustness (see Section 4).

Task 1: Dialog State Tracking A robust NLU and DST is the first step towards building a reliable dialog system. The dialog state is a summary of the entire conversation till the current turn. In a task-oriented system, it is represented in the form of *slot-value* pairs, where *slot* indicates the category/attribute of the user goal expressed in the utterance, and *value* is the corresponding information. For the evaluation metric, we report *joint goal accuracy*, which indicates the proportion of dialog turns where all the user’s search goal constraints are correctly identified (Mrksic et al., 2017). To specially study the NLU performance, we consider *intent classification*, which aims to automatically extract meaning from a natural language utterance in order to understand user’s goal (Hemphill et al., 1990; Zhu et al., 2019b).

Task 2: End-to-End Modeling The end-to-end (E2E) dialog models consider dialog history as input, and produce the natural language response. It jointly implements the dialog management (including DST and POL) and response generation (*i.e.*, NLG) components. Following Budzianowski et al. (2018), Inform, Success, and BLEU scores are reported. The first two metrics evaluate dialog task completion: Inform measures if the system provides a correct entity (inform rate), meanwhile Success measures the exact matching of answering all the requested information (success rate), and if the answered information matches users’ goal. BLEU evaluates how fluent the generated responses are compared to human-written responses. A combined score (Combined) is also reported using $\text{Combined} = (\text{Inform} + \text{Success}) \times 0.5 + \text{BLEU}$ as an overall quality measure, as suggested in (Budzianowski et al., 2018).

4 Robustness Diagnostic Checklist

Existing benchmarks assume a world of a “perfect” user who always provides precise, concise, and semantically unambiguous utterances. These goal-oriented dialog datasets are largely collected by crowd-sourcing, where a crowd-sourced worker enacts the part of a real user by following a set of template instructions provided for the task. This

method results in a dataset where most user utterances are straight-forward, stick to the goal and tend to leave out the variation/errors commonly found in real-world conversational data. To this end, we collect a suite of language variations to reveal the dialog sophistication in the real world, and measure the robustness of dialog models.

4.1 Checklist Tasks

Language Variations 🌈 It is well-known that humans communicate using language with fairly large variations such as different ways of expressions or personalized styles (Sacks et al., 1978), while template-based crowd-sourcing fails in covering the linguistic variations (Schegloff et al., 1977; Moore and Arar, 2019). Specifically, we consider four types of variations in RADDLE: (i) *Paraphrase* widely exists among different users, who may present restatements of the meaning of a text or message using other words. (ii) *Verbosity* describes a quality that users may express their intents using more words than needed. (iii) *Simplification* is a quality that users express their intents using fewer words to be concise. (iv) *Typos* often result from mistakes made in the typing. In Figure 1(b)-(e), we provide examples to illustrate these language variations.

Speech Errors 🗣️ It is desirable that dialog systems can leverage automatic speech recognition (ASR) techniques to serve the speech modality, as in Amazon Alexa. However, almost all dialog systems have typically assumed that the user input is written text, and hoped that the system would seamlessly integrate with speech inputs. Recently, it has been empirically shown in Gopalakrishnan et al. (2020) that dialog systems trained on written data is very sensitive to various types of synthetic and actual ASR hypotheses in the dialog history. To bring attention to this gap, RADDLE promotes speech robustness as an evaluation criterion. For example in Figure 1(f), “what’s available” can be transcribed as “once available” due to ASR deficiency, and a robust dialog system is expected to still correctly perceive user intents.

Unseen Entities 📖 Most existing DST methods are not designed to handle slot values that are not known to the tracker. The assumption that a predefined ontology exists for the dialog and one can enumerate all possible values for each slot is often not valid in real-world scenarios. Even if such lists or dictionaries exist, they can be very large in size

User : I want to tour a college in the center of town. what's available?
System : I have several different options available . did you want one with free entrance ?
User : Entrance fee is not an issue.,
System : Christ's College is free to enter.
User : Can i have the postcode and address for that college?,
System : Yes, it's located at saint andrew's street, postcode cb23bu. Is there any other information you need?
User : no, that's it. thanks.
System : thank you , goodbye .

(a) Standard dialog session

User : I want to take a tour in a college which lies in the center of town. Any recommendation?
System : I have several different options available . did you want one with free entrance ?
User : I don't have a problem paying the entrance fee.
System : Christ's College is free to enter.
User : Will you please give me postcode and address of that college?
System : Yes, it's located at saint andrew's street, postcode cb23bu. Is there any other information you need?
User : no, perfect. I appreciate that.
System : thank you , goodbye .

(b) Paraphrase

User : I want to tour a college. what do you have available? I know some college in the center area is awesome.
System : I have several different options available . did you want one with free entrance ?
User : I do not have any preference on the fee but does paying the fee give me something extra?
System : Christ's College is free to enter.
User : Okay, sounds good! I would like to visit that college to enjoy this afternoon, do you have its address and postcode?
System : Yes, it's located at saint andrew's street, postcode cb23bu. Is there any other information you need?
User : no, thanks.
System : thank you , goodbye .

(c) Verbosity

User : Any college available in the center of town?
System : I have several different options available . did you want one with free entrance ?
User : do not care about fee
System : Christ's College is free to enter.
User : Give me address and postcode.
System : Yes, it's located at saint andrew's street, postcode cb23bu. Is there any other information you need?
User : no, thanks.
System : thank you , goodbye .

(d) Simplification

User : want to tour a college in center. whats available
System : I have several different options available . did you want one with free entrance ?
User : I don't have a problem paying the entrance fee.
System : Christ's College is free to enter.
User : Plz tell me addr and postcode of that college.
System : Yes, it's located at saint andrew's street, postcode cb23bu. Is there any other information you need?
User : no, thats all. thx
System : thank you , goodbye .

(e) Typos

User : i want to tour college in the center of town once available
System : I have several different options available . did you want one with free entrance ?
User : entrance phi is not an issue.
System : Christ's College is free to enter.
User : Can i have the postcard and address for that college?,
System : Yes, it's located at saint andrew's street, postcode cb23bu. Is there any other information you need?
User : i know that's it uh thanks
System : thank you , goodbye .

(f) Speech errors

User : I want to tour a park in the Bellevue downtown . what's available?
System : I have several different options available . did you want one with free entrance ?
User : Entrance fee is not an issue.,
System : Downtown Park is a nice one and it is free to enter.
User : Can i have the postcode and address for that college?,
System : Yes, it's located at 10201 NE 4th St, postcode 98004. Is there any other information you need?
User : no, that's it. thanks.
System : thank you , goodbye .

(g) Unseen entities

User : I want to tour a college in the center of town. what's available?,
System : I have several different options available . did you want one with free entrance ?
User : today is a sunny day we should go outside for a walk.
System : Sorry I do not understand. Can you paraphrase? By the way, Christ's College is a good place and fee to enter.
User : Can i have the postcode and address for that college?,
System : Yes, it's located at 10201 NE 4th St, postcode 98004. Is there any other information you need?
User : no, that's it. thanks.
System : thank you , goodbye .

(h) Out-of-domain utterance

Figure 1: Illustration of different language perturbations in the robustness diagnostic checklist. The standard dialog example is shown in (a). Based on it, (b)-(e) are four types of language variations 🗨️, (f) shows speech error 🗨️, (e) shows unseen entities 🗨️, and (h) shows out-of-domain utterance 🗨️. In each case, some representative examples are highlighted in red text.

and highly dynamic (Xu and Hu, 2018). Therefore, *unseen entities* are common in dialogs, *i.e.*, entities that are not observed during training, but appear in the testing stage. In Figure 1(g), the entity *Bellevue downtown* is in the knowledge base but never appears in model training, a robust DST should be able to recognize it as a city/place, via generalizing from other similar entities learned during training.

Out-of-Domain Utterances 🗨️ Most deployed task-oriented dialog systems are built for a closed set of target domains. Thus, they are fragile when

dealing with out-of-domain (OOD) utterances (Lee and Shalyminov, 2019). Failure to detect OOD utterances often prevents the model from responding with an appropriate fallback action, hence leading to frustrating user experience. Therefore, it is important to endow task bots with the ability to detect OOD utterances for special handling (Larson et al., 2019). For example, in Figure 1(h), the user suggests an excursion to a task bot trained in college consulting, which is out of the bot's scope. The bot is expected to raise a flag to label the utterance as an outlier, and guides the user to focus on the

current domain.

4.2 Collection Protocols

The standard setting is sampled from MultiWOZ 2.0 (Budzianowski et al., 2018) but re-purposed in a few-shot learning setting.

The language variations corpus is created by workers on Amazon Mechanical Turks based on the standard corpus. To maximize the quality, we require workers in US locale and have a minimal previous approval rate of 90%. Assignments are constructed at the turn level. Given a user utterance and associated dialog history, workers are required to answer four questions, what are the paraphrase, typos, verbose, and simplified versions of the user utterance. Moreover, in each assignment, the workers are instructed to exactly mention the slot values in the answers if the given user utterance has them. We pay Turks 0.5\$ per assignment and each assignment can be finished in one to two minutes.

For the speech recognition errors setting, we employ the audio-level error simulation (Gopalakrishnan et al., 2020), which generates audio signals from texts, adds noise into the audio, and then decodes the audio with an ASR model to obtain hypotheses. In particular, we employ Microsoft Cognition text-to-speech service to synthesize audio signals. After injecting background noise into the audio signals, we use the speech recognition service to obtain a corpus of Word Error Rate (WER) of 30%.

For the `reminder` domain that is applied for unseen entity evaluation, we firstly simulate several dialogs as seed scenarios using an agenda-based simulator and then randomly replace the slots in the dialogs with new values. Similar to constructing the language variations corpus, we then hire workers to rewrite the corpus as diverse and realistic as possible. Finally, the out-of-domain corpus is developed following Lee and Shalymov (2019). We randomly choose 50% utterances in DSTC (Henderson et al., 2014) for the `Attraction` domain as the training set. For the test set, besides utterance from DSTC, we also introduce utterance from a diverse set of domains like `Stanford` (Eric and Manning, 2017), `Reddit`, `Twitter` (Sordani et al., 2015) to evaluate the capability of handling different out-of-domain utterances. A board of data researchers reviews all the collected data to ensure no ethical concerns in it.

5 Methods

5.1 Competitive Baselines

For baselines, we consider three representative methods, holding state-of-the-art positions on existing benchmarks such as MultiWoZ (Budzianowski et al., 2018).

DAMD (Zhang et al., 2020) is a state-of-the-art modular system, where each dialog module is implemented using a neural network, and the whole system is trained in an end-to-end manner.

GPT-2 represents a single multi-task learning model with impressive results on general language understanding and generation tasks. GPT-2 is an auto-regressive language model that leverages 12-24 layers of masked, multi-head self-attention Transformers. GPT-2 is pre-trained on extremely massive text data OpenWebText (Radford et al., 2019). It has demonstrated superior performance on characterizing human language data distribution and knowledge transfer. Given text prompts, GPT-2 can often generate fluent sentences. Its ancestral work GPT (with a smaller model size and less training data) has shown impressive results on language understanding tasks. In this paper, we consider GPT-2_{FT} as the approach of directly fine-tuning the pre-trained GPT-2 on a specific domain. Hence, GPT-2_{FT} can be viewed as SOLOIST without grounded pre-training, and serve as a strong baseline for both DST and E2E task.

SOLOIST represents recent model variants (Ham et al., 2020; Hosseini-Asl et al., 2020) to parameterize dialog system as a single auto-regressive model. SOLOIST subsumes different dialog modules (e.g. state tracker, dialog policy, response generator) into a single Transformer model. It has the similar capability with GPT-2 in understanding and generating natural language sentences but is pre-trained on large heterogeneous dialog corpora to gain additional capability of grounding text response in user goals and real-world knowledge for task completion (Peng et al., 2020a; Gao et al., 2020). For detailed description, please see Section A in Appendix.

5.2 Adversarially Robust SOLOIST

It is known that adversarial training can improve a model’s adversarial robustness, which refers to a model’s invariance to small (often imperceptible) perturbations of its inputs (*i.e.*, clean exam-

Model	Standard		Para.		Simp.		Typos		Verbo.		Speech ERR		Unseen		OOD		
	Avg.	Avg. _c	JGA ↑	C ↑	JGA ↑	C ↑	JGA ↑	C ↑	JGA ↑	C ↑	JGA ↑	C ↑	JGA ↑	IC ↑	JGA ↑	F1 ↑	
DAMD	-	-	14.18	48.99	6.75	44.13	5.78	42.93	5.33	42.58	7.08	42.56	9.1	45.94	-	-	-
GPT-2 _{FT}	47.46	46.53	40.52	67.36	31.36	62.72	28.82	59.44	22.31	54.15	30.40	54.16	31.41	65.95	28.28	51.29	47.37 83.86
SOLOIST	59.09	58.30	53.17	76.13	40.27	64.89	37.18	63.61	22.73	57.77	38.21	65.71	36.81	70.48	69.05	96.98	56.28 96.18
SOLOIST _{Adv}	61.03	60.14	55.47	79.06	42.11	71.13	38.28	69.89	23.30	63.17	40.02	69.36	39.02	72.33	69.56	98.79	55.03 89.94

Table 2: Overall results of baselines across all RADDLE tasks. C indicates the Combined metric, IC denotes intent classification accuracy. Avg. is averaged over all the tasks while Avg._c is averaged over all the robust checklist tasks. Para., Simp., Verbo. are short for Paraphrase, Simplification, and Verbosity. Note that it is not straightforward to directly apply DAMD to Unseen and OOD tasks since it requires extra annotations. As such, we omit results of DAMD on these two tasks.

ples) (Madry et al., 2017; Miyato et al., 2018; Liu et al., 2020; Li et al., 2020b). Adversarial examples are produced by adding perturbations on clean examples to fool the predictions of a trained model the most. Though fundamentally different, one may view adversarial examples as resembling the variations in natural language to some extent. Inspired by this idea, we propose an adversarially robust SOLOIST model, denoted as SOLOIST_{Adv}.

Specifically, for a dialog turn x drawn from the training dataset \mathcal{D} , and a neural model SOLOIST parameterized by θ , the *standard training* minimizes the empirical risk: $\min_{\theta} \mathbb{E}_{x \sim \mathcal{D}} \mathcal{L}_{\theta}(x)$, where $\mathcal{L}_{\theta}(x)$ is the SOLOIST learning objective defined in Appendix Section A. The key idea of *adversarial training* is to modify the objective by applying small perturbation δ to input word embeddings that maximize the adversarial loss: $\min_{\theta} \mathbb{E}_{x \sim \mathcal{D}} \max_{\delta} \mathcal{L}_{\theta}(x + \delta)$, where the inner maximization can be solved by running a number of projected gradient descent steps (Goodfellow et al., 2014; Bubeck, 2014). SOLOIST_{Adv} is trained in a hybrid manner that combines standard training and adversarial training. It augments the training dataset with adversarial examples that add perturbations in the word embedding space of original dialog turns, which improve the model’s robustness against noisy inputs that arguably covers language variations. In our experiments, SOLOIST_{Adv} employs adversarial training in both task-specific pre-training and fine-tuning stages.

5.3 Submission Details

Training We leverage the pre-trained checkpoints from the corresponding work, and fine-tune them on RADDLE. For SOLOIST_{Adv}, We apply 100k steps of adversarial training to the pre-trained checkpoints. Each domain is trained separately. We train our models with Adam with initial learning rate 5e-5 and batch size 1 for 20 epochs. We encourage subsequent submissions systems to devote

the same computation efforts in fine-tuning stage, e.g., up to one hour GPU time, for each model to ensure fair comparisons.

Evaluation The RADDLE benchmark follows the same evaluation model as GLUE (Wang et al., 2018) or Kaggle³. To evaluate a system on the benchmark, one must run the system on the provided test data for the tasks, then upload the results to the website <http://aka.ms/raddler> for scoring. The benchmark site shows per-task scores and a macro-average of those scores to determine a system’s position on the leaderboard. The website also provides fine- and coarse-grained results on the robustness diagnostic datasets. We will provide human evaluation services for top-ranked submissions on a quarterly basis. The human evaluation protocol follows Peng et al. (2020a) and Li et al. (2020c).

6 Benchmark Results

6.1 Overall Results

We first present the results of baseline methods across all tasks on the RADDLE benchmark in Table 2. As shown, GPT-2_{FT} fine-tuned with domain-specific dialog corpora outperforms the strong modular-based method DAMD. This highlights the efficacy of pre-trained language models. SOLOIST improves upon GPT-2_{FT} over 10 points in terms of average score, and consistently performs better than GPT-2_{FT} across all the tasks. These strong results indicate that large-scale task-specific pre-training on dialog corpora is crucial for effective and robust task adaptation. However, the performance of SOLOIST drops on robust checklist tasks. Benefiting from adversarial training, SOLOIST_{Adv} outperforms SOLOIST about 2 points.

³<https://www.kaggle.com/>

6.2 Robustness Diagnostic Checklist Results

Table 2 shows the overall performance of DST and E2E modeling under different variation settings.

Language Variations It is noticeable that all the models incur significant performance drops under each type of variation. Among all variation types, `Typos` has the most substantial impact on both `JGA` and `Combined` score resulting in 10 to 20 points of drop in performance. This is expected as misspelled keywords pose significant challenges for state tracking. The influence of other three types of variations are also prominent. The results reveal that existing SoTA dialog models trained on limited task-specific examples are not robust enough to handle various types of user utterances. Adversarial training improves robustness to language variations, boosting performance across all the language variations tasks.

Speech Errors We observe a clear degradation in all metrics for all models. This shows that during inference, models trained on textual data are sensitive and not robust to actual ASR hypotheses introduced in dialog history.

Unseen Entities Without task-specific pre-training, `GPT-2FT` only achieves less than 30% of `JGA` and 51.20 of dialog act accuracy even on a simple domain with most of the common entity values. `SOLOIST` performs significantly better than `GPT-2FT` by achieving 69.05% `JGA` and 96.98 dialog act accuracy but remains imperfect. `SOLOISTAdv` performs similar to `SOLOIST`, which is expected as adversarial training does not provides additional knowledge. These results imply that task-specific pre-training can improve the generalization capability of models but is still far from enough for production environments.

Out-of-Domain Utterances It is non-trivial for conventional modular-based dialog systems to handle OOD detection. It often requires an additional component to classify whether a user utterance as in-domain or not. As such, we omit the result of `DAMD` in our experiments. `GPT-2FT` achieves 83.96 F1 score while `SOLOIST` has 96.18 F1 score, which shows that task-specific pre-training can improve robustness of models to OOD utterances. It is interesting to observe that adversarial training hurts model’s performance on OOD detection. We conjecture that adversarial training enable models to tolerate disturbances on the inputs and thus yield

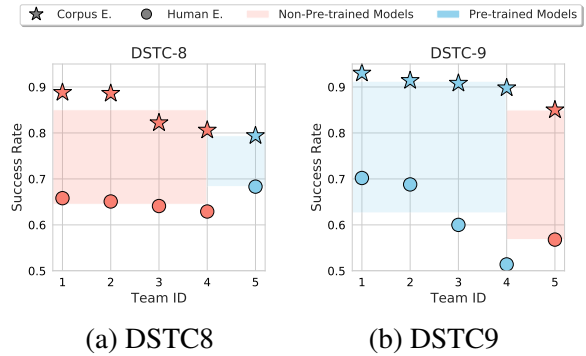


Figure 2: Corpus and human evaluation for different models in two recent Multi-domain Dialog Challenges: (a) DSTC8 and (b) DSTC9. The regions indicate the gap between human and corpus evaluations for different types of models. We observe that (i) In DSTC8, Team 5 is the winner, and the only submission adopting pre-trained GPT-2 models; The performance discrepancy between the corpus and human evaluation is significantly smaller than other teams using modular-based methods without pre-training. (ii) a general trend shifting from modular based systems to pre-trained end-to-end systems. (iii) a substantial drop in performance which indicates that pre-trained methods remain sensitive to noisy inputs.

more false positive predictions on this task.

Finally, it is worth pointing out some important trends in the dialog research community, based on the DSTC challenge (Kim et al., 2019; Gunasekara et al., 2020) in the last 2 years (Figure 2). In DSTC8 (Kim et al., 2019), the winning submission by Team 5 is the only one that uses pre-trained models (GPT-2). When moving from corpus evaluation to human evaluation, it exhibits the least performance drop relative to other submissions, which is strong evidence to demonstrate robustness of pre-trained models. By the time of DSTC9 (Gunasekara et al., 2020), the community have witnessed a general trend shift from modular systems to pre-trained end-to-end architectures. However, the significant performance gap between corpus evaluation and human evaluation indicates that pre-trained methods remain sensitive to noisy inputs. Such observations underscore the importance of robustness-oriented design and evaluation, for which RADDLE fills a major void.

7 Conclusion

We introduce RADDLE, a platform and collection of resources for evaluating and analyzing task-oriented dialog systems. We confirm (1) the utility of grounded pre-training and transfer learning methods in dialog systems: pre-training improves

generalization in a limited data setting, and (2) adversarial training improves robustness, but still leaves room for improvement. When evaluating these models on our diagnostic dataset, we find that they fail (often spectacularly) on many robustness test cases, suggesting possible avenues for future work. In summary, the question of how to design unified, efficient, robust models remains largely unexplored, and we believe that RADDLE can provide fertile soil for addressing this challenge.

Acknowledgement

We gratefully acknowledge the entire Project Philly team inside Microsoft, who provided the computing platform for our research. We also thank the anonymous reviewers whose suggestions helped clarify this work.

Ethical Considerations

The collection of our RADDLE dataset is consistent with the terms of use of any sources and the original authors' intellectual property and privacy rights. The dataset is collected with Amazon mechanical Turks, and each HIT requires up to two minutes to complete. The requested inputs are general language variations, and no privacy-related information is collected during data collection. Each HIT was paid 0.5 USD, with the hourly pay being 15% higher than the minimum wage requirements in our area. A board of data researchers has reviewed all the collected data to ensure no ethical concerns *e.g.*, toxic language and hate speech.

References

- Sébastien Bubeck. 2014. Convex optimization: Algorithms and complexity. *arXiv preprint arXiv:1405.4980*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. [Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026.
- Sam Coope, Tyler Farghly, Daniela Gerz, Ivan Vulic, and Matthew Henderson. 2020. [Span-convert: Few-shot span extraction for dialog with pretrained conversational representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 107–121. Association for Computational Linguistics.
- 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.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. [Unified language model pre-training for natural language understanding and generation](#). In *Advances in Neural Information Processing Systems*, pages 13042–13054.
- Mihail Eric and Christopher D Manning. 2017. Key-value retrieval networks for task-oriented dialogue. *arXiv preprint arXiv:1705.05414*.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2019. [Neural approaches to conversational ai](#). *Foundations and Trends® in Information Retrieval*, 13(2-3):127–298.
- Jianfeng Gao, Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayandeh, Lars Liden, and Heung-Yeung Shum. 2020. Robust conversational ai with grounded text generation. *arXiv preprint arXiv:2009.03457*.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.
- Karthik Gopalakrishnan, Behnam Hedayatnia, Longshaokan Wang, Yang Liu, and Dilek Hakkani-Tur. 2020. Are neural open-domain dialog systems robust to speech recognition errors in the dialog history? an empirical study. *arXiv preprint arXiv:2008.07683*.
- Chulaka Gunasekara, Seokhwan Kim, Luis Fernando D'Haro, Abhinav Rastogi, Yun-Nung Chen, Mihail Eric, Behnam Hedayatnia, Karthik Gopalakrishnan, Yang Liu, Chao-Wei Huang, et al. 2020. Overview of the ninth dialog system technology challenge: Dstc9. *arXiv preprint arXiv:2011.06486*.
- Donghoon Ham, Jeong-Gwan Lee, Youngsoo Jang, and Kee-Eung Kim. 2020. [End-to-end neural pipeline for goal-oriented dialogue systems using gpt-2](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 583–592.
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. [The ATIS spoken language systems pilot corpus](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Matthew Henderson, Iñigo Casanueva, Nikola Mrkšić, Pei-Hao Su, Tsung-Hsien Wen, and Ivan Vulic.

2020. [Convert: Efficient and accurate conversational representations from transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 2161–2174. 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.
- 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*.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. [Ctrl: A conditional transformer language model for controllable generation](#). *arXiv preprint arXiv:1909.05858*.
- Seokhwan Kim, Michel Galley, Chulaka Gunasekara, Sungjin Lee, Adam Atkinson, Baolin Peng, Hannes Schulz, Jianfeng Gao, Jinchao Li, Mahmoud Adada, et al. 2019. [The eighth dialog system technology challenge](#). *arXiv preprint arXiv:1911.06394*.
- Stefan Larson, Anish Mahendran, Joseph J. Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K. Kummerfeld, Kevin Leach, Michael A. Laurenzano, Lingjia Tang, and Jason Mars. 2019. [An evaluation dataset for intent classification and out-of-scope prediction](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1311–1316, Hong Kong, China. Association for Computational Linguistics.
- Sungjin Lee and Igor Shalyminov. 2019. Contextual out-of-domain utterance handling with counterfeit data augmentation. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7205–7209. IEEE.
- Sungjin Lee, Qi Zhu, Ryuichi Takanobu, Zheng Zhang, Yaoqin Zhang, Xiang Li, Jinchao Li, Baolin Peng, Xiujun Li, Minlie Huang, et al. 2019. [Convlab: Multi-domain end-to-end dialog system platform](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 64–69.
- Chunyuan Li, Xiang Gao, Yuan Li, Baolin Peng, Xiujun Li, Yizhe Zhang, and Jianfeng Gao. 2020a. [Optimus: Organizing sentences via pre-trained modeling of a latent space](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4678–4699, Online. Association for Computational Linguistics.
- Chunyuan Li, Xiujun Li, Lei Zhang, Baolin Peng, Mingyuan Zhou, and Jianfeng Gao. 2020b. Self-supervised pre-training with hard examples improves visual representations. *arXiv preprint arXiv:2012.13493*.
- Jinchao Li, Baolin Peng, Sungjin Lee, Jianfeng Gao, Ryuichi Takanobu, Qi Zhu, Minlie Huang, Hannes Schulz, Adam Atkinson, and Mahmoud Adada. 2020c. [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*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119.
- Xiaodong Liu, Hao Cheng, Pengcheng He, Weizhu Chen, Yu Wang, Hoifung Poon, and Jianfeng Gao. 2020. Adversarial training for large neural language models. *arXiv preprint arXiv:2004.08994*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. *arXiv preprint arXiv:1506.08909*.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- Shikib Mehri, Mihail Eric, and Dilek Hakkani-Tur. 2020. DialoGLUE: A natural language understanding benchmark for task-oriented dialogue. *arXiv preprint arXiv:2009.13570*.
- Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *T-PAMI*.
- Robert J Moore and Raphael Arar. 2019. *Conversational UX Design: A Practitioner’s Guide to the Natural Conversation Framework*. ACM.
- Nikola Mrksic, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve J Young. 2017. Neural belief tracker: Data-driven dialogue state tracking. In *ACL (1)*.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2019. Adversarial NLI: A new benchmark for natural language understanding. *arXiv preprint arXiv:1910.14599*.

- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayandeh, Lars Liden, and Jianfeng Gao. 2020a. [SOLOIST: few-shot task-oriented dialog with a single pre-trained auto-regressive model](#). *CoRR*, abs/2005.05298.
- Baolin Peng, Chenguang Zhu, Chunyuan Li, Xiujuan Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. 2020b. [Few-shot natural language generation for task-oriented dialog](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 172–182, Online. Association for Computational Linguistics.
- Baolin Peng, Chenguang Zhu, Michael Zeng, and Jianfeng Gao. 2020c. Data augmentation for spoken language understanding via pretrained models. *arXiv preprint arXiv:2004.13952*.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Harvey Sacks, Emanuel A Schegloff, and Gail Jefferson. 1978. A simplest systematics for the organization of turn taking for conversation. In *Studies in the organization of conversational interaction*. Elsevier.
- Emanuel A Schegloff, Gail Jefferson, and Harvey Sacks. 1977. The preference for self-correction in the organization of repair in conversation. *Language*.
- Pararth Shah, Dilek Hakkani-Tür, Gokhan Tür, Abhinav Rastogi, Ankur Bapna, Neha Nayak, and Larry Heck. 2018. Building a conversational agent overnight with dialogue self-play. *arXiv preprint arXiv:1801.04871*.
- Alessandro Sordani, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. A neural network approach to context-sensitive generation of conversational responses. *arXiv preprint arXiv:1506.06714*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gasic, Lina M Rojas Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. [A network-based end-to-end trainable task-oriented dialogue system](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 438–449.
- Chien-Sheng Wu, Steven Hoi, Richard Socher, and Caiming Xiong. 2020a. [Tod-bert: Pre-trained natural language understanding for task-oriented dialogues](#). *arXiv preprint arXiv:2004.06871*.
- Chien-Sheng Wu, Steven CH Hoi, Richard Socher, and Caiming Xiong. 2020b. [Tod-bert: Pre-trained natural language understanding for task-oriented dialogue](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 917–929.
- Qingyang Wu, Yichi Zhang, Yu Li, and Zhou Yu. 2019. [Alternating recurrent dialog model with large-scale pre-trained language models](#). *arXiv preprint arXiv:1910.03756*.
- Puyang Xu and Qi Hu. 2018. [An end-to-end approach for handling unknown slot values in dialogue state tracking](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1448–1457, Melbourne, Australia. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019. XLNet: Generalized autoregressive pretraining for language understanding. *NeurIPS*.
- Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020. [Task-oriented dialog systems that consider multiple appropriate responses under the same context](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9604–9611.
- Li Zhou, Jianfeng Gao, Di Li, and Heung-Yeung Shum. 2020. The design and implementation of xiaoice, an empathetic social chatbot. *Computational Linguistics*, 46(1):53–93.
- Chenguang Zhu. 2020. Boosting naturalness of language in task-oriented dialogues via adversarial training. *arXiv preprint arXiv:2004.14565*.
- Chenguang Zhu, Michael Zeng, and Xuedong Huang. 2019a. Multi-task learning for natural language generation in task-oriented dialogue. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1261–1266.
- Chenguang Zhu, Michael Zeng, and Xuedong Huang. 2019b. [Sim: A slot-independent neural model for dialogue state tracking](#). *arXiv preprint arXiv:1909.11833*.
- Qi Zhu, Zheng Zhang, Yan Fang, Xiang Li, Ryuichi Takano, Jinchao Li, Baolin Peng, Jianfeng Gao, Xiaoyan Zhu, and Minlie Huang. 2020. [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*, pages 142–149, Online. Association for Computational Linguistics.

A Background on SOLOIST

We review the SOLOIST (Peng et al., 2020a) for completeness. Each dialog turn is represented as:

$$\boldsymbol{x} = (\boldsymbol{s}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{r}), \quad (1)$$

where \boldsymbol{s} is the entire dialog history up to the current dialog turn, \boldsymbol{b} is the dialog belief state acquired from human annotation, \boldsymbol{c} is the DB state automatically retrieved from a database using \boldsymbol{b} , and \boldsymbol{r} is the delexicalized dialog response, from which the system response in natural language can be easily obtained with some automatic post-processing. In sum, each item in \boldsymbol{x} is by itself a sequence of tokens, the entire dialog turn can be viewed as a long sequence.

SOLOIST is a neural model parameterized by θ to characterize the sequence generation probability $p_{\theta}(\boldsymbol{x})$. It is pre-trained using publicly available heterogeneous dialog corpora with labels of belief states and DB states. The pre-trained model can be fine-tuned to any new task to generate responses grounded in task-specific user goals and a database. The pre-training and fine-tuning share the same multi-task objective for learning θ :

$$\mathcal{L}_{\theta} = \mathcal{L}_{\text{B}} + \mathcal{L}_{\text{R}} + \mathcal{L}_{\text{C}}, \quad (2)$$

where each task is described as follows:

Task 1: Belief Prediction For a belief state sequence of length T_b , we define the objective of predicting the belief state as:

$$\mathcal{L}_{\text{B}} = \log p(\boldsymbol{b}|\boldsymbol{s}) = \sum_{t=1}^{T_b} \log p_{\theta}(b_t|b_{<t}, \boldsymbol{s}), \quad (3)$$

where $b_{<t}$ indicates all tokens before t .

Task 2: Grounded Response Generation A delexicalized response of length T_r , $\boldsymbol{r} = [r_1, \dots, r_{T_r}]$, is generated by our model token-by-token from left to right, grounded in dialog history \boldsymbol{c} , belief state \boldsymbol{b} and DB state \boldsymbol{s} . The corresponding training objective is defined as

$$\begin{aligned} \mathcal{L}_{\text{R}} &= \log p(\boldsymbol{r}|\boldsymbol{c}, \boldsymbol{b}, \boldsymbol{s}) \\ &= \sum_{t=1}^{T_r} \log p_{\theta}(r_t|r_{<t}, \boldsymbol{c}, \boldsymbol{b}, \boldsymbol{s}). \end{aligned} \quad (4)$$

Task 3: Contrastive Objective A contrastive objective is employed to promote the matched items ($y = 1$ for positive samples \boldsymbol{x}) while driving

down the mismatched items ($y = 0$ for negative samples \boldsymbol{x}'). Since the the special token [EOS] attends all tokens in the sequence, the output feature on [EOS] is the fused representation of all items. We apply a binary classifier on top of the feature

$$\mathcal{L}_{\text{C}} = y \log(p_{\theta}(\boldsymbol{x})) + (1-y) \log(1 - p_{\theta}(\boldsymbol{x}')). \quad (5)$$

Please refer (Peng et al., 2020a) for more details.