

PRAL: A Tailored Pre-Training Model for Task-Oriented Dialog Generation

Jing Gu*

UC Davis

jkgu@ucdavis.edu

Qingyang Wu*

Columbia University

qingyang.wu@columbia.edu

Chongruo Wu

UC Davis

crwu@ucdavis.edu

Weiyan Shi

Columbia University

shi.weiyan@columbia.edu

Zhou Yu

Columbia University

zy2461@columbia.edu

Abstract

Large pre-trained language generation models such as GPT-2 have demonstrated their effectiveness as language priors by reaching state-of-the-art results in various language generation tasks. However, the performance of pre-trained models on task-oriented dialog tasks is still under-explored. We propose a Pre-trained Role Alternating Language model (PRAL), explicitly designed for task-oriented conversational systems. We design several techniques: start position randomization, knowledge distillation, and history discount to improve pre-training performance. In addition, we introduce a high quality large-scale task-oriented dialog pre-training dataset. We effectively adapt PRAL on three downstream tasks. With much less training data, PRAL outperforms or is on par with state-of-the-art models.

1 Introduction and Related Work

Current approaches to building task-oriented dialog systems still require a substantial amount of annotations and therefore are labor-intensive. On the other hand, large-scale pre-trained language models such as BERT (Devlin et al., 2019) and GPT (Radford et al., 2019) have achieved great success on various NLP tasks. There have been several attempts to apply these language models to dialog systems directly. For example, Transfer-Transfo (Wolf et al., 2019) fine-tuned GPT on the Persona-Chat dataset (Zhang et al., 2018b) and achieved the state-of-the-art performance on chitchat dialog generation. DialoGPT (Zhang et al., 2020) utilizes a large Reddit corpus to further pre-train GPT-2 (Zhang et al., 2020). All of these studies pointed to a promising direction towards building dialog systems with large-scale language models and less annotation.

However, these language models applied to dialog systems still have some limitations. First, further pre-training language models for dialog systems requires a considerable amount of training data. Small pre-training dialog datasets would not provide a large amount of commonsense knowledge needed for dialog generation. However, a diverse collection of high-quality dialog datasets is difficult to obtain. Besides, these language models usually do not consider dialog feature in their structures.

To tackle these issues, we propose Pre-trained Role Alternating Language model (PRAL), a language model designed explicitly for dialog generation. To begin with, we collect and process 13 dialog datasets, ranging from TV transcripts to pizza ordering dialogs, to enrich the pre-training data with high-quality dialog corpora. Second, we adopt ARDM proposed in Wu et al. (2019) and use two separate GPT-2 to model the two speakers in the dialog. Next, we apply Start Position Randomization (SPR) to cope with the variable lengths in dialogs, which prevents the language model from binding the position index with the text information. Additionally, we utilize a Teacher model to perform knowledge distillation and incorporate common sense knowledge into the dialog generation. Finally, we re-weigh each utterance with discount factors and emphasize on the later part in a dialog to better incorporate contextual information.

In summary, we propose PRAL and design several effective techniques to improve the dialog model pre-training. Our pre-trained model improves the success rate on CamRest676 and MultiWOZ dataset, and the coherence and diversity scores on PersuasionForGood. Our model is data-efficient and use 10x less than SOLOIST and 1000x less than DialoGPT in terms of training data size. We also process and present a collection of high-

* Equal contribution

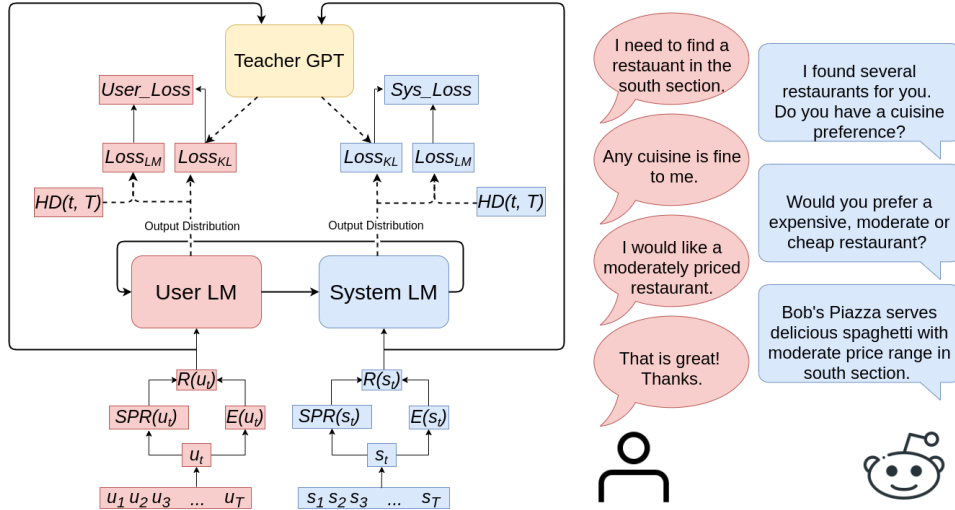


Figure 1: An overview of PRAL’s architecture. PRAL has separate language models for each speaker. The representation of user utterance u_t or system u_s is from word embedding E and the randomized position embedding SPR . $HD(t, T)$ is the history discount weight of each utterance. Teacher GPT provides supervision for the two language models. $Loss_{LM}$ and $Loss_{KL}$ denote the losses of the language models and the KL divergence.

Dataset Statistics	
# Datasets	13
# Dialogues	142,298
Avg. turns per dialogue	12.66
Avg. tokens per turn	11.78
Avg. tokens per dialogue	149.25
Total unique tokens	108,106

Table 1: Statistics of our dataset

quality dialog datasets suitable for pre-training large-scale language models for dialog systems.

2 PretrainDial Dataset for Pre-training

Large clean dialog datasets are difficult to find. Therefore, we constructed *PretrainDial*, a large-scale multi-domain dialog corpus for dialog pre-training. *PretrainDial* is a large-scale pre-training dataset and can only be collected from existing dialogs. We carefully selected 13 existing dialog corpora, ranging from chitchat such TV transcripts to task-oriented dialogs, and design a sophisticated text processing pipeline. Table 1 shows the statistics of the *PretrainDial* dataset. Please check appendix A for more details about the dataset statistic and the text processing pipeline.

3 Methods

In this section, we will first briefly introduce ARDM, our base dialog model, and then describe a set of techniques proposed in PRAL. Figure 1 shows the main structure of PRAL.

3.1 Alternating Roles Dialog Model

The basic idea behind ARDM (Wu et al., 2019) is to simultaneously model the user and system with two separate GPT-2 to capture the different language styles among different speakers. A dialog can be considered as a sequence of utterances $d = \{u_1, s_1, u_2, s_2, \dots, u_T, s_T\}$, where T is the total number of turns. We use p_u and p_s to represent the probability of the user utterance and system utterance. The dialog distribution is defined as:

$$p(d) = \prod_{t=1}^T p_u(u_t | u_{<t}, s_{<t}) p_s(s_t | u_{\leq t}, s_{<t}) \quad (1)$$

However, ARDM does not contain prior knowledge about dialog. In contrast, PRAL is designed for dialog system and absorbs abundant dialog knowledge during the pre-training process. To further improve ARDM or other dialog generation models, we propose three effective techniques to improve pre-training efficiency.

3.2 Start Position Randomization

We propose to randomize the start position to improve pre-training model’s quality. Transformer-based language models use position embedding to encode the location information for each token. It supports a fixed maximum position, and the position index always starts from 0. However, since most dialogs contain less than 1024 tokens, most vectors in the positional embedding would remain

zero and not update during pre-training. Besides, position embedding should only provide location information. However, the fixed start position will bond specific text with a particular position index. For example, “hi” is bonded with index 0 as “hi” usually appears at the beginning of the dialog. Therefore, the model is likely to overfit on the first several positional embeddings.

To address this issue, we propose to perform Start Position Randomization (SPR). L stands for the total number of tokens in a dialog, and the maximum start position index is $1024 - L$. We randomize the start position to be any number between 0 to $1024 - L$. It would disentangle the positional information from the textual meaning and force the model to update all the positional embeddings.

3.3 Teacher GPT

Neural networks models suffer from catastrophic forgetting (Kirkpatrick et al., 2016). Since we have finetuned GPT-2 with a new dialog corpus, the updated model is at risk in forgetting the prior knowledge from the original GPT-2. Teacher Model is used to calculate the distillation loss (Hinton et al., 2015) between the fixed GPT and our two language models. It constrains the language model from generating a token distribution that is too different from the original token distribution. The Teacher Model has two functions. First it avoids language model from catastrophic forgetting the knowledge in the original GPT-2 weights (Kirkpatrick et al., 2016). Secondly, when GPT-2 Large is used as Teacher Model, it imparts knowledge into our language models. The ablation study in table 2a validates the the functions.

3.4 History Discount

In dialog generation, historical utterances closer to current utterance should have a more significant impact on the generation than the ones that are further. Because in human conversations, we tend to prioritize local coherence over distant history coherence as well. Therefore, we introduce discount factor γ to re-weigh the importance of each utterance based on the turn number. For a dialog with a total number of T utterances and its current utterance index to be t , we weigh the language model loss with γ^{T-t} . By incorporating the discount factor γ , the model focus more on recent history in the generation process.

3.5 Optimization

We use a language modeling loss to optimize our model, shown in Equation 2.

$$Loss_{LM} = \sum_{t=1}^T \gamma^{T-t} \sum_{l=1}^{L_t-1} CE(P_{tl}, G_{t(l+1)}) \quad (2)$$

CE denotes the cross-entropy loss. T is the total number of utterances in a dialog, and L_t is the total number of tokens in the t^{th} utterance. For the loss of each utterance t in the dialog, it is weighed by the discount factor described in section 3.4. We combine loss from all words as the cross-entropy between the output probability distribution $P_{t(l+1)}$ and the ground truth $G_{t(l+1)}$.

The final loss combines the language model loss and KL divergence loss:

$$Loss = Loss_{LM} + \alpha KL(p, p^{constraint}) \quad (3)$$

The factor α is used to expedite model convergence and it decreases exponentially as the number of iterations increases, i.e. $\alpha = \alpha_0 \lambda^{iter}$.

4 Experiments

We pre-train PRAL on *PretrainDial*. We use GPT-2 large as the Teacher model. We use AdamW optimizer with warm-up steps as 10 percent of the training step. The learning rate is set to be 1×10^{-4} . For the calculation of loss, we set α_0 to be 0.1 and set λ to be 0.9999. The discount factor γ is set to be 0.95. To show the generalizability, we fine-tune PRAL on three downstream dialog generation tasks, CamRest676, MultiWOZ and PersuasionForGood, as is shown in Table. 2. Refer to Appendix B for more experiment setting.

CamRest676 (Rojas-Barahona et al., 2016) is a dialog dataset for restaurant recommendation containing 680 dialogs. We use BLEU-4 metrics to measure the quality of generated sentences, and Success F1 to evaluate the responses on specific slots, such as address, phone, postcode. Sequicity is the state-of-the-art method in task-oriented dialog tasks that requires annotations. PRAL beat all other models, including a concurrent work SOLOIST (Peng et al., 2020) on both BLEU-4 and Success F1. It is worth noting that PRAL does not need any annotation. SOLOIST and DialogPT have a close performance with our model. However, SOLOIST uses around 1 Million dialogues, DialogPT uses around 147 million dialogues, meanwhile we only use around 142K dialogues, which

Model	BLEU-4	Success F1
Sequicity	21.4	0.852
Sequicity (w/o RL)	22.9	0.821
GPT-2-finetune	21.8	0.851
DialoGPT	25.2	0.861
SOLOIST	25.5	0.871
ARDM	26.2	0.864
PRAL	27.2	0.874
- w/ Teacher GPT(small)	26.9	0.869
- w/o Teacher GPT	25.0	0.865
- w/o loss discount	27.0	0.867
- w/o SPR	26.6	0.869

(a) Results on CamRest676 dataset.

Model	Supervision		BLEU-4	Inform	Success
	Dialog State	Dialog Act			
Human	-	-	-	0.989	0.965
Baseline	✓	×	18.9	0.825	0.729
HDSA	✓	✓	23.6	0.877	0.734
LaRL	✓	×	12.8	0.828	0.792
SOLOIST	✓	×	18.0	0.896	0.793
ARDM	×	×	20.6	0.874	0.728
PRAL	×	×	21.2	0.899	0.798

(b) Results on MultiWOZ dataset

	Perplexity ↓	BLEU-1 ↑	BLEU-2 ↑	Fluency ↑	Logic ↑	Coherence ↑	Diversity ↑	Overall ↑	Avg. Donation ↑
ARDM	10.1	16.5	6.44	0.39	0.41	0.37	0.27	0.18	0.62
PRAL	10.3	17.3	10.9	0.61	0.59	0.63	0.73	0.82	0.99

(c) PersuasionForGood. Automatic Evaluation and Human Evaluation Results

Table 2: Evaluation on three datasets

is a thousand times less. This further shows PRAL is data-efficient.

Ablation studies on CamRest676 shows that the Teacher GPT plays the most important role. The fact that PRAL with Teacher GPT (Small) in table 2a outperforms PRAL without Teacher GPT (Small) shows the importance of the knowledge in the original model weights. When using GPT-2 Large as Teacher Model, the performance is better than that of using GPT-2 small, which validates the effect of knowledge distillation.

MultiWOZ (Budzianowski et al., 2018) contains around 10k dialogues covering various domains. We evaluate the models with on BLEU-4, Inform Rate, and Success Rate which measures if the system provides the requested information. PRAL outperforms the attention seq2seq model which is used as the baseline in Multiwoz (Budzianowski et al., 2018) in all metrics. Without using any annotation, PRAL also outperforms or achieve comparable results with HDSA (Budzianowski et al., 2018), LaRL (Zhao and Kawahara, 2019) and SOLOIST. Except for HDSA which requires both dialog state and dialog act, PRAL achieves a better BLEU score than all other models. PRAL outperforms ARDM in all metrics, which further validates the effectiveness of the pre-training process.

PersuasionForGood We also evaluate our method on PersuasionForood (Wang et al., 2019), where a persuader tries to persuade users to donate money to children. There are a total of 1,017 dialogues. Although not a traditional task-oriented dialog bench-

mark, it is a good benchmark for human evaluation. Automatic metrics evaluation is efficient but could fail to capture the text quality on a deeper and complicated level. We choose this task also because it benefits children. Unlike CamRest676 and Multiwoz, the language in PersuasionForGood dataset is so diverse that BLEU-4 scores of all of the models are too low to be scientific metrics. Therefore, we use BLEU-1 and BLEU-2 instead. Our model achieves a significantly higher score on BLUE metrics, especially on BLEU-2 (63% up). In human evaluation, we ask evaluators that how much they are willing to donate after the conversation and acquire their ratings in terms of fluency, logic, coherence, and diversity. The result suggests that PRAL outperforms ARDM on all the metrics. For human evaluation details, please refer to Appendices C.

Case studies show some linguistic problems in ARDM, such as repetition and unnaturalness. Meanwhile, with pre-training, PRAL is more natural and persuasive. Please refer to Appendices D for an example of PRAL and ARDM.

5 Conclusion

We propose PRAL, a large pre-trained dialog system for task-oriented generation. We incorporated methods that are designed for large dialog system into PRAL with good performances on three downstream tasks. The model generates more fluent, coherent, diverse, and logical dialogs according to human evaluation results. We also release a high-quality dialog dataset for the pre-training process.

References

- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Ultes Stefan, Ramadan Osman, and Milica Gašić. 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 (EMNLP)*.
- Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Ben Goodrich, Daniel Duckworth, Semih Yavuz, Amit Dubey, Kyu-Young Kim, and Andy Cedilnik. 2019. **Taskmaster-1: Toward a realistic and diverse dialog dataset**. 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 4516–4525, Hong Kong, China. Association for Computational Linguistics.
- Chit-Chat Challenge. a. Chitchat dataset. <https://github.com/BYU-PCCL/chitchat-dataset>.
- Kaggle Challenge. b. Friends corpus. <https://www.kaggle.com/vinayvk/friends-series-data-set>.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. 2011. Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In *CMCL@ACL*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Layla El Asri, Hannes Schulz, Shikhar Sharma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul Mehrotra, and Kaheer Suleman. 2017. **Frames: a corpus for adding memory to goal-oriented dialogue systems**. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 207–219, Saarbrücken, Germany. Association for Computational Linguistics.
- Joachim Fainberg, Ben Krause, Mihai Dobre, Marco Damonte, Emmanuel Kahembwe, Daniel Duma, Bonnie L. Webber, and Federico Fancellu. 2018. Talking to myself: self-dialogues as data for conversational agents. *ArXiv*, abs/1809.06641.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. **Distilling the knowledge in a neural network**.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. **The curious case of neural text degeneration**. In *International Conference on Learning Representations*.
- James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2016. **Overcoming catastrophic forgetting in neural networks**. *CoRR*, abs/1612.00796.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. **Dailydialog: A manually labelled multi-turn dialogue dataset**. In *IJCNLP*.
- Ilya Loshchilov and Frank Hutter. 2019. **Decoupled weight decay regularization**. In *International Conference on Learning Representations*.
- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayan-deh, Lars Liden, and Jianfeng Gao. 2020. **Soloist: Few-shot task-oriented dialog with a single pre-trained auto-regressive model**. *arXiv preprint arXiv:2005.05298*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. **Language models are unsupervised multitask learners**.
- Filip Radlinski, Krisztian Balog, Bill Byrne, and Karthik Krishnamoorthi. 2019. **Coached conversational preference elicitation: A case study in understanding movie preferences**. In *SIGDIAL 2019*.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2019. **Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset**. *ArXiv*, abs/1909.05855.
- Reddit. **Reddit corpus**. <https://zissou.infosci.cornell.edu/convokit/documentation/subreddit.html>.
- Lina Maria Rojas-Barahona, Milica Gaić, Nikola Mrksic, Pei hao Su, Stefan Ultes, Tsung-Hsien Wen, Steve J. Young, and David Vandyke. 2016. **A network-based end-to-end trainable task-oriented dialogue system**. In *EACL*.
- Xuwei Wang, Weiyan Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. **Persuasion for good: Towards a personalized persuasive dialogue system for social good**. In *ACL*.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. **Transfertransfo: A transfer learning approach for neural network based conversational agents**. *CoRR*, abs/1901.08149.
- Qingyang Wu, Yichi Zhang, Yu Li, and Zhou Yu. 2019. **Alternating roles dialog model with large-scale pre-trained language models**.
- Justine Zhang, Cristian Danescu-Niculescu-Mizil, Lucas Dixon, Yiqing Hua, Nithum Thain, and Dario Taraborelli. 2018a. **Conversations gone awry: Detecting early signs of conversational failure**. In *ACL*.

- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018b. [Personalizing dialogue agents: I have a dog, do you have pets too?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.
- Yizhe Zhang, Siqu Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. [Dialogpt: Large-scale generative pre-training for conversational response generation.](#) In *ACL, system demonstration*.
- Tianyu Zhao and Tatsuya Kawahara. 2019. [Effective incorporation of speaker information in utterance encoding in dialog.](#)

A Dataset sources

Our dataset contains high-quality dialogues which are selected from other 13 datasets listed in Table 3. PretrainDial is a large-scale pre-training dataset and can only be collected from existing dialogs. Due to the page limit as a short paper, we didn’t elaborate on the process in the paper. First, we collected dialog datasets that are commonly used in recent years. Then we filtered out the datasets with various standards such as content appropriateness. For example, we filtered “Conversations Gone Awry” Dataset because the conversation involves necessary background knowledge. Then, we process the text in the selected datasets. This step is essential since these datasets contain unnecessary noise, especially for datasets that contain raw text such as Friends dataset. The processing includes: (1) We replaced less informative appeared entity. For example, replace a long URL link with the word “URL”. (2) Delete meaningless repetition. (3) Delete responses that are not written in English. (4) Delete offensive language. (5). In some datasets such as Reddit, the conversation involves more than two people, so we extract a complete conversation flow involving only two people. Note there are more detailed process steps. We cannot describe all of them. We will release the text processing script, which we believe is helpful for the community when collecting dialog datasets.

B Experiment Setting Detail

B.1 Training Details

We initialize PRAL with a large pre-trained language model GPT-2 small with 117M parameters (Radford et al., 2019). We follow a special format in GPT-2 as the “trigger” so that the model can zero-shot dialog response generation, by prefixing the user role token “A:” or “B:”, and suffixing the end of utterance token “\n\n\n”. We first pre-train PRAL on *PretrainDial* and then further fine-tune PRAL on the specific task dataset. We apply AdamW optimizer (Loshchilov and Hutter, 2019), and the number of warm-up ratio is set to 0.1. Learning rate is 1×10^{-4} in the pre-training process and 3×10^{-5} in fine-tune process. The dropout rate is set to 0.1 for all tasks. For the calculation of loss in the pre-training process, we set α_0 to be 0.1 and set λ to be 0.9999. The discount factor γ is set to be 0.95.

B.2 Decoding Details

In the downstream task, we decode utterances by nucleus sampling (Holtzman et al., 2020) with different hyper-parameters (top-p, top-k). We also vary the temperature of $T < 1$ to find the best setting for the specific down-stream dialog task. We use nucleus sampling for all methods. In CamRest676 task, we set top-p 0.2 and temperature 0.7 for our model. For MultiWOZ task, we set the top-p to 0.2 and the temperature to 0.7. In PersuasionForGood task, to generate diverse responses, we use a top-p of 0.9 and a temperature of 0.7.

C Human Evaluation Detail

Twenty people participated in the human evaluation. ARDM is the state-of-the-art model for PersuasionForGood task. Each person will have ten conversations with PRAL and ARDM in random orders, five conversations for each model. 1) For the donation task, the participants will be asked, “How much will you donate after talking to the bot? Please choose from 0-2\$”. 2) For fluency, logic, and coherence, the participants will be asked, “Which one do you think is more fluent/logic/coherent?” and choose the model with better performance on the corresponding metric. 3) For diversity, participants compare performances across dialogs, so they will be asked, “Which bot do you think generates more diverse responses?” after talking with each model for five times.

D Persuasion Example

Table 4 and Table 5 show an PRAL example and an ARDM example.

E Ethical Statement

As knowledge in language model comes from training corpus, we used a strict pipeline to avoid offensive and other inappropriate text. We believe such a selective training source could contribute a more unbiased dialog system. However, the Teacher Model could be a potential source of biased knowledge, although we notice that is not the case in our experiments. In the future, we will also explore how to select the knowledge to import from Teacher Model.

1	CCPE-M dataset (Radlinski et al., 2019)
2	ChitChat-Dataset (Challenge, a)
3	Conversations Gone Awry Dataset(Wiki) (Zhang et al., 2018a)
4	Conversations Gone Awry Dataset(CMV) (Zhang et al., 2018a)
5	Cornell Movie-Dialogs Corpus (Danescu-Niculescu-Mizil and Lee, 2011)
6	DailyDialog (Li et al., 2017)
7	Frames Dataset (El Asri et al., 2017)
8	PersonaChat ConvAI2 Dataset (?)
9	The Schema-Guided Dialogue Dataset (Rastogi et al., 2019)
10	Self-dialogue Corpus (Fainberg et al., 2018)
11	Friends Series Dataset (Challenge, b)
12	Taskmaster-1 (Byrne et al., 2019)
13	Reddit Corpus (Reddit)

Table 3: 13 Dataset Sources

Role	Utterance
SYS	Good morning, how are you doing today?
USR	Good!
SYS	Have you ever heard of Save the Children?
USR	I have not heard of it. Would you like to give an introduction?
SYS	It is a charity that helps children in need around the world.
USR	Cool
SYS	I would like to give a little bit more information.
USR	Sure. Go ahead.
SYS	Save the Children is an international non-governmental organization that promotes children’s rights, provides relief and helps support children in developing countries.
USR	It is a good organization. I am glad to know it.
SYS	Yes, they are a very professional organization. Would you be interested in donating some of your task money to this charity?
USR	Sure. I would like to donate some money to it.
SYS	Thank you so much!
USR	no problem.
SYS	How much do you like to donate to the charity now? Your donation will be directly deducted from your task payment.
USR	I would like to give \$2 to them.
SYS	Thank you so much!
USR	You are welcome
SYS	I hope you have a great day!
USR	You too.

Table 4: An example conversation between a machine persuader trained by PRAL and a human persuadee.

Role	Utterance
SYS	Good morning.
USR	Good!
SYS	Have you ever heard of Save the Children?
USR	I have not heard of it. Would you like to give an introduction?
SYS	It is a charity that helps children.
USR	Cool
SYS	Do you wanna know more about it?.
USR	Sure. Go ahead.
SYS	It is a charity that helps children.
USR	I know. But what is this organization really about?
SYS	Save the Children is an international non-governmental organization that promotes children's rights, provides relief and helps support children in developing countries.
USR	It is a good organization. I would like to make some contributions.
SYS	Would you be interested in donating some money to this charity?
USR	Sure. I think I can do that.
SYS	How much do you like to donate to the charity now? Your donation will be directly deducted from your task payment.
USR	I would like to give \$1 to them.
SYS	Thank you so much! Bye!
USR	Bye!

Table 5: An example conversation between a machine persuader trained by ARDM and a human persuadee.