

user could be provided stories tailored for his/her preference;

- **Tutorial:** While reading an instruction book could be tedious and time-consuming, a chat bot could efficiently walk an user through the process.

As shown in Figure 1, for our task, a good conversation should have the following characteristics:

1. **Coverage:** The chat bot should try to convey as much information in the passage as possible, instead of mumbling about irrelevant information;
2. **Coherence:** The chat bot’s response should be coherent to the user’s utterance, making the user feel that his/her questions are followed and addressed.

In summary, we propose a novel framework which consists of the following two chat bot modules: 1) Teacher bot, which attempts to transfer the information in an input passage to a user through conversation; and 2) Student bot, which responds to Teacher bot to form coherent conversations during training. The two bots are trained in a two-phase manner: In Phase 1, we pre-tune the two chat bots on *Wizard of Wikipedia* (Dinan et al., 2019) dataset, enabling both bots with the basic ability of conversing over a passage. In Phase 2, we fine-tune Teacher bot through self-play with Student bot, guided by reinforcement rewards. In this process, we enhance Teacher bot to be more informative while maintaining the ability to coherently address human users. Specifically, the fine-tuning phase is unsupervised, i.e. Teacher bot could be transferred to various domains or corpora without additional annotated dialogue datasets.

Our contributions include: 1) A novel task of information-acquisition-oriented dialogue system; 2) A novel unsupervised learning framework which enables a teacher bot to carry out informative and coherent conversations with human users for information acquisition purpose; 3) Extensive experiments with human evaluation demonstrate the effectiveness of our proposed approach.

2 Approach

In order to obtain an informative and attentive teaching dialogue system, we propose a framework that consists of two chat bots in different roles, and leverage both supervised learning and unsupervised reinforcement learning, as illustrated in Figure 2. The unsupervised reinforcement learning enables the system to be fine-tuned on other text corpus

where no annotation or dialogue data is required.

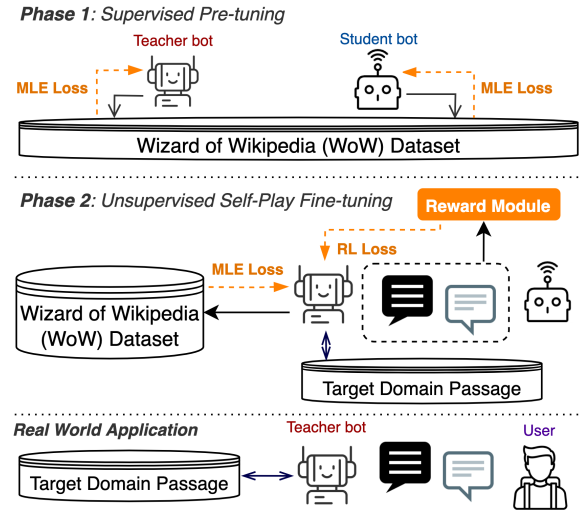


Figure 2: Our two-phase training framework

2.1 Model Architecture

Given a passage P , the conversation between Teacher bot X and Student bot Y can be denoted as a sequence of turns $C = \{U_1^X, U_1^Y, \dots, U_N^X, U_N^Y\}$, where N is the number of turns in the conversation. In order to mimic our use case, Teacher bot has access to P whereas Student bot does not.

Teacher bot X aims at transmitting the information in P to the student. At the n th turn, X takes as input P and the conversation history $H_n^Y = \{U_1^X, U_1^Y, \dots, U_{n-1}^X, U_{n-1}^Y\}$, and outputs U_n^X . Teacher bot X adopts DoHA (Prabhumoye et al., 2021), a pre-trained model for document-grounded text generation, and is tuned in supervised phase and unsupervised self-play phase.

In order to fine-tune Teacher bot X with reinforcement learning on full conversations, as a practical approach, we train a **Student bot** Y to carry on conversations with X . Student bot Y takes the conversation history $H_n^Y = \{U_1^X, U_1^Y, \dots, U_n^X\}$ as input, and output U_n^Y . It adopts BART (Lewis et al., 2020) model.

2.2 Phase 1: Supervised Pre-Tuning

This phase trains Teacher bot X to initialize and carry out conversations based on a given passage P , and trains Student bot Y to respond appropriately to X . To this end, we pre-tune both X and Y on the *Wizard of Wikipedia* (WoW) dataset. WoW was chosen as the pre-tuning dataset because of its two characteristics: 1) **Open-domain:** WoW contains conversations on a broad range of topics and domains across Wikipedia, thus the pre-tuned Teacher

and Student bots have greater potentials to be successfully transferred to other domains during the fine-tuning stage; 2) Content-grounded: in WoW, the teaching bot’s utterances are grounded on passages, which is similar to our task. We present the gold passage to Teacher bot directly, though, different from the WoW’s original setting (Dinan et al., 2019) where Teacher bot searches a large corpus for supporting passages.

We optimize the maximum-likelihood objective for Teacher bot by minimizing the following loss:

$$L_{mle}^X = - \sum_{n=1}^N \sum_{m=1}^{M_n} \log(p(x_m | x_1, \dots, x_{m-1}, H_n^X, P))$$

where N is the total number of turns in the conversation, $\{x_1, \dots, x_{M_n}\}$ is Teacher bot’s response at the n_{th} turn, M_n is the number of words in U_n^X . The loss function for Student bot, L_{mle}^Y , is similar to L_{mle}^X , with the exception of not including a passage P as input.

$$L_{mle}^Y = - \sum_{n=1}^N \sum_{m=1}^{M_n} \log(p(y_m | y_1, \dots, y_{m-1}, H_n^Y))$$

where $\{y_1, \dots, y_{M_n}\}$ is Student bot’s response at the n_{th} turn.

2.3 Phase 2: Unsupervised Self-Play Fine-tuning

In this phase, we aim at improving Teacher bot’s ability to present informative and coherent conversations. This is achieved by reinforcement learning on Teacher bot with the help of Student bot, and could be applied in a novel target domain even where dialogue dataset is absent. We adopt a self-play approach, i.e. we let Teacher bot and Student bot chat with each other over a passage in the target domain to generate multiple turns of conversations. In this fine-tuning phase, we keep Student bot frozen, and reward Teacher bot when the generated conversation achieves higher scores. In order to reduce the variance of the gradient estimate, we apply self-critic reinforcement learning (Rennie et al., 2017). Specifically, at each turn, we let Teacher bot generate two separate utterances: 1) U^s , which is sampled from the model, i.e. $x_m^s \sim p(x | x_1^s, \dots, x_{m-1}^s)$, and 2) U^* , which is obtained by greedy search, i.e. $x_m^* = \arg \max_w p(x | x_1^*, \dots, x_{m-1}^*)$. We optimize

the model by minimizing the following RL loss:

$$L_{rl}^X = - \sum_{n=1}^N (R(U_n^s) - R(U_n^*)) \sum_{m=1}^{M_n} \log(p(x_m^s | x_1^s, \dots, x_{m-1}^s, H_n^X, P))$$

where $R()$ is the reward function, which we will cover in Section 2.4, and P is a passage from the target domain corpus.

If not taking into account language modeling, optimizing RL loss alone would lead Teacher bot to generate inarticulate and even grammatically incorrect utterances. To keep the fluency of Teacher bots, we optimize a combined loss L^X consisting of RL loss L_{rl}^X on the new target domain data and MLE loss L_{mle}^X on the pre-tuning dialogue dataset, so the language style acquired during the pre-tuning phase would not get lost during RL fine-tuning:

$$L^X = \gamma L_{rl}^X + (1 - \gamma) L_{mle}^X$$

where $\gamma \in (0, 1)$ is a scaling factor accounting for the emphasis on L_{rl}^X and L_{mle}^X . We note that while L_{mle}^X should be obtained on an annotated content-grounded dialogue dataset (e.g. WoW), L_{rl}^X could be obtained on any target domain passage corpus even without dialogue data. This enables our approach to be transferred to an unsupervised text corpus.

2.4 Reward Functions

2.4.1 Coverage

We define the coverage reward of a Teacher bot’s utterance U^X as:

$$R_{cov} = \text{ROUGE}_1(P, H + U^X) - \text{ROUGE}_1(P, H)$$

where $\text{Rouge}_1(P, H)$ is the Rouge-1 F1 (Lin, 2004) score of the conversation history H to the input passage P . Intuitively, this function favors utterances that cover more information in the passage and have less overlap with the conversation history.

2.4.2 Coherence

Dialogue coherence datasets We explore neural coherence scoring models trained on two open-domain dialogue coherence datasets:

1. *WoW-coherence dataset* We reuse the WoW dataset to heuristically build a dialogue coherence classification dataset. Specifically, for each multi-turn dialogue in WoW, we label the ground truth response to its conversation history as *coherent* response, and all later responses in the same dialogue as *incoherent* responses.

2. *InferConv dataset* (Dziri et al., 2019) This is an open-domain dialogue coherence classification dataset built from PersonaChat conversational data (Zhang et al., 2018b). The dataset casts a response as the hypothesis and the conversation history as the premise, thus convert dialogue coherence evaluation into an NLI task. The dataset classifies the relationship between the response and the conversation history into three categories: *entailed*, *neutral* and *contradict*. Table 2 summarizes statistics of these datasets.

Coherence scoring models Based on the same pre-trained model BERT (Devlin et al., 2019), we train two different coherence scoring models on the two dialogue coherence classification datasets respectively. Both models take the concatenation (with [SEP]) of the conversation history and a candidate response as input, and minimize the cross entropy loss between the predicted label and the gold label. We use different methods to attain the coherence reward R_{coh} from the two models.

For model *WoW-coherence*, we define the coherence reward with softmax:

$$R_{coh} = \frac{e^{o_c}}{e^{o_c} + e^{o_i}}$$

where o_c and o_i are the logits for *coherent* and *incoherent* labels in the output layer. For model *InferConv*, we observe some responses labeled as *neutral* are appropriate responses but are not closely related to conversation history (e.g. “*That’s interesting!*”), we thus heuristically assign constant scores s_e , s_n and s_c as coherence reward R_{coh} when the response is predicted as *entailed*, *neutral* and *contradict*. In the remainder of the paper, we use *WoW-coherence* as the default coherence model, and compare it with *InferConv* in Section 4.

2.4.3 Mixed Reward

The coverage and coherence rewards are combined with a hyper-parameter β , yielding the final reward:

$$R = \beta R_{cov} + (1 - \beta) R_{coh}$$

3 Experimental Settings

We proceed by describing our datasets, comparison systems and evaluation metrics. We then show the performance of our proposed approach compared to state-of-the-art in §4.

Pre-Tuning	WoW-Train	WoW-Valid	
#Utterances	166,787	17,715	
#Dialogues	18,430	1,948	
#Words/utterance	16.6	16.6	
#Words/passage	110.3	109.8	
Fine-Tuning	Wikipedia	CNN-DM	Paper Abs.
#Passages	50,000	50,000	22,512
#Words/passage	111.7	129.8	149.3
Test Set	Wikipedia	CNN-DM	Paper Abs.
#Passages	1,000	1,000	500
#Words/passage	112.7	129.9	148.1

Table 1: Datasets statistics.

3.1 Datasets

Wizard of Wikipedia (Dinan et al., 2019) contains a total of 22,311 human-human conversations crowd-sourced via Amazon’s Mechanical Turk. The conversations are grounded in Wikipedia passages covering a wide range of topics: *e-book*, *toga party*, *armadillo*, etc. Both Teacher and Student bots are pre-tuned on the *WoW* dataset during Phase 1. Different from *WoW*’s original setting, we present the gold passage to the Teacher bot directly, instead of searching a large corpus for supporting passages. This allows us to focus less on retrieval and more on creating a Teacher bot to deliver informative and attentive dialogues.

We consider knowledge sources of various sorts as Teacher bot’s target domain during fine-tuning. CNN/DailyMail contains a large collection of online news articles with an average of 781 tokens per article (See et al., 2017). The full content of the article cannot be conveyed in a short conversation. Thus, we use the first 130 tokens of each article as a supporting passage, assuming it covers the most important content of the news article.

Academic papers have become an omnipresent source of knowledge. We create our own dataset containing papers published in recent years (2017–2021) at major venues, including ACL, EMNLP, NAACL, EACL, Findings and ICLR conferences. Similarly, we use paper abstracts as supporting passages instead of full articles. Moreover, we include Wikipedia passages from the *WoW* dataset, without conversations, as another source of knowledge. The *CNN-DM*, *Paper Abstracts* and *Wikipedia* datasets are used in Phase 2 of unsupervised self-play fine-tuning. Statistics of these datasets are summarized in Table 1.

3.2 Comparison Models

Our baseline Teacher bot builds on the state-of-the-art content-grounded dialogue generation model:

WoW-Coh	Coherent Resp.		Incoherent	All
Train	74,092	233,142	307,234	
Valid	3,939	12,362	16,301	
Test	3,865	12,098	15,963	

InferConv	Entail	Neutral	Contradict	All
Train	218,181	579,434	261,984	1,059,599
Valid	28,072	12,242	9,780	50,094
Test	3,119	1,087	1,360	5,566

Table 2: Details of the datasets used to train our neural coherence scoring models.

Coherence	
3	TeacherBot provides coherent responses to the evaluator’s input.
2	TeacherBot provides largely coherent responses (with minor coherency issues) to the evaluator’s input.
1	TeacherBot does not respond properly to the evaluator’s input.
Readability	
3	TeacherBot’s responses are easy to read, containing no grammatical or semantic errors.
2	TeacherBot’s responses read smoothly but may contain 1-2 grammatical or semantic errors.
1	TeacherBot’s responses contain >2 grammatical or semantic errors, or are nonsensical.
Overall Quality	
An initial score of 3 is given to a dialogue, then 1 point is deducted for each of the following issues, with a minimum score of 0.	
<ul style="list-style-type: none"> • Uninformative, i.e. it provides <2 correct answers during QA. • Incoherent, i.e. the average coherence score is <2 points. • Low readability, i.e. the average readability score is <2.5 points. • Any other issues that could lead to an ineffective conversation e.g. words are repeated between turns. 	

Table 3: A scoring rubric provided to human evaluators.

DOHA (Prabhumoye et al., 2021). It includes two improvements to the architectures of pre-trained encoder-decoder models (Lewis et al., 2020): building context-driven representation of the supporting document, and enabling document-headed attention to acquire information from the document. DOHA has demonstrated strong performance in document-grounded generation. All DOHA models are pre-tuned on the WoW dataset.

Our FULL Teacher bot is created to converse in an informative and coherent manner. It extends DOHA by incorporating both coverage and coherence rewards in unsupervised self-play fine-tuning. Additionally, we ablate FULL model by removing each of the two rewards: +COV uses only the coverage reward for fine-tuning, i.e. setting $\beta = 1$ in our reward function (§2.4). +COH utilizes only the WoW-coherence reward, i.e. setting $\beta = 0$. Please refer to appendix for more implementation details and hyper-parameters.

3.3 Evaluation Metrics

We investigate a wide range of metrics to evaluate Teacher bot’s performance. *Objective metrics* measure the content coverage and coherence of Teacher bot’s utterances. *Subjective metrics*, devised with human-in-the-loop, provide a holistic evaluation of a conversation, focusing on its overall effectiveness

and various aspects of linguistic quality.

Objective Metrics. Teacher bot converses with Student bot over a passage for three turns. That is, Teacher bot initiates the dialogue and provides two responses to Student bot. We objectively evaluate Teacher bot’s utterances in terms of information coverage and coherence as follows.

- ROUGE (Lin, 2004) is one of the most widely used metrics for measuring information coverage. We consider three variants in this study: R-1, R-2 and R-L, which respectively measure the overlap of unigrams, bigrams and the longest common subsequence between the given passage and Teacher bot’s utterances.

- QA_{F1} and QA_{CONF} are two variants of SummaQA (Scialom et al., 2019), a question answering-based evaluation metric. If a conversation is rich in information, it could be used as a surrogate for the passage to answering important questions. To this end, SummaQA generates Cloze-style questions from a passage by masking out entities, then employs a QA system to answer those questions based on a conversation. A higher QA performance suggests the conversation has better coverage. Particularly, QA_{F1} reports the F1 score for question answering; QA_{CONF} measures the confidence of the QA system in predicting answers.

- *WoW-Coherence* and *InferConv* are neural coherence scoring models (§2.4.2) repurposed for evaluation. These models quantitatively assess if Teacher bot has produced a coherent response given the conversation history, or not.

- DPR_{RELV} provides a new perspective on dialogue coherence evaluation (Zhang et al., 2018a). It draws on the Dense Passage Retriever model (DPR; Karpukhin et al., 2020) to predict if a Teacher bot’s response is relevant to Student bot’s input. A higher relevance score means the input and response share the same topic, suggesting a coherent conversation.

Subjective Metrics. We recruit 24 human evaluators to interact with Teacher bots. Each evaluator is asked to converse with bots over four passages. For each passage, the evaluator chats with four different Teacher bots for three turns, where Teacher bot initiates the conversation and responds twice to the evaluator’s input. We randomly select 48 passages for evaluation, i.e., 16 passages from each of the three test sets. To evaluate conversations produced from *Paper Abstracts*, we require evaluators, 8 in total, to be either PhD students or have

Dataset	Model	Coverage Metrics					Coherence Metrics			Subjective Metrics				Avg Len
		R-1	R-2	R-L	QA _{CONF}	QA _{F1}	WoW-Coh	InferCohv	DPR _{RELv}	QA _{HUMAN}	Coh	Read	Overall	
Wikipedia	DoHA	48.09	41.27	44.31	19.76	19.39	0.503	0.550	0.555	30.0	2.07	2.79	1.84	15.51
	+COV	74.62	71.66	72.25	30.90	34.38	0.307	0.466	0.543	45.63	1.85	2.85	2.12	28.11
	+COH	44.87	35.87	39.64	18.96	17.17	0.807	0.694	0.578	33.12	2.42	2.81	2.34	17.11
	FULL	62.78	58.83	60.61	25.94	27.01	0.617	0.630	0.556	38.74	2.26	2.88	2.37	20.69
CNN-DM	DoHA	38.89	30.09	32.91	15.90	15.80	0.521	0.567	0.538	28.12	2.28	2.67	1.87	16.29
	+COV	81.52	78.52	73.73	30.98	38.46	0.253	0.381	0.525	58.59	2.14	2.79	2.37	36.08
	+COH	30.45	18.61	24.0	13.31	11.31	0.845	0.692	0.561	40.15	2.40	2.48	1.93	16.23
	FULL	65.77	60.45	57.83	25.5	30.36	0.604	0.692	0.559	52.5	2.57	2.71	2.53	29.76
Papers	DoHA	36.27	28.20	30.60	10.29	5.34	0.565	0.452	0.557	30.5	1.58	2.55	1.4	15.61
	+COV	72.63	69.69	49.18	20.32	17.61	0.271	0.141	0.529	57.5	1.73	2.77	1.96	33.55
	+COH	32.96	21.24	26.21	8.18	3.95	0.806	0.547	0.576	25.47	1.82	2.26	1.25	16.3
	FULL	59.65	54.22	47.88	15.46	13.81	0.766	0.501	0.560	51.76	2.16	2.53	2.09	27.37

Table 4: We compare Teacher bots based on naive DoHA(Prabhumoye et al., 2021) model to variants fine-tuned using different reward functions, Avg len refers to average utterance length.

obtained a PhD degree. For fair comparison, we shuffle and hide the order of Teacher bots presented to evaluators. Human evaluators were suggested to feed the same or similar inputs across Teacher bots on the same passage whenever possible. Throughout the conversation, the passage was not shown to the evaluators. After the conversation, human evaluators were asked to complete the following evaluation tasks:

- **QA_{HUMAN}**: Five sentences are randomly selected from each passage and one important entity is masked out in each sentence. The evaluators are presented with each corrupted sentence and asked if the sentence could be recovered by referencing the conversation with Teacher bot. We report the ratio of sentences that could be correctly recovered.

- **Linguistic Quality**: We ask human evaluators to rate each conversation along three dimensions: *Coherence*: Does Teacher bot provide coherent responses to the evaluator’s input? *Readability*: Are Teacher bot’s utterances easy to read, containing no grammatical or semantic errors? *Overall Quality*: How will the conversation score in terms of informativeness, coherence, readability and all aspects considered? The scoring rubric provided to human evaluators is shown in Table 3.

4 Objective Results

Results on Test Sets. Table 4 presents objective evaluation results obtained for various Teacher bots on three test sets: *CNN-DM*, *Paper Abstracts* and *Wikipedia*. We observe that our FULL Teacher bot is able to substantially outperform the baseline system DOHA on all datasets and across all objective metrics. It strikes a fine balance between delivering information-rich conversations and ensuring those conversations are coherent and attentive. Further, we find that optimizing a single reward, whether it be coverage or coherence, produces suboptimal results. For instance, +COV tends to produce longer




Figure 3: ROUGE gain and utterance length tend to decrease as the number of turns increases

utterances than other variants. It improves information coverage, but yields low coherence scores, leading to a performance even inferior to the baseline DOHA. Our findings suggest that it is important for the reinforcer L_{rl}^X to learn with both coverage or coherence rewards.

Trading off Coverage for Coherence. In Figure 4, we plot the learning curves of coverage and coherence scores when the reinforcer adopts a single reward (+COV, +COH) or both (FULL). We use R_{cov} and R_{coh} to approximate coverage and coherence scores. These plots are generated using 50 validation instances from the *Paper Abstracts* dataset. We observe that with only the coverage reward (+COV), Teacher bot tends to aggressively copy content from the passage, while disregarding the conversation history. This inevitably leads to incoherent conversations. Conversely, +COH can improve on coherence, but falls short on delivering informative conversations. Finally, our FULL Teacher bot trades off coverage for substantially higher coherence, thus achieving a significant improvement over the baseline DOHA model.

Trading off Coverage for Coherence. We are curious to know the amount of information brought by each utterance produced by Teacher bot. To this end, we define information gain $\mathcal{IG}(\cdot)$ as the

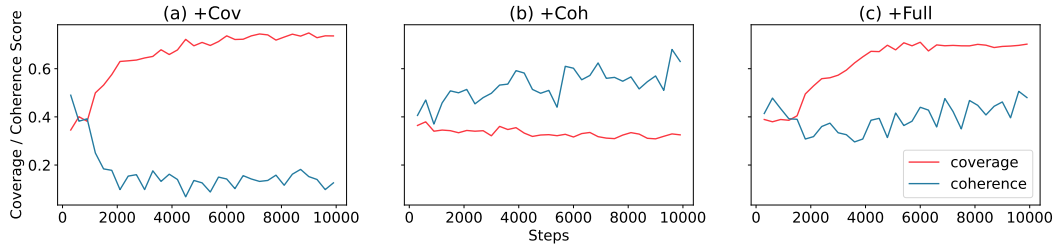


Figure 4: Coverage and coherence scores when applying different reward scores

Dataset	Model	Coverage		Coherence		
		QA _{CONF}	QA _{F1}	WoW-Coh	IC	DPR
Wikipedia	WoW	18.96	17.17	0.807	0.694	0.578
	InferCov	20.41	19.61	0.647	0.676	0.561
CNN-DM	WoW	13.31	11.31	0.845	0.692	0.561
	InferCov	14.95	14.07	0.719	0.728	0.550
Papers	WoW	8.18	3.95	0.806	0.547	0.576
	InferCov	9.47	4.31	0.737	0.554	0.558

Table 5: Teacher bot’s performance when using different coherence score models during fine-tuning

improvement of ROUGE scores brought by an utterance U :

$$IG(U) = \text{ROUGE}(P, H + U) - \text{ROUGE}(P, H),$$

where P is the supporting passage, and H represents the conversation history. We consider three ROUGE variants, R-1, R-2 and R-L, respectively. Figure 3 illustrates the gain of information for each of the three turns. The average ROUGE gain is reported for each turn, using conversations produced for the *Paper Abstracts* dataset. We observe that there is a general tendency across turns that information gain is decreasing. This is in part because that at the beginning of a conversation, Teacher bot has no constraints regarding content selection, it could rephrase any content selected from the supporting passage to initiate a dialogue. In subsequent turns, Teacher bot has to exercise caution in response generation considering both the conversation history and overall coherence of the conversation. Consequently, we find that the average length of the utterances also decreases in subsequent turns.

A Comparison of Coherence Scoring Models. We compare two Teacher bots fine-tuned only with coherence reward from different coherence scoring models (i.e. *WoW-coherence* and *InferConv*). We demonstrate their objective results in Table 5. According to the results, Teacher bot fine-tuned with *InferConv* achieves slightly better coverage metrics. However, in terms of coherence metrics, Teacher bot fine-tuned with *WoW-coherence* model generally achieves better performance. Based on this

observation, *WoW-coherence* scoring model better measures coherence in conversations.

5 Subjective Results

We demonstrate the subjective evaluation results in Table 4 and have the following observations:

1. For question answering, +COV achieves the best performance on all three datasets. This again proves that the coverage reward helps make the conversation more informative;
2. For coherence scores, +COH achieves the best performance on *Wikipedia*. However, on the other two datasets it was outperformed by FULL.
3. For readability scores, on *CNN-DM* and *Paper Abstracts*, +COH achieves the lowest performance while +COV achieves the highest.
4. For overall scores, FULL demonstrates the best performance. This suggests FULL delivers conversations that are more balanced in coverage, coherence and readability.
5. The *Paper Abstracts* corpus is the most challenging among all the test corpora, as our Teacher bots generally show worse performance in coherence, readability and overall scores. We found passages in *Paper Abstracts* contain volumes of professional vocabularies thus are more complicated for people to understand. In addition, it’s also more difficult for Student bot to respond appropriately during self-play fine-tuning (See examples in [appendix](#)). As a result, transferring Teacher bots to this domain is more challenging.

A Case Study. We show a few Teacher bots’ responses to users in Table 6. After analyzing the cases, we have the following observations:

1. The coverage reward encourages Teacher bots to directly copy content from the input passage,

Example A (Passage from Wikipedia)	
Passage	The Concorde is a British–French turbojet-powered supersonic passenger airliner ... It had a maximum speed over twice the speed of sound, at Mach 2.04, ... First flown in 1969, Concorde entered service in 1976 and operated for 27 years. ...
User	What is the plane famous for?
Bot Cov	First flown in 1969, Concorde entered service in 1976 and operated for 27 years.
Bot Coh	Concorde is famous for the speed of sound at Mach 2.04.
Bot Full	Concorde is famous for having a maximum speed over twice the speed of sound, at Mach 2.04.
Example B (Passage from CNN News)	
Passage	The US Food and Drug Administration, ... Molnupiravir is for the treatment of mild-to-moderate coronavirus disease in adults with positive results of direct SARS-CoV-2 viral testing ... This is the second Covid-19 antiviral pill authorized for ill people to take at home, before they get sick enough to be hospitalized ...
User	Tell me more about the drug.
Bot Cov	This is the second Covid-19 antiviral pill authorized for ill people to take at home, before they get sick enough to be hospitalized.
Bot Coh	The drug is for people with mild to moderate coronavirus disease.
Bot Full	Molnupiravir is for the treatment of mild-to moderate coronavirus disease in adults with positive results of direct SARS-CoV-2 viral testing.
Example C (Passage from Paper Abstracts)	
Passage	... In this paper we describe a novel framework to guide a supervised, extractive summarization system with question-answering rewards ... Our results compare favorably with those reported by strong summarization baselines as evaluated by automatic metrics and human assessors ...
User	What did you propose?
Bot Cov	Our results compare favorably with those reported by strong summarization baselines as evaluated by automatic metrics and human assessors.
Bot Coh	A proposal to use the extractive summarizer to answer questions.
Bot Full	We propose a novel framework to guide a supervised, extractive summarization system with question-answering rewards.

Table 6: A few example responses from different Teacher bots. In example A, +COH and FULL abstractively generates *Concord is famous for ...* to make the response more coherent to user’s question. However, the underlined part in +COH’s statement misrepresent the input passage and is inaccurate. In example B, all responses seem coherent because of the open-ended question.

while the coherence reward encourages abstractively generating new content: As shown in Example A, +COV directly extracts a part of the original passage as response, regardless of the user’s question, while +COH and FULL abstractively rewrite the response to make it more coherent.

2. Putting too much weight on coherence reward could make Teacher bot become so abstractive that it misrepresents the original passage and lead to incoherence and semantic/grammar errors. (See +COH’s response in example A and C) This explains +COH’s low coherence and readability scores on *CNN-DM* and *Paper Abstracts*. This observation suggests the necessity to carefully choose the weight for coherence rewards, and to coupling coherence reward with coverage reward, which could make the chat bot less abstractive.

3. Generally, user utterances could be classified into two categories: **Information-seeking queries** which request certain information (e.g. the user’s utterance in example A); **Open statements** which do not have specific requests (e.g. the user’s utterance in example B). We found evaluators tend to give high coherence scores to response to open statements, as they could be addressed by a wider range of responses.

6 Related Work

Content-Grounded Dialogue Generation Content grounded dialogue generation is the task of using the information provided in external content (e.g. a passage, etc.) to guide dialogue generation. Compared to previous research, our task has the following novelties. 1) Compared to content-grounded information-retrieval-oriented dialogue such as doc2dial (Feng et al., 2020) and ABCD (Chen et al., 2021) where the chat bot responds to user query in a passive way, we expect our chat bots to convey knowledge proactively. 2) Compared to chit chat-oriented dialogue such as (Zhou et al., 2018; Dinan et al., 2019; Komeili et al., 2021; Xu et al., 2021), our task is more focused on extensive conversation in a particular topic, and aims at helping the end user acquire knowledge or information from a given passage. 3) Contrasted to chat bots that are applied in a single domain (Zhou et al., 2018; Moghe et al., 2018; Xu et al., 2019), our chat bot could be transferred to other domains through self-talk based fine-tuning.

Another line of research works focus on content grounded text generation models (Prabhumoye et al., 2021; Zhao et al., 2020). Compared with ordinary text generation models (e.g. BART (Lewis et al., 2020)), these models are specifically designed to model external content as an additional input, and achieve better performance on content grounded dialogue generation tasks including *CMU DoG* (Zhao et al., 2020) and *Wizard of Wikipedia* (Dinan et al., 2019).

RL in Text Generation Reinforcement learning has been applied in various natural language generation tasks, including image caption (Rennie et al., 2017), automatic summarization (Paulus et al., 2018), machine translation (Kang et al., 2020) and poem generation (Yang et al., 2019). Specifically, when applying reinforcement learning in dialogue generation (Li et al., 2016; Zhao et al., 2019; Shi et al., 2019; Yamazaki and Aizawa, 2021; Liu et al.,

2020), self-play is often used to enable scoring multi-turn dialogues. Compared to previous dialogue generation research using RL and self-play, our two-phase framework enables transferring the teacher bot to other domains by optimizing a mixed reward of coverage and coherence.

Educational Dialogue Systems and Conversational QA There have been research works applying dialogue systems for educational purposes. Some chat bots are for language practice. (Ruan et al., 2021; Stewart and File, 2007; Huang et al., 2017) Others are specially designed for education in a single domain or task, e.g. moral education (Peng et al., 2019), educational debate (Yuan et al., 2008). Compared with previous educational dialogue systems, our system is for information acquisition without domain restriction. Our task is also related to conversational question answering (CQA), e.g. (Kočiský et al., 2018; Rajpurkar et al., 2016; Joshi et al., 2017; Zellers et al., 2018). However, most existing CQA systems passively response to user queries in single turn conversations, while our system actively engage with users in multi-turn conversations.

7 Conclusion

We propose an information-acquisition-oriented dialogue system that transfers information and knowledge in passages to users through conversations. An unsupervised self-talk approach is introduced leveraging novel rewards to enable Teacher bots to deliver informative and attentive conversations. Experiments with automatic and human evaluations demonstrate the effectiveness of our approach. Some interesting future directions include extending the conversations to be based on a set of documents and specializing our dialogue systems for in specific domain, e.g. patient education.

8 Ethical Considerations

Our models are pre-tuned on *Wizard of Wikipedia* dataset and fine-tuned on three corpora: *Wikipedia*, *CNN-DailyMail* and *Paper-Abstracts* (Abstracts of papers from ACL, EMNLP, NAACL, EACL, Findings and ICLR submissions from 2017 to 2021). All the datasets used in this paper are publicly available. Moreover, we did not use full-length Wikipedia or CNN-Daily Mail news articles in our experiments, but tailored versions of 100-150 words. This is because a full length Wikipedia/CNN-Daily Mail article may contain too much content to be covered in a short conversation.

As described in (Maynez et al., 2020; Kryscinski et al., 2020; Lebanoff et al., 2020; Zhou et al., 2021), current state of the art neural conditional text models can output hallucinated content unfaithfully to the input text, which impedes the safe deployment of the models. We note that our Teacher bots may also generate utterances that are not supported by the input passage.

References

- Derek Chen, Howard Chen, Yi Yang, Alexander Lin, and Zhou Yu. 2021. [Action-based conversations dataset: A corpus for building more in-depth task-oriented dialogue systems](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3002–3017, Online. 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, Minneapolis, Minnesota. Association for Computational Linguistics.
- John Dewey. 1923. *Democracy and education: An introduction to the philosophy of education*. macmillan.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of Wikipedia: Knowledge-powered conversational agents. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Nouha Dziri, Ehsan Kamaloo, Kory Mathewson, and Osmar Zaiane. 2019. [Evaluating coherence in dialogue systems using entailment](#). 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 3806–3812, Minneapolis, Minnesota. Association for Computational Linguistics.
- Song Feng, Hui Wan, Chulaka Gunasekara, Siva Patel, Sachindra Joshi, and Luis Lastras. 2020. [doc2dial: A goal-oriented document-grounded dialogue dataset](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8118–8128, Online. Association for Computational Linguistics.
- Roberta Michnick Golinkoff, Erika Hoff, Meredith L Rowe, Catherine S Tamis-LeMonda, and Kathy Hirsh-Pasek. 2019. Language matters: Denying the existence of the 30-million-word gap has serious consequences. *Child development*, 90(3):985–992.
- Jin-Xia Huang, Kyung-Soon Lee, Oh-Woog Kwon, and Young-Kil Kim. 2017. [A chatbot for a dialogue-based second language learning system](#), pages 151–156.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Xiaomian Kang, Yang Zhao, Jiajun Zhang, and Chengqing Zong. 2020. [Dynamic context selection for document-level neural machine translation via reinforcement learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2242–2254, Online. Association for Computational Linguistics.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. [The NarrativeQA reading comprehension challenge](#). *Transactions of the Association for Computational Linguistics*, 6:317–328.
- Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2021. [Internet-augmented dialogue generation](#). *CoRR*, abs/2107.07566.
- Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. [Evaluating the factual consistency of abstractive text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9332–9346, Online. Association for Computational Linguistics.
- Logan Lebanoff, John Muchovej, Franck Dernoncourt, Doo Soon Kim, Lidan Wang, Walter Chang, and Fei Liu. 2020. [Understanding points of correspondence between sentences for abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 191–198, Online. Association for Computational Linguistics.
- Rosemary Lever and Monique Sénéchal. 2011. Discussing stories: On how a dialogic reading intervention improves kindergartners’ oral narrative construction. *Journal of experimental child psychology*, 108(1):1–24.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [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*, pages 7871–7880, Online. Association for Computational Linguistics.

- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016. [Deep reinforcement learning for dialogue generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Austin, Texas. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Qian Liu, Yihong Chen, Bei Chen, Jian-Guang Lou, Zixuan Chen, Bin Zhou, and Dongmei Zhang. 2020. [You impress me: Dialogue generation via mutual persona perception](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1417–1427, Online. Association for Computational Linguistics.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Nikita Moghe, Siddhartha Arora, Suman Banerjee, and Mitesh M. Khapra. 2018. [Towards exploiting background knowledge for building conversation systems](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2322–2332, Brussels, Belgium. Association for Computational Linguistics.
- Suzanne E Mol, Adriana G Bus, Maria T De Jong, and Daisy JH Smeets. 2008. Added value of dialogic parent–child book readings: A meta-analysis. *Early education and development*, 19(1):7–26.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. [A deep reinforced model for abstractive summarization](#). In *International Conference on Learning Representations (ICLR)*.
- Yan Peng, Penghe Chen, Yu Lu, Qinggang Meng, Qi Xu, and Shengquan Yu. 2019. A task-oriented dialogue system for moral education. In *International Conference on Artificial Intelligence in Education*, pages 392–397. Springer.
- Shrimai Prabhumoye, Kazuma Hashimoto, Yingbo Zhou, Alan W Black, and Ruslan Salakhutdinov. 2021. [Focused attention improves document-grounded generation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4274–4287, Online. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. [Self-critical sequence training for image captioning](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1179–1195.
- Sherry Ruan, Liwei Jiang, Qianyao Xu, Zhiyuan Liu, G. M. Davis, Emma Brunskill, and J. Landay. 2021. Englishbot: An ai-powered conversational system for second language learning. *26th International Conference on Intelligent User Interfaces (IUI)*.
- Thomas Scialom, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2019. [Answers unite! unsupervised metrics for reinforced summarization models](#). 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 3246–3256, Hong Kong, China. Association for Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Mike Sharples. 2005. *Learning as conversation transforming education in the mobile age*. Citeseer.
- Weiyang Shi, Kun Qian, Xuewei Wang, and Zhou Yu. 2019. [How to build user simulators to train RL-based dialog systems](#). 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 1990–2000, Hong Kong, China. Association for Computational Linguistics.
- Iain Stewart and Portia File. 2007. [Let’s chat: A conversational dialogue system for second language practice](#). *Computer Assisted Language Learning*, 20.
- Hu Xu, Bing Liu, Lei Shu, and Philip S Yu. 2019. Review conversational reading comprehension. *arXiv preprint arXiv:1902.00821*.
- Jing Xu, Arthur Szlam, and Jason Weston. 2021. [Beyond goldfish memory: Long-term open-domain conversation](#). *CoRR*, abs/2107.07567.
- Takato Yamazaki and Akiko Aizawa. 2021. [Phrase-level action reinforcement learning for neural dialog response generation](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 5028–5038, Online. Association for Computational Linguistics.

- Zhichao Yang, Pengshan Cai, Yansong Feng, Fei Li, Weijiang Feng, Elena Suet-Ying Chiu, and Hong Yu. 2019. [Generating classical Chinese poems from vernacular Chinese](#). 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 6155–6164, Hong Kong, China. Association for Computational Linguistics.
- Tangming Yuan, David Moore, and Alec Grierson. 2008. A human-computer dialogue system for educational debate: A computational dialectics approach. *International Journal of Artificial Intelligence in Education*, 18(1):3–26.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. [SWAG: A large-scale adversarial dataset for grounded commonsense inference](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104, Brussels, Belgium. Association for Computational Linguistics.
- Hainan Zhang, Yanyan Lan, Jiafeng Guo, Jun Xu, and Xueqi Cheng. 2018a. Reinforcing coherence for sequence model in dialogue generation. In *IJCAI*, pages 4567–4573.
- 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.
- Tiancheng Zhao, Kaige Xie, and Maxine Eskenazi. 2019. [Rethinking action spaces for reinforcement learning in end-to-end dialog agents with latent variable models](#). 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 1208–1218, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xueliang Zhao, Wei Wu, Chongyang Tao, Can Xu, Dongyan Zhao, and Rui Yan. 2020. [Low-resource knowledge-grounded dialogue generation](#). In *International Conference on Learning Representations*.
- Chunting Zhou, Graham Neubig, Jiatao Gu, Mona Diab, Francisco Guzmán, Luke Zettlemoyer, and Marjan Ghazvininejad. 2021. [Detecting hallucinated content in conditional neural sequence generation](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1393–1404, Online. Association for Computational Linguistics.
- Kangyan Zhou, Shrimai Prabhumoye, and Alan W Black. 2018. [A dataset for document grounded conversations](#). In *Proceedings of the 2018 Conference*
- on Empirical Methods in Natural Language Processing*, pages 708–713, Brussels, Belgium. Association for Computational Linguistics.

9 Appendix

9.1 Implementation Details

9.1.1 Pre-tuning

Key hyper-parameters for both models.

- source max len: 1024
- target max len: 128
- batch size: 8
- train epoch: 3
- learning rate: $2e-5$
- Both Teacher and Student bots adopt the initialized weights of `bart-base` from hugging-face.
- Both DoHA models and BART models are based on the implementation presented in the DoHA paper (Prabhumoye et al., 2021).

Note we initialize our models with `bart-base` instead of `bart-large` as our self-play fine-tuning is very computational intensive and time consuming. With our current setting, the self-play fine-tuning takes about 2.5 days on one single NVIDIA Tesla V100 GPU.

9.1.2 Self-Play Reinforced Fine-Tuning

Mixed Loss. Optimizing our mixed loss is achieved by intermittently minimizing the MLE loss for a batches and then minimizing the RL loss for b batches. In other word, the parameter γ in Section 2.3 is determined by a and b . Specifically, in each fine-tuning step, the parameters of Teacher bot are updated for 4 times, once over an RL batch, 3 times over MLE batches.

Key Hyper-Parameters. We apply the same set of basic hyper-parameters for all Teacher bots during all fine-tuning process:

- MLE batch size: 8
- RL batch size: 5
- maximum coverage score: 0.5
- train epoch: 3
- learning rate: $1e-6$
- Fine-tuning steps: 10,000
- β for three reinforced Teacher bots: FULL- $\beta = 0.7$, +COV- $\beta = 1.0$, +COH- $\beta = 0.0$

9.2 Coherence Scoring Models

Both *WoW-coherence* and *InferConv* are trained based on `bert-base-cased`.

Key hyper-parameters for both models.

- max length: 256
- batch size: 32
- learning rate $2e-5$
- epochs 3

For *InferConv* classification model, the constant scores are $s_e = 1.0$, $s_n = 0.2$ and $s_c = 0.0$

The *WoW-coherence* and *InferConv* classification model achieves 82.1% and 88.4% accuracy on respective test sets.

All other hyper-parameter settings for Teacher bot, Student bot and coherence scoring models are based on the system’s default setting. All our experiments were run on servers with Nvidia A100 and V100 GPUs.

9.3 Dataset Details

The passages in train/validation/test set of our *Wikipedia* corpus are randomly sampled from passages in train/validation/test set of *WoW* respectively. Similarly, passages in train/validation/test set of our *CNN-DM* corpus are randomly sampled from the train/validation/test set of the original *CNN-DM* dataset. For *Paper Abstracts*, we randomly distribute all collected paper abstracts into train/validation/test sets.

9.4 Rewards function Details

During our exploration, we have explored multiple variations of coverage reward functions:

1. Reward score is gained at the end of each turn of conversation, and it is calculated the ROUGE score improvement of the teacher bot’s utterance (As we applied in the paper);
2. Reward score is gained at the end of each turn of conversation, and it is calculated the ROUGE score improvement of both the teacher bot and student bot’s utterances;
3. Reward score is gained at the end of the entire conversation, and is calculated as the ROUGE score improvement of all teacher bot’s utterances.

According to our experience, the performance of 1 is similar than 2. Specifically, we notice in most cases, our student bot contribute marginally in terms of ROUGE score improvement. In addition, in practice, we observe 1 shows better performance than 3.

9.5 Human Evaluation Details

All human evaluation passages are collected randomly, the length of the *Wikipedia* and *CNN* passages are tailored to 100-150 words to conform to the length of passages in the fine-tuning stage. During conversations, we suggest our evaluators to use

utterances relevant to the topic, so that their utterances could be appropriately addressed by Teacher bots referencing content in the passage. All human evaluators we recruit have at least a bachelor degree, each evaluator is rewarded with a \$20 gift card for participation.

The evaluation results are collected using Google Colab, an example evaluation page is available [through this link](#) (Personal information are anonymized during reviewing stage).

Note that in our human evaluation, Teacher bot chat about a passage with two evaluators. We use Pearson correlation coefficient to measure Inter-annotator agreement, the average Pearson correlation coefficient among each pair of evaluators is 0.16, implying weakly positive correlation. We note different evaluators use different utterances to chat with the same Teacher bot, thus it is reasonable that the conversation with the same Teacher bot over the same passage have large variance.

9.6 Examples

We provide more examples of conversations in self-play in Table 7 and human evaluations in Table 8.

Example D (Passage from Wikipedia)	
Passage	For 4,000 years China has been a nation of farmers. By the time the People’s Republic of China was established in 1949, virtually all arable land was under cultivation; irrigation and drainage systems constructed centuries earlier and intensive farming practices already produced relatively high yields. But little prime virgin land was available to support population growth and economic development. However, after a decline in production as a result of the Great Leap Forward (1958–60), agricultural reforms implemented in the 1980s increased yields and promised even greater future production from existing cultivated land.
Dialogue	<p>[Teacher bot]: China has been a nation of farmers for 4,000 years!</p> <p>[Student bot]: Wow, that’s a long time. What do they grow?</p> <p>[Teacher bot]: They grow a variety of crops, including chicken, fish, fruits, and vegetables.</p> <p>[Student bot]: That sounds like a lot of food. What else can you tell me about China?</p> <p>[Teacher bot]: Well, by the time the People’s Republic of China was established in 1949, virtually all arable land was under cultivation.</p>
Example E (Passage from CNN-DailyMail)	
Passage	Everywhere you look there are two-wheeled drivers waiting ominously for passengers to hop on board. They’re the backbone of public transport in Uganda and the fastest way to get around the capital city, Kampala. The boda boda motorbike taxi is a staple used by all sectors of society. Traditionally, the usual way of catching a ride was to venture into any street corner packed with tens of boda bodas or simply waiting for one to pass by. But now, a new local startup are bringing this classic mode of transport into the technological age by providing the ability to hail a boda boda at the click of a smartphone. "Bodas are the main thing getting people from A to B," says Alastair Sussock, co-founder of SafeBoda ...
Dialogue	<p>[Teacher bot]: Ugandan start-up SafeBoda are bringing this classic mode of transport into the technological age by providing the ability to hail a boda boda at the click of a smartphone.</p> <p>[Student bot]: I’ve never heard of them. What are they?</p> <p>[Teacher bot]: They are the backbone of public transport in Uganda and the fastest way to get around the capital city, Kampala.</p> <p>[Student bot]: That’s really cool. How long have they been around?</p> <p>[Teacher bot]: Well they have been around for a long time. Traditionally, the usual way of catching a ride was to venture into any street corner packed with tens of boda bodas or simply waiting for one to pass by.</p>
Example F (Passage from Paper Abstracts)	
Passage	Existing open-domain question answering (QA) models are not suitable for real-time usage because they need to process several long documents on-demand for every input query, which is computationally prohibitive. In this paper, we introduce query-agnostic indexable representations of document phrases that can drastically speed up open-domain QA. In particular, our dense-sparse phrase encoding effectively captures syntactic, semantic, and lexical information of the phrases and eliminates the pipeline filtering of context documents. Leveraging strategies for optimizing training and inference time, our model can be trained and deployed even in a single 4-GPU server. Moreover, by representing phrases as pointers to their start and end tokens, our model indexes phrases in the entire English Wikipedia (up to 60 billion phrases) using under 2TB. Our experiments on SQuAD-Open show that our model is on par with or more accurate than previous models with 6000x reduced computational cost, which translates into at least 68x faster end-to-end inference benchmark on CPUs ...
Dialogue	<p>[Teacher bot]: We introduce query-agnostic indexable representations of document phrases that can drastically speed up open-domain QA.</p> <p>[Student bot]: What is QA?</p> <p>[Teacher bot]: It is a process where our dense-sparse phrase encoding effectively captures syntactic, semantic, and lexical information of the phrases and eliminates the pipeline filtering of context documents.</p> <p>[Student bot]: What are the basic concepts of QA?</p> <p>[Teacher bot]: Well the basic concepts are that we represent phrases as pointers to their start and end tokens, our model indexes phrases in the entire English Wikipedia (up to 60 billion phrases) using under 2TB.</p>

Table 7: Conversation examples of Teacher bot fine-tuned with both rewards chat with Student bot during self-play. Example F demonstrates an example where Student bot’s question seems too broad to be responded using content in the passage.

Example G (Passage from Wikipedia)	
Passage	Interstellar is a 2014 epic science fiction film co-written, directed and produced by Christopher Nolan. It stars Matthew McConaughey, Anne Hathaway, Jessica Chastain, Bill Irwin, Ellen Burstyn, and Michael Caine. Set in a dystopian future where humanity is struggling to survive, the film follows a group of astronauts who travel through a wormhole near Saturn in search of a new home for mankind. Paramount Pictures, Warner Bros. Pictures, and Legendary Pictures co-financed the film. Cinematographer Hoyte van Hoytema shot it on 35 mm in the Panavision anamorphic format and IMAX 70 mm. Interstellar uses extensive practical and miniature effects and the company Double Negative created additional digital effects. Interstellar premiered on October 26, 2014, in Los Angeles, California. In the United States, it was first released on film stock, expanding to venues using digital projectors. The film had a worldwide gross of over \$677 million (and \$701 million with subsequent re-releases), making it the tenth-highest-grossing film of 2014.
Dialogue	<p>[Teacher bot]: Interstellar is a 2014 epic science fiction film directed and produced by Christopher Nolan.</p> <p>[Human user]: Who stars in the film?</p> <p>[Teacher bot]: The film stars Matthew McConaughey, Anne Hathaway, Jessica Chastain, Bill Irwin, Ellen Burstyn, and Michael Caine.</p> <p>[Human user]: What story does it tell?</p> <p>[Teacher bot]: The film follows a group of astronauts who travel through a wormhole near Saturn.</p>
Example H (Passage from CNN News)	
Passage	The US Food and Drug Administration on Monday granted full approval to the Pfizer/BioNTech Covid-19 vaccine for people age 16 and older. This is the first coronavirus vaccine approved by the FDA, and is expected to open the door to more vaccine mandates. The vaccine will be marketed as Comirnaty, the FDA said in its announcement on Monday. The Pfizer/BioNTech vaccine has been authorized for emergency use in the United States since mid-December for people age 16 and older, and in May, the authorization was extended to those 12 and older. FDA says it's working as fast as possible to fully approve vaccines, as urgency rises amid Covid surge FDA says it's working as fast as possible to fully approve vaccines, as urgency rises amid Covid surge ...
Dialogue	<p>[Teacher bot]: The Pfizer/BioNTech vaccine has been authorized for emergency use in the United States since mid-December for people age 16 and older.</p> <p>[Human user]: Give me more details.</p> <p>[Teacher bot]: This is the first coronavirus vaccine approved by the FDA, and is expected to open the door to more vaccine mandates.</p> <p>[Human user]: So, what will happen next?</p> <p>[Teacher bot]: The Pfizer/BioNTech Covid-19 vaccine for people age 16 and older will be marketed as Comirnaty, the FDA said in its announcement on Monday.</p>
Example I (Passage from Paper Abstracts)	
Passage	We introduce a new language representation model called BERT, which stands for Bidirectional Encoder Representations from Transformers. Unlike recent language representation models, BERT is designed to pre-train deep bidirectional representations from unlabeled text by jointly conditioning on both left and right context in all layers. As a result, the pre-trained BERT model can be fine-tuned with just one additional output layer to create state-of-the-art models for a wide range of tasks, such as question answering and language inference, without substantial task-specific architecture modifications. BERT is conceptually simple and empirically powerful. It obtains new state-of-the-art results on eleven natural language processing tasks, including pushing the GLUE score to 80.5% (7.7% point absolute improvement), MultiNLI accuracy to 86.7% (4.6% absolute improvement), SQuAD v1.1 question answering Test F1 to 93.2 (1.5 point absolute improvement) and SQuAD v2.0 Test F1 to 83.1 (5.1 point absolute improvement).
Dialogue	<p>[Teacher bot]: We introduce BERT, which stands for Bidirectional Encoder Representations from Transformers.</p> <p>[Human user]: What did you propose?</p> <p>[Teacher bot]: A new language representation model called BERT which is conceptually simple and empirically powerful.</p> <p>[Human user]: How is the result?</p> <p>[Teacher bot]: The pre-trained BERT model can be fine-tuned with just one additional additional output layer to create state-of-the-art models for a wide range of tasks, such as question answering and language inference, without substantial task-specific architecture modifications.</p>

Table 8: Conversation examples of Teacher bot fine-tuned with both rewards chat with human evaluators.