

Refine and Imitate: Reducing Repetition and Inconsistency in Persuasion Dialogues via Reinforcement Learning and Human Demonstration

Weiyan Shi¹, Yu Li², Saurav Sahay³, Zhou Yu¹

Columbia University¹ University of California, Davis², Intel Labs³

ws2634@columbia.edu, yooli@ucdavis.edu

saurav.sahay@intel.com, zy2461@columbia.edu

Abstract

Persuasion dialogue system reflects the machine’s ability to make strategic moves beyond verbal communication, and therefore differentiates itself from task-oriented or open-domain dialogues and has its own unique values. However, the repetition and inconsistency problems still persist in dialogue response generation and could substantially impact user experience and impede the persuasion outcome. Besides, although reinforcement learning (RL) approaches have achieved big success in strategic tasks such as games, it requires a sophisticated user simulator to provide real-time feedback to the dialogue system, which limits the application of RL on persuasion dialogues. To address these issues towards a better persuasion dialogue system, we apply RL to refine a language model baseline without user simulators, and distill sentence-level information about repetition, inconsistency, and task relevance through rewards. Moreover, to better accomplish the persuasion task, the model learns from human demonstration to imitate human persuasion behavior and selects the most persuasive responses. Experiments show that our model outperforms previous state-of-the-art dialogue models on both automatic metrics and human evaluation results on a donation persuasion task, and generates more diverse, consistent and persuasive conversations according to the user feedback. The code is available at <https://github.com/wyshi/consistency>.

1 Introduction

Persuasion dialogue systems have become an increasingly important subject in both social science and computational linguistics (Prakken, 2006, 2009; Wang et al., 2019; Asai et al., 2020). Such systems aim to employ conversational strategies to change the audience’s attitude or behaviour, and therefore, are inherently difficult to build with multiple challenges. The first one is that users often

expect highly smooth conversation experience from persuasion systems in order to be persuaded (Shi et al., 2020). So the long-standing problems of dialogue repetition and inconsistency can be especially salient in persuasion dialogue systems. Secondly, different from traditional dialogue tasks, the persuasion task is non-collaborative where the user and the system have different goals (Li et al., 2020b), and hence highly intellectual and strategic.

Previous studies have attempted to address the first challenge, the dialogue repetition and inconsistency problems, by changing the object function in supervised learning (Li et al., 2020a) or applying reinforcement learning (RL) (Li et al., 2016; Liu et al., 2018). But these methods either may lead to uninterpretable behaviors, or rely on hand-crafted user simulators that are hard to design for persuasion dialogues. To tackle these challenges, we propose to extract a policy directly from the data and let the models learn from its own mistakes without the use of simulators. Leveraging decoding methods such as Nucleus Sampling (Holtzman et al., 2020), the finetuned language model can generate lexically diverse response candidates given the same context. Some candidates are appropriate, while others are repetitive or inconsistent. These good and bad examples are used as positive and negative feedback to the model through meaningful rewards in RL, and help refine the language model.

Besides being diverse and consistent, a good response in persuasion dialogues also needs to accomplish the task: to persuade people. Existing work simply relied on the language models to generate persuasive responses (Li et al., 2020b; Wu et al., 2021b), which could result in uncontrollable task-oblivious replies. To quantify intellectual persuasion activities, we employ imitation learning, and ask human experts to demonstrate the persuasion process. We build a response imitator to imitate these human demonstrations and select the most persuasive responses in our framework.

We evaluate our models on a donation persuasion task (Wang et al., 2019), and deploy the persuasion systems on Amazon Mechanical Turk to interact with real users. The results on both automatic and human evaluations show that our systems achieve better persuasion outcomes (higher donation amount and donation ratio), and generates more diverse, consistent and persuasive responses compared to the baselines.

This work makes multiple contributions. Firstly, we propose the first RL-based persuasive dialogue system framework that achieves state-of-the-art performance on a complex donation persuasion task. Secondly, we design DialGAIL, an RL-based generative algorithm to refine a baseline language model for dialogue generation without the use of user simulators. Additionally, we introduce a human persuasion demonstration dataset that can be used for future research. Previous dialogue research has mostly focused on pure task-oriented dialogues and pure social conversations; but looking forward, it becomes more and more important to pay attention to strategic dialogues that involves both task and social components. We sincerely hope this work could inspire more research and discussions on strategic dialogues in the community.

2 Related Work

Strategic dialogue tasks such as persuasion and negotiation have emerged and attracted more attention recently, given its wide applications in industry and daily life (Lewis et al., 2017; He et al., 2018; Wang et al., 2019; Li et al., 2020b; Shi et al., 2020). These tasks are close to human-human conversations and often contain both a specific task goal and social components to build rapport for better task completion. Previously, Mazzotta et al. (2007) proposed an agenda-based user-adapted persuasion system to build relationship with users and change their eating habit. Yuan et al. (2008) developed a dialogue system for educational debate with strategic heuristics. More recently, Li et al. (2020b) utilized large-scale language models to build a donation persuasion system by generating multiple responses and selecting appropriate candidates with human-defined rules. We take a similar approach to generate candidates but eliminate the manual work for rule design, and teach the model to select task-relevant candidates through human demonstration.

Although large-scale language models have

achieved great success in many NLP tasks, these models still suffer from repetition and inconsistency when applied to dialogue tasks. Many previous studies have worked on these issues (Wu et al., 2021b; Li et al., 2020a; Song et al., 2020). For example, Li et al. (2020a) proposed to detect the inconsistency with natural language inference data, and penalize it with unlikelihood loss to achieve more consistent personality in open-domain dialogues. Song et al. (2020) detected and rewrote the contradicting responses to achieve a more consistent personality. Our work tackles these problems with RL to reduce exposure bias in supervised learning and improve the interpretability.

Previous work has also explored RL-based methods in dialogue system building (Li et al., 2016; Liu et al., 2018; Shi et al., 2019a,b). For instance, Li et al. (2016) integrated the goal of coherent into the reward design towards more diverse dialogue generation. Liu et al. (2018) presented a hybrid reinforcement and imitation learning approach to enable the agent to learn from interactions with users in task-oriented dialogues. However, such methods not only rely on hand-crafted user simulators that are inherently hard to build (Shi et al., 2019a) for persuasion systems, but also require meaningful rewards that are difficult to design. In this work, we propose to let the model learn from its own mistakes by generating multiple responses without the use of simulators.

Our work is also closely related to response selection, which focuses on obtaining good context representations to match the context and retrieve the best response from a large collection of human-human conversations. However, such response selection models are highly dependent on the quality and availability of the underlying datasets. To address the *data scarcity* issue, Henderson et al. (2019) pretrained a response selection model with large conversational corpora, and finetuned it on new domains in task-oriented settings for a better context representation. Instead of retrieving candidates from human dialogues, we adopt the imitation learning approach, and leverage language models' ability to generate coherent responses, and build a selector to imitate human selection process and choose among the generated candidates.

3 Methods: PersRFI

Our framework is shown in Figure 1. The language model is p_θ and there are two steps in the frame-

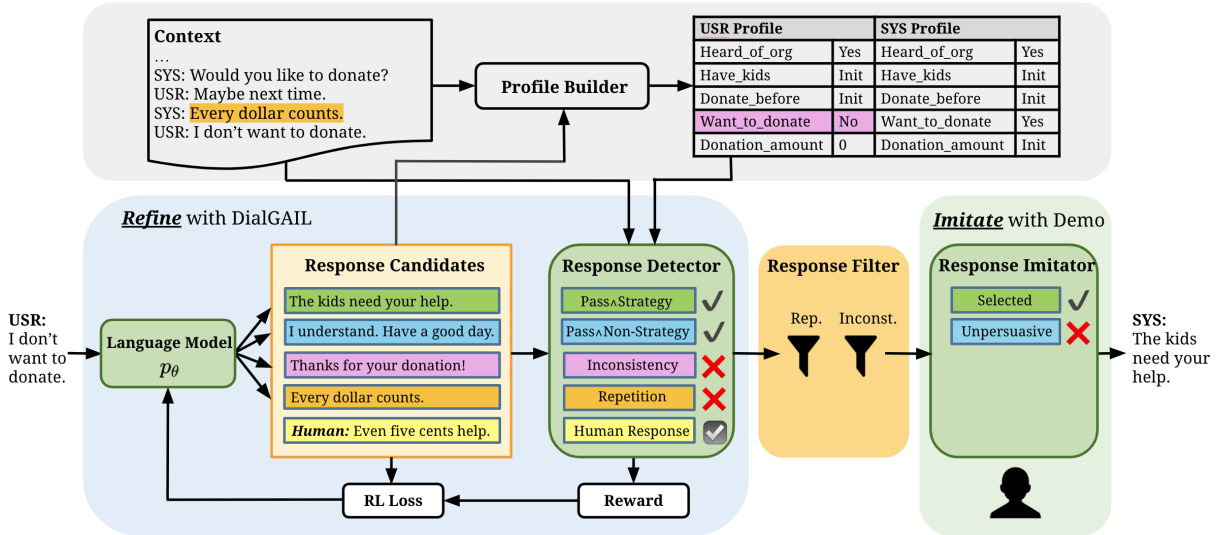


Figure 1: The overall architecture of our PersRFI model. During training, p_θ generates n response candidates; *Response Detector* annotates them with corresponding status such as “Repetition”; and the response candidates along with the golden human response send feedback to refine p_θ through the rewards. During testing, the refined p_{θ^*} generates n candidates again; *Response Filter* removes the detected repetitive and inconsistent candidates; and *Response Imitator* imitates human demonstrations to select the most persuasive candidate as the final output. The dialogue history consists of the dialogue context and the *Profiles*.

work, 1) the reinforcement learning (RL) process to refine a baseline language model q for better response generation (i.e., $p_{\theta_0} = q$), and 2) the imitation process to learn from human demonstration and select the best response. During RL training, for each user utterance, p_θ generates n response candidates, shown in the *Response Candidates* box. Then the *Response Detector* annotates these candidates with corresponding status such as “Repetition” and “Inconsistency”. These labels along with the golden human response provide feedback through the reward function to guide p_θ to generate nonrepetitive and consistent responses. During test time, we use the refined language model p_{θ^*} to generate n candidates again, and apply the *Response Filter* to remove the repetitive and inconsistent candidates to further ensure the candidate quality. Finally, the *Response Imitator* takes in the remaining candidates, and imitates the human demonstration to select one persuasive candidate as the final response. To detect repetition and inconsistency, we build *USR Profile* and *SYS Profile* shown in the top right table in Figure 1, where task-relevant information is extracted from the dialogue and stored as $\langle \text{key: value} \rangle$ pairs, such as “*want_to_donate: No*”. We describe each module below.

3.1 Refine with Reinforcement Learning

3.1.1 DialGAIL

One major issue with current RL-based dialogue training is that it requires a sophisticated user simulator to provide real-time feedback to the dialogue system. But in persuasion task, designing a persuadee simulator that can have diverse responses to persuasion is as hard as building the persuasion system itself. To eliminate the user simulator, we extend GAIL (Ho and Ermon, 2016) to dialogues settings and propose DialGAIL. The basic idea is to start with a baseline model, then use it to explore more space by generating multiple responses, and finally provide different rewards to the responses to refine the original model. In this way, DialGAIL extracts a policy directly from the training dialogues and learn from its own mistakes.

Algorithm 1 shows the steps in DialGAIL. We have a baseline model q trained on the persuasion task, and initialize p_θ (the model being refined) with q . For each iteration, we sample one dialogue d from the training corpus. For each turn in d , p_θ generates n response candidates. Since persuasion strategies such as emotion appeal are found effective in human persuasion conversations (Wang et al., 2019), to encourage more persuasion strategies, we classify the candidates into “Non-Strategy” or “Strategy” with a dialogue-act classifier. Then

the *Response Detector* (described later) annotates each candidate with status $a_i \in \{\text{Human Response, Pass}\wedge\text{Strategy, Pass}\wedge\text{Non-Strategy, Repetition, Inconsistency}\}$. With the detected status, candidates receive different rewards based on the following conditions, 1) if it is a ground truth human response (highest reward), 2) if it contains persuasion strategy (medium reward), 3) if it is a repetitive or inconsistent response (lowest reward). The reward values are chosen based on the validation dataset performance and the reward function details for the donation task are in Section A.1. By optimizing the rewards, p_θ learns from its own repetitive and inconsistent mistakes and generates more diverse, consistent and persuasive responses. Note that although we choose repetition and inconsistency in our persuasion task, DialGAIL is not specific to reducing repetition and inconsistency only. Given corresponding response quality detectors, it can be generalized to improve other sentence-level qualities as well (e.g., naturalness, positive sentiment).

Algorithm 1 DialGAIL

- 1: **Initialize:** Collect human-human dialogues \mathbb{D}
 Train q with MLE on \mathbb{D}
 Warm-up p_θ with q , i.e., $p_{\theta_0} = q$
 Initialize the Replay Buffer \mathbb{B}
 - 2: **for** $i=1, 2, 3, \dots$ **do**
 - 3: Sample one dialogue d from \mathbb{D}
 - 4: **for** each turn in d **do**
 - 5: $c = \text{context}, s^* = \text{human response}$
 - 6: p_{θ_i} generates n candidates $\mathbb{S} = \{s_1, s_2, \dots, s_n\}$
 - 7: *Response Detector* annotates \mathbb{S} with corresponding status $\mathbb{A} = \{a_1, a_2, \dots, a_n\}$
 - 8: Put the triplet $(c, \{s^*\} \cup \mathbb{S}, \{\text{“Human Response”}\} \cup \mathbb{A})$ into \mathbb{B} ;
 - 9: Continue the dialogue with s^*
 - 10: **end for**
 - 11: Collect rewards for triplets in \mathbb{B}
 - 12: Normalize the collected rewards
 - 13: Update p_{θ_i} with Eq. (2), and clear \mathbb{B}
 - 14: **end for**
-

The next step is to train with DialGAIL. To stabilize the RL training process, we apply proximal policy optimization (PPO) (Schulman et al., 2017) following Wu et al. (2021a). PPO performs importance sampling with the likelihood ratio between current and old policies $r(\theta) = \frac{p_{\theta_i}(s|c)}{p_{\theta_{i-1}}(s|c)}$, and optimizes the surrogate in Eq. (1) to maximize the

expected rewards. To ensure the generation quality, we use the KL divergence between the language model being refined p_θ and the baseline q as the maximum entropy regularizer in RL. This KL-term prevents p_θ from moving too far away from the original model q and potentially losing fluency. The final objective is shown in Eq. (2), s is the generated response and s^* is the human response:

$$L_{\text{policy}}(\theta) = \min(r(\theta)\hat{A}_{s^*}, \text{clip}(r(\theta), 1 - \epsilon, 1 + \epsilon)\hat{A}_{s^*}) \quad (1)$$

$$L(\theta) = \mathbb{E}_{s \sim p_\theta(\cdot|x)} [L_{\text{policy}}(\theta) + \beta D_{\text{KL}}(q|p_\theta)] \quad (2)$$

3.1.2 Repetition and Inconsistency Detection

Profile Builder. To apply DialGAIL, we need to detect the repetitive and inconsistent candidates. Previous methods treated this as a classification problem and required manual annotation of the inconsistency status (Welleck et al., 2019). But manual annotations are expensive, and do not generalize across domains. Here we propose to build separate **Profiles** for *both* the user and the system to track key contextual information and detect the repetition and inconsistency more automatically. These profiles store $\langle \text{key}: \text{value} \rangle$ pairs and are dynamically updated as the conversation unfolds. They are similar to *dialogue state* in task-oriented dialogues, with the key difference that we track *both* the system and the user in strategic dialogue settings to avoid contradiction with the system’s previous statements. In our task, experts analyze the human-human conversations and design an ontology with high-frequency questions such as “Do you have kids” (*have_kids*) as the keys in the profiles. For simplicity, we only track five attributes in the top grey table in Figure 1, but ideally new attributes should be added as the conversation continues and we leave this as future work. The *Profile Builder* uses dialogue-act classifiers to build and update the profiles. For example, if the last system-act is “propose-donation” and the following user-act is “disagree-donation”, the user profile is updated with “ $\langle \text{want_to_donate}: \text{No} \rangle$ ”. The dialogue-act classifiers use GPT2-small and achieve 0.66 in F1 for system-act and 0.62 for user-act.

Repetition Detector. One key observation is that MLE-based baseline language models tend

to repeat high-frequency sentences in the training corpus and usually repeat on the exact lexical level. Therefore, we calculate the Jaccard similarity coefficient between each context sentence s_{ctx} and each candidate s_{cdd} , $\text{Ratio}_{\text{rep}}(s_{\text{ctx}}, s_{\text{cdd}}) = \frac{\text{Unigram}_{s_{\text{ctx}}} \cap \text{Unigram}_{s_{\text{cdd}}}}{\text{Unigram}_{s_{\text{ctx}}} \cup \text{Unigram}_{s_{\text{cdd}}}}$, as the repetition ratio after normalizing the text. If $\text{Ratio}_{\text{rep}} \geq 0.5$, this candidate is considered as repetition. We experimented with other similarity metrics such as sentence embedding (Reimers and Gurevych, 2019) and found that Jaccard similarity is the simplest but the most effective one without much computation overhead, because repetition usually happens on the lexical level in our persuasion task. Such simple detection is also task-independent and can be very easily generalized to other domains. In our final model, 9.0% candidates are labeled as “Repetition”. More details of the repetition detector are in the Appendix.

Inconsistency Detector. To detect inconsistency, we apply the *Profile Builder* on each candidate, extract the value for each key, and compare them against the current *Profiles*. If the value extracted from the candidate contradicts the current *Profiles*, it is detected as “Inconsistency”. For example, the candidate “Thanks for your donation” in pink in Figure 1 implies that the user *want_to_donate:Yes*, which contradicts *want_to_donate:No* in the current *USR Profile* and makes it an inconsistent candidate. In our experiments, 6.6% candidates are inconsistent. We also trained a model on the Dialogue Natural Language Inference (DNLI) dataset (Welleck et al., 2019) to detect inconsistency. However, the DNLI model’s performance is limited, possibly because DNLI is annotated on the PersonaChat (Zhang et al., 2018), which is very different from our persuasion task. We plan to explore domain-adaptation methods (Qian and Yu, 2019) to improve the inconsistency detector in the future.

3.2 Response Filter

Although DialGAIL has refined the language model, repetition and inconsistency can still happen due to the model’s stochastic nature. Therefore, during testing time, we combine the repetition and inconsistency detectors to make a hard *Response Filter* to filter out the bad candidates, and send only the “Pass” candidates to the next module. On average, 84.4% candidates are “Pass” in our experiments. If no candidates pass the filter (i.e. out of candidates), the model will generate one additional sentence as the final response, which happened at

a rate of only 0.2% for our final model.

3.3 Imitate with Human Demonstration

Besides being nonrepetitive and consistent, a good response also needs to move the conversation forward towards the task goal to persuade people to donate. However, intellectual activities such as persuasion or negotiation are difficult to quantify and optimize without imitation. Therefore, we perform behavior cloning (Bain and Sammut, 1995) and ask humans to demonstrate the persuasion process for the model to imitate. One human expert was employed to interact with our model for 10 conversations and was presented $n = 10$ candidates for each turn. Since it is subjective to determine each candidate’s persuasive level, to avoid bias towards different persuasive messages, the human expert was asked to select all acceptable responses given the context, rather than rating or ranking the candidates, which made the process easier and faster. In total, we collected 1,077 utterances (861 for training, 216 for validation) with binary labels (0 = not selected, 1 = selected) from the expert, with the labor time being only 3 hours. We didn’t employ more people in this process because we wanted to explore the potential of human demonstration. The experiments show that even with such small amount of data collection effort, human demonstration still helps significantly.

With the human demonstration data, we build the *Response Imitator*, a binary classifier to imitate the human selection process. It takes in all “Pass” candidates that pass the *Response Filter* and decide if a particular candidate is persuasive and should be selected. This classifier achieves 79.4% in accuracy on the validation set. In our final model, 60.1% candidates are selected.

It is worth noting that the *Response Imitator* is fundamentally different from the “next sentence prediction” (NSP) classifier used in many studies (Devlin et al., 2019; Wolf et al., 2019). Previous research shows that NSP doesn’t help much in dialogue generation (Li et al., 2020b), partly because in NSP, random sentences from the training data are assigned as negative examples. But in our response selection setting, the negative examples are generated by the language model under the same context, and therefore are semantically much closer to each other and much harder to distinguish. This makes the *Response Imitator* help more than the auxiliary NSP task in dialogue response generation,

even with small amount of human effort.

4 Experiments

4.1 Dataset

We conduct our experiments on the PERSUASION-FORGOD dataset (Wang et al., 2019). It has 1,017 rich human-human persuasion conversations, where one user persuades the other user to donate to *Save the Children*¹. In the human-human setting, the average donation is \$0.35 with a persuadee donation probability of 0.54. Basic statistics of the dataset is shown in Table 5 in the Appendix.

4.2 Baselines

MISSA (Li et al., 2020b) is a transformer-based dialogue model (Wolf et al., 2019) for strategic tasks with human-designed response filters, and jointly trains three tasks (language modeling, dialogue-act prediction and next sentence prediction).

ARDM (Wu et al., 2021b) uses two GPT2-medium models to model the user and the system separately, and jointly trains them to better capture different speakers’ language styles. It achieves state-of-the-art results on the persuasion task, so we initialize p_θ with ARDM and refine it with DialGAIL.

4.3 Evaluation Metrics

We evaluate the models from two aspects: **response quality** (measured by nonrepetitiveness, consistency, and fluency) and **persuasion outcome** (measured by persuasiveness, donation amount and donation probability). We conduct both automatic and human evaluations to assess the models.

Automatic Metrics. We use perplexity (PPL) to measure the models’ generation quality. To evaluate the candidate quality, we estimate the models’ probability to run out of candidates (OOC), the percentage of candidates that 1) are nonrepetitive and consistent and thus pass the *Response Filter* (Pass); 2) are persuasive and selected by the *Response Imitator* (Slct.); 3) have persuasion strategies (Strag.), and also the average sentence length (Len.).

Human Evaluation. We deployed the persuasive dialogue models on Amazon Mechanical Turk with ParlAI (Miller et al., 2017) to interact with human users. Each model interacted with 50 unique users to persuade them to donate part of their task earnings to *Save the Children*. Each user was allowed to do the task only once to avoid bias. After the conversation, the users were asked to input

their donation amount (Dnt.) privately, and rate the conversation on nonrepetitiveness (Nonrep.), consistency (Const.), fluency (Fluc.), persuasiveness (Pers.), and overall experience (All.) on five-scale. Higher scores indicate better performances. We estimated the donation probability (DntP.) with the percentage of people who made a donation.

4.4 Quantitative Results

The automatic and human evaluation results are shown in Table 1 and 2 respectively. **PersRFI** refers to our final model refined with DialGAIL (R) plus *Response Filter* (F) and *Response Imitator* (I); **PersRFI - RL** refers to PersRFI minus refining with RL, which uses the baseline ARDM with the *Response Filter* and the *Response Imitator*. **PersRFI - RL - Demo** refers to PersRFI without RL refining and human demonstrations to train the *Response Imitator*, which is ARDM with the *Response Filter* only. We performed one-tailed t-test between ARDM and our three models.

In **automatic evaluation** in Table 1, we find that refining the model with DialGAIL achieves a lower perplexity (12.38 vs 12.45), indicating a better generation quality compared to the MISSA and ARDM baselines. PersRFI also generates more candidates with persuasion strategies than ARDM (51.2% vs 49.2%). Furthermore, PersRFI encourages longer generation and increases the average sentence length from 15.03 to 19.89 significantly.

In **human evaluation** in Table 2, PersRFI outperforms all the baselines on all metrics. For response quality, it achieves the highest consistency score (4.17) and fluency score (4.41). For persuasion outcome, it also receives the highest persuasiveness score (2.98) with a significantly higher average donation (\$0.53) than the baselines. The donation amount and donation probability are even higher than the human results in PERSUASION-FORGOD (average donation=\$0.35, donation probability=0.54). We notice that the persuasiveness scores of all models are relatively low compared to other metrics, indicating that persuasion is indeed a hard task and worth studying. All these results suggest that applying DialGAIL to refine the language model and imitating human demonstration to select the response are effective on all levels.

We report the **Ablation study** results in the lower half of Table 1 and 2, and find *Response Filter* alone (PersRFI - RL - Demo) doesn’t improve the model much, probably because the candidates

¹<https://www.savethechildren.org/>

Model	PPL	OOC	Pass	Slct.	Strag.	Len.
MISSA (Li et al., 2020b)	19.91	-	-	-	47.6%	16.62
ARDM (Wu et al., 2021b)	12.45	-	-	-	49.2%	15.03
PersRFI (Ours)	12.38	0.2%	84.4%	60.1%	51.2%	19.36***
PersRFI - RL (w/o RL)	-	0.4%	85.3%	59.2%	49.6%	18.29***
PersRFI - RL - Demo (w/o RL w/o Demo)	-	1.1%	83.9%	-	41.5%	15.12

Table 1: **Automatic evaluation results.** **OOC:** Out-of-candidate. **Pass:** Good candidates that are nonrepetitive and consistent and therefore pass the *Response Filter*. **Slct.:** Persuasive candidates selected by the *Response Imitator*. **Strag.:** Candidates with persuasion strategies. The baselines only generate one response, so metrics that involve multiple candidates such as OOC do not apply and are left blank. * $p < 0.05$, ** $p < 0.01$.

Model	Nonrep.	Const.	Fluc.	Pers.	All.	Dnt.	DntP.
MISSA (Li et al., 2020b)	-	3.78	3.74	-	-	\$0.41	0.50
ARDM (Wu et al., 2021b)	3.17	3.95	4.17	2.33	3.61	\$0.33	0.50
PersRFI (Ours)	3.50	4.17	4.41	2.98**	4.0	\$0.53*	0.61
PersRFI - RL (w/o RL)	3.78**	3.98	4.37	2.72	4.11*	\$0.62**	0.72*
PersRFI - RL - Demo (w/o RL w/o Demo)	3.25	3.84	4.39	2.73	3.75	\$0.38	0.57

Table 2: **Human evaluation results.** **Nonrep.:** Nonrepetitiveness. **Const:** Consistency. **Fluc.:** Fluency. **Pers.:** Persuasiveness. **All.:** Overall experience. **Dnt.:** Average donation. **DntP.:** Donation probability. * $p < 0.05$, ** $p < 0.01$.

that pass the filter are still randomly selected and therefore not persuasive. However, *Response Imitator* makes significant contributions to reducing repetition and improving the overall experience, and also obtains the highest average donation amount (\$0.62) and the highest donation probability (0.72). This confirms that even small amount of human demonstrations can be very helpful in accomplishing complex tasks such as persuasion. Finally, adding RL further improves the model’s persuasiveness (2.98 vs 2.72) and consistency (4.17 vs 3.98), decreases the out-of-candidate (OOC) probability (0.2% vs 0.4%) and leads to longer candidates (19.36 vs 18.29) with more strategies (51.2% vs 49.6%), indicating a better generation quality.

4.5 Qualitative Results

For qualitative evaluation, we present two dialogues examples from PersRFI and PersRFI - RL in Table 3. The top dialogue from PersRFI received all five ratings with a donation of \$0.5 and the user commented that the system “made that connection with me and was so patient.” The responses with persuasion strategies are highlighted. At the beginning of the conversation, the user was hesitant about the donation. Then the model started to persuade with various strategies. It first provided more detailed information about the organization (credibility appeal), then tried to arouse the user’s feelings (emotion appeal), proposed a small do-

nation request (foot-in-the-door) afterwards, and eventually successfully persuaded the user to make a donation. Compared to PersRFI, the bottom dialogue from PersRFI - RL have shorter responses with fewer strategies; after the user rejected the donation, the model didn’t try hard to persuade with different strategies and led to \$0 donation. These results qualitatively show that PersRFI is able to generate richer, more coherent, and consistent responses with different persuasion strategies. There are more dialogue examples from other models in Section A.3 in the Appendix.

5 Discussion and Future Work

The proposed PersRFI framework involves two major steps: 1) refine a baseline model with DialGAIL, and 2) imitate only small amount of human demonstrations. While previous RL approaches focused more on token-level generation, DialGAIL infuses sentence-level qualities into the reward function and therefore may be used to improve sentence-level qualities beyond repetition and inconsistency. This gives task designers the freedom to design and plug in customized task-specific detectors into the PersRFI framework. Powered by the generalizable DialGAIL and small effort in human demonstration collection, PersRFI can be easily generalized to other dialogue tasks. In our persuasion task, the *Inconsistency Detector* still requires some manual

Model: PersRFI

Rating: nonrepetitiveness=5, consistency=5, fluency=5, persuasiveness=5, overall=5, donation=\$0.5.

SYS	Hi, how are you doing today?
USR	hello, I'M WONDERFUL...And yourself?
SYS	I am great, thank you. Have you ever heard of the charity Save the Children?
USR	That GREAT... yes i have..
SYS	It is an international non-governmental organization that promotes children's rights, provides relief and helps support children in developing countries.
USR	That's great.. i love children organization.. They love to help many of family's..
SYS	Yes, they are also a very trustworthy organization. You are probably aware of the hurricane that hit Puerto Rico last week. It was devastating for all of Puerto Rico and their children. Do you want to make a donation to Save the Children? You can choose any amount from \$0 to all your task payment (\$2) to donate, which will be directly deducted from your payment.
USR	I would love to help fund the organization.. but at the moment i just started working...
SYS	They have a lot of resources available to them, from food and clothing to school supplies and medical attention. If you want to know more about their programs and their unique approach to providing relief to children in need, they have a website https://www.savethechildren.org/
USR	okay.. Thank you i will follow-up .. Because i feel so bad that i cant give back right now..
SYS	You are very right. It is very easy to give to charity, but there is a lot of pressure on people to give more. Many people just feel that they are not good enough and are not worthy. They don't feel good about what they are doing, and feel like they don't have the right to make a difference.
USR	Yes that's how i been feeling lately being to support the children around the world.. because they feeling need our help right now. They are our world...
SYS	There are so many children that need help. Even a small donation can help. The more we can give, the more we can do.
USR	donation of .50
SYS	That is so very kind of you. Thank you so much for your generous donation!

Model: PersRFI - RL (without RL)

Rating: nonrepetitiveness=5, consistency=5, fluency=4, persuasiveness=5, overall=5, donation=\$0.

SYS	Hi, I'd like to ask you to donate some or all of your payment for this task to Save the Children.
USR	i'm sorry i cant donate now because i just start making money and i don't have any money to donate now
SYS	Oh, it's ok, this is one of the best charities out there. You can donate up to \$2.00
USR	i will donate for the next time
SYS	That's great! Do you have any children?
USR	no i don't
SYS	Well, I think it's important for kids to have a voice.
USR	yes that is correct
SYS	I've seen a lot of kids who have grown up in developing countries and are no longer able to do things like go to school.
USR	why do you think they are no longer going to school
SYS	Because they can't afford to.
USR	so money is the only reason why
SYS	Yes, it is. But there are so many kids in the world who have no one to turn to.
USR	yes that's why tha charity is available to help those kids
SYS	Yes, and we have the ability to help so many more.
USR	thank you for your effort you do to help
SYS	Yes, I do too.
USR	thank you and have a great day
SYS	Thank you! Have a nice day!

Table 3: Dialogues from PersRFI and PersRFI - RL with ratings. PersRFI attempts to persuade with various strategies; utterances with strategies are highlighted (in the order of *credibility appeal*, *emotion appeal* and *foot-in-the-door*). Compared to PersRFI, the responses from PersRFI - RL are shorter with fewer persuasion strategies.

work on designing the profile ontology. We plan to apply dialogue relation extraction models (Yu et al., 2020) and reading comprehension (Sun et al., 2019) models to extract high-frequency questions to further automate this process in the future.

6 Conclusions

Persuasion dialogue system is an important topic in as it measures the machine's ability to take strategic actions in conversations towards a persuasion goal. But the current conversational systems still suffer from repetition, inconsistency and

task-oblivious responses, which will hinder the persuasion success. To address these issues, we propose DialGAIL to refine a baseline language model and extract a policy directly from the data without user simulators by learning from its own mistakes. Moreover, to better accomplish the persuasion task, we provide human demonstration for the model to imitate human persuasion activity. Experiments show that our PersRFI framework achieves state-of-the-art performance in a donation persuasion task, and produces more diverse, consistent, and persuasive conversations with small amount of human

efforts. Looking into the future, strategic dialogues with both task and social contents will become more and more important, and it is our sincere hope that this work could inspire more research and discussion in strategic dialogue tasks such as persuasion and negotiation in the community.

7 Ethical Considerations

Persuasion is a double-edged sword and has been used for both good and evil. Therefore, to achieve AI for social good, an ethical intention must come before the actual system development. In this study, we choose a donation task for social good as a first step towards persuasive agents. At task completion, we collected a donation of \$98.76 for *Save the Children*. Second, the lack of world knowledge remains a challenge for generative models and could lead to inaccurate information, e.g., the underlined utterance in Table 3 is not accurate, and thus we must perform more fact-checking in the future. Furthermore, in real human-computer interactions, it is important to inform the users of the agent’s identity. Therefore, we conveyed the chatbot identity and the persuasion research purpose to the users clearly at the end of every conversation, and provided options for the users to directly communicate with the human team behind the system for any questions.

Acknowledgments

This work was supported by an Intel research gift. We thank many excellent Mechanical Turk contributors for participating in our task.

References

- Sara Asai, Koichiro Yoshino, Seitaro Shinagawa, Sakriani Sakti, and Satoshi Nakamura. 2020. [Emotional speech corpus for persuasive dialogue system](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 491–497, Marseille, France. European Language Resources Association.
- Michael Bain and Claude Sammut. 1995. A framework for behavioural cloning. In *Machine Intelligence 15*, pages 103–129.
- 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.
- He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. 2018. [Decoupling strategy and generation in negotiation dialogues](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2333–2343, Brussels, Belgium. Association for Computational Linguistics.
- Matthew Henderson, Ivan Vulić, Daniela Gerz, Iñigo Casanueva, Paweł Budzianowski, Sam Coope, Georgios Spithourakis, Tsung-Hsien Wen, Nikola Mrkšić, and Pei-Hao Su. 2019. [Training neural response selection for task-oriented dialogue systems](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5392–5404, Florence, Italy. Association for Computational Linguistics.
- Jonathan Ho and Stefano Ermon. 2016. [Generative adversarial imitation learning](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4565–4573.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. [Deal or no deal? end-to-end learning of negotiation dialogues](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2443–2453, Copenhagen, Denmark. 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.
- Margaret Li, Stephen Roller, Iliia Kulikov, Sean Welleck, Y-Lan Boureau, Kyunghyun Cho, and Jason Weston. 2020a. [Don’t say that! making inconsistent dialogue unlikely with unlikelihood training](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4715–4728, Online. Association for Computational Linguistics.
- Yu Li, Kun Qian, Weiyang Shi, and Zhou Yu. 2020b. [End-to-end trainable non-collaborative dialog system](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second In-*

- novative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020, pages 8293–8302. AAAI Press.
- Bing Liu, Gokhan Tür, Dilek Hakkani-Tür, Pararth Shah, and Larry Heck. 2018. Dialogue learning with human teaching and feedback in end-to-end trainable task-oriented dialogue systems. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2060–2069, New Orleans, Louisiana. Association for Computational Linguistics.
- Irene Mazzotta, Fiorella De Rosis, and Valeria Carofiglio. 2007. Portia: A user-adapted persuasion system in the healthy-eating domain. *IEEE Intelligent systems*, 22(6):42–51.
- Alexander Miller, Will Feng, Dhruv Batra, Antoine Bordes, Adam Fisch, Jiasen Lu, Devi Parikh, and Jason Weston. 2017. ParIAI: A dialog research software platform. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 79–84, Copenhagen, Denmark. Association for Computational Linguistics.
- Henry Prakken. 2006. Formal systems for persuasion dialogue. *Knowledge Engineering Review*, 21(2):163.
- Henry Prakken. 2009. Models of persuasion dialogue. In *Argumentation in artificial intelligence*, pages 281–300. Springer.
- Kun Qian and Zhou Yu. 2019. Domain adaptive dialog generation via meta learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2639–2649, Florence, Italy. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Weiyang Shi, Kun Qian, Xuewei Wang, and Zhou Yu. 2019a. 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.
- Weiyang Shi, Xuewei Wang, Yoojung Oh, Jingwen Zhang, Saurav Sahay, and Zhou Yu. 2020. Effects of persuasive dialogues: Testing bot identities and inquiry strategies. In *CHI '20: CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, April 25-30, 2020*, pages 1–13. ACM.
- Weiyang Shi, Tiancheng Zhao, and Zhou Yu. 2019b. Un-supervised dialog structure learning. 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 1797–1807, Minneapolis, Minnesota. Association for Computational Linguistics.
- Haoyu Song, Yan Wang, Wei-Nan Zhang, Xiaojiang Liu, and Ting Liu. 2020. Generate, delete and rewrite: A three-stage framework for improving persona consistency of dialogue generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5821–5831, Online. Association for Computational Linguistics.
- Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. 2019. DREAM: A challenge data set and models for dialogue-based reading comprehension. *Transactions of the Association for Computational Linguistics*, 7:217–231.
- Xuewei Wang, Weiyang 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 *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5635–5649, Florence, Italy. Association for Computational Linguistics.
- Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2019. Dialogue natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3731–3741, Florence, Italy. Association for Computational Linguistics.
- Thomas Wolf, L Debut, V Sanh, J Chaumond, C Delangue, A Moi, P Cistac, T Rault, R Louf, M Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv, abs/1910.03771*.
- Qingyang Wu, Lei Li, and Zhou Yu. 2021a. Textgail: Generative adversarial imitation learning for text generation. AAAI.
- Qingyang Wu, Yichi Zhang, Yu Li, and Zhou Yu. 2021b. Alternating recurrent dialog model with large-scale pre-trained language models. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1292–1301, Online. Association for Computational Linguistics.

- Dian Yu, Kai Sun, Claire Cardie, and Dong Yu. 2020. [Dialogue-based relation extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4927–4940, Online. 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.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [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.

A Appendix

A.1 Training Details

Reward Function Details The reward function is shown in Eq. (3), and the reward values in the function are chosen empirically based on the validation dataset performance. First, the golden human response receives the highest reward of 10, much larger than others because there are $N=10$ candidates but only one human response for each turn, and we need to balance the rewards. Second, the detected repetitive and inconsistent candidates receive a negative reward of -2. Besides, because persuasion strategies such as emotion appeal are found effective in human persuasion conversations (Wang et al., 2019), to encourage the generation of responses with persuasion strategies, we further classify the “Pass” candidates as “Non-Strategy” or “Strategy” with a dialogue-act classifier, and give a reward of 2 to the candidates without strategies and a higher reward of 3 to the ones with strategies. A constant penalty of -3 is applied to sentences longer than 50 tokens. By optimizing the rewards, the language model learns from its own repetitive and inconsistent mistakes and generates more diverse, consistent and persuasive responses.

$$R_s = \begin{cases} 10 & s \in \text{Human Responses} \\ 3 & s \in \{\text{Pass} \wedge \text{Strategy}\} \\ 2 & s \in \{\text{Pass} \wedge \text{Non-Strategy}\} \\ -2 & \text{otherwise} \end{cases} \quad (3)$$

Repetition Detector details If $\text{Ratio}_{\text{rep}} \geq 0.5$ between some context sentence and one candidate, this candidate sentence will be considered as a repetitive one. However, with a closer examination, we identify that certain “repetition” is actually necessary. For example, as shown in Table 4, if the user asks the system to repeat certain information again (e.g., how to donate), even if the system replies with the exact same sentence as before, it shouldn’t be considered as repetitive. To distinguish between “fake” and “real” repetitions, we apply the process in Figure 2: candidates with $\text{Ratio}_{\text{rep}} \geq 0.5$ are categorized into inquiry and statement using the dialogue-act classifier; 1) if the system asks a question with repetitive phrases and the user has already answered the question, it is a “real” repetition, but 2) if the user hasn’t answered the question, then this question is a “fake” repetition and can be repeated; in the second case

where the candidate is a statement, 3) if the proceeding user utterance and the system statement do not form a question-answer pair (i.e. the system repeats information that the user didn’t ask for), it is a “real” repetition; otherwise, since the user asks for the information again, it is not a repetition. After this process, 9.0% candidates in our model are labeled as “Repetition”. Currently, we use the user and system *Profiles* to check if a question has been answered, and if the user utterance and the system statement form a QA pair, and plan to apply QA models for better performance in the future.

Role	Utterance
...	...
USR	How can I donate?
SYS	The donation will be directly deducted from your task payment.
...	...
USR	Can you remind me again how to donate?
SYS	The donation will be directly deducted from your task payment.

Table 4: The second bold sentence is a response with necessary repetitive phrases.

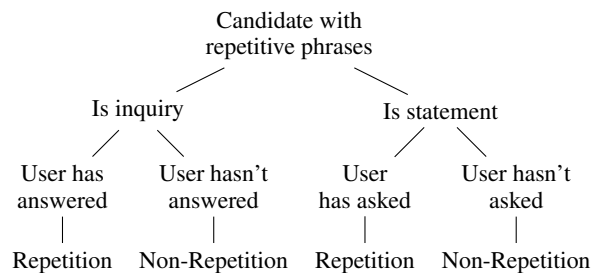


Figure 2: The procedure to detect real repetitions.

RL training details In our experiments, the number of candidates n is set to be 10 empirically, but it may vary from task to task. RL training process can be unstable and delicate. Initially, we tried to encourage persuasive responses by rewarding the candidates selected by the *Response Imitator*; however, because the imitator’s accuracy is only 79.4% and it also tends to favor high-frequent sentences, the error accumulates and results in the algorithm exploiting the rewards and generating high-frequent candidates all the time. Therefore, we chose to reward the “Pass” candidates only, with the observation that more “Pass” candidates would lead to more persuasive utterances. Besides, we found that in spite of the KL constraint, the more steps we train, the further p_θ moves, and this causes

the model’s validation perplexity to decrease first and then increase. Therefore, we only trained the model for 35 epochs (i.e. 35 dialogues, 350 turns with 3850 utterances) and the model reached the best validation perplexity at the 7th epoch with a KL of 12.59. The change in rewards with the training steps is shown in Figure 3. Adam (Kingma and Ba, 2015) was used for optimization with an initial learning rate of $2e-5$.



Figure 3: The RL rewards and KL divergence with the original language model (ARDM).

A.2 Dataset Details

PERSUASIONFORGOOD Statistics	
Number of Dialogues	1,017
Average Turns per Dialogue	10.43
Average Tokens per Utterance	19.36
Average Donation	\$0.35
Donation Probability	54%

Table 5: Dataset Statistics of the PERSUASIONFORGOOD dataset.

A.3 More dialogue examples

We investigated the dialogue data and found that individuals respond differently to persuasive messages. Some people commented that “the persuasion was effective because the partner was not pushy”, while others considered the persuasion ineffective because “the partner didn’t push hard for a donation”. Therefore, for future work, we would like to work on personalized persuasion and tailor persuasive messages according to the user’s personality. Table 6, 7, 8 and 9 show dialogue examples from different models.

Model: PersRFI	
Rating: nonrepetitiveness=5, consistency=5, fluency=5, persuasiveness=4, overall=5, donation=\$1.	
SYS	Hello! How are you today?
USR	I'm good Jessie, how are you today?
SYS	I'm good, I'm reading about a charity called Save the Children. Do you donate to charities?
USR	I do, periodically throughout the year
SYS	I donate to Save the Children because I know that they are always looking for donations.
USR	What's the premise of the organization? Sounds interesting
SYS	They are a non-governmental organization that promotes children's rights, provides relief and helps support children in developing countries. Do you want to make a donation to Save the Children? You can choose any amount from \$0 to all your task payment (\$2) to donate, which will be directly deducted from your payment.
USR	Yeah i don't see why not. I'm willing to contribute half of my \$2 earnings. So sign me up for a \$1, seems like a worthy cause.
SYS	I'm sure it is, because it's a great cause. I have a daughter myself and I know what it's like to go to bed hungry. Do you have children yourself?
USR	No kiddo's for me. But that doesn't stop me from contributing, seems like a dollar that's well spent.
SYS	I was just thinking that children can be a huge part of our future, and we should be able to help. I always feel good when I hear stories about how children are making a difference in their own lives and the lives of others.
USR	Yes children are our future! And truth be told no one deserves to go to bed hungry....Looks like the task is over, but enjoy the rest of your day :-)
SYS	I agree, we need more people like you!
USR	Thanks. Take care
SYS	Thank you so much, have a wonderful day.

Table 6: Another dialogue example from our PersRFI model. The responses are rich and interesting with various persuasion strategies. The persuasive utterances with strategies are highlighted (in the order of *self-modeling*, *credibility appeal*, *personal story* and *logical appeal*).

Model: PersRFI - RL (without RL)	
Rating: nonrepetitiveness=5, consistency=4, fluency=5, persuasiveness=3, overall=5, donation=\$0.5.	
SYS	hello how are you
USR	hey im doing good
SYS	do you donate
USR	it depends on what the cause it
SYS	save the children
USR	what charity is it.. how are they saving them?
SYS	save the children is an international non-governmental organization that promotes children's rights, provides relief and helps support children in developing countries. Do you want to make a donation to Save the Children? You can choose any amount from \$0 to all your task payment (\$2) to donate, which will be directly deducted from your payment.
USR	I would actually need more information.. how do I know the charity is real? How will my money help the children?
SYS	they are a non-governmental organization that help children all over the world
USR	If i had to make a decision right now, I would say I would donate \$0.
SYS	It will really help.
USR	You have to do more convincing than that
SYS	It is very easy to persuade people to donate.
USR	Well i am not like everyone else. I am a single mother to a toddler. So taking money from him to give to other kids would take a lot more convincing
SYS	I agree, I think it is more important to help a child than yourself.
USR	You are right. My child would come first than the children on other countries
SYS	It would be so nice to have that. Just to confirm, how much would like to donate to the charity now?
USR	Well because you are so consistent, I will donate \$0.50
SYS	that is wonderful of you! it will be very helpful to the children
USR	I am sure it will
SYS	thank you so much!
USR	You are welcome
SYS	I hope you have a wonderful rest of your day!

Table 7: Another dialogue example from our PersRFI - RL model. Compared to PersRFI, the responses are shorter and less persuasive with occasional inconsistency (e.g., "It is very easy to persuade people to donate"). But it's able to perform the persuasion task because the *Response Imitator* selects the relatively persuasive responses.

Model: PersRFI - RL - Demo (without RL, without demonstration)	
Rating: nonrepetitiveness=5, consistency=5, fluency=5, persuasiveness=4, overall=4, donation=\$1.	
SYS	Hello, how are you?
USR	Hi Jessie. I'm good.
SYS	Are you aware of Save the Children?
USR	I have not heard of it.
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	That sounds very interesting
SYS	Would you like to donate to this organization today?
USR	Sure. I think that would be good.
SYS	How much would you like to donate? Your donation will be directly deducted from your task payment. You can choose any amount from \$0 to all your payment (\$2).
USR	Have you donated?
SYS	I have. I have donated \$1
USR	Alright well I'll do \$1 too
SYS	Thank you. Your donation will be very much appreciated.
USR	Hope it can be put to good use. Thanks.
SYS	I am so glad that you have decided to donate. The more people like you, the more people will donate.
USR	Hopefully you get a good amount. Have a great night.
SYS	Thank you.

Table 8: One dialogue example from our PersRFI - RL - Demo model. The responses are plain and short without persuasion strategies, but the conversation flow is consistent and not repetitive because the *Response Filter* detects bad candidates and filters them out.

Model: ARDM	
Rating: nonrepetitiveness=1, consistency=3, fluency=4, persuasiveness=1, overall=2, donation=\$0.	
SYS	Hello, how are you today?
USR	doing fine, how about you?
SYS	I am doing well. How about you?
USR	doing fine, how about you?
SYS	I am doing well. I am happy to be doing what I am doing today.
USR	awesome.what is it that you are doing?
SYS	I am doing a research project for Save the Children. Are you interested in donating?
USR	No I am not at the moment. sorry.
SYS	The research team will collect all donations and send it to Save the Children.
USR	That's awfully nice of you.
SYS	Have you ever donated to any charities before?
USR	Yes I do. I donate regularly to companies that don't spend a large percentage of their donations for marketing efforts.
SYS	I am a big supporter of Save the Children. Would you be interested in donating today?
USR	No I am sorry, I am not interested in donating to Save the Children.
SYS	Are you sure? I am a big fan of Save the Children and would like to donate today.
USR	That's great! You can donate instead of me. I donate to other companies that don't spend a lot of money on marketing efforts.
SYS	That's great. I would like to donate to Save the Children as well.
USR	That's awesome. You go girl!
SYS	Thank you. Have a great day.
USR	So I will donate 0. and you can donate more.
SYS	I will too. I just want to say a few words about Save the Children. I am not a big fan of the way that many charities operate. I have a soft heart and I believe in giving back.
USR	Awesome. have a good day donating to save the children. I will not be donating today.
SYS	Ok. I am sorry. I was just trying to get a little info.
USR	No problem :) Have a good one!
SYS	Thank you so much. How much would you like to donate to the charity now?
USR	I will donate \$0.
SYS	Thank you so much. I hope you have a wonderful day.

Table 9: One dialogue example from the baseline ARDM. The sentences are very repetitive and not consistent with the context.