

Generating Persona Consistent Dialogues by Exploiting Natural Language Inference

Haoyu Song, Wei-Nan Zhang, Jingwen Hu, Ting Liu
 Research Center for Social Computing and Information Retrieval
 Harbin Institute of Technology, Heilongjiang Province, China
 {hysong, wnzhang, jwhu, tliu}@ir.hit.edu.cn

Abstract

Consistency is one of the major challenges faced by dialogue agents. A human-like dialogue agent should not only respond naturally, but also maintain a consistent persona. In this paper, we exploit the advantages of natural language inference (NLI) technique to address the issue of generating persona consistent dialogues. Different from existing work that re-ranks the retrieved responses through an NLI model, we cast the task as a reinforcement learning problem and propose to exploit the NLI signals from response-persona pairs as rewards for the process of dialogue generation. Specifically, our generator employs an attention-based encoder-decoder to generate persona-based responses. Our evaluator consists of two components: an adversarially trained naturalness module and an NLI based consistency module. Moreover, we use another well-performed NLI model in the evaluation of persona-consistency. Experimental results on both human and automatic metrics, including the model-based consistency evaluation, demonstrate that the proposed approach outperforms strong generative baselines, especially in the persona-consistency of generated responses.

Introduction

Despite the recent success of dialogue generation in open-domain by training from large volumes of human-to-human interaction data (Shang, Lu, and Li 2015; Serban et al. 2016; Li et al. 2017; Zhu et al. 2019), conversing to a dialogue agent is still in its infancy, and one major issue for these data-driven models is the lack of a consistent persona (Vinyals and Le 2015; Li et al. 2016a; Zhang et al. 2018; Song et al. 2019a). Figure 1 shows how consistency affects the quality of dialogues.

One practical approach to increase the consistency of a dialogue agent is to explicitly define a set of personal facts describing the characters (the *personas*) of the agent and learn to generate responses that reflect the predefined personas (Zhang et al. 2018). However, due to the lack of consistency modeling and the maximum-likelihood estimation (MLE) objective function, these persona-based models still face the inconsistency issue (Welleck et al. 2019).

Persona (for Dialogue Agent)

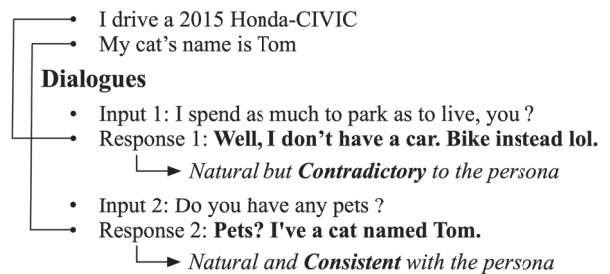


Figure 1: Naturalness is an important attribute of dialogue responses. In persona-based dialogue generation, the consistency with persona is another essential factor to consider. An ideal response should be not only natural but also consistent with the persona.

Natural Language Inference (NLI) learns a mapping between a sentence pair and an entailment category. Taking advantages of the NLI techniques in natural language understanding (Bowman et al. 2015), the detection of persona-consistency can be modeled as an NLI task (Welleck et al. 2019), which assigns a label of *entailment*, *neutral*, or *contradiction* to an “(utterance, persona)” pair. Meanwhile, existing persona-based dialogue models are limited by their loss functions. For these deep generative models, it is difficult to design a differentiable training objective to exploit the NLI based consistency detection method. Besides designing a differentiable training objective, reinforcement learning (RL) offers another solution to this problem, which back-propagates the reward signals to guide the generator.

In this paper, different from re-ranking the archived responses (Welleck et al. 2019), we take advantages of the NLI techniques in guiding the generation of persona consistent dialogues. Specifically, we propose a system trained using reinforcement learning. Our model has one evaluator with two modules and one generator. The evaluator consists of a naturalness module and a consistency module. The naturalness module is trained adversarially for higher accuracy. As for the consistency module, we use an NLI styled classifier

to detect the consistency between responses and personas. We further employ two different NLI classifiers in our experiments to investigate the role of NLI signals. The generator, which is a persona-based attentive Seq2Seq model (Zhang et al. 2018), takes message and persona texts as input and generates responses that reflect the persona texts. Note that more advanced generative models such as MASS (Song et al. 2019b) can also be exploited as our generator.

We summarize the contributions as follows:

- We propose an RL framework for persona consistent dialogue generation, thus addressing the challenge of training objective need to be differentiable in persona-based dialogue models.
- To the best of our knowledge, this is the first work that exploits NLI techniques to enhance the generation of persona consistent dialogues.
- Evaluations are carried out both quantitatively and qualitatively, and experimental results show that our model outperforms strong baselines, especially in terms of persona-consistency.

Related Work

Persona-based Dialogue Generation In open-domain dialogue generation, Zhang et al. (2018) initiate a new line of research (the persona-based dialogue) by introducing the Persona-Chat dataset, with explicit persona information in each dialogue session. They further propose two generative models, *persona-Seq2Seq* and *Generative Profile Memory Network*, to incorporate persona texts into responses. In the persona-based scenario, a model is associated with a persona, which is composed of several persona texts (See the top two sentences in Figure 1). A response is then generated using both the input message and the assigned persona. Following this line, Yavuz et al. (2018) apply the *DeepCopy* model in the persona-based dialogue generation. These works have laid a solid foundation for this area. Through attention or copy, generated responses can reflect the predefined persona. However, the loss functions in these models do not take the consistency issue into account, and inconsistency is still a problem to be addressed in the existing approaches (Welleck et al. 2019).

Natural Language Inference The task of Natural Language Inference (NLI) is to learn a function $f_{NLI}(p, h) \rightarrow \{E, N, C\}$, where p and h denote *premise* and *hypothesis* respectively. The output E, N and C represent *entailment*, *neutral* and *contradiction* between the *premise* and *hypothesis*. Since the release of large scale corpus SNLI (Bowman et al. 2015), deep neural network methods have made promising progress (Chen et al. 2017; Gong, Luo, and Zhang 2018; Kim et al. 2019). Welleck et al. (2019) model the detection of dialogue consistency as an NLI task and propose the DNLI (Dialogue NLI) dataset, which is similar to the SNLI but in the domain of persona-based dialogue. Further, they verify the effectiveness of using NLI model to re-rank candidate responses in a retrieval-based dialogue model. Compared with the retrieval-based model, the responses from

generative models are not limited to the given dataset. Moreover, exploiting consistency detection method in deep generative dialogue models has not been explored yet.

Reinforcement Learning In recent years, deep reinforcement learning has been widely applied in natural language processing, such as machine translation (Wu et al. 2018), visual question generation (Fan et al. 2018), paraphrase generation (Li et al. 2018), anaphora resolution (Yin et al. 2018) etc. The advantage of reinforcement learning lies in that it does not need a differentiable objective function. In open-domain dialogue generation, Li et al. (2016b) manually defined three rewards and use reinforcement learning to train the dialogue agent. Further, Li et al. (2017) apply adversarial learning method (Yu et al. 2017) for dialogue generation and propose the *REGS* model. This model shows its strength in the naturalness of generated responses. However, natural responses can also be inconsistent, especially in the persona-based scenario (as shown in Figure 1).

Proposed Approach

Problem Definition

Our goal is to learn a generative model G to deliver persona consistent dialogues, which can be formally defined as: given an input message X , and a set of persona texts $P = \{P_1, P_2, \dots, P_n\}$, to generate a response \hat{Y} , based on both input message X and the persona text set P , i.e., $\mathcal{G}(X, P) = \hat{Y}$. Moreover, \hat{Y} should be consistent with the persona text set P , which means the NLI category between \hat{Y} and any P_i should be *entailment* or *neutral*, rather than *contradiction*, i.e., $\forall P_i \in P, NLI(\hat{Y}, P_i) \in \{E, N\}$, where E and N denote *entailment* and *neutral* respectively.

Evaluator

The proposed reinforcement learning framework consists of two components: an evaluator and a sequence generator, as illustrated in Figure 2.

As forementioned, an ideal response should be not only natural but also consistent with personas. Therefore, we consider these two attributes of responses while training the generator. More concretely, whether a response is as natural as from human (*natural* or *unnatural*) and whether a response is consistent with predefined personas (*entailment*, *neutral* or *contradiction*). These two attributes are independent of each other, and an ideal response Y^* should satisfy:

$$Y^* \in \text{Natural} \cap \text{Entailment}. \quad (1)$$

The key idea is to encourage the generator to generate responses that satisfy Formula (1). We use the policy gradient method in reinforcement learning to train the generator. We will discuss this in detail later.

Notice that our evaluator consists of two modules, rather than one jointly trained module, which is due to the difference between the two attributes. For the naturalness module, it can benefit from the adversarial training scheme (Yu et al. 2017; Li et al. 2017), as naturalness is reflected in the training data. Naturalness as a submodule in the evaluator can

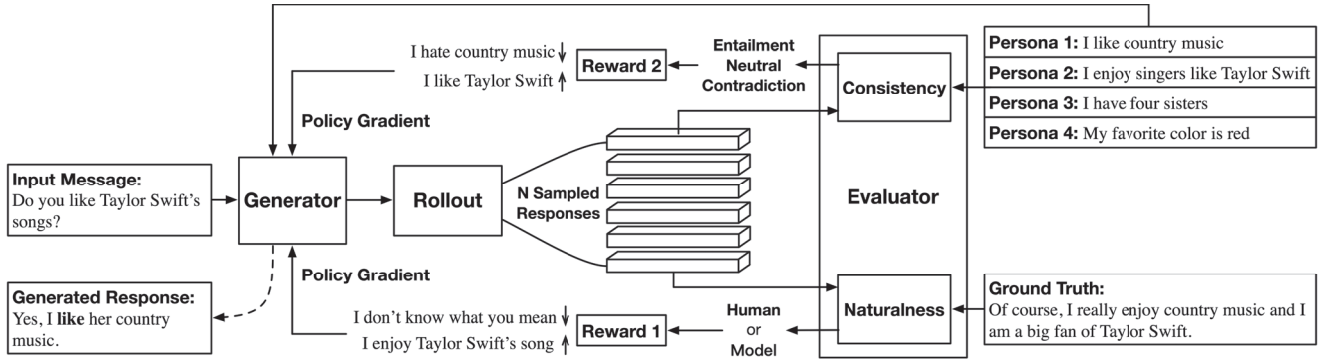


Figure 2: The overall framework of our model, which mainly consists of a generator and an evaluator. The dashed connection only appears in the generation process. The \uparrow and \downarrow denote that the generation of a response is encouraged and discouraged by the reward signals respectively.

achieve higher accuracy with adversarial training. In contrast, no labels are available in the training process to improve the performance of natural language inference.

Naturalness Module The naturalness module E_N is proposed to distinguish between human responses and model generated ones. As the generator is updating during the training process, new examples from models are generated. Therefore, we shall update the E_N .

It is safe to assume that responses from humans are always more natural than the ones from models. From this observation, we take responses from the training data as positive examples and responses from the generator G as negative examples. E_N guides the sequence generator G to predict responses closer to the examples from the training data, which is more natural.

In more detail, the naturalness module E_N is a binary classifier that takes response Y or \hat{Y} as input¹ and produces a softmax distribution over two classes, indicating whether the response is from human (natural) or model (unnatural). The input is encoded into a vector representation using a bidirectional GRU encoder, which is then fed to a highway architecture, followed by a fully-connected layer with two-class softmax function.

The objective function of E_N is to minimize the cross-entropy between the ground truth label and the predicted probability. And the reward from E_N is:

$$R_1 = E_N^+, \quad (2)$$

where E_N^+ is the output probability of \hat{Y} from the human.

Consistency Module The consistency module E_C is an NLI classifier. We introduce this module to detect the consistency in dialogues by distinguishing $\{entailment, neutral, contradiction\}$ between generated responses and the persona texts. Recent NLI models (Conneau et al. 2017; Chen et al. 2017; Gong, Luo, and Zhang 2018; Kim et

¹In our experiments, we found that the way of taking $\{X, Y\}$ as input to E_N didn't bring significant performance improvements in the accuracy of E_N , so we choose the more straightforward way.

al. 2019) are usually trained on large-scale datasets like SNLI (Bowman et al. 2015). The domain adaption problem could lead to a performance gap. Therefore, a better dataset for our task is the recently released DNLI (Welleck et al. 2019), which is in the persona-based dialogue domain.

The consistency module E_C is not updated in the adversarial training process of E_N . Due to the assumption that responses from humans are natural, E_N can always get positive examples (the human responses from training data) and negative examples (the generated responses from G) during the adversarial training process. However, as exemplified in Figure 1, a natural response does not necessarily entail persona texts and vice versa. Due to this difference, E_C cannot be iteratively updated like E_N .

In addition to exploiting NLI models in dialogue generation, another issue worth exploring is how the performance of different NLI signals affects the quality of dialogue generation. Thus in our experiments, we apply two NLI classifiers with performance differences:

- **Base Model** We use the GRU to learn the sentence representations and then put them into a multilayer perceptron (MLP) classifier. The MLP has a hidden layer with *tanh* activation and a softmax output layer in our experiments. For training, we use a multi-class cross-entropy loss. In the following sections, we abbreviate this model as E_{base} .
- **Finetuned BERT** With multilayer bidirectional Transformers (Vaswani et al. 2017), BERT (Devlin et al. 2018) has achieved state-of-the-art results on various natural language understanding tasks, including NLI. We finetune the $BERT_{base}$ model on the DNLI dataset and achieve best results compared with several other reported results on this dataset. In the following sections, we abbreviate this model as E_{bert} .

Finally, we can get the three-class confidences from the output layer of a consistency module. The reward from E_C can formulate as:

$$R_2 = \max_i E_i - \max_j C_j, i, j \in \{1, 2, \dots, |P|\}, \quad (3)$$

where E is the confidence for *entailment* and C is the confidence for *contradiction*. Index i (or j) denotes the confi-

dence is calculated between \hat{Y} and P_i (or P_j). This reward is designed to encourage *entailment* and discourage *contradiction* between \hat{Y} and P .

Reinforcement Learning

We formalize the persona consistent dialogue generation problem as a reinforcement learning task. That is, we train a generator G to produce a response $\hat{Y}_{1:t} = \{y_1, y_2, \dots, y_t\}$, where y_i represents a word in the vocabulary. At each timestep t , the state s_t is the current produced word $(y_1, y_2, \dots, y_{t-1})$, and the action a is the next selected word y_t . The policy model $G(Y_t|Y_{1:t-1})$ defines the probability that selecting the t -th word depending on the previously generated words, which is the current state.

Sequence Generator Our generator G_θ takes a form similar to Seq2Seq model with attention mechanism. The only difference is that we prepend persona texts to the input sequence, i.e., $X = \forall P_i \in P || X$, where $||$ denotes the concatenation. The same strategy is also applied to the generative model in Zhang et al. (2018).

Reward Estimation In reinforcement learning, the training objective is to maximize the accumulated future rewards. We encourage the generator to generate responses that are close to human and consistent with the predefined persona. Based on rewards from the naturalness module and consistency module, the final reward function is:

$$\mathcal{R}_{E_\phi}(\hat{Y}|X, P) = \lambda R_1 + (1 - \lambda)R_2, \quad (4)$$

where E_ϕ is our evaluator with the parameter ϕ . We train the generator $G_\theta(y_t|Y_{1:t-1})$ to generate a response from the initial state s_0 to maximize its expected final reward:

$$\mathcal{J}(\theta) = \sum_{t=1}^T G_\theta(\hat{y}_t|s_{t-1}) \cdot Q_{E_\phi}^{G_\theta}(\hat{y}_t, s_{t-1}), \quad (5)$$

where $Q_{E_\phi}^{G_\theta}(\hat{y}_t, s_{t-1})$ is the action-value function at timestep t . When there is a finished response $\hat{Y}_{1:T}$, the evaluator can provide a reward by Eq. (4) for the action-value function:

$$Q_{E_\phi}^{G_\theta}(\hat{y}_T, s_{T-1}) = \mathcal{R}_{E_\phi}(\hat{Y}_{1:T}|X, P). \quad (6)$$

Rollout Policy Our evaluator is trained to predict based on a complete sequence. Thus the reward from Eq. (6) can only be used for the final states in a response (the generation of a response must be finished), which will hurt the effectiveness of training the generator.

To evaluate the action-value Q at an intermediate state s_t ($t < T$), a common strategy is to apply rollout policy, such as Monte Carlo search, to sample the last $T - t$ words for the partially decoded response $\hat{Y}_{1:t}$ (Yu et al. 2017; Li et al. 2017; Fan et al. 2018). When applying a rollout policy, the model keeps sampling words from the generative model until the decoding is finished. This process is repeated for N times, and the average reward of the sampled responses by Eq. (4) is used as the action-value for the state s_t :

$$Q_{E_\phi}^{G_\theta}(\hat{y}_{t+1}, \hat{s}_t) = \frac{1}{N} \sum_{i=1}^N \mathcal{R}_{E_\phi}(\text{rollout}_{G_\theta}^i(\hat{Y}_{1:t+1})|X, P). \quad (7)$$

Algorithm 1 Sketch of the training procedure

Requires: generator G_θ , evaluator E_N and E_C , dialogue corpus \mathcal{S} , nli dataset \mathcal{L} .

- 1: Randomly initialize G_θ , E_N and E_C .
 - 2: Pretrain G_θ using MLE on \mathcal{S} .
 - 3: Pretrain E_N using negative samples from G_θ by minimizing the cross-entropy loss.
 - 4: Pretrain E_C on \mathcal{L} accordingly.
 - 5: **for** number of training iterations **do**
 - 6: **for** G-steps **do**
 - 7: Sample \hat{Y} from G_θ
 - 8: **for** t in 1 : T **do**
 - 9: Compute $Q(\hat{y}_t, s_{t-1})$ by Eq. (7)
 - 10: Update G_θ via Policy Gradient by Eq. (8)
 - 11: **for** teacherforce-steps **do**
 - 12: Update G_θ via MLE
 - 13: **for** E_N -steps **do**
 - 14: Sample \hat{Y} from new G_θ and sample Y from \mathcal{S}
 - 15: Update E_N via cross-entropy loss
 - 16: **return** G_θ
-

With rollout policy, the gradient of Eq. (5) can be solved by policy gradient method:

$$\nabla_\theta \mathcal{J}(\theta) = \sum_{t=1}^T \mathbb{E}[\nabla_\theta \log G_\theta(\hat{y}_t|s_{t-1}) \cdot Q_{E_\phi}^{G_\theta}(\hat{y}_t, s_{t-1})], \quad (8)$$

and the expectation \mathbb{E} can be approximated by sampling methods.

When N are large enough, MC search leads to a reasonable estimate of the sentence rewards. However, this comes at a high computational cost. When N decreases for the balance of computational time, the diversity of the sampled responses are affected, which could lead to a poor estimate. Therefore, we propose a different rollout policy: 1. at step t , the model first generates a t words' subsequence with beam search. 2. at step $t + 1$, the model keeps N different words with the top probabilities. 3. after step $t + 1$, the model continues to generate words with a sampling-based method for the partially decoded sequences with $t + 1$ words.

We apply this rollout policy for a balance of the computational time and the sample diversity. In this way, we can get diverse samples, even when N is relatively small.

Adversarial Training

As forementioned, the E_N needs adversarial training to get higher accuracy. Algorithm 1 shows the overall training process of the proposed approach, including the adversarial training of E_N . In Eq. (8), the ground-truth responses are not directly exposed to the generator in the training process. Practically, updating the generator G_θ only using the gradients from Eq. (8) leads to unstable training. The same issue is also reported in Li et al. (2017). To alleviate this issue, we follow Sutskever, Vinyals, and Le (2014) and use the Teacher Force strategy to train G_θ , via MLE loss together with rewards from the evaluator.

Experiments

Datasets

Persona-Chat We perform persona-based dialogue generation experiments on the Persona-Chat dataset (Zhang et al. 2018). The conversations are obtained from crowdworkers who were randomly paired and asked to act the part of a given persona. Also, the persona is created by another set of crowdworkers. This dataset contains 164,356 utterances in 10,981 dialogues and has a set of 1,155 personas, each consisting of four or five persona texts. The testing set contains 1,000 dialogues (15,119 utterances) and 200 never seen before personas. We set aside 968 dialogues (15,705 utterances) together with its personas from the training set for validation. The final data has 10,000/968/1,000 dialogues for train/validate/test².

As reported in Zhang et al. (2018), pretraining on larger datasets would yield better results. Thus we use another two million input-response pairs from OpenSubtitles to pretrain all models in our experiments, and we report this instead.

DNLI The recently released Dialogue Natural Language Inference dataset (Welleck et al. 2019) offers a new domain for NLI models. DNLI mainly consists of *utterance-persona* pairs, which are labeled as entailment (E), neutral (N), or contradiction (C). This dataset has 310,110/16,500/16,500 pairs for train/validate/test. Due to the length limit, we show other statistics of the DNLI dataset in the appendix.

Baselines

In the persona-based dialogue generation area, to the best of our knowledge, no previous work has explicitly modeled the consistency issue. To evaluate our model, we compared the proposed approach with the following strong models:

- **S2SA** Seq2Seq is a generative dialogue model with the context attention mechanism (Shang, Lu, and Li 2015). This is the only model **without** persona information.
- **Transformer** Transformer is one of the state-of-the-art sequence transduction models (Vaswani et al. 2017). We concatenate persona texts to the message as its input.
- **REGS** Reward for Every Generation Step is an adversarially trained model with Monte Carlo search for response generation (Li et al. 2017). We regard persona texts as dialogue context while training this model.
- **Per-S2S** This is a Seq2Seq model that prepends all persona texts to the input message (Zhang et al. 2018).
- **GPMN** Generative Profile Memory Network is a generative model that encodes persona as individual memory representations in a memory network (Zhang et al. 2018).
- **DeepCopy** DeepCopy is a hierarchical pointer network, which extends the pointer-generator network to copy tokens from relevant persona texts (Yavuz et al. 2018).

To make the following sections more concise, we abbreviate the proposed Reinforcement Learning based Consistent

²Note that the test set in ConvAI2 is different from the test set in Zhang et al. (2018) and is not publicly available.

Dialogue Generation approach as **RCDG**. Considering we have two different implementations (E_{base} and E_{bert}) of the consistency module E_C , we use $RCDG_{base}$ and $RCDG_{bert}$ to denote implemented with E_{base} and E_{bert} , respectively.

Experimental Settings

For the generator, both encoder and decoder are two-layer GRU with a hidden size 500. Embeddings of size 300 are randomly initialized and updated during training. Vocabulary size is 18,300, and other tokens are replaced with the *UNK* token. Encoder and decoder share the same vocabularies and embeddings. The model parameters are optimized using Adam with an initial learning rate of 0.0003. Learning rate decay is 0.98. Training minibatch size is 32. We set λ to 0.4 and N to 5. We implement the model in *OpenNMT-py*.

Evaluation Metrics

Consistency Evaluation First, we evaluate the persona-consistency of different models. Dziri et al. (2019) has shown that entailment techniques can be used as a surrogate for human judgment in evaluating dialogue consistency. Following this work, we employ NLI model to automatically evaluate the persona-consistency of the generated responses. For a generated response \hat{Y} and a set of persona texts $P = \{P_1, P_2, \dots, P_n\}$, an NLI model can assign an entailment category l_i to each (\hat{Y}, P_i) pair, where $l_i \in \{E, N, C\}$. Then we simulate the human evaluator in deciding the entailment category between \hat{Y} and P by:

$$NLI(\hat{Y}, P) = \begin{cases} E & \text{if } E \in L \\ C & \text{elif } C \in L \\ N & \text{otherwise} \end{cases} \quad (9)$$

where $L = \{l_1, l_2, \dots, l_n\}$.

Considering we have used BERT as a consistency evaluator in the training process, it is not fair to use the same model again for evaluation. Thus we introduce another well-performed NLI model DIIN (Gong, Luo, and Zhang 2018), as a third party, to evaluate all the dialogue models.

Dialogue Quality Evaluation Second, the quality of generated dialogues is also an essential factor to consider. We evaluate the dialogue quality of different models with the following metrics:

- **Perplexity** Following Zhang et al. (2018), we use perplexity (ppl) to measure the fluency of responses. Lower perplexity means better fluency.
- **Embedding metrics** Following Serban et al. (2016), we use Embedding Average (Ave.), Embedding Greedy (Grd.), and Embedding Extrema (Ext.) as evaluation metrics. These metrics are based on word embeddings, and they measure the relevance of a response regarding a target response. We use GoogleNews 300D word vectors.
- **Distinct** Following Li et al. (2015), we calculate the token ratios of distinct bigrams (Distinct-2, abbreviated as Dst. for convenience). We use this metric to measure how diverse the responses are.

Model		Dev	Test
Welleck <i>et al.</i> 2019	InferSent	85.82	85.68
	ESIM	86.31	88.20
Models in this work	DIIN	86.72	88.84
	E_{base}	80.48	81.26
	E_{bert}	87.67	89.14

Table 1: Accuracy of different models on the DNLI dataset.

Model	Entail.(%)	Contr.(%)
Human	48.00	1.16*
S2SA	8.37	12.94
GPMN	12.98	11.53
Per-S2S	13.27	12.19
REGS	14.08	10.83
Transformer	14.20	9.00
DeepCopy	14.62	12.17
RCDG _{base}	18.71 (28.0%)	5.93 (34.1%)
RCDG _{bert}	19.07 (30.4%)	5.56 (38.2%)

Table 2: NLI model-based persona-consistency evaluation results. *Entail.* denotes entailment (the higher the better). *Contr.* denotes contradiction (the lower the better). Best results are in bold, and the percentages in the parentheses are improvements regarding baselines’ best results. * We show some contradiction examples of Human in the appendix.

Human Evaluations In addition to the automatic evaluations, we also recruit five well-educated human judges to evaluate the generated responses.

Quantitatively evaluating the persona-consistency in generative models is a non-trivial task for humans. One major challenge is that the majority of the responses are neutral regarding the persona texts. As we can see in the first row of Table 2, even in the test set of Persona-Chat (from *Human*), half of the responses are neutral regarding the personas. This is plausible because many conversations in the real world do not ground on personas, such as greeting and question. With the limited sample size, we did not get statistically significant results in human evaluation when quantitatively evaluating the persona-consistency: the human judges labeled most of the sampled responses neutral.

Instead, we exploit human evaluations to verify the effectiveness of the model-based evaluation. Responses from all models are divided into three categories, and we randomly sample 150 response-persona pairs from each category. The judges are instructed to give a 5-scale score to each pair: **0**: definitely contradiction; **1**: potential contradiction; **2**: definitely neutral; **3**: potential entailment; **4**: definitely entailment. Note that the judges evaluate samples from each category (predicted by the DIIN), rather than from each model.

For dialogue quality, the evaluation is conducted following the usual practice. We sample 100 responses from each model and randomly shuffle them for judges. The five judges rate each response with a 3-scale criteria: **0**: persona contra-

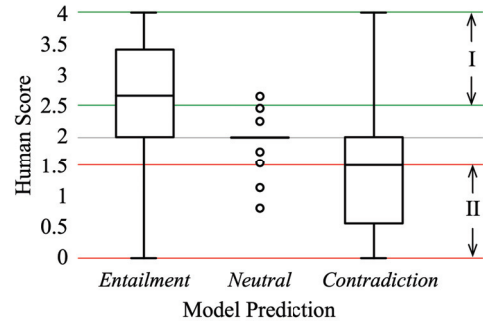


Figure 3: Boxplot of the human scores for consistency versus the model prediction categories. Three categories of model predictions are on the horizontal axis. With an average score greater or equal than 2.5, the area I is the score interval that is likely to be *Entailment*. Similarly, area II is likely to be *Contradiction*. This figure shows the correlation between human scoring and model prediction.

diction, irrelevant to the input, or grammatically broken; **1**: the response reply to the message, but is not informative; **2**: the response is relevant and informative.

Results of Consistency

Table 1 shows the performance of different models on the DNLI dataset. The first two rows of results are reprinted from Welleck et al. (2019). We implement the other three models. DIIN is the model for persona-consistency evaluation. The last two models (E_{base} and E_{bert}) are two different implementations of the consistency module.

Automatic Results We report the model-based persona-consistency evaluation results in Table 2. With the explicit modeling of persona-consistency and reinforcement learning, our approach achieves the highest entailment percentage and a much lower contradiction percentage, compared with all other baselines.

The last two rows in Table 2 are the results of our approach, with different implementations of the consistency module. Both of them outperform other baselines significantly. Our RCDG_{bert} gets better results, but this comes with higher computational costs, compared with our RCDG_{base}. The results could be interpreted to mean that the NLI signals work well, regardless of the NLI model structure.

Human Validation The human evaluation scores of each category are depicted in Figure 3. For the entailment category, more than half of the samples get an average score in the interval I, which means three judges agree that the sample is likely to be entailment or two judges agree and one of them is confident. For the neutral category, most samples get an average score of 2, and there are only a few outliers. This leads to the overlapping of the boxplot quartile lines. Figure 3 shows that the model-based evaluation is in a relatively good agreement with human evaluation. We have done a preliminary experiment in the evaluation of consistency, while a full study is beyond the scope of this paper.

Model	ppl	Ave.	Grd.	Ext.	Dst.
DeepCopy	42.8	62.1	43.2	45.1	863
Per-S2S	36.3	61.5	45.1	42.5	719
S2SA	34.8	59.8	41.9	43.5	473
GPMN	34.3	65.3	45.7	43.2	741
REGS	33.6	64.3	44.2	44.8	1009
Transformer	28.1	63.4	43.9	43.6	1505
RCDG _{base}	30.2	66.7	46.9	46.4	1289
RCDG _{bert}	29.9	66.9	47.2	46.8	1275

Table 3: Automatic results, and Dst. is scaled by 10^{-4} .

Model	0	1	2	Avg	\mathcal{K}
S2SA	0.378	0.406	0.216	0.838	0.54
GPMN	0.250	0.446	0.304	1.054	0.46
Per-S2S	0.224	0.482	0.294	1.068	0.45
REGS	0.242	0.440	0.318	1.076	0.42
DeepCopy	0.224	0.450	0.326	1.102	0.48
Transformer	0.212	0.458	0.330	1.118	0.43
RCDG _{base}	0.182	0.440	0.378	1.196	0.50
RCDG _{bert}	0.180	0.436	0.384	1.204	0.47

Table 4: The results of human evaluation for response quality, together with the Fleiss Kappa (\mathcal{K}). The \mathcal{K} coefficient between 0.41 and 0.6 means moderate agreement.

Results of Dialogue Quality

We first report the automatic evaluation results of dialogue quality in Table 3. Our methods are the best in the three embedding metrics, which indicates that our generated responses are most relevant to the ground truth. As our model is designed to address naturalness and consistency issues effectively, these results are within expectation. We notice that Transformer gets the best results in perplexity and distinct-2. It could be interpreted to mean that Transformer has a better language model compared with all other RNN based models. This also inspires us to use more advanced sequence models as our generator in future work. Except for the Transformer, our methods perform best in these RNN-based models.

We report the human evaluation results in Table 4. Our model has the highest ratio of 2, which means our generated responses are of higher quality. The Transformer also performs well in human evaluation, but it gets many 1 points. One reason could be that this model generates more questions rather than declarative sentences, which makes human judges feel less informative.

Ablation Study

As the proposed model achieves better performance than previous approaches, we conduct an analysis to gain further insight on how the integration of different modules helps the response generation. We report the results in Table 5.

As we can see, the performance of the vanilla generator is not outstanding. With the help of the adversarially trained naturalness module E_N , the dialogue quality

Model	ppl	Ave.	Dst.	E.(%)	C.(%)
Generator	34.5	63.8	823	13.23	12.08
+ E_N	32.3	64.8	1078	14.01	11.32
+ E_{base}	40.9	60.3	639	16.98	7.52
+ E_{bert}	41.3	59.9	517	17.13	7.19
+ $E_N + E_{base}$	30.2	66.7	1289	18.71	5.93
+ $E_N + E_{bert}$	29.9	66.9	1275	19.07	5.56

Table 5: Ablation test of different modules.

Persona Input	
Transformer	I'm a 12 years old girl Hi there! I am age 24, how about you?
REGS	30. I want to be a baseball announcer.
GPMN	I'm 30 years old.
DeepCopy	I am 32 and older than you.
Ours	I'm 12 years old. I am 12 years old.
Persona Input	
Transformer	I only listen to country music What sort of music do you listen to?
REGS	What kind of music do you like?
GPMN	I love country . Especially tyler farr.
DeepCopy	I love rap . Tupac is my favorite. You?
Ours	My sort of music pays well. My favorite is country , what about you?

Table 6: Sampled dialogues from different models.

is improved. Meanwhile, if we directly apply reinforcement learning without the naturalness module, although the consistency of the generated response is improved, the quality has decreased significantly, as shown in the $+E_{base}$ and $+E_{bert}$. When we integrate the naturalness module and the consistency module, the performance achieves the best.

Finally, we show some generated examples in Table 6.

Conclusion and Future Work

In this paper, we consider modeling the persona-consistency in open-domain dialogue generation by exploiting natural language inference. To this end, we cast the task as a reinforcement learning problem and leverage natural language inference signals in the deep generative model. We demonstrate the effectiveness of our approach in comparison with several baselines by experiments on the Persona-Chat dataset. In the future, we plan to apply our model to larger scale datasets. Furthermore, we plan to use more advanced generators in our approach to achieve higher performance.

Acknowledgments

The paper is supported by the National Natural Science Foundation of China under Grant No.61772153. In addition, we want to acknowledge the Heilongjiang Province Art Planning Project 2019C027 and the Heilongjiang Province Social Science Research Project 18TQB100. We also want to thank all the anonymous reviewers for their comments.

References

- Bowman, S. R.; Angeli, G.; Potts, C.; and Manning, C. D. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 632–642.
- Chen, Q.; Zhu, X.; Ling, Z.-H.; Wei, S.; Jiang, H.; and Inkpen, D. 2017. Enhanced lstm for natural language inference. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1657–1668.
- Conneau, A.; Kiela, D.; Schwenk, H.; Barrault, L.; and Bordes, A. 2017. Supervised learning of universal sentence representations from natural language inference data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 670–680.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dziri, N.; Kamaloo, E.; Mathewson, K.; and Zaiane, O. R. 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)*, 3806–3812.
- Fan, Z.; Wei, Z.; Wang, S.; Liu, Y.; and Huang, X. 2018. A reinforcement learning framework for natural question generation using bi-discriminators. In *Proceedings of the 27th International Conference on Computational Linguistics*, 1763–1774.
- Gong, Y.; Luo, H.; and Zhang, J. 2018. Natural language inference over interaction space. In *International Conference on Learning Representations*.
- Kim, S.; Hong, J.-H.; Kang, I.; and Kwak, N. 2019. Semantic sentence matching with densely-connected recurrent and co-attentive information. In *AAAI*.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, B. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.
- Li, J.; Galley, M.; Brockett, C.; Spithourakis, G. P.; Gao, J.; and Dolan, B. 2016a. A persona-based neural conversation model. *arXiv preprint arXiv:1603.06155*.
- Li, J.; Monroe, W.; Ritter, A.; Jurafsky, D.; Galley, M.; and Gao, J. 2016b. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics.
- Li, J.; Monroe, W.; Shi, T.; Jean, S.; Ritter, A.; and Jurafsky, D. 2017. Adversarial learning for neural dialogue generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2157–2169.
- Li, Z.; Jiang, X.; Shang, L.; and Li, H. 2018. Paraphrase generation with deep reinforcement learning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3865–3878.
- Serban, I. V.; Sordani, A.; Bengio, Y.; Courville, A. C.; and Pineau, J. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*, volume 16, 3776–3784.
- Shang, L.; Lu, Z.; and Li, H. 2015. Neural responding machine for short-text conversation. *arXiv preprint arXiv:1503.02364*.
- Song, H.; Zhang, W.-N.; Cui, Y.; Wang, D.; and Liu, T. 2019a. Exploiting persona information for diverse generation of conversational responses. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 5190–5196.
- Song, K.; Tan, X.; Qin, T.; Lu, J.; and Liu, T.-Y. 2019b. MASS: Masked sequence to sequence pre-training for language generation. In Chaudhuri, K., and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning Research*, 5926–5936. Long Beach, California, USA: PMLR.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, 3104–3112.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Vinyals, O., and Le, Q. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Welleck, S.; Weston, J.; Szlam, A.; and Cho, K. 2019. Dialogue natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3731–3741. Florence, Italy: Association for Computational Linguistics.
- Wu, L.; Tian, F.; Qin, T.; Lai, J.; and Liu, T.-Y. 2018. A study of reinforcement learning for neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3612–3621.
- Yavuz, S.; Rastogi, A.; Chao, G.-I.; and Hakkani-Tür, D. 2018. Deepcopy: Grounded response generation with hierarchical pointer networks. *Advances in neural information processing systems*.
- Yin, Q.; Zhang, Y.; Zhang, W.-N.; Liu, T.; and Wang, W. Y. 2018. Deep reinforcement learning for chinese zero pronoun resolution. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 569–578.
- Yu, L.; Zhang, W.; Wang, J.; and Yu, Y. 2017. Seqgan: Sequence generative adversarial nets with policy gradient. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Zhang, S.; Dinan, E.; Urbanek, J.; Szlam, A.; Kiela, D.; and Weston, J. 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)*, 2204–2213.
- Zhu, Q.; Cui, L.; Zhang, W.-N.; Wei, F.; and Liu, T. 2019. Retrieval-enhanced adversarial training for neural response generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3763–3773.