

Counterfactual Off-Policy Training for Neural Dialogue Generation

Qingfu Zhu[‡], Weinan Zhang^{‡*}, Ting Liu[‡], William Yang Wang[‡]

[‡]Harbin Institute of Technology, Harbin, China

[‡]University of California, Santa Barbara, USA

{qfzhu, wnzhang, tliu}@ir.hit.edu.cn

william@cs.ucsb.edu

Abstract

Open-domain dialogue generation suffers from the data insufficiency problem due to the vast size of potential responses. In this paper, we propose to explore potential responses by counterfactual reasoning. Given an observed response, the counterfactual reasoning model automatically infers the outcome of an alternative policy that could have been taken. The resulting counterfactual response synthesized in hindsight is of higher quality than the response synthesized from scratch. Training on the counterfactual responses under the adversarial learning framework helps to explore the high-reward area of the potential response space. An empirical study on the DailyDialog dataset shows that our approach significantly outperforms the HRED model as well as the conventional adversarial learning approaches.

1 Introduction

Open-domain dialogue generation (Shang et al., 2015a; Vinyals and Le, 2015; Sordoni et al., 2015a) intends to produce coherent responses given dialogue history. Nevertheless, it suffers from data insufficiency problem as there may exist many potential responses for a given dialogue history (Li et al., 2016). An ideal way of exploring the potential responses is to train the model by chatting with real users, which is usually time-consuming and labor-intensive in practice. Although replacing a real user with a user simulator could address the issue, the simulator only roughly approximates real user statistics, and its development process is costly (Su et al., 2016).

In contrast, humans could independently reason potential responses based on past experiences from the true environment. Having observed a response, one might naturally ask himself or herself: “What




Figure 1: An example of a counterfactual response, which is a potential response inferred in hindsight from given observed response.

would happen if I respond differently, while everything else in the environment remains the same.” Answering the question will result in a potential response (as an example in Figure 1), and it is beneficial for improving future decision making (Roese, 1997). The potential response inferred in hindsight is called a *counterfactual* response, where the concept “counterfactual” describes the posterior process of reasoning the outcome of alternative actions (i.e., a different responding policy) that could have been taken while keeping everything else unchanged (Buesing et al., 2019).

Motivated by this, we propose a counterfactual off-policy training (COPT) approach to explore potential responses. Building upon the adversarial learning framework, COPT casts a dialogue generator as a structural causal model (SCM), which describes a generation process with two ingredients: *scenarios* and *causal mechanisms* (Wright, 1920; Buesing et al., 2019). The scenario is a random noise variable that captures all unobserved yet relevant aspects of the environment, i.e., user profiles. The causal mechanism is a deterministic function that takes a scenario and dialogue history as input and outputs a response. In this way, reasoning a counterfactual response in an observed response’s environment can be achieved by feeding the scenario of the observed response into the

*Corresponding author.

causal mechanism. After generating the counterfactual response, the generator will receive a reward from a discriminator and optimize itself accordingly.

Intuitively, a counterfactual response is synthesized by grounding the model in the scenario where an observed response occurs, rather than the scenario sampled from scratch as standard adversarial learning-based approaches. This improves the quality of the synthesized responses and subsequently benefits the model that learns from the synthesis. To verify the effectiveness of our approach, we conduct experiments on the public available DailyDialog dataset (Li et al., 2017b). Experimental results show that our approach significantly outperforms previous adversarial learning-based approaches in both automatic and human evaluations. The contributions of this paper are summarized as follows:

- We connect the concept of counterfactual reasoning with the dialogue generation by casting the dialogue generation model as a structural causal model.
- Our counterfactual response is of higher quality than the response synthesized from scratch in standard adversarial learning-based dialogue generation model.
- Our approach is model-agnostic and can be applied to any adversarial learning-based dialogue generation model.

2 Related Work

Dialogue Generation Data-driven dialogue systems can be roughly divided into two categories: retrieval-based (Leuski et al., 2006; Ji et al., 2014; Yan et al., 2016) and generation-based (Shang et al., 2015b; Sordoni et al., 2015b; Vinyals and Le, 2015). Responses of retrieval-based methods come from a fixed candidate response set and thus are incapable of being customized. The generation-based methods can create new responses, but the vanilla sequence to sequence model tends to produce generic responses (Li et al., 2016).

One way to address the generic response problem is by introducing external knowledge, such as keywords (Mou et al., 2016; Zhu et al., 2019b), topics (Xing et al., 2017), persona information (Zhang et al., 2019; Song et al., 2019), and retrieved candidate responses (Song et al., 2018; Wu et al., 2019; Zhu et al., 2019a). Another way is to optimize

the architecture of networks. There are two architectures widely employed in this research line: the variational auto-encoder (Bowman et al., 2016; Zhao et al., 2017) and the generative adversarial network (Goodfellow et al., 2014; Li et al., 2017a; Zhang et al., 2018; Xu et al., 2018; Tuan and Lee, 2019). Our approach falls into the latter category. The differences between our approach and other adversarial learning-based approaches are as follows. First, we cast the dialogue generation model as an SCM to explore potential responses in the environment where observed responses occur. Second, we learn on counterfactual responses that inferred from the SCM. Third, a pre-trained behavior policy is involved during the generation process, making our approach an off-policy algorithm and benefits the exploration of potential responses.

Counterfactual Reasoning The counterfactual reasoning is a concept derived from psychology. It describes the human capacity to learn from experience by reasoning the outcome of an alternative action that could have been taken (Pearl and Mackenzie, 2018). Combined with the SCM, counterfactual reasoning improves the performance of policy evaluation in reinforcement learning (Buesing et al., 2019; Oberst and Sontag, 2019). In the area of NLP, counterfactual reasoning in previous work is mainly used for data augmentation (Qin et al., 2019; Fu et al., 2020; Kaushik et al., 2020), which rewrites the original data given a counterfactual label or condition. In this paper, we connect the concept of counterfactual reasoning with the dialogue generation and are the first to cast a generation model as an SCM under the adversarial learning framework.

3 Method

We cast a dialogue generation model as an SCM to explore potential responses by counterfactual reasoning during the training process. We will first review the concept of the SCM (Sec. 3.2), and then introduce our COPT approach (Sec. 3.3).

3.1 Notation

We use capital letters for random variables (e.g., V), lowercase letters for instances of random variables (e.g., v), and bold letters for vectors (e.g., $\mathbf{V} = \{V_1, \dots, V_N\}$). During the training process, we denote the response generated by COPT as *counterfactual response*. In contrast, the response of standard adversarial learning-based dialogue




Figure 2: An example of an SCM and an intervention. **Left:** An SCM with random variables V , scenarios U , and causal mechanisms F represented by colored squares. **Right:** A new SCM after taking an intervention on the left SCM. The original causal mechanism $f_2(V_1, U_2)$ (denoted by the orange square) is replaced by $f_2^T(V_1, U_2)$ (denoted by the purple square).

generation (i.e., REGS Li et al., 2017a) is denoted as *standard response*.

3.2 Background: Structural Causal Model

A structural causal model over random variables $V = \{V_1, \dots, V_N\}$ consists of independent noise random variables $U = \{U_1, \dots, U_N\}$ with distribution P_U and deterministic functions $F = \{f_1, \dots, f_N\}$ such that $V_i = f_i(PA_i, U_i)$, where $PA_i \subset V$ are the parents of V_i in a given DAG (Buesing et al., 2019). U is called *scenarios*, and F is called *causal mechanisms*. Figure 2 (Left) shows an example of an SCM. Each random variable V_i is determined by its parents in V , U_i , and f_i , e.g., $V_2 = f_2(V_1, U_2)$.

During the training process, we cast a dialogue generation model as an SCM over two random variables: dialogue history X and response Y . This is achieved by converting the conditional distribution $P(Y|X)$ into a deterministic function $Y = f_\pi(X, U)$ (for more details see Sec. 3.3). The scenario U is a random noise variable that captures all unobserved yet relevant properties, like user profiles. The causal mechanism is denoted as f_π to highlight the role of the policy (parameters) π of the model. The dialogue generation SCM makes it possible to sample counterfactual responses in the scenario where observed responses occur. This improves the quality of synthesized responses and subsequently helps the model to explore the high-reward area of the potential response space in the training process.

Intervention in SCM Given an SCM, an intervention T is defined as the replacement of some causal mechanisms. Figure 2 shows an example of intervention. The original causal mechanism $f_2(V_1, U_2)$ in the left SCM is replaced with $f_2^T(V_1, U_2)$, resulting in a new SCM in the right. Accordingly, intervention in our dialogue gener-

ation SCM denotes the update of the policy. For instance, the update from the behavior policy μ that generates observed responses to the target policy π that we aim to learn is the intervention of replacing $f_\mu(X, U)$ with $f_\pi(X, U)$.

Counterfactual Reasoning in SCM Given an SCM and observed a variable $V_i = v_i$, counterfactual reasoning answers the question: “*What the variable V_i would have been if I take an intervention T while remaining everything else unchanged*”. In this way, generating a counterfactual response can be seen as querying: “*Having observed a response $Y = y$, what the response Y would have been if I take an intervention by following the target policy π , rather than the behavior policy μ that generates the observed responses*”.

Typically, counterfactual reasoning answers the question by the following steps (as Figure 3):

- Observed $Y = y$ when $X = x$, infer the scenario u in hindsight from $Y = f_\mu(X, U)$.
- Take an intervention by replacing the causal mechanism $f_\mu(X, U)$ with $f_\pi(X, U)$.
- Reason a counterfactual response $\hat{y} = f_\pi(x, u)$ by the resulting new SCM.

In the following sections, we denote an observed response from the training set as Y and a model-generated response as \hat{Y} .

3.3 Counterfactual Off-Policy Training

Our COPT approach is model-agnostic and can be applied to any adversarial learning-based dialogue generation model. Without loss of generality, we take the combination of COPT and the reward for every generation step (REGS) model (Li et al., 2017a) as an example in this section. It consists of two main components: a generator G and a discriminator D .

Generator The generator G is a sequence to sequence (Seq2Seq) model (Sutskever et al., 2014) equipped with the attention mechanism (Bahdanau et al., 2015; Luong et al., 2015). During the encoding process, G reads the dialogue history into hidden states using an encoder LSTM (Hochreiter and Schmidhuber, 1997):

$$H_i = \text{LSTM}(X_i, H_{i-1}), \quad (1)$$

where X_i is the i -th word of the dialogue history, and H_i denotes the corresponding hidden state.

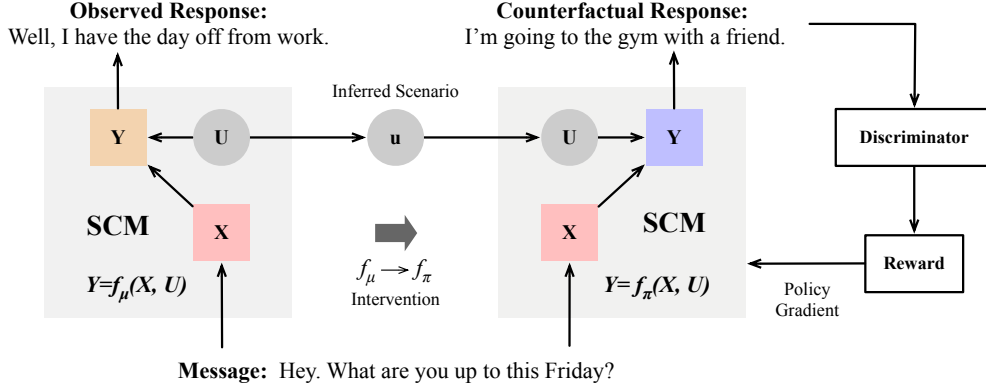


Figure 3: The architecture of our COPT approach. π is the target policy that we aim to learn. μ is the behavior policy that generates observed responses. First, we infer the scenario \mathbf{u} where the observed response occurs. Then we update the policy from μ to π , which can be seen as an intervention on the left SCM and results in the right SCM. Then, the counterfactual response is reasoned in the inferred scenario \mathbf{u} by the causal mechanism $Y = f_\pi(X, U)$.

At the j -th decoding time step, the hidden states are summarized into a context vector C_j by the attention mechanism. Subsequently, G predicts the distribution of the next word over the vocabulary by a decoder LSTM:

$$S_j = \text{LSTM}([e(\hat{Y}_{j-1}), C_j], S_{j-1}), \quad (2)$$

$$P_j^\pi(\hat{Y}_j | \mathbf{X}, \hat{\mathbf{Y}}_{1:j-1}) = \text{softmax}(S_j \cdot \mathbf{O}), \quad (3)$$

where the bracket $[\cdot, \cdot]$ denotes concatenation, and $e(\cdot)$ denotes the embedding of a word. S_j is the j -th hidden state of the decoder LSTM. \hat{Y}_{j-1} is the word generated in the previous time step. \mathbf{O} is the output matrix. We use the superscript in P_j^π to highlight the role of the policy (parameters of G).

Adversarial learning-based dialogue generation model is optimized according to the reward of responses sampled from $P_j^\pi(\hat{Y}_j | \mathbf{X}, \hat{\mathbf{Y}}_{1:j-1}) \in \mathbb{R}^{|\mathcal{V}|}$ (abbreviated as P_j^π in the following), where $|\mathcal{V}|$ is the vocabulary size. Using the Gumbel-Max Trick (Luce, 2012), the sampling process can be achieved by:

$$\hat{Y}_j = \arg \max_{\hat{y}_j} (\log P_j^\pi + U_j), \quad (4)$$

where the element of U_j follows the standard Gumbel distribution. In this way, the generator turns into a Gumbel-Max SCM (Oberst and Sontag, 2019), whose scenarios and causal mechanisms are represented by U_j and Equation 4, respectively.

From the perspective of the SCM, each response is generated in a scenario. For instance, a standard response is produced in a scenario sampled from scratch. In contrast, the scenario for a counterfactual response is inferred from an observed response $\mathbf{y} = \{y_j | y_j = \arg \max_{y_j} (\log p_j^* + \mathbf{u}_j)\}$, where

$*$ is the user's policy that generates the observed response in the true environment. However, the user's policy is not available in practice, which hinders the posterior inference of the scenario. To this end, we introduce a behavior policy μ instead and learn it by minimizing the MLE loss on observed responses. In this way, an observed response can be seen as generating in a scenario \mathbf{u}_j while following the policy μ : $y_j = \arg \max_{y_j} (\log p_j^\mu + \mathbf{u}_j)$.

According to Oberst and Sontag (2019), there are two ways to infer the scenario \mathbf{u}_j in hindsight from $y_j = \arg \max_{y_j} (\log p_j^\mu + \mathbf{u}_j)$ given y_j and μ . One way is the rejection sampling, which samples \mathbf{u}_j from the standard Gumbel distribution and rejects those where $y_j \neq \arg \max_{y_j} (\log p_j^\mu + \mathbf{u}_j)$. The other way of the posterior inference makes use of the properties of the shifted Gumbel $\mathbf{g} = \log p_j^\mu + \mathbf{u}_j$: the maximum of \mathbf{g} follows the standard Gumbel distribution and is independent with the argmax of \mathbf{g} (Maddison et al., 2014). Therefore, \mathbf{g} can be obtained by first sampling a maximum and then sampling the remaining elements truncated at the maximum. And \mathbf{u}_j is subsequently computed by subtracting $\log p_j^\mu$ from \mathbf{g} . We employ the second method to infer the scenario in COPT because it is more time-efficient than rejection sampling¹.

Given the scenario inferred from the observed response, COPT reasons the counterfactual response by feeding the dialogue history and the scenario into the SCM (Equation 4). Then the discriminator evaluates the counterfactual response and returns

¹In our experiments, using the second sampling method takes 0.79 seconds training on a batch when the batch size is 64. In contrast, the rejection sampling takes roughly 1.45 hours, making it hard to be used in practice.

a reward to the generator. Note that the counterfactual response and the SCM are utilized for the training process. During the inference process, responses are generated in the same way as the standard adversarial learning-based dialogue generation (beam search or sampling from P_j^π , we use the former in our approach) because the observed response is not available.

Discriminator The discriminator D provides a reward for each generation step. It takes as input the dialogue history \mathbf{X} , the word \tilde{Y}_j produced in the current generation step, and the prefix $\tilde{\mathbf{Y}}_{1:j-1}$ in previous steps, where $\tilde{\mathbf{Y}} \in \{\mathbf{Y}, \hat{\mathbf{Y}}\}$ can be either an observed response or a model-generated response. The output reward $D(\tilde{Y}_j|\mathbf{X}, \tilde{\mathbf{Y}}_{1:j-1})$ is the probability that \tilde{Y}_j is human-generated. Concretely, D first reads \mathbf{X} and $\tilde{\mathbf{Y}}_{1:j}$ with an encoder-decoder model. Then, it computes the reward by a Multi-Layer Perceptron (MLP), which takes as input the last hidden state of the decoder.

Adversarial Learning We train G and D under the adversarial learning framework, where G tries to fool D by generating human-like responses while D aims to distinguish between model-generated and human-generated (the observed) responses. Since a response is a sequence of discrete tokens, we pass by the gradient of D to G using the policy gradient algorithm. In this way, G converts into an agent whose partially generated response and parameters define a state and a policy, respectively. At each generation step, the agent takes an action by producing a word and observes a reward from D to update its policy.

Note that there are two policies in COPT: the target policy that we aim to learn and the behavior policy used for the reasoning of scenarios. The behavior policy is pre-trained and then froze during adversarial learning because it aims to maximize the likelihood of a fixed set of observed responses. Introducing the behavior policy makes COPT an off-policy approach because the counterfactual response, from which the target policy learns, is not entirely based on the target policy itself.

The goal of the generator is to minimize the negative expected reward: $J_G(\theta) = -\mathbb{E}_{\hat{\mathbf{Y}}_{1:j} \sim G} D(\hat{Y}_j|\mathbf{X}, \hat{\mathbf{Y}}_{1:j-1})$, where θ is the parameters of π . The gradient of θ can be derived by the

Algorithm 1 Counterfactual Off-Policy Training

- 1: Pre-train π and μ with MLE loss;
 - 2: Pre-train D on positive instances sampled from observed responses, and negative instances generated by pre-trained π ;
 - 3: **for** epoch in number of epochs **do**
 - 4: **for** g in g-steps **do**
 - 5: Infer \mathbf{u} from an observed response;
 - 6: Generate a counterfactual response in \mathbf{u} ;
 - 7: Optimize θ according to Equation 5;
 - 8: **end for**
 - 9: **for** d in d-steps **do**
 - 10: Sample positive instances from observed responses;
 - 11: Sample negative instances from responses generated by π ;
 - 12: Update ϕ according to Equation 6;
 - 13: **end for**
 - 14: **end for**
-

likelihood ratio trick (Williams, 1992):

$$\begin{aligned} \nabla J_G(\theta) = & -\mathbb{E}_{\hat{\mathbf{Y}}_{1:j} \sim G} D(\hat{Y}_j|\mathbf{X}, \hat{\mathbf{Y}}_{1:j-1}) \\ & \cdot \nabla \log G_\pi(\hat{Y}_j|\mathbf{X}, \hat{\mathbf{Y}}_{1:j-1}), \end{aligned} \quad (5)$$

where $G_\pi(\hat{Y}_j|\mathbf{X}, \hat{\mathbf{Y}}_{1:j-1})$ is the probability of generating \hat{Y}_j with the policy π given \mathbf{X} and $\hat{\mathbf{Y}}_{1:j-1}$.

The discriminator distinguishes between observed responses and model-generated responses. This is achieved by minimizing the following loss:

$$\begin{aligned} J_D(\phi) = & -\mathbb{E}_{\mathbf{Y}_{1:j} \sim \mathcal{S}} \log D(Y_j|\mathbf{X}, \mathbf{Y}_{1:j-1}) \quad (6) \\ & -\mathbb{E}_{\hat{\mathbf{Y}}_{1:j} \sim G} \log(1 - D(\hat{Y}_j|\mathbf{X}, \hat{\mathbf{Y}}_{1:j-1})), \end{aligned}$$

where ϕ is the parameters of D . As a positive instance, $\mathbf{Y}_{1:j}$ is a prefix randomly sampled from observed response set \mathcal{S} . A negative instance $\hat{\mathbf{Y}}_{1:j}$ for training D is a prefix of a standard response, rather than a counterfactual response. This is because the latter is of higher quality than the former (as shown in Sec. 4.7).

Pre-training Initialized with different parameters, π and μ are pre-trained on the training set with MLE loss. The pre-training of D depends on the specific model that COPT applied to. For example, REGS pre-trains D on the prefix of a response. In contrast, the discriminator of StepGAN (Tuan and Lee, 2019) is randomly initialized during the adversarial learning process. The overall algorithm of COPT is summarized as Algorithm 1.

| | |
|------------------------------|--------|
| Training Dialogues | 11,118 |
| Validation Dialogues | 1,000 |
| Test Dialogues | 1,000 |
| Average Tokens Per Dialogue | 114.7 |
| Average Tokens Per Utterance | 14.6 |

Table 1: Statistics of the DailyDialog dataset.

4 Experiments

4.1 Data

The experiments are conducted on the DailyDialog dataset (Li et al., 2017b).² It is a multi-turn dialogue dataset and covers various topics of daily life. The dataset has already been divided into training, validation, and test sets, as shown in Table 1. Given a dialogue that consists of K utterances, we divide it into $K-1$ instances. Each instance has at most three continuous utterances. The last utterance is the response, and the previous utterances are concatenated as the dialogue history.

4.2 Baselines

We compare COPT with the following dialogue generation models:

- HRED (Serban et al., 2016): The hierarchical recurrent encoder-decoder. An implementation by Park et al. (2018) is available³.
- REGS (Li et al., 2017a): Reward for every generation step. Its discriminator is trained on partially generated responses to provide a reward for each generation step.
- DPGAN (Xu et al., 2018): The diversity-promoting GAN introduces a language model based discriminator to encourage the generation of informative responses.⁴
- StepGAN (Tuan and Lee, 2019): The step-wise GAN trains the discriminator by maximizing the average of state-action values of observed responses. During the adversarial learning process, the discriminator assigns scores for every generation step in the same way as REGS.

Distinct from previous approaches, COPT casts a dialogue generation model as an SCM and trains it

²<http://yanran.li/dailydialog>

³<https://github.com/ctr4si/A-Hierarchical-Latent-Structure-for-Variational-Conversation-Modeling>

⁴<https://github.com/lancopku/DPGAN>

| Model | Time (s/epoch) |
|--------------|----------------|
| HRED | 84 |
| DPGAN | 608 |
| REGS | 912 |
| REGS+COPT | 1,215 |
| StepGAN | 951 |
| StepGAN+COPT | 1,244 |

Table 2: The average training time (in seconds per epoch) on a single GPU.

on counterfactual responses. It is model-agnostic and can be applied to any adversarial learning-based dialogue generation model, such as REGS, DPGAN, and StepGAN.

4.3 Training Details

We implement REGS, StepGAN, and their variants with COPT using OpenNMT (Klein et al., 2017), an open-source framework for building sequence to sequence models. We manually tune the parameters according to the perplexity on the validation set. The vocabulary consists of the most frequent 10,000 words. Including more words (up to 17,438, the total number of DailyDialog vocabulary) observes no improvement but takes more time for training. We use 300 dimensional GloVe (Pennington et al., 2014) vectors to initialize word embeddings. Both the encoder and the decoder are a two-layer LSTM in G and a single layer LSTM in D . The number of hidden units is 500.

During the adversarial learning process, we use the ADAM algorithm to alternately optimize G and D for one batch and five batches. The batch size is 64. We have tested the learning rate from 1e-6 to 1e-3. REGS+COPT and StepGAN+COPT achieve the best performance on 1e-5. The number of parameters for all the baselines is in a range of 21M to 26M. Equipping an adversarial learning baseline with COPT will introduce extra parameters with the same amount of the generator’s parameters. Contributed by the behavior policy, the parameters are learned by pre-training, and COPT will not increase the number of trainable parameters in adversarial learning. Table 2 shows the average training time. COPT may increase the training time due to the posterior inference of scenarios. But it facilitates the exploration of the high-reward area of the potential response space and subsequently improves the quality of responses.

| Model | Dist-1 | Dist-2 | BLEU-1 | BLEU-2 | BLEU-3 | BLEU-4 |
|------------------------------|--------------|--------------|--------------|-------------|-------------|-------------|
| HRED (Serban et al., 2016) | 0.011 | 0.045 | 33.01 | 4.54 | 1.05 | 0.29 |
| DPGAN (Xu et al., 2018) | 0.002 | 0.008 | 31.61 | 3.74 | 0.43 | 0.06 |
| REGS (Li et al., 2017a) | 0.021 | 0.097 | 38.39 | 6.85 | 2.00 | 0.68 |
| REGS+COPT | 0.026 | 0.116 | 39.80 | 7.67 | 2.34 | 0.77 |
| StepGAN (Tuan and Lee, 2019) | 0.018 | 0.097 | 34.76 | 6.51 | 1.97 | 0.60 |
| StepGAN+COPT | 0.026 | 0.117 | 43.83 | 8.43 | 2.66 | 0.95 |

Table 3: Automatic evaluation results of distinct-1 (Dist-1), distinct-2 (Dist-2), and BLEU scores.

| | |
|------------------|--|
| Dialogue History | A: This is our counter-offer. B: We are sincere to make business with you, but the difference between your counter-offer and our price is too great. |
| HRED | A: We can do that, but you can take a look at this one. |
| DPGAN | A: Think that the will be, but he said that they had a think we have a good job in a. Me, I am sure that I can be able to have a look at the first course. |
| REGS | A: If you are willing to withdraw money, we can offer you a private vehicle. |
| REGS+COPT | A: If you are willing to pay a commission fee, we can offer you \$5 % discount. |
| StepGAN | A: What do you mean by that? |
| StepGAN+COPT | A: In that case, we can offer you a discount. |

Table 4: An example of generated responses given dialogue history between person A and B.

4.4 Evaluation Metrics

Automatic Evaluation We evaluate the diversity and the relevance of generated responses using *distinct* (Li et al., 2016) and *BLEU* (Papineni et al., 2002), respectively. The *distinct-k* is the number of distinct *k*-grams normalized by the number of words of responses. Since BLEU might correlate weakly with human judgments of quality in the single-reference setting (Liu et al., 2016), we use the multi-reference DailyDialog test set (Gupta et al., 2019), where each instance is augmented with four human-written diverse responses.⁵

Human Evaluation The human evaluation is conducted on 200 instances randomly sampled from the test set. We create a project on Amazon Mechanical Turk (Buhrmester et al., 2016) (AMT) and employ five AMT workers to give a preference between two responses generated by our approach and a baseline.⁶ To maintain the quality of the evaluation, the task is visible to workers whose approve rate is greater than 95%, and the number of approved is greater than 500.

⁵<https://github.com/prakharguptaz/multirefeval>

⁶<https://requester.mturk.com/>

4.5 Results

Table 3 shows the results of automatic evaluation. Both REGS and StepGAN outperform HRED in distinct-1 and distinct-2, indicating that adversarial learning is beneficial for improving the diversity of responses. There is no increase in DPGAN compared with HRED in our experiments. We believe this is because the scale of the DailyDialog dataset is not large enough for sufficiently training the language model based discriminator. For the same reason, COPT is not added to DPGAN. After introducing COPT, both distinct-1 and distinct-2 in REGS and StepGAN further increase, and the improvement is significant (t-test, $p < 0.01$). This suggests that COPT is model-agnostic to adversarial learning-based approaches and helps to promote the diversity. In terms of BLEU in Table 3, both REGS and StepGAN achieve higher BLEU scores with COPT, and the improvements of BLEU-1 and BLEU-2 are significant ($p < 0.05$). This demonstrates the effectiveness of COPT in improving the relevance of responses. The less significant result of BLEU-3 and BLEU-4 is mainly due to the sparsity of tri-grams and four-grams, which are harder to be covered by references than uni-grams and bi-grams.

The human evaluation results are shown in Table 5. Our approach is clearly preferred as it

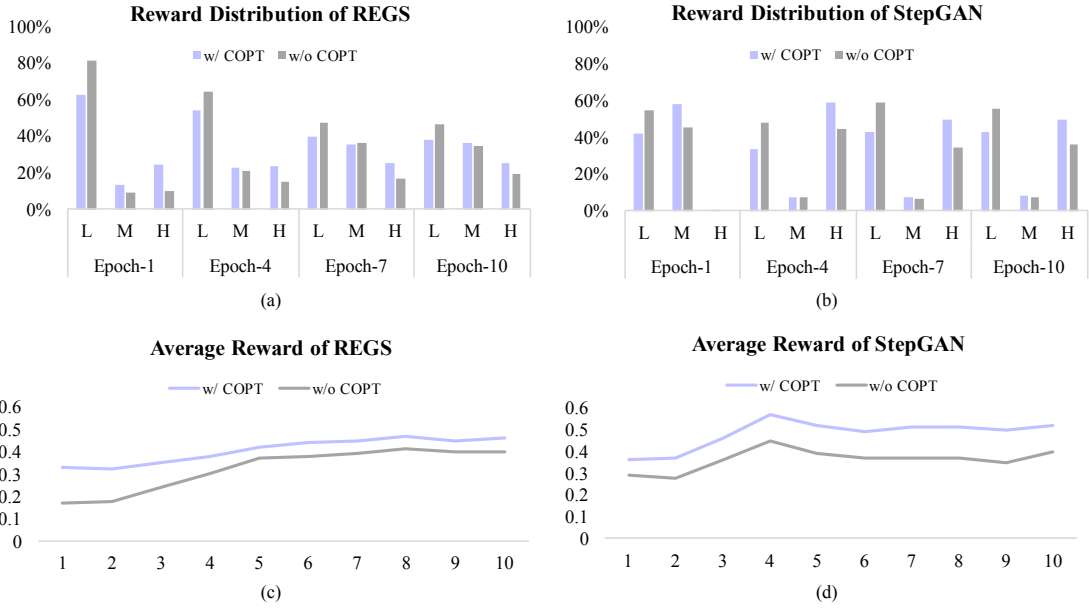


Figure 4: Reward distribution and the average reward of the counterfactual response (w/ COPT) and the standard response (w/o COPT). The y-axis in (a) and (b) is the percentage, and the x-axis corresponds to three reward intervals in different epochs, including: **L**ow reward interval [0.00, 0.33], **M**iddle reward interval (0.33, 0.66], and **H**igh reward interval (0.66, 1.00]. The y-axis in (c) and (d) is the reward, and the x-axis corresponds to epochs.

| | Win | Tie | Lose |
|--------------------------|-------|-------|-------|
| REGS+COPT vs. HRED | 63.15 | 13.20 | 23.65 |
| REGS+COPT vs. REGS | 32.28 | 43.81 | 23.92 |
| StepGAN+COPT vs. HRED | 65.95 | 15.38 | 18.67 |
| StepGAN+COPT vs. StepGAN | 45.69 | 20.91 | 33.40 |

Table 5: Wins, losses, and ties (in %) of our approach against baselines based on the human evaluation.

has more winning instances than losing instances ($p < 0.01$). The results indicate that COPT helps improve the quality of responses. Following Zhou et al. (2018) and Ke et al. (2018), we measure the agreement of annotators using inter-rater consistency. The percentage of instances that at least three annotators have the same preference (3/5 agreement) is 84.18%. The percentage for 4/5 agreement is 46.89%.

4.6 Case Study

Table 4 shows an example of responses generated by baselines and our approach. The response of DPGAN sometimes is not fluent and can be very long. We believe this is also because the scale of the DailyDialog dataset is not enough for the language model discriminator. The response of HRED is not as informative as that of our approach. Its first part is generic, and what the pronoun “that” refers to is not clear. The response of StepGAN is not infor-

mativ enough as well. In contrast, the response of REGS is quite informative, but its content is not entirely relevant to the dialogue history. After introducing COPT, the responses of REGS+COPT and StepGAN+COPT propose offering a discount to address Person B’s concern of the price, which is both informative and relevant.

4.7 Analysis

To further analyze COPT’s effectiveness in exploring the high-reward area of the potential response space during the training process, we compare the reward of a counterfactual response and a standard response on the same 10,000 randomly sampled training instances. However, the comparison between the two types of responses could be biased if their rewards are computed by different discriminators. Besides, the quality of responses is determined not only by the way they generated (with or without COPT) but also by the generator. To focus on the analysis of COPT and eliminate the bias between generators and discriminators, we generate and evaluate the two types of responses using an identical generator and its corresponding discriminator. Here, we use REGS+COPT and StepGAN+COPT as testbeds because they could generate both the two types of responses.

Figure 4 shows the distribution of rewards and

the average reward. The percentage of counterfactual responses in the high reward interval (0.66, 1.00] is higher than that of standard responses. Meanwhile, counterfactual responses generated with COPT achieve a higher average than standard responses. The results demonstrate the effectiveness of the counterfactual response in exploring the high-reward area of the potential response space during the training process. Note that the distribution and the average between different epochs are not comparable due to the update of the discriminator as the training processes.

5 Conclusion

We propose a model-agnostic approach, COPT, that can be applied to any adversarial learning-based dialogue generation models. In contrast to existing approaches, it learns on counterfactual responses inferred from the structural causal model, taking advantage of observed responses. This helps the model to explore the high-reward area of the potential response space. Experiments show that the COPT significantly improves the quality of the generated responses, which demonstrates the effectiveness of this approach.

Acknowledgments

The authors would like to thank all the anonymous reviewers for their insightful comments. This paper is supported by the National Natural Science Foundation of China under Grant No. 62076081, No. 61772153, and No. 61936010. The author from UCSB is not supported by any of the projects above.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the Third International Conference on Learning Representations*.

Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. [Generating sentences from a continuous space](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21.

Lars Buesing, Theophane Weber, Yori Zwols, Nicolas Heess, Sebastien Racaniere, Arthur Guez, and Jean-Baptiste Lespiau. 2019. Woulda, coulda, shoulda:

Counterfactually-guided policy search. In *Proceedings of the Seventh International Conference on Learning Representations*.

- Michael Buhrmester, Tracy Kwang, and Samuel D Gosling. 2016. Amazon’s mechanical turk: A new source of inexpensive, yet high-quality data?
- Tsu-Jui Fu, Xin Wang, Matthew Peterson, Scott Grafton, Miguel Eckstein, and William Yang Wang. 2020. Counterfactual vision-and-language navigation via adversarial path sampling. In *Proceedings of the 16th European Conference on Computer Vision*.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Proceedings of the Twenty-eighth Conference on Neural Information Processing Systems*, pages 2672–2680.
- Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi, and Jeffrey Bigham. 2019. [Investigating evaluation of open-domain dialogue systems with human generated multiple references](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 379–391, Stockholm, Sweden. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Journal of Neural computation*, 9(8):1735–1780.
- Zongcheng Ji, Zhengdong Lu, and Hang Li. 2014. An information retrieval approach to short text conversation. *arXiv preprint arXiv:1408.6988*.
- Divyansh Kaushik, Eduard Hovy, and Zachary C Lipton. 2020. Learning the difference that makes a difference with counterfactually-augmented data. In *Proceedings of the Eighth International Conference on Learning Representations*.
- Pei Ke, Jian Guan, Minlie Huang, and Xiaoyan Zhu. 2018. [Generating informative responses with controlled sentence function](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1499–1508, Melbourne, Australia. Association for Computational Linguistics.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72.
- Anton Leuski, Ronakkumar Patel, David Traum, and Brandon Kennedy. 2006. [Building effective question answering characters](#). In *Proceedings of the 7th SIGdial Workshop on Discourse and Dialogue*, pages 18–27.

- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017a. [Adversarial learning for neural dialogue generation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2157–2169.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017b. Dailydialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. [How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- R Duncan Luce. 2012. *Individual choice behavior: A theoretical analysis*. Courier Corporation.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective approaches to attention-based neural machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421.
- Chris J Maddison, Daniel Tarlow, and Tom Minka. 2014. A* sampling. In *Advances in Neural Information Processing Systems*, pages 3086–3094.
- Lili Mou, Yiping Song, Rui Yan, Ge Li, Lu Zhang, and Zhi Jin. 2016. [Sequence to backward and forward sequences: A content-introducing approach to generative short-text conversation](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3349–3358.
- Michael Oberst and David Sontag. 2019. Counterfactual off-policy evaluation with gumbel-max structural causal models. In *International Conference on Machine Learning*, pages 4881–4890.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Yookoon Park, Jaemin Cho, and Gunhee Kim. 2018. [A hierarchical latent structure for variational conversation modeling](#). 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 1792–1801, New Orleans, Louisiana. Association for Computational Linguistics.
- Judea Pearl and Dana Mackenzie. 2018. *The book of why: the new science of cause and effect*. Basic Books.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Lianhui Qin, Antoine Bosselut, Ari Holtzman, Chandra Bhagavatula, Elizabeth Clark, and Yejin Choi. 2019. [Counterfactual story reasoning and generation](#). 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 5043–5053, Hong Kong, China. Association for Computational Linguistics.
- Neal J Roese. 1997. Counterfactual thinking. *Psychological bulletin*, 121(1):133.
- Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015a. [Neural responding machine for short-text conversation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1577–1586, Beijing, China. Association for Computational Linguistics.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015b. [Neural responding machine for short-text conversation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1)*, pages 1577–1586.
- Haoyu Song, Wei-Nan Zhang, Yiming Cui, Dong Wang, and Ting Liu. 2019. Exploiting persona information for diverse generation of conversational responses. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*.
- Yiping Song, Rui Yan, Cheng-Te Li, Jian-Yun Nie, Ming Zhang, and Dongyan Zhao. 2018. An ensemble of retrieval-based and generation-based human-computer conversation systems. In *Proceedings of*

- the 27th International Joint Conference on Artificial Intelligence and the 23rd European Conference on Artificial Intelligence.*
- Alessandro Sordani, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015a. [A neural network approach to context-sensitive generation of conversational responses](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 196–205, Denver, Colorado. Association for Computational Linguistics.
- Alessandro Sordani, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015b. [A neural network approach to context-sensitive generation of conversational responses](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 196–205.
- Pei-Hao Su, Milica Gašić, Nikola Mrkšić, Lina M. Rojas-Barahona, Stefan Ultes, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2016. [Online active reward learning for policy optimisation in spoken dialogue systems](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2431–2441, Berlin, Germany. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Proceedings of the Twenty-Eighth Conference on Neural Information Processing Systems*, pages 3104–3112.
- Yi-Lin Tuan and Hung-Yi Lee. 2019. Improving conditional sequence generative adversarial networks by stepwise evaluation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(4):788–798.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Journal of Machine Learning*, 3(8):229–256.
- Sewall Wright. 1920. The relative importance of heredity and environment in determining the piebald pattern of guinea-pigs. *Proceedings of the National Academy of Sciences of the United States of America*, 6(6):320.
- Yu Wu, Furu Wei, Shaohan Huang, Zhoujun Li, and Ming Zhou. 2019. Response generation by context-aware prototype editing. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, volume 17, pages 3351–3357.
- Jingjing Xu, Xuancheng Ren, Junyang Lin, and Xu Sun. 2018. [Diversity-promoting GAN: A cross-entropy based generative adversarial network for diversified text generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3940–3949.
- Rui Yan, Yiping Song, and Hua Wu. 2016. Learning to respond with deep neural networks for retrieval-based human-computer conversation system. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 55–64.
- Wei-Nan Zhang, Qingfu Zhu, Yifa Wang, Yanyan Zhao, and Ting Liu. 2019. Neural personalized response generation as domain adaptation. *World Wide Web*, 22(4):1427–1446.
- Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, and Bill Dolan. 2018. Generating informative and diverse conversational responses via adversarial information maximization. In *Proceedings of the Thirty-Second Conference on Neural Information Processing Systems*, pages 1810–1820.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. [Learning discourse-level diversity for neural dialog models using conditional variational autoencoders](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664, Vancouver, Canada. Association for Computational Linguistics.
- Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. Commonsense knowledge aware conversation generation with graph attention. In *the 27th International Joint Conference on Artificial Intelligence and the 23rd European Conference on Artificial Intelligence*, pages 4623–4629.
- Qingfu Zhu, Lei Cui, Wei-Nan Zhang, Furu Wei, and Ting Liu. 2019a. [Retrieval-enhanced adversarial training for neural response generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3763–3773, Florence, Italy. Association for Computational Linguistics.
- Qingfu Zhu, Weinan Zhang, Lei Cui, and Ting Liu. 2019b. Order-sensitive keywords based response generation in open-domain conversational systems. *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)*, 19(2):1–18.