

Controllable Dialogue Generation with Disentangled Multi-grained Style Specification and Attribute Consistency Reward

Zhe Hu^{1*}, Zhiwei Cao^{2*}, Hou Pong Chan³,
Jiachen Liu¹, Xinyan Xiao¹, Jinsong Su², Hua Wu¹

¹ Baidu Inc, Beijing, China ² Xiamen University, China

³ University of Macau, Macau SAR, China

¹{huzhe01, liujiachen, xiaoxinyan, wu_hua}@baidu.com
²lines@stu.xmu.edu.cn, jssu@xmu.edu.cn ³hpchan@um.edu.mo

Abstract

Controllable text generation is an appealing but challenging task, which allows users to specify particular attributes of the generated outputs. In this paper, we propose a controllable dialogue generation model to steer response generation under multi-attribute constraints. Specifically, we define and categorize the commonly-used control attributes into global and local ones, which possess different granularities of effects on response generation. Then, we significantly extend the conventional seq2seq framework by introducing a novel two-stage decoder, which first uses a *multi-grained style specification layer* to impose the stylistic constraints and determine word-level control states of responses based on the attributes, and then employs a *response generation layer* to generate final responses maintaining both semantic relevancy to the contexts and fidelity to the attributes. Furthermore, we train our model with an attribute consistency reward to promote response control with explicit supervision signals. Extensive experiments and in-depth analyses on two datasets indicate that our model can significantly outperform competitive baselines in terms of response quality, content diversity and controllability.

1 Introduction

As a long-standing task in natural language processing, dialogue generation aims to automatically produce responses given input contexts. In this aspect, the dominant methods are neural sequence-to-sequence (seq2seq) models (Cho et al. 2014b; Vinyals and Le 2015; Shang, Lu, and Li 2015) trained to maximize the log-likelihood over responses in an end-to-end fashion. However, such generated responses not only lack controllability and interpretability (Hua and Wang 2019), but also tend to be boring with genericness and repetitiveness (Li et al. 2016a).

One important reason for these defects stems from the fact that the above models neglect the one-to-many relationship between context and response (Zhang et al. 2018a; Xu et al. 2019). As shown in Figure 1, for the same context, there exist multiple valid responses corresponding to different attributes. Generally, the conventional methods maximizing the likelihood of responses given contexts are unable

*These authors contributed equally. Work done while Zhiwei Cao was an intern at Baidu.

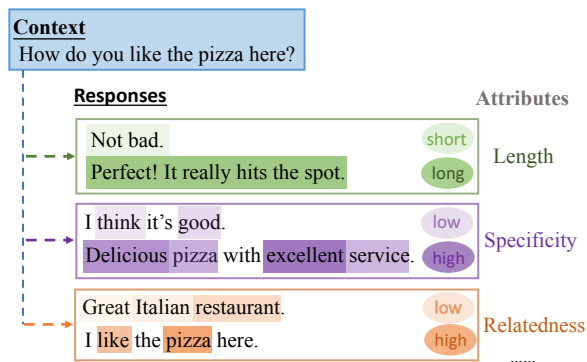


Figure 1: There exist multiple responses with different control attributes for the same input context. Length controls response from a global perspective, whereas **Specificity** and **Relatedness** can be directly reflected on each response token. Color brightness indicates the corresponding attribute values.

to explicitly learn the correspondence relationship between response and attributes. Thus, incorporating explicit control into the generation is crucial to tackle the above defects.

To achieve this goal, many efforts have been devoted to exploring control variables for dialogue generation (Zhao, Zhao, and Eskenazi 2017; See et al. 2019; Zhang et al. 2018a). However, these studies mainly focus on leveraging a single attribute to control a specific aspect, which is unsuitable for real applications that involve multiple attributes. Recently, Xu et al. (2019) propose a memory-enhanced seq2seq model to govern response generation with multiple variables. Nevertheless, there still exist three drawbacks in these studies: 1) They equally consider all attributes, but as shown in Figure 1, different attributes impact generation with varying effects (e.g., some attributes control response globally whereas some attributes possess the fine-grained influence on each response token). This hinders the model flexibility to accurately reflect the attributes on outputs. 2) Controllable generation involves a complicated disentanglement process, where the model is required to generate responses maintaining both relevancy to the contexts and fidelity to the attributes, especially under the multi-attribute constraints. However, existing dialogue models couple style specification and response generation altogether in a single

module, which leads to low interpretability and controllability. 3) Current methods are usually trained with the maximum likelihood objective and only learn weak connections between the control attributes and responses, thus often generating outputs inconformable to the attributes.

In this paper, we propose **CRAYON**, a framework to generate **C**ontrollable **R**esponse with **m**ulti-**g**rained **s**tyle specification and **a**tttribute **c**ONSistency reward. We consider important dialogue attributes including *specificity*, *sentiment*, *response-relatedness*, *question-asking* and *response length*. We further classify these attributes into two categories based on their properties: *global attributes* affecting the generation of responses from an overall perspective, and *local attributes* influencing the generation of each response word. Such classification enables our model to more flexibly and accurately control response generation at different levels.

To tackle the second drawback, we separate the control states and semantic states by dividing the generation process into two steps of style specification and surface generation, which further improves the model controllability and interpretability. Specifically, as a significant extension of conventional seq2seq method (Cho et al. 2014b), our model is equipped with a novel two-stage controlled decoder: 1) a **multi-grained style specification layer** first imposes *stylistic constraints* and generates a sequence of word-level control states based on the attributes, and 2) a **response generation layer** then handles *semantic requirements* on relevancy and produces a final response. To the best of our knowledge, our work is the first attempt that applies word-level style specification with multi-grained control to achieve better disentanglement for controllable generation.

Furthermore, we apply reinforcement learning (RL) with Markov Decision Process to optimize the model towards dedicated reward functions. During this process, we design reward functions that explicitly encourage the generated responses to satisfy the attribute constraints. By introducing direct supervision signals on attribute fidelity, our model is able to generate more diverse responses with better controllability.

We carry out experiments on two dialogue generation datasets, Persona-Chat and DailyDialog. Automatic and human evaluations show that our model significantly outperforms both controllable and non-controllable baselines towards response quality, content diversity and controllability, demonstrating the ability to disentangle the complex controllable generation under multi-attribute constraints.

2 Our Model

2.1 Task Formulation and Model Overview

Given an input context $\mathbf{x} = (x_1, \dots, x_N)$, our model aims to generate a response $\mathbf{y} = (y_1, \dots, y_M)$ that also satisfies the control attributes \mathbf{z} . During training, the model jointly learns *response generation* and *attribute prediction*. By doing so, our model is able to generate responses in both scenarios whether control attributes are explicitly given by users or automatically inferred from contexts. We will specify the details later.




Figure 2: The overview of CRAYON. \mathbf{x} , \mathbf{y} and \mathbf{z} denotes the input context, response and control attributes respectively. The attributes can be either provided by the user or automatically inferred from the contexts.

As shown in Figure 2, our model is based on an encoder-decoder framework, which mainly consists of three components: 1) a **context encoder** (§ 2.2) first converts an input context into a sequence of hidden states; 2) an **attribute predictor** (§ 2.3) predicts the control attributes \mathbf{z} according to the input context; 3) a **two-stage controlled decoder** (§ 2.4) takes as inputs the encoder outputs and attributes, and generates a final response in a controllable manner.

2.2 Context Encoder

The context encoder maps input contexts into hidden representations using a bi-directional GRU network (Cho et al. 2014a). Formally, given an input \mathbf{x} , the encoder generates a sequence of hidden states $\{\mathbf{h}_i^x\}_{i=1}^N$, which will be used for both attribute predictor and initial states of decoder layers.

2.3 Attribute Predictor

We design an attribute predictor to predict each attribute given the input context. This enables our model to generate responses with proper attributes when they are not provided during inference. Specifically, taking the last context hidden state \mathbf{h}_N^x as input, we predict the attributes as follows:

$$P(z_j|\mathbf{x}) = \text{softmax}(\text{MLP}_j(\mathbf{h}_N^x)), \quad (1)$$

where z_j is the j -th attribute, and $\text{MLP}_j(*)$ denotes a multi-layer feed-forward network for z_j .

Inspired by Lian et al. (2019), we further incorporate the posterior distribution $P'(z_j|\mathbf{x}, \mathbf{y})$ to leverage both input context and response, and train the predictor by reducing the prediction divergence between its prior $P(z_j|\mathbf{x})$ and posterior distributions to improve the prediction performance. Please refer to the supplementary material for the details.

2.4 Two-stage Controlled Decoder

Figure 3 shows the basic architecture of our two-stage controlled decoder. We first use a multi-grained style specification layer to generate a sequence of word-level control states based on the control attributes. Then, we stack a response generation layer to produce final response with both input context and control states. By separating the control states and semantic states, our model is able to address the complicated disentanglement and generate responses maintaining both semantic relevancy to the contexts and fidelity to the attributes.




Figure 3: Our two-stage controlled decoder. The *multi-grained style specification layer* first generates a sequence of control states $\{h_t^z\}$ based on the attributes, and meanwhile predicts local attribute values $\{v_{kt}^*\}$ for each local attribute (Eq. 4). The *response generation layer* then produces the final response based on the control states and context.

Multi-grained Style Specification Layer We design a novel multi-grained style specification layer to disentangle the *stylistic information* of the response based on the given attributes. With attributes z as input, we first introduce an attribute embedding layer to obtain their embeddings. Then, we concatenate all local attribute embeddings and all global ones into two control vectors: e^{z^l} and e^{z^g} , respectively. Based on these two vectors, we finally generate the control state for each response word.

Specifically, given the local control vector e^{z^l} , we first use a GRU network to calculate a sequence of local control states $\{h_t^{z^l}\}$:

$$h_t^{z^l} = \text{GRU}(h_{t-1}^{z^l}, k_t), \quad (2)$$

$$k_t = e^{z^l} \odot \sigma(\text{MLP}([h_{t-1}^{z^l}; e^{z^l}])), \quad (3)$$

where $\sigma(\ast)$ is a sigmoid nonlinear transformation and \odot represents element-wise multiplication to dynamically inject attributes into each token. Then, we concatenate each local control state with the global control vector e^{z^g} to form the final word-level control state, i.e., $h_t^z = [h_t^{z^l}; e^{z^g}]$. Through the above operations, each control state is governed by both local and global control vectors, e^{z^l} and e^{z^g} , where e^{z^l} is dynamically imposed at each time step, while e^{z^g} statically impacts the control state from a global perspective.

Particularly, to enhance the representations of the control states, we further introduce an auxiliary task of local style prediction to predict local attribute values for each response token. We define the local style prediction loss as

$$\mathcal{L}_{l\text{-style}} = - \sum_{k=1}^{|\mathcal{Z}^l|} \sum_{t=1}^M \log P(v_{kt}^* | h_t^{z^l}), \quad (4)$$

where $|\mathcal{Z}^l|$ is the number of local attributes, and v_{kt}^* represents the ground-truth value of the k -th local attribute for the t -th response word. We will introduce the label construction of v_{kt}^* in Section 4.1.

Response Generation Layer On the top of the multi-grained style specification layer, we adopt a response generation layer based on another GRU network to handle the *semantic requirements* and generate the final response.

Concretely, at the t -th step, the response generation layer consumes control state h_t^z and the previous generated token y_{t-1} to calculate the semantic hidden state h_t^r :

$$h_t^r = \text{GRU}(h_{t-1}^r, \tanh(W_w y_{t-1} + W_z h_t^z)), \quad (5)$$

where W_* are trainable parameters. We also leverage the attention mechanism (Bahdanau, Cho, and Bengio 2014) over the input context to compute a context vector c_t , and then calculate the probability of the next generated word y_t :

$$P(y_t | y_{1:t-1}) = \text{softmax}(W_g [h_t^r; c_t] + b_g), \quad (6)$$

$$c_t = \text{ATT}(H^x, h_t^r), \quad (7)$$

where W_g and b_g are trainable parameters, and $\text{ATT}(\ast)$ represents the attention operation.

3 Model Training

Our model is first trained with maximum likelihood (ML) objective. Since ML objective does not provide direct supervision on attribute fidelity, we further employ policy-based reinforcement learning (RL) with an *attribute consistency reward* to continuously train our model. By doing so, we expect our model to generate responses with better fidelity to the control attributes.

3.1 ML Training Objective

Our ML training objective mainly includes: negative log-likelihood loss (\mathcal{L}_{nl}), local style prediction loss ($\mathcal{L}_{l\text{-style}}$), constrained bag-of-words loss ($\mathcal{L}_{c\text{-bow}}$), and attribute prediction loss ($\mathcal{L}_{\text{attr}}$)¹:

¹Details of attribute prediction loss are provided in supplementary.

$$\mathcal{L}_{\text{ml}} = \mathcal{L}_{\text{nll}} + \alpha \cdot \mathcal{L}_{\text{1-style}} + \beta \cdot \mathcal{L}_{\text{c-bow}} + \gamma \cdot \mathcal{L}_{\text{attr}}, \quad (8)$$

where α , β and γ are balancing coefficients.

Specifically, we adopt the negative log-likelihood loss for response generation:

$$\mathcal{L}_{\text{nll}} = - \sum_{(\mathbf{x}, \mathbf{y}, \mathbf{z}) \in D} \log P(\mathbf{y}|\mathbf{x}, \mathbf{z}; \theta), \quad (9)$$

where \mathbf{y} , \mathbf{x} and \mathbf{z} are response, context and control attributes, and θ is model parameter set.

Inspired by Zhao, Zhao, and Eskenazi (2017), we further design a novel constrained bag-of-words loss to improve the model interpretability and controllability:

$$\mathcal{L}_{\text{c-bow}} = -\mathbf{E}_{\mathbf{z} \sim P(\mathbf{z}|\mathbf{x})} \sum_{t=1}^M \log P_b(y_t|\mathbf{x}, \mathbf{z}), \quad (10)$$

where $P_b(\ast)$ is an MLP with softmax, which transforms \mathbf{h}_N^x and e_z to a probability distribution with the dimension same as vocabulary size V . The constrained bag-of-words loss discards word orders and facilitates the model to capture the global semantics of the target response. Also, compared with the original BOW loss, the C-BOW loss optimizes the model to ground the control variables with the explicit semantic information corresponding to the attributes, which enhances the model interpretability from probabilistic perspective, as mentioned in (Fu, Feng, and Cunningham 2019).

3.2 RL with Attribute Consistency

We apply the self-critical policy gradient training algorithm (Rennie et al. 2017) to use discrete metrics as RL rewards:

$$\mathcal{L}_{\text{rl}} = -(\mathbf{R}(\mathbf{y}^s) - \mathbf{R}(\hat{\mathbf{y}})) \log P(\mathbf{y}^s|\mathbf{x}, \mathbf{z}; \theta), \quad (11)$$

where \mathbf{y}^s is a sampled response obtained by sampling words from $P(\mathbf{y}^s|\mathbf{x}, \mathbf{z}; \theta)$, and $\hat{\mathbf{y}}$ is a self-critical baseline yielded by greedily selecting words that maximize the output probability at each time step. $R(\ast)$ is the reward function.

Reward Function. To encourage responses to satisfy the control attributes, we design an attribute consistency reward. Specifically, for discrete attributes such as *Question-asking*, we give a reward of 1 if the response conforms the attribute, or give 0 otherwise. For continuous attributes such as *Specificity*, we first quantize the value into discrete bins, and measure the reverse distance between the response bin value \hat{z} and the attribute bin value z^* as the reward: $1 - |\hat{z} - z^*|/(\# - 1)$, where $\#$ is the number of bins. The final reward is written as $R(\mathbf{y}) = \sum_{z_i} R_i(\mathbf{y})$.

4 Experiments

4.1 Experimental Setups

Datasets and Preprocessing We conduct experiments on *Persona-Chat* and *DailyDialog* datasets. *Persona-Chat* (Zhang et al. 2018b) is a conversation dataset grounded

Dataset	Train	Val.	Test
Persona-Chat	122,343	14,602	14,056
DailyDialog	69,107	6,458	6,128

Table 1: Statistics on the datasets for experiments.

on personas, where each participant is assigned with a persona profile serving as background knowledge. We prepend persona texts to dialogue history as the input context. *DailyDialog* (Li et al. 2017b) is a multi-turn chit-chat dataset containing conversations about daily life. We follow the pre-processing as Bao et al. (2020), and the data statistics are summarized in Table 5.

We exploit five control attributes in our experiments: 1) **Specificity (Spe.):** following Zhang et al. (2018a) and See et al. (2019), we calculate the specificity value based on the normalized inverse response frequency (NIDF), which is then discretized into 3 bins; 2) **Sentiment (Sent.):** we run Stanford CoreNLP (Manning et al. 2014) to annotate sentiment results, which can be labeled as “*positive*”, “*neutral*” or “*negative*”; 3) **Response-relatedness (Rel.):** following See et al. (2019), we compute response-relatedness based on the cosine similarity between the embeddings of response and the last utterance, and discretize the value into 3 bins; 4) **Question-asking (Q-A):** we also consider question-asking as implemented in See et al. (2019). This binary feature is set to “*True*” if and only if at least one word in $\{how, what, when, where, which, who, whom, whose, why, ?\}$ appear in the response; and 5) **Length (Len.):** we quantize the response length into 3 bins to represent different size ranges (Fan, Grangier, and Auli 2018). Furthermore, we categorize *Sentiment*, *Length* and *Question-asking* as global attributes, and *Specificity* and *Response-relatedness* as local attributes based on their properties. Details of attribute categorization are provided in the supplementary.

For local prediction loss as in Equation 4, we construct local attribute labels for each response word. For *Specificity*, we discretize the NIDF score of each token into 6 bins. For *Response-relatedness*, we compute the cosine similarity between the word embedding and the embedding of the last utterance, and then quantize this similarity into 6 bins.

Baselines and Settings We first compare our model with the following *non-controllable baselines*: 1) SEQ2SEQ (Bahdanau, Cho, and Bengio 2014): the standard sequence-to-sequence model with attention mechanism; 2) TRANSFORMER (Vaswani et al. 2017): the standard transformer-based model which has been proved effective for text generation tasks; 3) CVAE (Zhao, Zhao, and Eskenazi 2017): a conditional variational autoencoder that captures the discourse-level diversity; 4) PER-CVAE (Song et al. 2019): a memory-augmented architecture with incorporation of explicit persona texts. Besides, we consider *controllable baselines* including 1) GTMNES2S (Xu et al. 2019): a sequence-to-sequence architecture with a goal tracking memory network²; 2) CT (See et al. 2019; Keskar

²Since the source code of GTMNES2S are not released, we implement their method. Details are provided in supplementary.

	Persona-Chat					DailyDialog				
	PPL \downarrow	BLEU-1 \uparrow	BLEU-2	Dist.1 \uparrow	Dist.2	PPL \downarrow	BLEU-1 \uparrow	BLEU-2	Dist.1 \uparrow	Dist.2
<i>Non-controllable Comparisons</i>										
SEQ2SEQ	30.83	19.95	3.26	1.63	13.34	30.08	19.59	2.07	3.55	23.09
TRANSFORMER	32.08	18.34	2.44	1.57	11.78	29.44	18.22	1.92	4.10	22.81
CVAE*	†	16.67	1.97	2.07	15.39	39.68	14.79	1.07	4.05	26.52
PER-CVAE*	40.91	17.14	1.97	2.74	22.35	-	-	-	-	-
<i>System Setting (the attributes are not provided and need to be predicted)</i>										
CT-APPEND	31.66	19.10	3.03	1.72	15.18	34.18	19.06	1.94	3.85	25.69
CT-EMB	33.53	19.57	3.19	1.68	13.98	35.79	18.40	2.04	3.82	26.48
GTMNES2S*	31.36	18.84	3.24	1.03	14.66	29.05	17.48	2.01	2.91	22.79
CRAYON (ours)	27.80	21.07	3.55	1.77	16.15	26.80	21.91	2.74	4.32	27.23
CRAYON + RL	27.77	20.64	3.40	2.05	18.49	27.50	21.66	2.84	5.02	31.33
<i>Oracle Setting (the true attributes are provided)</i>										
CT-APPEND	25.26	23.81	4.55	1.87	17.10	26.80	23.20	2.82	3.86	26.41
CT-EMB	26.08	23.85	4.72	2.06	17.94	26.94	22.28	2.74	3.92	27.55
GTMNES2S*	27.79	20.44	3.77	1.06	14.68	26.83	19.97	2.48	2.25	19.36
CRAYON (ours)	21.97	25.87	5.34	2.20	20.85	21.52	25.29	3.73	4.86	32.24
CRAYON + RL	21.87	26.00	5.50	2.46	22.83	21.87	25.01	3.86	5.55	35.58

Table 2: Experimental results on Persona-Chat and DailyDialog datasets. The best scores are in bold. †: perplexity is very unstable due to the sampling process. *: results obtained by running code released by their authors or our implementation. Our best model variants are significantly better than all comparisons ($p < 0.01$, Welch’s t -test) on perplexity and BLEU, except for BLEU-2 under the system setting on Persona-Chat. The *Non-controllable* and *System Setting* are comparable where only contexts are available during test time.

et al. 2019): the conditional training method which directly incorporates control attributes into inputs. For CT, we report the performance of two variants: CT-APPEND that appends the attributes to the input contexts and CT-EMB which concatenates the embeddings of control attributes to the decoder’s input at every step. We apply the same attributes as our model for controllable baselines.

We employ both *system setting* and *oracle setting* for controllable models. Under the system setting, the attributes are not given and need to be predicted by the model based on contexts, and we focus on evaluating the response quality. Under the oracle setting, the true attributes are provided, and we mainly focus on evaluating the model controllability. The training details are given in supplementary.

4.2 Automatic Evaluation

We adopt perplexity (PPL) to evaluate response fluency and BLEU-1/2 (Papineni et al. 2002) to evaluate relevance. For diversity, we employ Distinct-1/2 (Dist.1/2) (Li et al. 2016a) to calculate the ratio of distinct uni-grams or bi-grams. The results are shown in Table 2.

Comparison with Non-controllable Baselines. Under the system setting, where the same inputs (only contexts) are leveraged, our model achieves remarkably lower perplexity than all non-controllable baselines on both datasets. These results indicate that our model is able to generate more fluent responses. As for relevance, our model significantly outperforms the baselines in terms of BLEU-1 and BLEU-2.³ Furthermore, our RL variant produces more diverse responses with larger distinct-1/2. Although PER-CAVE achieves high distinct scores, its perplexity and BLEU scores are very low,

³We cannot conduct significant test on Dist.1/2 since we compute the inter-distinct on total generated words as implemented in (Li et al. 2016a).




Figure 4: Ablation results of our model variants.

indicating a low response quality. By contrast, our model can achieve both high response quality and diversity. All these improvements demonstrate that by incorporating important control attributes, our model can produce more appropriate and diverse responses.

Comparison with Controllable Baselines. Moreover, compared to CT and GTMNES2S that directly incorporate attributes into input contexts and rely on one single module to deal with both stylistic constraints and semantic requirements altogether, our model surpasses these baselines towards all aspects under both system and oracle settings. The results prove the effectiveness of our model on disentangling the complicated controllable generation under multi-attribute constraints. Furthermore, after applying RL training, the distinct scores are improved, showing that introducing explicit training signals on attributes benefits the model to generate more diverse responses. This is consistent with our motivation to tackle the one-to-many mapping problem via explicit control. Overall, the above results demonstrate that our model can achieve better response quality and di-

	Read.	Coh.	Rich.	Overall.
SEQ2SEQ	3.13	3.10	2.85	3.14
PER-CVAE	2.74	2.54	2.62	2.77
CT-APPEND	2.76	2.59	2.69	2.81
CT-EMB	2.59	2.63	2.57	2.59
GTMNES2S	2.74	2.38	3.13	2.77
CRAYON (ours)	3.29	3.30*	3.20	3.41*
CRAYON + RL	3.17	3.15	3.36*	3.38*

Table 3: Human evaluation with scores on a scale of 1 to 5 (best). *: significantly better than all comparisons ($p < 0.005$, Welch’s t -test). The Krippendorff’s α values for all aspects exceed 0.4, indicating general consensus to intermediate agreement.

versity than controllable baselines.

Ablation Study on Model Variants. We further analyze our model variants to quantify the contributions of various components under the system setting. As shown in Figure 4, CRAYON achieves the best performance towards both response quality and diversity. Meanwhile, removing C-BOW loss (\mathcal{L}_{c-bow}) and local style prediction loss ($\mathcal{L}_{l-style}$) leads to performance degradation. In particular, removing local style prediction loss brings significant decreases to response diversity. The results imply the effectiveness of these two losses. Finally, we consider the variant without multi-grained control, where we treat global attributes as local ones with identical labels for all tokens, and the results also drop. This result indicates that incorporating attributes in different granularities indeed help the model to effectively steer the generation process.

4.3 Human Evaluation

We also conduct human evaluation to analyze the response quality. Concretely, we select 100 random samples from Persona-Chat test set, and ask three proficient English speakers to rate on a Likert scale of 1 to 5 (best) regarding *readability* (*Read.*), *coherence* (*Coh.*), *content richness* (*Rich.*) and *overall quality* (*Overall.*).⁴ As shown in Table 3, our model variant achieves the highest scores on all aspects. After applying the attribute consistency reward, the readability and coherence slightly decrease. We find that introducing RL training makes the responses sometimes contain grammatical errors, but meanwhile the model generates more diverse and informative responses with better content richness. Notably, CT-EMB and CT-APPEND produce low results. We hypothesize that directly adding all attributes to inputs brings complex disentanglement issues and the model is hard to learn the correct relation between attributes and responses.

4.4 Analysis on Model Controllability

Accuracy of Control. We study the control accuracy in Table 4. Specifically, we compute the attribute values of each generated response with the same classifiers as in data preprocessing and compare the values with the true input attributes. First, our model trained with ML objectives outperforms CT and GTMNES2S on both datasets. This shows

⁴Detailed guidelines are given in the supplementary.

	Q-A.	Len.	Sent.	Rel.	Spe.
<i>Persona-Chat</i>					
CT-APPEND	92.57	81.12	76.15	63.48	54.49
CT-EMB	96.44	80.69	74.59	63.74	58.78
GTMNES2S	97.15	70.16	74.90	67.58	62.43
CRAYON (ours)	98.32	85.21	75.09	69.99	64.46
<i>w/o</i> $\mathcal{L}_{l-style}$	97.52	82.10	72.30	66.19	52.61
CRAYON + RL	98.73	85.31	84.28	73.87	75.10
<i>DailyDialog</i>					
CT-APPEND	92.54	76.41	72.05	71.63	52.10
CT-EMB	91.78	81.69	75.92	73.68	56.76
GTMNES2S	93.88	67.38	72.98	74.38	60.94
CRAYON (ours)	96.48	88.35	76.62	76.59	58.19
<i>w/o</i> $\mathcal{L}_{l-style}$	96.33	83.56	75.77	74.76	55.65
CRAYON + RL	98.17	90.29	82.17	79.65	67.40

Table 4: Control Accuracy (%) of each attribute.




Figure 5: Contributions of each single control attributes.

that our multi-grained style specification layer can accurately reflect the stylistic constraints and help to produce desired responses that are conformable to the attributes. Second, removing local style prediction loss leads to performance decreases, indicating the effectiveness of this auxiliary task on strengthening the style specification. Third, the observation of the normal accuracy on *Specificity* is consistent with prior work (See et al. 2019). We speculate that *Specificity* is a relatively implicit attribute, and the small-scale training data is insufficient for the model to learn the mapping between the attribute and response. Nonetheless, the results in Figure 5 empirically show that the model can still produce high-quality responses with this attribute. Fourth, after applying RL with attribute consistency reward, the results of all attributes are improved, especially for *Specificity*, proving that our model can generate responses with better fidelity to the attributes by introducing explicit supervision signals.

Ablation on Single Attribute. We study the contribution of each attribute in Figure 5. We consider one single attribute to control the response generation at each time, and present the corresponding perplexity and distinct-2. As shown in the left figure, among all attributes, *Length* and *Question-asking* are effective on both datasets. While *Sentiment* is quite useful on DailyDialog, it does not contribute much on Persona-Chat. In the right figure, we can conclude that *Specificity* is more relevant to the response diversity, while *Length* and *Response-relatedness* are less relevant. This indicates that introducing proper attributes is critical. In addition, compared with the results in Table 2, incorporating all attributes can significantly improve the model performance towards both response fluency and diversity, further proving the im-

<p>Context: R1: How are you this evening ?</p> <p>Response: ① Q-A=False, Len.=1, Sent.=Negative, Rel.=1, Spe.=0 Generated: <i>Tired</i>. Watching a movie on the weekends . ② Q-A=False, Len.=1, Sent.=Neutral, Rel.=0, Spe.=0 Generated: <i>Not bad</i> . Just listening to rap music , cooking some pizza . ③ Q-A=True, Len.=2, Sent.=Positive, Rel.=0, Spe.=2 Generated: I am <i>doing very well</i> . Just finished eating sushi . Where are you from ?</p>
<p>Context: R1: Hey guess what ? I am a jewish boy and still enjoy jazz music. R2: Cool ! I am jewish too ! I play pop music , i am the lead singer of bb2 .</p> <p>Response: ① Q-A=False, Len.=1, Sent.=Positive, Rel.=1, Spe.=0 Generated: That is <i>pretty cool</i> . I love music too ! ② Q-A=False, Len.=2, Sent.=Positive, Rel.=1, Spe.= 2 Generated: <i>Awesome</i> ! I am a jazz singer . I love music and jazz music mostly .</p>
<p>Context: R1: Good evening , what are you doing tonight ? R2: Hi, how are you ? I just got in from working on my car .</p> <p>Response: ① Q-A=False, Len.=1, Sent.=Positive, Rel.=0, Spe.=1 Generated: <i>Great</i> , i am doing great tonight . Enjoying some romance books . ② Q-A=False, Len.=1, Sent.=Negative, Rel.=0, Spe.= 0 Generated: I am <i>tired</i>, getting ready to go .</p>

Figure 6: Sample responses with corresponding attributes, generated by our model with RL on Persona-Chat test set. Words with large specificity scores are in boldface and color-coded.

portance of introducing controls into dialogue generation.

4.5 Case Study

We show sample outputs of our model with the corresponding attributes in Figure 7. Our model is able to generate proper responses with desired attributes. Specifically, given a negative sentiment, our model correctly generates “*tired*”, while for the positive sentiment our model generates phrases such as “*very well*” and “*pretty cool*”. Besides, when a high specificity is given, our model produces response words such as “*sushi*”, “*jazz*” and “*romance*”. With respect to *Question-asking*, the model learns to correctly use “?”. Overall, our model can properly reflect the stylistic constraints and generate responses conformable to the controls. This proves the effectiveness of our model regarding controllability by incorporating fine-grained controls into the two-stage controlled decoder.

5 Related Work

Dialogue Generation. Neural response generation models are mostly based on seq2seq framework (Sutskever, Vinyals,

and Le 2014; Li et al. 2016a). To improve the quality of response and address problems such as generic and safe response (Sutskever, Vinyals, and Le 2014; Mou et al. 2016), many extensions under the encoder-decoder framework have been proposed. For instance, maximum mutual information objective (Li et al. 2016a) or diverse beam search (Mou et al. 2016) are utilized to address the generic response issue during decoding. Besides, some work also tackles this problem during model training. For example, adversarial learning (Li et al. 2017a; Zhang et al. 2018c) or reinforcement learning (Li et al. 2016b; Yao et al. 2016; Xu, Wu, and Wu 2018; Saleh et al. 2020) based methods could directly improve the quality of responses. Some studies also adopt latent variables to capture the response variation and control the generation (Zhao, Zhao, and Eskenazi 2017; Gao et al. 2019a), yet these latent variables are difficult to explain and hard to control the generation with specific attributes. Sankar and Ravi (2019) leverage discrete attributes with reinforcement learning to promote response diversity. However, RL is only used for dialogue attribute prediction without direct supervision on responses.

Controllable Generation. Our work is also in line with controllable generation. Recent work incorporates control attributes such as specificity (Takayama and Arase 2020), topic (Baheti et al. 2018), dialogue acts (Xu, Wu, and Wu 2018), phrase (Wu et al. 2020), and style (Wang et al. 2017; Gao et al. 2019b). Hedayatnia et al. (2020) use a dialogue policy to control responses at the turn and sentence levels with grounded knowledge. Xu et al. (2019) propose a memory-enhanced seq2seq model with a multi-attribute controlling mechanism. Gupta et al. (2020) control dialogue generation based on the semantic frames of retrieved exemplars to improve coherency.

Conditional training (Keskar et al. 2019) and weighted decoding (See et al. 2019) are commonly used for controllable generation. However, these methods couple style specification and response realization in a single module, which makes it hard to deal with the complex disentanglement introduced by multiple constraints. Our work is different from the above methods in the following aspects: 1) We divide control attributes into global and local ones to facilitate more flexible controls in different granularities; 2) We adopt multi-grained style specification and response generation to address complicated disentanglement; 3) We design an attribute consistency reward to introduce the direct training signals on control and promote response fidelity to attributes.

6 Conclusion

In this paper, we have presented CRAYON, a controllable dialogue generation model with a novel two-stage controlled decoder and an attribute consistency reward to steer the generation process in a way of multi-grained controls. Both automatic and human evaluations indicate that our proposed model can address the complicated disentanglement and generate high-quality responses with better fidelity to controls. In the future, we plan to extend our method to other generation tasks and incorporate pre-trained models.

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv preprint arXiv:1409.0473*.
- Baheti, A.; Ritter, A.; Li, J.; and Dolan, B. 2018. Generating More Interesting Responses in Neural Conversation Models with Distributional Constraints. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3970–3980. Brussels, Belgium: Association for Computational Linguistics.
- Bao, S.; He, H.; Wang, F.; Wu, H.; and Wang, H. 2020. PLATO: Pre-trained Dialogue Generation Model with Discrete Latent Variable. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 85–96. Online: Association for Computational Linguistics.
- Cho, K.; Van Merriënboer, B.; Bahdanau, D.; and Bengio, Y. 2014a. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Cho, K.; van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014b. Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1724–1734. Doha, Qatar: Association for Computational Linguistics.
- Fan, A.; Grangier, D.; and Auli, M. 2018. Controllable Abstractive Summarization. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, 45–54. Melbourne, Australia: Association for Computational Linguistics.
- Fu, Y.; Feng, Y.; and Cunningham, J. P. 2019. Paraphrase Generation with Latent Bag of Words. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Gao, J.; Bi, W.; Liu, X.; Li, J.; and Shi, S. 2019a. Generating Multiple Diverse Responses for Short-Text Conversation. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*, volume 33, 6383–6390.
- Gao, X.; Zhang, Y.; Lee, S.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, B. 2019b. Structuring Latent Spaces for Stylized Response 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)*, 1814–1823. Hong Kong, China: Association for Computational Linguistics.
- Gupta, P.; Bigham, J. P.; Tsvetkov, Y.; and Pavel, A. 2020. Controlling Dialogue Generation with Semantic Exemplars. *arXiv preprint arXiv:2008.09075*.
- Hedayatnia, B.; Gopalakrishnan, K.; Kim, S.; Liu, Y.; Eric, M.; and Hakkani-Tur, D. 2020. Policy-Driven Neural Response Generation for Knowledge-Grounded Dialog Systems. In *Proceedings of the 13th International Conference on Natural Language Generation*, 412–421. Dublin, Ireland: Association for Computational Linguistics.
- Hua, X.; and Wang, L. 2019. Sentence-Level Content Planning and Style Specification for Neural Text 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)*, 591–602. Hong Kong, China: Association for Computational Linguistics.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical Reparametrization with Gumbel-Softmax. In *Proceedings International Conference on Learning Representations 2017*. Toulon, France: OpenReviews.net.
- Keskar, N. S.; McCann, B.; Varshney, L. R.; Xiong, C.; and Socher, R. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Klein, G.; Kim, Y.; Deng, Y.; Senellart, J.; and Rush, A. 2017. OpenNMT: Open-Source Toolkit for Neural Machine Translation. In *Proceedings of ACL 2017, System Demonstrations*, 67–72. Vancouver, Canada: Association for Computational Linguistics.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, B. 2016a. A diversity-promoting objective function for neural conversation models. In *2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL HLT 2016 - Proceedings of the Conference*, 110–119. ISBN 9781941643914.
- 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*, 1192–1202. Austin, Texas: Association for Computational Linguistics.
- Li, J.; Monroe, W.; Shi, T.; Jean, S.; Ritter, A.; and Jurafsky, D. 2017a. Adversarial Learning for Neural Dialogue Generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2157–2169. Copenhagen, Denmark: Association for Computational Linguistics.
- Li, Y.; Su, H.; Shen, X.; Li, W.; Cao, Z.; and Niu, S. 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)*, 986–995. Taipei, Taiwan: Asian Federation of Natural Language Processing.
- Lian, R.; Xie, M.; Wang, F.; Peng, J.; and Wu, H. 2019. Learning to select knowledge for response generation in dialog systems. *arXiv preprint arXiv:1902.04911*.
- Manning, C.; Surdeanu, M.; Bauer, J.; Finkel, J.; Bethard, S.; and McClosky, D. 2014. The Stanford CoreNLP Natural Language Processing Toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 55–60. Baltimore, Maryland: Association for Computational Linguistics.

- Mou, L.; Song, Y.; Yan, R.; Li, G.; Zhang, L.; and Jin, Z. 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*, 3349–3358. Osaka, Japan: The COLING 2016 Organizing Committee. ISBN 9784879747020.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 311–318. Philadelphia, Pennsylvania, USA: Association for Computational Linguistics.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543. Doha, Qatar: Association for Computational Linguistics.
- Rennie, S. J.; Marcheret, E.; Mroueh, Y.; Ross, J.; and Goel, V. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7008–7024.
- Saleh, A.; Jaques, N.; Ghandeharioun, A.; Shen, J.; and Picard, R. 2020. Hierarchical reinforcement learning for open-domain dialog. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 8741–8748.
- Sankar, C.; and Ravi, S. 2019. Deep Reinforcement Learning For Modeling Chit-Chat Dialog With Discrete Attributes. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, 1–10. Stockholm, Sweden: Association for Computational Linguistics.
- See, A.; Roller, S.; Kiela, D.; and Weston, J. 2019. What makes a good conversation? How controllable attributes affect human judgments. 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)*, 1702–1723. Minneapolis, Minnesota: Association for Computational Linguistics.
- Shang, L.; Lu, Z.; and Li, H. 2015. 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)*, 1577–1586. Beijing, China: Association for Computational Linguistics.
- Song, H.; Zhang, W.-N.; Cui, Y.; Wang, D.; and Liu, T. 2019. Exploiting Persona Information for Diverse Generation of Conversational Responses. 5190–5196.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to Sequence Learning with Neural Networks. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems*, volume 27, 3104–3112. Curran Associates, Inc.
- Takayama, J.; and Arase, Y. 2020. Consistent Response Generation with Controlled Specificity. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 4418–4427. Online: Association for Computational Linguistics.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. 30: 5998–6008.
- Vinyals, O.; and Le, Q. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Wang, D.; Jojic, N.; Brockett, C.; and Nyberg, E. 2017. Steering Output Style and Topic in Neural Response Generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2140–2150. Copenhagen, Denmark: Association for Computational Linguistics.
- Wu, Z.; Galley, M.; Brockett, C.; Zhang, Y.; Gao, X.; Quirk, C.; Koncel-Kedziorski, R.; Gao, J.; Hajishirzi, H.; Ostendorf, M.; et al. 2020. A controllable model of grounded response generation. *arXiv preprint arXiv:2005.00613*.
- Xu, C.; Wu, W.; Tao, C.; Hu, H.; Schuerman, M.; and Wang, Y. 2019. Neural Response Generation with Meta-words. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 5416–5426. Florence, Italy: Association for Computational Linguistics.
- Xu, C.; Wu, W.; and Wu, Y. 2018. Towards explainable and controllable open domain dialogue generation with dialogue acts. *arXiv preprint arXiv:1807.07255*.
- Yao, K.; Peng, B.; Zweig, G.; and Wong, K.-F. 2016. An attentional neural conversation model with improved specificity. *arXiv preprint arXiv:1606.01292*.
- Zhang, R.; Guo, J.; Fan, Y.; Lan, Y.; Xu, J.; and Cheng, X. 2018a. Learning to Control the Specificity in Neural Response Generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1108–1117. Melbourne, Australia: Association for Computational Linguistics.
- Zhang, S.; Dinan, E.; Urbanek, J.; Szlam, A.; Kiela, D.; and Weston, J. 2018b. Personalizing Dialogue Agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2204–2213. Melbourne, Australia: Association for Computational Linguistics.
- Zhang, Y.; Galley, M.; Gao, J.; Gan, Z.; Li, X.; Brockett, C.; and Dolan, B. 2018c. Generating Informative and Diverse Conversational Responses via Adversarial Information Maximization. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31, 1810–1820. Curran Associates, Inc.
- Zhao, T.; Zhao, R.; and Eskenazi, M. 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)*, 654–664. Vancouver, Canada: Association for Computational Linguistics.

A Additional Experiment Details

Preprocessing. We conduct experiments on *Persona-Chat* (Zhang et al. 2018b) and *DailyDialog* (Li et al. 2017b) datasets. We follow the preprocessing as Bao et al. (2020). The data statistics are summarized in Table 5. We further filter the samples with a reference length shorter than 3. Our vocabulary size is 14,119 on Persona-Chat and 20,014 on DailyDialog.

Dataset	Train	Val.	Test
Persona-Chat	122,343	14,602	14,056
DailyDialog	69,107	6,458	6,128

Table 5: Statistics on the datasets for experiments.

Training Details. We implement our model using OpenNMT (Klein et al. 2017). We use a two-layer bidirectional GRU for both encoder and decoder with the 300-dimensional hidden size (150 per direction). For Transformer, we apply a 4-layer encoder and decoder with 6 attention heads and the dimension of 300. We initialize word embeddings with GloVe (Pennington, Socher, and Manning 2014) and fine-tuned them during training. The dimension of attribute embeddings is 300, which is randomly initialized. The batch size is 64. We apply the Adam optimization (Kingma and Ba 2014) with warm-up steps of 500 and the maximal learning rate of $5e-4$. We implement early stopping based on the perplexity on validation set. Also we set dropout with a retention probability of 0.9 to prevent overfitting. All parameters are tuned from the validation set. To make the model more robust to attributes under system setting, we apply schedule sampling during training: we use gold control attributes with a probability of 80%, and use predicted attributes from the predictor with a probability of 20%. During the RL stage, since only incorporating RL loss leads to a degradation of response fluency, we apply a combination of RL loss and NLL loss.

For PER-CAVE, we use their implementations.⁵ For GTMNES2S, since the source code is not released, we reproduced the model according to their paper. We apply the same attributes as our model, and set the hidden size and number of layers the same as our model.

All experiments are trained on NVIDIA Tesla V100 GPUs. For our model, it takes approximately 2 hours to converge in the ML training stage, and 10 hours in the RL training stage. When selecting the model checkpoint, we choose the one based on the validation perplexity. The perplexity score on Persona-Chat validation set is 22.43 and on DailyDialog is 20.33.

B Categorization of Attributes

We categorize *Sentiment*, *Length* and *Question-asking* as global attributes, and *Specificity* and *Response-relatedness* as local attributes based on the property whether the attribute can be directly reflected on each token or it affects the generation from an overall perspective and cannot be directly

⁵<https://github.com/vsharecodes/percvae>

reflected on each token. In particular, we consider *Specificity* as a local attribute since the specificity of the whole response strongly depends on the specificity score of each word. Besides, we consider *Response-relatedness* as a local attribute because 1) the relatedness of response and context can be reflected on each response token (the sentence embeddings are calculated as the weighted sum of word embeddings) and 2) See et al. (2019) find that weighted decoding is more effective for response-relatedness, where they consider the similarity between each response token and context as the decoding feature. Thus, we incorporate the token-level relatedness into local attribute prediction to bring a performance gain.

C Attribute Predictor

We design an attribute predictor to predict each attribute given input context. This benefits our model to generate responses with proper attributes when they are not provided. Particularly, inspired by Lian et al. (2019), we train the predictor by reducing the prediction divergence between its prior distribution and posterior distribution.

Prior Distribution. Taking the last context hidden state \mathbf{h}_N^x as input, we define the prior distribution as follows:

$$P(z_j|\mathbf{x}) = \text{softmax}(\text{MLP}_j^{\text{prior}}(\mathbf{h}_N^x)), \quad (12)$$

where z_j is the j -th attribute, and $\text{MLP}_j^{\text{prior}}(*)$ denotes a multi-layer feed-forward network for z_j .

Posterior Attribute Distribution. Unlike the prior distribution solely based on input context, the calculation of the posterior distribution involves both input context and response. Specifically, we first use the same context encoder to learn the semantic representation of the response, forming the response hidden states $\{\mathbf{h}_t^y\}_{t=1}^M$. Then, we define the posterior distribution $P'(z_j|\mathbf{x}, \mathbf{y})$ as follows:

$$P'(z_j|\mathbf{x}, \mathbf{y}) = \text{softmax}(\text{MLP}_j^{\text{post}}([\mathbf{h}_N^x; \mathbf{h}_M^y])), \quad (13)$$

where $[\cdot]$ represents concatenation. Compared with the prior distribution, the posterior one better fits the true distributions since responses can also be utilized.

During training, we use gold control attributes with a probability of 80%, and use predicted attributes from the predictor with a probability of 20% as the inputs for the decoder. Since discrete control attributes are non-differentiable for gradient backpropagation, we apply the *Gumbel-Softmax Reparameterization* trick (Jang, Gu, and Poole 2017) to sample control attributes.

Training Objective. To effectively train the attribute predictor, we define a comprehensive training objective including the attribute prediction loss (\mathcal{L}_{acc}) and the prediction divergence loss (\mathcal{L}_{kl}):

$$\mathcal{L}_{\text{attr}} = \lambda_1 \mathcal{L}_{\text{acc}} + \lambda_2 \mathcal{L}_{\text{kl}}, \quad (14)$$

where λ_* are hyper-parameters used to balance the preference among the three losses. We set λ_1 as 1.0 and λ_2 as 0 in the first 1,000 steps, and then change λ_2 to 1.0.

\mathcal{L}_{acc} are used to directly train the predictor. Formally, we define the loss in the following way:

$$\mathcal{L}_{\text{acc}} = -\log \sum_j P'(z_j|\mathbf{x}, \mathbf{y}), \quad (15)$$

where e_z is the concatenation of all attribute embeddings.

Besides, \mathcal{L}_{kl} is approximated as the Kullback-Leibler divergence between the posterior distribution and the prior one:

$$\mathcal{L}_{\text{kl}} = \sum_j \text{KL}(P'(z_j|\mathbf{x}, \mathbf{y}), P(z_j|\mathbf{x})), \quad (16)$$

where we use $\text{KL}(\ast)$ to denote the Kullback-Leibler divergence function. During training, the model learns to enforce the prior and posterior distributions to be as close as possible by minimizing \mathcal{L}_{kl} . Note that the attribute prediction loss is beneficial for learning the posterior distribution of the attribute predictor, which in turn helps the prior distribution to approach the same distribution. Thus, our model can sample attributes to generate desirable responses even without input attributes.

D Human Evaluation Guideline

In our survey, each human annotator is presented with 100 samples given the persona and dialogue history. The annotators are asked to evaluate 6 responses for every sample, on a scale of 1 to 5 for the following aspects. 1 means the worst and 5 means the best. To avoid bias, we anonymize the models and shuffle the outputs to the annotators. More details are in Table 6.

- **Readability:** whether the response is fluent, complete, grammatically correct and can be understood;
- **Coherence:** whether the response is relevant with the dialogue context and consistent with the dialogue history or background knowledge;
- **Content richness:** whether the response is informative, interesting and encourages you to continue the conversation;
- **Overall quality:** this is a general assessment that whether you think it is a good response or not.

E Sample Output

We further include sample outputs of our model on Persona-Chat test set from Figure 7 and Table 7 to Table 10.

Readability:	
1	Not readable, contains fragments, missing components, or serious grammar errors
3	Contains relatively minor grammatical errors, not very fluent, but understandable
5	Correct Grammar, very fluent and complete
Coherence:	
1	Not relevant with the contexts, or inconsistent with dialogue history or background knowledge
3	Relevant to the contexts, but with minor conflicts to the dialogue history or background knowledge
5	Completely coherent and relevant to the dialogue contexts and background knowledge.
Content richness:	
1	Very generic or boring. Do not want to continue this conversation
3	Contains some information, but somewhat not interesting or informative
5	Interesting, informative and you want to continue the conversation.
Overall quality:	
1	Not a valid response
3	Can be a response, but contains some language errors or not informative
5	A good response.

Table 6: Explanations on human evaluation aspect scales.

<p>Context: <i>R1:</i> Hello there hows it going ? <i>R2:</i> Hi, i am Joanna. How are you tonight ? <i>R1:</i> Hey i am Brianna, i am okay .</p> <p>Response: ① Q-A=False, Len.=0, Sent.=Positive, Rel.=1, Spe.=0 <i>Generated:</i> That is great ! ② Q-A=False, Len.=1, Sent.=Positive, Rel.=1, Spe.=0 <i>Generated:</i> Oh that is great ! What kind of music do you enjoy ? ④ Q-A=False, Len.=2, Sent.=Positive, Rel.=1, Spe.=0 <i>Generated:</i> Great . Where are you from ? i live in South Carolina . you ever been to South Carolina ?</p>
<p>Context: <i>R1:</i> How is it going ? What are you up to ? <i>R2:</i> Good , just reading a book about traveling .</p> <p>Response: ① Q-A=True, Len.=2, Sent.=Positive, Rel.=0, Spe.=0 <i>Generated:</i> Nice ! I love reading book . How many books are you reading ? ② Q-A=True, Len.=2, Sent.=Positive, Rel.=1, Spe.= 2 <i>Generated:</i> That is good . I'm always reading about the bible .</p>
<p>Context: <i>R1:</i> Hi , how are you today ?</p> <p>Response: ① Q-A=False, Len.=1, Sent.=Positive, Rel.=1, Spe.=2 <i>Generated:</i> I am good . Just listening to nirvana music . ② Q-A=False, Len.=1, Sent.=Neutral, Rel.=1, Spe.=2 <i>Generated:</i> Just fine and eating some pizza . ② Q-A=False, Len.=1, Sent.=Negative, Rel.=1, Spe.=2 <i>Generated:</i> Tired . My job is a musician .</p>
<p>Context: <i>R1:</i> Hi ! How are you doing today ? <i>R2:</i> I am wonderful thank you . I am listening to classical music .</p> <p>Response: ① Q-A=False, Len.=1, Sent.=Neutral, Rel.=0, Spe.=1 <i>Generated:</i> I am a big singer . I am a violent person . ② Q-A=False, Len.=1, Sent.=Neutral, Rel.=2, Spe.=1 <i>Generated:</i> My favorite genre of music ! My mom loves classical music !</p>

Figure 7: Sample responses with different attributes, generated by our model with RL on Persona-Chat test set.

<p>Persona i do not drive because i live in new york i like jazz music my favorite tv show is the office i do not eat anything sweet .</p> <p>Context B: hello , how are you tonight ? are you watching the office A: hi ! i am well , thanks . i actually just got back from a hike . B: that sounds wonderful , i live in the city and do not have hiking trails . A: i am in ohio where there are lots of wooded trails . B: ohio is lovely , we went hiked blue hen and buttermilk falls once . A: yes , very nice . my 2 sons love to hike buttermilk falls . B: that hill though , lol . do you like sweets ? i do not . A: i try to not eat many sweets to set a good example for my boys . B: i do not like them that much . do you have any hobbies ? A: i am single mom so other than hiking , a lot of running my boys around to sports</p>
<p>Per-CVAE i ' m a bit of a vegan . i can not wait for that .</p> <p>Seq2Seq i enjoy running too .</p> <p>CT-append that is good . i ' m on road trips .</p> <p>CT-emb oh cool . i love to hike . i have only 2 kids , but i ' m a UNK too !</p> <p>GTMNES2S that is good . i have long hair and dyed my hair purple !</p> <p>Transformer it is a great place to go</p> <p>Ours that is fun . i like to learn about history and watching tv .</p> <p>Ours+RL that ' s a great hobby , i also enjoy listening to music while playing</p>

Table 7: Sample outputs on persona dataset.

<p>Persona i like to play racquetball i enjoy oldies music i am looking forward to retiring in five years i love trying barbeque restaurants i am an administrative assistant</p> <p>Context B: good morning ! where are you from ? A: from texas . a large house ranch was where i was raised . B: ah , good old texas ! best barbecue i ever had in my life was in texas . A: do you like horses ? i love working with horses . B: i have never tried but i am retiring in 5 years so maybe i will . A: what is your favorite musician ? mine is john denver . B: chuck berry is the best ever ! what kind of hobbies do you do ? A: one of my hobbies is working out every day . you ?</p>
<p>Per-CVAE whats your favorite season ? mine and me .</p> <p>Seq2Seq i like to cook barbeque and i love to eat barbeque .</p> <p>CT-append i love to take photos , and you ?</p> <p>CT-emb i have to to my house , UNK . and go to UNK</p> <p>GTMNES2S i play piano and dance . how often do you play with your aww ? that's cool . i like a little than i do .</p> <p>Transformer i am in california .</p> <p>Ours my favorite hobby is racquetball . and i try to go to the gym .</p> <p>Ours+RL i enjoy playing racquetball and playing volleyball .</p>

Table 8: Sample outputs on persona dataset.

Persona

i have been at the same job for a years
i do not get along with my dad
i am a nurse
i hate cats

Context

B: hi ! i am headed to work at the hospital this morning . you ?
A: just relaxing . i work retail , so not as exciting as the hospital .
B: i have been a nurse at this hospital for 35 years so not too exciting
A: i bet you have some crazy stories though .
B: the worst was when i treated a lady dressed like a cat . cats repulse me
A: i am not a cat lover either . i would rather deal with dogs .
B: my dad has 5 dogs and i do not enjoy him so dogs bother me too
A: i went to a friends wedding last week where they used their dog as a ring barer
B: haha that is silly . at least it was not a disgusting cat
A: very true . do you have a favorite band you like to listen to while working ?

Per-CVAE

i love the musical troupe

Seq2Seq

i like listening to the radio

CT-append

i like to listen to my nephew . we are a handful .

CT-emb

i love to dance . i have a band , its a UNK , the UNK

GTMNES2S

i like to listen to music when i play video games , i do have a couple of them .

Transformer

oh nice . i hope to find a family soon .

Ours

not really . i listen to music and love to listen to nickelback while i take it

Ours+RL

not really much but i like jazz . have you ever been to a concert ?

Table 9: Sample outputs on persona dataset.

Persona

i was arrested for stealing
i want children badly but not right time in my life
my parents were always gone when i was growing up
i want to be a fashion designer
my best friend took my boyfriend in high school

Context

B: hi how are you ? i am in jail right now .
A: i am good . . . not in jail , lol .
B: i stole a water bottle from the store and ended up here .
A: oh . i am a managerial consultant in wisconsin .
B: wow ! i have always wanted to be a fashion designer .
A: maybe once you get out of jail you can be .
B: yeah but my ex boyfriend from high school is dating my best friend now .
A: oh that stinks . i am single .
B: i am too , i think it is my parents fault
A: what type of music do you listen to ?
B: i really like pop . you ?
A: i like rock and pop . metallica is my favorite .

Per-CVAE

i guess . i like to be taller lol

Seq2Seq

i like all kinds of music . what kind of music do you like ?

CT-append

i like rap i listen to pop too . i listen to hip hop

CT-emb

i like rock . i am actually UNK to do , but i have a couple of my boyfriend at school .

GTMNES2S

i like simple that is not all i do sometimes , but i never bet i could go fishing !

Transformer

that is very true !

Ours

i like metallica . they are my favorite .

Ours+RL

i love metallica ! i even listen to them every day .

Table 10: Sample outputs on persona dataset.