

Learning towards Selective Data Augmentation for Dialogue Generation

Xiuying Chen^{1*}, Mingzhe Li^{2*}, Jiayi Zhang³,
Xiaoqiang Xia³, Chen Wei³, Jianwei Cui³, Xin Gao^{1†}, Xiangliang Zhang⁴, Rui Yan^{5†}

¹Computational Bioscience Research Center, KAUST

²Ant Group

³Xiaomi AI Lab

⁴University of Notre Dame

⁵Gaoling School of Artificial Intelligence, Renmin University of China

xiuying.chen@kaust.edu.sa, limingzhe.lmz@antgroup.com

Abstract

As it is cumbersome and expensive to acquire a huge amount of data for training neural dialog models, data augmentation is proposed to effectively utilize existing training samples. However, current data augmentation techniques on the dialog generation task mostly augment all cases in the training dataset without considering the intrinsic attributes between different cases. We argue that not all cases are beneficial for augmentation task, and the cases suitable for augmentation should obey the following two attributes: (1) low-quality (the dialog model cannot generate a high-quality response for the case), (2) representative (the case should represent the property of the whole dataset). Herein, we explore this idea by proposing a *Selective Data Augmentation framework* (SDA) for the response generation task. SDA employs a dual adversarial network to select the lowest quality and most representative data points for augmentation in one stage. Extensive experiments conducted on two publicly available datasets, *i.e.*, DailyDialog and OpenSubtitles, show that our framework can improve the response generation performance with respect to various metrics.

Introduction

Open-domain dialogue generation is becoming a research hotspot in the community of natural language processing due to its penitential applications (Li et al. 2019; Chen et al. 2021b). Generally, in the paradigm of deep neural networks, a large quantity of training data is required for facilitating the convergence of these models. As such, a data augmentation framework that can generate reliable training cases is the crux of building a robust dialogue model.

As shown in Figure 1(a), existing data augmentation methods for the dialog generation task mainly investigate different ways to augment all data samples without considering their distinct attributes. For example, Hou et al. (2018) augmented each case by leveraging other cases with similar semantic meaning in the training dataset, and Li et al. (2019) generated diversified versions for each query and response

in an adversarial style. However, we argue that in practice, the attributes of the training cases vary, thus, not all cases are necessary for augmentation. The augmentation of dull responses such as “I don’t know” and noisy samples with unpaired queries and responses even brings harm to the model. Taking one step further, we assume that whether each case is beneficial for augmentation should be examined from two aspects. From the generation quality aspect, the generation model may perform relatively well in some cases, for example, the cases with safe answers. Correspondingly, it is redundant and sometimes harmful to augment these cases (Csaky, Purgai, and Recski 2019). Thus, we should only focus on part of the data where the model fails to adapt to (*low-quality*). From the dataset attribute side, the quality of user-generated training data varies greatly, and noisy samples frequently appear (Cai et al. 2020). Hence, we should augment the representative cases that reflect a larger set of their properties (*representative*), instead of some noisy samples that do not represent the general attribute of the whole dataset. This is also inspired by a previous study (Schröder and Niekler 2020), which shows that training on representative cases can increase the quality of the resulting model.

Based on this assumption, in this work, we propose a novel *Selective Data Augmentation* framework, namely SDA, to accurately select the most informative data points from the training dataset by simultaneously considering the generation quality and representativeness. The overview is illustrated in Figure 1(b). The dialog selector is required to select the samples maximizing the distance between generated responses and original responses (*low-quality*) while minimizing the distance between selected samples and original samples (*representative*).

Concretely, we use a dual generative adversarial (Dual-GAN) framework to assist the dialog selector in the distance measurement between deep feature representations. From the generation quality side, a discriminator tries to discriminate between the generated response and the ground-truth response, while the dialog selector aims to trick the discriminator. If the generated responses cannot fool the discriminator, then the selected samples have low quality. From the representativeness side, we measure the distance by the re-

*These authors contributed equally.

†Corresponding authors.

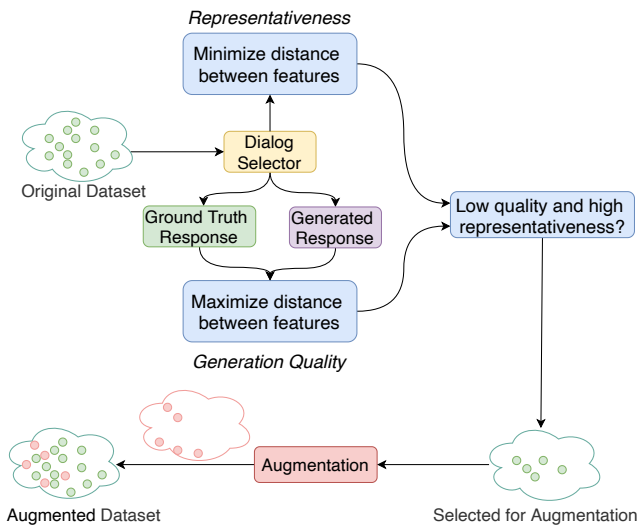


Figure 1: Our goal is to simultaneously select the lowest quality and the most representative cases in the training dataset for augmentation (best viewed in color).

construction process. If the selected samples successfully reconstruct the original data, then the selected cases have high representativeness. Concretely, the samples selected by the dialog selector are sent to a variational autoencoder (VAE), which embeds the features of selected samples into the same latent space and then reconstructs them. The reconstructed features are fed to the representativeness discriminator, which tries to discriminate between the original samples and the reconstructed samples. If the selected samples successfully fool the discriminator, then the selected samples have high representativeness. In this way, the dialog selector is encouraged to take both generation quality and representativeness into consideration during data selection.

Our main contributions can be summarized as follows: (1) We propose the selective data augmentation task, which aims to select suitable training cases for augmentation. (2) We propose a dual adversarial framework for *Selective Data Augmentation* (SDA), which can simultaneously learn to select the lowest quality and most representative data points for augmentation. (3) Extensive experiments conducted on two public dialog datasets show that our approach can improve the dialog generation performance. We also show the universality of our framework for the story generation task.

Related Work

Dialog Generation. Existing approaches to improve neural dialogue generation models mainly target building more powerful learning systems, using extra information such as conversation topics (Zou et al. 2021), persona profile (Chan et al. 2019), user emotions (Song et al. 2019), out-sourcing knowledge (Li et al. 2021), or pretrained models (Tuan et al. 2021). Another popular framework for dialogue generation concentrates on using VAE (Zhao, Zhao, and Eskénazi 2017), in which a latent variable is introduced to benefit the dialogue model with more diverse response generation. As

the GAN framework facilitates training the generator,

Data Augmentation. In the paradigm of deep learning, data augmentation is an effective way to boost the performance of neural models. To name a few, Kurata, Xiang, and Zhou (2016) proposed to generate labeled data with the decoder LSTM based on the perturbed encoded vector for the semantic slot filling task. Andreas (2020) designed a simple data augmentation protocol aimed at providing a compositional inductive bias in conditional and unconditional sequence models. Kobayashi (2018) and Wu et al. (2019) showed that contextual augmentation using label-conditional language models helps text classification tasks. In terms of dialog generation task, Li et al. (2019) proposed a generative model to augment existing data, where the CVAE was employed as the generator to output more training data with diversified expressions. Louvan and Magnini (2020) proposed Lightweight augmentation, a set of word-span and sentence-level augmentation methods for low-resource slot filling and intent classification. The most recent work was proposed by Cai et al. (2020), where they proposed to augment and reweight all learning samples.

Methodology

Open-domain dialogue generation involves generating a response $R^i = (r_1^i, \dots, r_j^i, \dots, r_m^i)$ for a user-issued query $Q^i = (q_1^i, \dots, q_k^i, \dots, q_{m'}^i)$, where r_j^i refers to the j -th word in the response in i -th case, and q_k^i denotes the k -th word in the query in i -th case. m and m' are the word length of a response and a query, respectively. The entire dialogue system is trained under D , i.e., maximizing the $P(R^i|Q^i)$, where $D = \{(Q^i, R^i)\}_{i=1}^N$ is the dataset and N refers to the number of training query-response pairs. For the data augmentation task, the original dataset D is increased to $D' = \{(Q^i, R^i)\}_{i=1}^{N'}$, where N' is the data size after augmentation. In our selective data augmentation task, we aim to select suitable cases suitable for augmentation and increase the data size from N to N' . Correspondingly, the response generation changes from $\text{argmax}P(R|Q, D)$ to $\text{argmax}P(R|Q, D')$.

Overall, the *Dialog Selector* assigns select weights to existing samples. To select the lowest quality and most representative cases, we propose two discriminators to assist this process, as shown in Figure 2. Firstly, a *Generation Quality Discriminator* (GQD) discriminates between the ground-truth response and the generated response. The dialog selector will assign high weights to cases that cannot fool GQD. Secondly, to examine the representativeness of the selected samples, a reconstruction and a discrimination process are employed. The intuition is that if the selected cases can successfully reconstruct the original data, then the selected cases are representative. Concretely, a reconstructor first embeds the selected samples into the same latent space and then reconstructs them. A *Representativeness Discriminator* (RD) is then required to classify whether the input belongs to the original samples or the selected samples. Dialog selector will assign high weights to cases that fool RD.

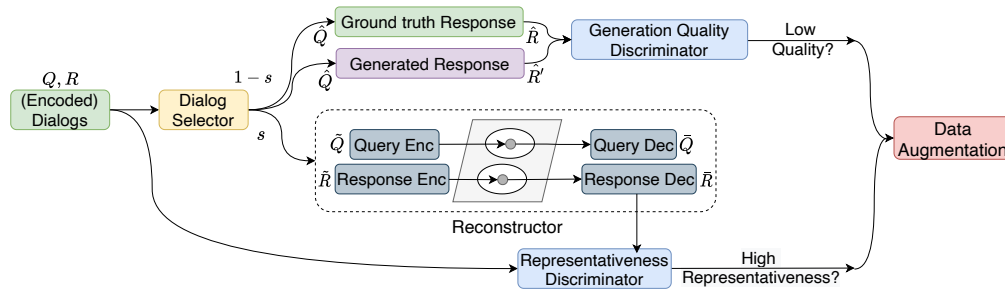


Figure 2: Major components of our approach. The *dialog selector* selects samples that will be examined by two discriminators. The *generation quality discriminator* examines the generation quality of the selected cases, and the *representativeness discriminator* examines the representativeness of the selected samples. The samples with low quality and high representativeness, *i.e.*, high s score, will be selected for data augmentation.

Dialog Selector

We first employ a bi-directional recurrent neural network (Bi-RNN) to model the temporal interactions between words in the query and response, denoted by $Q = \{h_1^{q_i}, \dots, h_{m'}^{q_i}\}$ and $R = \{h_1^{r_i}, \dots, h_m^{r_i}\}$, respectively. i denotes the sample index. The final hidden state $h_{m'}^{q_i}$ and $h_m^{r_i}$ denotes the overall representation for the query and response. The dialog selector adopts a simple architecture, consisting of a 5-layer multi-layer perceptron (MLP) with Xavier initialization, to map the input feature to a score:

$$s^i = \sigma(\text{MLP}_a([h_{m'}^{q_i}; h_m^{r_i}])), \quad (1)$$

where $[\cdot]$ denotes the concatenation operation and σ denotes the sigmoid function.

In the next subsections, we will propose dual discriminators to assist the selection process. As preparation, the original representation Q and R are weighted using these scores, and we use \hat{R} to denote this process:

$$\hat{R}^i = (1 - s^i)R^i, \tilde{R}^i = s^i R^i. \quad (2)$$

\hat{R}^i is employed for quality discrimination, and \tilde{R}^i is used for representativeness discrimination. Note that we use $1 - s^i$ and s^i as the weights for quality and representativeness branches, respectively, to ensure the optimization of these two terms in the same direction. Notations for \hat{Q}^i and \tilde{Q}^i are similar.

To prevent the selector from assigning equal importance to all data points, we employ a length regularizer loss \mathcal{L}_{LR} to limit the number of selected elements, and use a determinantal point process (DPP) loss \mathcal{L}_{dpp} (Szegedy et al. 2015) to ensure the diversity of selected data points:

$$\mathcal{L}_{\text{LR}} = \left\| \sigma - \frac{1}{N} \sum_{i=1}^N s^i \right\|_2, \mathcal{L}_{\text{dpp}} = -\log(P(\mathbf{s})). \quad (3)$$

For \mathcal{L}_{LR} , σ represents the percentage of cases for subset selection. For \mathcal{L}_{dpp} , $P(\mathbf{s})$ is a probability that DPP assigns to the select \mathbf{s} . We compute $P(\mathbf{s}; L) = \frac{\det(L(\mathbf{s}))}{\det(L+I)}$, where L is an $N \times N$ similarity matrix between every case, I is an identity matrix, and $L(\mathbf{s})$ is a smaller square matrix, cut down from L given \mathbf{s} . For i -th case and j -th case, the pairwise similarity values are defined as $L_{i,j} = s^i s^j [h_{m'}^{q_i}; h_m^{r_i}] [h_{m'}^{q_j}; h_m^{r_j}]$.

Generation Quality Discriminator

Generation quality discriminator (GQD) aims to evaluate whether the generated responses are feasible for a given query. We achieve this by measuring the matching degree between query-response pairs in an adversarial fashion. The weighted ground truth query-response pair is treated as a positive case, while the query with the generated response pair is the negative case. Concretely, for the positive pair, we concatenate the weighted ground-truth response \hat{R}^i with the weighted query \hat{Q}^i . Then, a fully-connected neural network with a sigmoid activation function is utilized to flatten the feature matrix, resulting in the final matching score $m_g^i \in (0, 1)$. The matching score m_f^i between the negative instance \tilde{R}^i and query \tilde{Q}^i is also calculated as the above-mentioned procedure, except that we have a dimension transformation on the generated response to align it with \tilde{Q}^i . Note that our framework does not rely on specific response generation models, and in our case, we employ LSTM-based RNN as the generator.

In the paradigm of GAN, the training objective of GQD is to maximize the matching score of positive instances and minimize the negative ones, while the dialog selector is optimized by aiming to maximize the matching score the generated response:

$$\mathcal{L}_D = -\sum_{i=1}^N (\log(1 - m_f^i) + \log(m_g^i)), \quad (4)$$

$$\mathcal{L}_G = -\sum_{i=1}^N (\log(m_f^i)). \quad (5)$$

The dialog selector will learn to assign high weights, *i.e.*, $1 - s^i$, to samples that are difficult for GQD to identify, which leads to a low s^i score. In other words, the cases that obtain high s^i scores have low quality.

Reconstructor

In the next two subsections, we introduce our representativeness selection process, where the samples that can be used to reconstruct the original dataset are selected. This is inspired by Mahasseni, Lam, and Todorovic (2017), where they address the problem of finding representative images in a video. Their key idea is to find a subset of images that can

be used to reconstruct the whole video. In this work, we extend this ideology from video level to dataset level to find representation cases instead of images. Since dialog data is paired, we use the query to illustrate this process.

Our reconstructor takes the form of VAE, which is commonly used to effectively learn feature representations (Zhao, Zhao, and Eskénazi 2017). VAE defines a posterior distribution over the observed data, given an unobserved latent variable. Overall, VAE consists of an encoder and a decoder. The encoder maps the weighted query \tilde{Q}^i to a latent space e , and the decoder reconstructs the query from e .

Concretely, the encoder computes posterior distributions $q_\theta(e|\tilde{Q}^i)$, where the latent representations e is sampled. The reconstruction process can be formulated as $p_\theta(\tilde{Q}^i|e)$, representing the probability of generating input \tilde{Q}^i conditioned on e . Herein θ represents the parameters of the above encoders and reconstruction decoder. Because of the intractable integral of the marginal likelihood $p_\theta(\tilde{Q}^i)$, the posterior $q_\theta(e|\tilde{Q}^i)$ is simulated by variational approximation $q_\phi(e|\tilde{Q}^i)$, where ϕ is the parameters for q . When learning the VAE, the objective is to maximize the variational lower bound of $\log p_\theta(\tilde{Q}^i)$:

$$\mathcal{L}_{VAE}^q = \text{KL}(q_\phi(e|\tilde{Q}^i)||p_\theta(e)) - \mathbb{E}_{q_\phi(e|\tilde{Q}^i)}[\log p_\theta(\tilde{Q}^i|e)],$$

where the KL denotes KL-divergence, the regularization for encouraging the approximated posterior $q_\phi(e|\tilde{Q}^i)$ to be close to the prior $p_\theta(e)$, *i.e.*, standard Gaussian distribution. $\mathbb{E}[\cdot]$ is the reconstruction loss conditioned on the approximated posterior $q_\phi(e|\tilde{Q}^i)$.

We denote the reconstructed query as \bar{Q}^i , and reconstructed response as \bar{R}^i .

Representativeness Discriminator

The discrimination processes for query and response are similar, and we use the query to illustrate this process. Representativeness discriminator (RD) takes \bar{Q}^i and \tilde{Q}^i as input and aims to classify them into two distinct classes (*i.e.*, selected or original). RD adopts the same architecture in GQD except that it does not have the dimension transformation. We omit the details here due to limited space. RD aims to maximize the correct matching result, while the dialog selector aims to select cases that can fool RD. If RD cannot distinguish the selected cases from the original one, the dialogs with high s^i scores are seen to have good representativeness to the dataset. Hence, the dialog selector will learn to assign high s^i score to the representative case to fool RD.

Experiment

Experiment Setup

Datasets. Following Cai et al. (2020), we conduct experiments on two English conversation datasets: (1) *DailyDialog* (Li et al. 2017), a collection of real-world dialogues widely used in open-domain dialogue generation. This is a multi-turn dataset, and we treat each turn as a training pair in this work. The overlapping pairs are removed from the dataset. (2) *OpenSubtitles* (Lison and

Tiedemann 2016), a group of human-human conversations converted from movie transcripts. We split the *DailyDialog* dataset to 54,889/6,005/5,700, and *OpenSubtitles* to 64,000/8,000/8,000.

Implementation Details. (1) *Hyperparameter setting:* We implement our models in TensorFlow on an NVIDIA GTX 1080 Ti GPU. We truncate the input dialog to 20 words, the minimum decoding step is 10, and the maximum step is 30. The default σ in Equation 1 is set to 0.6 except in the augmentation percentage analysis. The batch size is set to 16, and we limit the vocabulary size to 50K. (2) *Optimization techniques:* We employ a set of techniques to deal with the posterior collapsed problem in VAE (Bowman et al. 2016) including Bag-Of-Words (BOW) and KL annealing. We increase the kl loss coefficient by 0.5 every 10,000 batches. Readers can refer to work by (Zhao, Zhao, and Eskénazi 2017) for details. For the GANs in our framework, we train the discriminator for one step every five steps for the generator, since it is it would be harder for generation than classification. The generators and discriminators are adversarially trained until GQD cannot discriminate between ground-truth and generated responses and RD is not able to distinguish between the summary and original datasets. The framework comes to convergence in less than an hour. (3) *Augmentation details:* We select 60% cases with the highest scores for augmentation if not specified, based on the experiment result on the validation dataset. For the selected cases, we employ the back-translation technique (Sennrich, Haddow, and Birch 2016) to augment them by ten times following Li et al. (2019). We choose the back-translations since it provides more diverse augmented text with different structures while preserving the meaning of the original text (Einolghozati et al. 2019; Chen et al. 2021a). Our evaluation metrics include distinctness (Li et al. 2016), BLEU (Papineni et al. 2002), and embedding metrics (Gu et al. 2019).

Baselines. We compare our model on following classic generation structure: (1) **SEQ2SEQ** (Bahdanau, Cho, and Bengio 2015): a sequence-to-sequence model with attention mechanisms. (2) **CVAE** (Zhao, Zhao, and Eskénazi 2017): a latent variable model using conditional variational auto-encoder, trained with KL-annealing and a BOW loss. (3) **Transformer** (Vaswani et al. 2017): an encoder-decoder architecture relying solely on the attention mechanisms. (4) **GPT-2** (Radford et al. 2019): a large-scale pre-trained language model, which is finetuned by the full training dataset. We also compare our approach with native augmentation, previous data augmentation, or instance weighting methods: (1) **Random:** we randomly select 60% data for augmentation, to compare with our selective augmentation method. Comparisons with different augmentation percentages can be found in the discussion section. (2) **Calibration** (Shang et al. 2018): a calibration network measures the quality of data samples and enables weighted training for dialogue generation. (3) **CVAE-GAN** (Li et al. 2019): a model that combines CVAE and GAN for augmentation. (4) **Manipulation** (Cai et al. 2020): it augments all the cases in the training process and reweights them.

	Models	Dist-1	Dist-2	Dist-3	BLEU-1	BLEU-2	BLEU-3	BLEU-4	Avg	Ext	Gre
(a)	SEQ2SEQ	1.36	5.98	11.19	12.85	2.26	1.12	0.93	76.56	42.98	62.24
	SEQ2SEQ (★)	1.63	7.24	15.01	14.65	2.73	1.27	0.81	77.83	43.60	63.56
	CVAE	1.88	6.50	12.08	11.38	2.08	1.10	0.75	74.18	40.71	62.49
	CVAE (★)	3.26	13.57	22.41	13.49	3.31	1.30	0.92	75.72	42.30	63.68
	Transformer	1.19	6.21	15.13	13.29	2.28	1.13	0.76	76.38	44.07	62.73
	Transformer (★)	2.92	15.20	23.70	14.17	2.52	1.33	0.85	77.46	45.69	64.20
(b)	GPT-2	2.16	7.44	16.15	15.27	2.84	1.66	0.78	78.27	45.39	64.17
	GPT-2 (★)	2.57	9.06	19.54	16.29	3.23	1.54	0.81	79.15	46.23	64.65
	SEQ2SEQ	1.37	2.22	6.62	10.03	1.57	1.01	0.84	61.88	45.34	50.45
	SEQ2SEQ (★)	1.56	3.94	5.83	10.78	2.00	1.29	0.97	62.36	46.24	51.14
	CVAE	0.70	2.22	5.92	9.90	1.87	1.07	0.92	65.37	49.60	53.63
	CVAE (★)	1.88	4.29	9.40	11.15	2.09	1.18	0.93	67.64	50.74	54.72
(b)	Transformer	1.57	3.28	6.39	8.76	2.35	1.21	0.87	66.91	44.40	54.18
	Transformer (★)	2.92	7.38	10.14	10.35	2.60	1.49	0.91	68.04	45.99	54.96
	GPT-2	3.12	4.32	7.29	10.97	3.30	2.15	1.15	67.28	48.60	55.07
	GPT-2 (★)	3.51	5.36	8.59	11.63	3.56	2.45	1.17	68.37	49.18	55.50

Table 1: Automatic evaluation results (%) on (a) DailyDialog and (b) OpenSubtitles. “★” denotes that the model is trained using our proposed framework. The metrics Average, Extrema, and Greedy are abbreviated as Avg, Ext, and Gre, respectively. Numbers in bold mean that the improvement to the best baseline is statistically significant (a two-tailed paired t-test with p-value <0.01).

	Dist-1	Dist-2	Dist-3	BLEU-1	BLEU-2	BLEU-3	BLEU-4	Avg	Ext	Gre
Full model	1.63	7.24	15.01	14.65	2.73	1.27	0.81	77.83	43.60	63.56
<i>w/o selective augmentation</i>	1.42	5.03	12.98	13.31	2.30	1.15	0.68	76.26	42.45	62.54
<i>w/o quality discriminator</i>	1.59	5.55	13.73	13.26	2.55	1.18	0.96	77.32	42.78	62.67
<i>w/o representativeness discriminator</i>	1.55	6.23	13.21	14.02	2.66	1.25	0.73	76.71	42.83	63.92

Table 2: Ablation test of our model (%) on DailyDialog, which is instantiated on the naive SEQ2SEQ. Numbers in bold mean that the improvements to the ablation models are statistically significant.

Main Results

Automatic evaluation. We instantiate our framework on a number of classic dialog generation models including SEQ2SEQ, CVAE, Transformer, and GPT-2. The automatic evaluation results are shown in Table 1. It can be seen that our model outperforms vanilla baselines on almost all automatic metrics. The improvements are consistent across both datasets, demonstrating the superiority and general applicability of our framework.

In addition, we compare our model with the existing augmentation methods. We select SEQ2SEQ as the response generation model following since all compared models are constructed on this classic baseline (Cai et al. 2020). Not surprisingly, as shown in Table 3, our framework outperforms most of the baseline methods. Concretely, SDA outperforms Random baseline in all metrics, demonstrating that selection is necessary to improve the performance of data augmentation. CVAE-GAN augments each case in the training dataset, and Manipulation augments every case in each training step, while our model only augments 60% data and achieves better performance. This demonstrates that selective data augmentation is more effective and efficient, outperforming data augmentation methods that require generating more augmented cases. The statistical significance of observed differences between the performance of two runs is tested using a two-tailed paired t-test for $\alpha = 0.01$.

Human Evaluation. We also employ a human evaluation on Amazon Mechanical Turk. For better annotation quality,

we employ three annotators and require the annotators to have at least 99% approval rate with at least 1,000 approved HITs. These annotators are hired to evaluate the quality of generated responses on DailyDialog dataset, where the evaluation is conducted in a double-blind fashion. Totally, 200 randomly sampled responses generated by each model are rated by each annotator with two different aspects, *i.e.*, *readability* and *informativeness*. Criteria are scored from 1 to 3, *i.e.*, bad, normal, and good. The results of the human evaluation are listed in Table 4. Our model significantly outperforms most of the baselines in terms of all the metrics. Particularly, our model increases informativeness by approximately 2.4% over Manipulation. The kappa statistics is 0.42 and 0.45 for readability and informativeness, respectively, which indicates moderate agreement between annotators. We also show a representative case from DailyDialog in Table 5, It can be seen that our model can generate a more diverse and interesting response that describes in detail how it feels to have a lover.

Discussions

Ablation Study. We also list the results of the ablation study in Table 2, aiming to investigate the influence of different modules in our proposed model. It can be seen that the performance of all metrics drops if we direct augment all the training data without selection. This demonstrates that selection is important for augmentation. We also find that the Dist-1 score drops by 2.45% and 4.91% after GQD and the

	Models	Dist-1	Dist-2	Dist-3	BLEU-1	BLEU-2	BLEU-3	BLEU-4	Avg	Ext	Gre
(a)	Random	1.38	5.24	11.70	13.25	2.17	1.13	0.73	77.02	43.12	62.68
	Calibration (Shang et al. 2018)	1.53	5.96	11.77	13.09	2.28	1.02	0.75	77.15	42.94	62.77
	CVAE-GAN (Li et al. 2019)	1.54	5.63	13.50	14.00	2.59	1.24	0.98	77.21	43.19	62.96
	Manipulation (Cai et al. 2020)	1.58	6.42	14.52	14.26	2.87	1.16	0.95	77.53	43.32	63.12
	SDA	1.63	7.24	15.01	14.65	2.92	1.27	0.81	77.83	43.60	63.56
(b)	Random	1.40	2.46	5.76	10.05	1.22	1.03	0.93	61.97	45.51	50.72
	Calibration (Shang et al. 2018)	1.43	2.58	5.82	10.20	1.23	1.08	0.68	62.03	45.57	50.83
	CVAE-GAN (Li et al. 2019)	1.49	2.83	5.07	10.26	1.28	1.17	0.87	62.28	45.74	50.76
	Manipulation (Cai et al. 2020)	1.41	3.40	5.93	10.37	1.58	1.24	0.94	62.29	46.00	50.22
	SDA	1.56	3.94	5.83	10.78	2.00	1.29	0.97	62.36	46.24	51.14

Table 3: Performance (%) of our approach instantiated on the naive SEQ2SEQ and the baseline approaches on (a) DailyDialog and (b) OpenSubtitles. Numbers in bold mean that the improvement to the best baseline is statistically significant.




Figure 3: (a) Generation quality of selected and unselected cases. (b) Reconstruction performance of selected and unselected cases. (c) Histogram of different augmentation frequencies.

Model	Readability	Informativeness
Calibration	1.63	1.68
CVAE-GAN	1.85	1.81
Manipulation	1.91	2.07
SDA	2.01	2.12

Table 4: Human evaluation on two aspects: Readability and informativeness.

RD are removed, respectively. This indicates that the jointly selected cases from the quality and representativeness aspects help generate more diverse and accurate responses.

Analysis of Selected Samples. In this subsection, we examine whether the model successfully selects the lowest quality and most representative cases for augmentation. We calculate the BLEU scores of selected and unselected cases in the response generation task and response reconstruction task. From Figure 3(a) and Figure 3(b) we can see that the selected cases have lower BLEU scores in terms of the generation quality and higher scores in the reconstruction task. This demonstrates that the model needs to be polished to generate better responses for the selected cases. In the meantime, the selected data itself is not noise data and represents the overall data distribution. To further glean the insights regarding which samples are favored by the augmentation model, we also list examples with different augmentation scores in Figure 3(c). We notice that samples frequently

-I got a ticket yesterday.	
ground-truth	Really? What did you get one for?
CVAE-GAN	Is that right?? Is that possible ?
Manipulation	88 yuan, please.
SDA	Really? How much is it?
- What do you mean? You have a lover? .	
ground-truth	A fiance.
CVAE-GAN	You've had a lot of your own! lover!
Manipulation	No, I'm serious.
SDA	Yeah, she's so different, she is the sun!

Table 5: Responses generated by baselines and our model. The top case is selected from DailyDialog, and the bottom case is from OpenSubtitles.

augmented by SDA are more reliable and meaningful context, where the response is closely related to the query and leads to a new topic. While for the dialog pairs self augmented, they contain universal and safe content such as "I don't know" or "I'd forgotten about it".

Visualization of Dual Training. To visualize the select process, we draw the loss curve of the generation quality discriminator (\mathcal{L}_D in Equation 4) and the response generator evaluator (\mathcal{L}_G in Equation 4) in Figure 4(a), and show the accuracy of GQD in Figure 4(b). When the training begins, the loss of the GQD and RGE fluctuates from time to time, as well as the accuracy curve, which verify the adversarial training. After several steps, the training converges, and




Figure 4: (a) Loss curve of the quality discriminator and generator evaluator. (b) Accuracy curve of the quality discriminator. (c) Relationship between the selective percentage for augmentation and the embedding scores. Blue denotes our model, and orange denotes the Random model.

	Models	BLEU-1	BLEU-2	BLEU-3	BLEU-4	Extrema	Average	Greedy
RocStories	CVAE	25.81	9.69	3.60	1.48	51.35	56.32	60.11
	Seq2Seq	23.24	8.96	3.40	1.50	51.33	56.49	60.07
	Transformer	25.52	5.96	3.54	1.45	51.29	56.37	60.06
	GPT-2	30.21	11.08	3.64	1.53	51.72	58.49	60.35
	GPT-2(★)	30.96	11.46	3.84	1.52	52.27	58.95	61.20

Table 6: Automatic evaluation results on RocStories for storytelling. Numbers in bold mean that the improvement to the best baseline is statistically significant (t-test with p-value <0.01).

the accuracy of RD stays around 50%, which means GQD cannot distinguish between the generated response and the ground truth one. In other words, the model successfully assigns low s scores, *i.e.*, high $(1-s)$, to the cases with high-quality generated responses so that GQD cannot perform better than a random guess. The accuracy curve of the RD is similar to that of GQD, which proves that our model assigns high s scores to the most representative cases so that RD cannot distinguish between the reconstructed and the original cases.

Impact of Augmented Data Scale. For previous experiments, the percentage of cases for augmentation is set to 60%. In this subsection, we change this percentage to study what is the influence of scale for augmentation and whether selective augmentation is still beneficial under different selection percentages. We also select the random baseline model for better comparison, where the cases for augmentation are randomly sampled. The result on the DailyDialog test dataset is shown in Figure 4(c). For Random baseline, its performance generally improves with the augmentation percentage. This result shows that random augmentation will benefit the dialog generation task, and the more cases are augmented, the better performance will be obtained. However, this is not true for selective augmentation. It can be seen that to begin with, the embedding scores of SDA increase fast with the selective percentage for augmentation. After the percentage reaches 60%, the growth stops, and when the percentage increases from 80% to 100%, there is even a drop in the performance. Similar performance is also observed on the OpenSubtitles dataset. This demonstrates that it only benefits the model if we select the proper cases

in the dataset for augmentation, otherwise, augmenting some cases brings harm to the model.

Universality of our framework. In addition, we test the generalization ability of our framework on the story generation task. RocStories dataset (Mostafazadeh et al. 2016) consists of 98,163 high-quality hand-crafted stories, which capture causal and temporal commonsense relations of daily events. Each story paragraph contains 5 sentences with an average of 43 words. Following the previous work (Yu et al. 2021), we split the dataset into 8:1:1 for training, validation, and test, and use BLEU as the evaluation metric. As can be seen from Table 6, equipped with augmentation data, our method outperforms GPT-2 by 2.4%, 3.4%, and 1.4% on RocStories in terms of BLEU-1, BLEU-2, and Greedy, respectively, which proves the superiority of our model. This experiment also demonstrates that our framework does not rely on a specific task, and can be extended to various text generation scenarios.

Conclusion and Broader Impacts

In this paper, we propose a selective data augmentation framework to improve the performance of dialogue models. We propose a dual adversarial network to select data for augmentation from the quality and representativeness aspects. One is to examine whether the case is of low generation quality, and the other one is whether the case is representative of the dataset. Experiments conducted on three public datasets demonstrate the effectiveness of our framework. In the future, we would like to explore the effectiveness of selective data augmentation on more generation tasks.

Acknowledgments

We would like to thank the anonymous reviewers for their constructive comments. This work was supported by the SDAIA-KAUST Center of Excellence in Data Science and Artificial Intelligence (SDAIA-KAUST AI). This publication is based upon work supported by the King Abdullah University of Science and Technology (KAUST) Office of Research Administration (ORA) under Award No FCC/1/1976-44-01, FCC/1/1976-45-01, URF/1/4663-01-01, RGC/3/4816-01-01, and BAS/1/1635-01-01. This work was also supported by NSFC Grant No. 62122089 and CCF-Tencent Rhino-Bird Open Research Fund.

References

- Andreas, J. 2020. Good-Enough Compositional Data Augmentation. In *Proc. of ACL*.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. *Proc. of ICLR*.
- Bowman, S.; Vilnis, L.; Vinyals, O.; Dai, A.; Jozefowicz, R.; and Bengio, S. 2016. Generating Sentences from a Continuous Space. In *Proc. of CoNLL*.
- Cai, H.; Chen, H.; Song, Y.; Zhang, C.; Zhao, X.; and Yin, D. 2020. Data Manipulation: Towards Effective Instance Learning for Neural Dialogue Generation via Learning to Augment and Reweight. In *Proc. of ACL*.
- Chan, Z.; Li, J.; Yang, X.; Chen, X.; Hu, W.; Zhao, D.; and Yan, R. 2019. Modeling Personalization in Continuous Space for Response Generation via Augmented Wasserstein Autoencoders. In *Proc. of EMNLP*.
- Chen, J.; Tam, D.; Raffel, C.; Bansal, M.; and Yang, D. 2021a. An Empirical Survey of Data Augmentation for Limited Data Learning in NLP. *arXiv preprint arXiv:2106.07499*.
- Chen, X.; Cui, Z.; Zhang, J.; Wei, C.; Cui, J.; Wang, B.; Zhao, D.; and Yan, R. 2021b. Reasoning in dialog: Improving response generation by context reading comprehension. In *Proc. of AACL*, volume 35, 12683–12691.
- Csaky, R.; Purgai, P.; and Recski, G. 2019. Improving Neural Conversational Models with Entropy-Based Data Filtering. In *Proc. of ACL*.
- Einolghozati, A.; Gupta, S.; Mohit, M.; and Shah, R. 2019. Improving robustness of task oriented dialog systems. *Proc. of NeurIPS*.
- Gu, X.; Cho, K.; Ha, J.-W.; and Kim, S. 2019. DialogWAE: Multimodal Response Generation with Conditional Wasserstein Auto-Encoder. In *Proc. of ICLR*.
- Hou, Y.; Liu, Y.; Che, W.; and Liu, T. 2018. Sequence-to-Sequence Data Augmentation for Dialogue Language Understanding. *Proc. of COLING*.
- Kobayashi, S. 2018. Contextual Augmentation: Data Augmentation by Words with Paradigmatic Relations. In *Proc. of ACL*.
- Kurata, G.; Xiang, B.; and Zhou, B. 2016. Labeled Data Generation with Encoder-Decoder LSTM for Semantic Slot Filling. In *Proc. of Interspeech*.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, W. 2016. A Diversity-Promoting Objective Function for Neural Conversation Models. *Proc. NAACL*.
- Li, J.; Qiu, L.; Tang, B.; Chen, D.; Zhao, D.; and Yan, R. 2019. Insufficient Data Can Also Rock! Learning to Converse Using Smaller Data with Augmentation. In *Proc. of AAAI*.
- Li, J.; Xia, Y.; Yan, R.; Sun, H.; Zhao, D.; and Liu, T.-Y. 2021. Stylized Dialogue Generation with Multi-Pass Dual Learning. *Proc. of NeurIPS*.
- Li, Y.; Su, H.; Shen, X.; Li, W.; Cao, Z.; and Niu, S. 2017. DailyDialog: A Manually Labelled Multi-turn Dialogue Dataset. *Proc. of IJCNLP*.
- Lison, P.; and Tiedemann, J. 2016. OpenSubtitles2016: Extracting Large Parallel Corpora from Movie and TV Subtitles. In *Proc. of LREC*.
- Louvan, S.; and Magnini, B. 2020. Simple is Better! Lightweight Data Augmentation for Low Resource Slot Filling and Intent Classification. In *Proceedings of the 34th Pacific Asia Conference on Language, Information and Computation*.
- Mahasseni, B.; Lam, M.; and Todorovic, S. 2017. Unsupervised Video Summarization with Adversarial LSTM Networks. *Proc. of CVPR*.
- Mostafazadeh, N.; Chambers, N.; He, X.; Parikh, D.; Batra, D.; Vanderwende, L.; Kohli, P.; and Allen, J. 2016. A corpus and cloze evaluation for deeper understanding of common-sense stories. In *Proc. of AACL*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proc. of ACL*.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*.
- Schröder, C.; and Niekler, A. 2020. A Survey of Active Learning for Text Classification using Deep Neural Networks. *ArXiv*.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Improving Neural Machine Translation Models with Monolingual Data. In *Proc. of ACL*.
- Shang, M.; Fu, Z.; Peng, N.; Feng, Y.; Zhao, D.; and Yan, R. 2018. Learning to Converse with Noisy Data: Generation with Calibration. In *Proc. of IJCAI*.
- Song, Z.; Zheng, X.; Liu, L.; Xu, M.; and Huang, X.-J. 2019. Generating responses with a specific emotion in dialog. In *Proc. of ACL*.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *Proc. of CVPR*.
- Tuan, Y.-L.; Pryor, C.; Chen, W.; Getoor, L.; and Wang, W. Y. 2021. Local Explanation of Dialogue Response Generation. *Proc. of NeurIPS*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Proc. of NIPS*.

- Wu, X.; Lv, S.; Zang, L.; Han, J.; and Hu, S. 2019. Conditional bert contextual augmentation. In *International Conference on Computational Science*.
- Yu, M.-H.; Li, J.; Chan, Z.; Zhao, D.; and Yan, R. 2021. Content Learning with Structure-Aware Writing: A Graph-Infused Dual Conditional Variational Autoencoder for Automatic Storytelling. In *Proc. of AAAI*.
- Zhao, T.; Zhao, R.; and Eskénazi, M. 2017. Learning Discourse-level Diversity for Neural Dialog Models using Conditional Variational Autoencoders. In *Proc. of ACL*.
- Zou, Y.; Zhao, L.; Kang, Y.; Lin, J.; Peng, M.; Jiang, Z.; Sun, C.; Zhang, Q.; Huang, X.; and Liu, X. 2021. Topic-Oriented Spoken Dialogue Summarization for Customer Service with Saliency-Aware Topic Modeling. In *Proc. of AAAI*.