

# Adversarial Learning on the Latent Space for Diverse Dialog Generation

Kashif Khan<sup>1\*</sup>, Gaurav Sahu<sup>1\*</sup>, Vikash Balasubramanian<sup>1\*</sup>, Lili Mou<sup>2</sup>, Olga Vechtomova<sup>1</sup>

<sup>1</sup>University of Waterloo

<sup>2</sup>Dept. Computing Science, University of Alberta  
Alberta Machine Intelligence Institute (Amii)

{kashif.khan, gaurav.sahu, v9balasu, ovechtom}@uwaterloo.ca  
doublepower.mou@gmail.com

## Abstract

Generating relevant responses in a dialog is challenging, and requires not only proper modeling of context in the conversation, but also being able to generate fluent sentences during inference. In this paper, we propose a two-step framework based on generative adversarial nets for generating conditioned responses. Our model first learns a meaningful representation of sentences by autoencoding, and then learns to map an input query to the response representation, which is in turn decoded as a response sentence. Both quantitative and qualitative evaluations show that our model generates more fluent, relevant, and diverse responses than existing state-of-the-art methods.<sup>1</sup>

## 1 Introduction

Dialog generation is a challenging problem because it not only requires us to model the context in a conversation but also to exploit it to generate a relevant and fluent response. A dialog generation system can be divided into two parts: 1) encoding the context of the conversation, and 2) generating a response conditioned on the given context. A generated response is considered to be “good” if it is meaningful, fluent, and most importantly, relevant to the given context.

With the advancement of deep learning, sequence-to-sequence (Seq2Seq) models (Sutskever et al., 2014) are adopted for dialog systems to encode conversational context and generate a response. However, they suffer from the problem of generic utterance generation, e.g., always generating “I don’t know” (Serban et al., 2016; Li et al., 2016). One possible explanation (Wei et al., 2019) is the high uncertainty in dialog generation. A plausible response is analogous to a “mode” of a continuous distribution, and the response distribution is thus multimodal. However, the decoder of a Seq2Seq model is trained by cross-entropy loss, which is equivalent to minimizing the KL divergence between the target and predicted distributions. The asymmetric nature of KL divergence makes the learned distribution wide-spreading, analogous to the mode-averaging problem for continuous variables.

Variational encoder-decoders (Serban et al., 2017; Bahuleyan et al., 2018; Zhao et al., 2017) and Wasserstein encoder-decoders (Bahuleyan et al., 2019) adopt probabilistic modeling to encourage diversity in responses. However, their decoders are also trained by cross-entropy loss against the target sequence, still making the model generate generic utterances.

In this paper, we propose an approach that uses adversarial learning in the latent space for dialog generation. We first train a variational autoencoder (VAE) (Kingma and Welling, 2014) on sentences, and then apply a generative adversarial network (GAN) on the latent space of the VAE. At inference time, we obtain the latent representation of the response from the generator of the GAN and decode it using the VAE’s decoder. In this way, we can benefit from the mode-capturing property of GANs (Mao et al., 2019; Thanh-Tung et al., 2019). Also, our GAN is trained on the latent space, and techniques like Gumbel-Softmax and reinforcement-learning (RL) are not necessary, which largely simplifies the training procedure. We further introduce a mean squared error (MSE) auxiliary loss to our adversarial module,

\*Equal contribution.

<sup>1</sup>The code is available at [https://github.com/vikigenius/conditional\\_text\\_generation](https://github.com/vikigenius/conditional_text_generation)  
This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

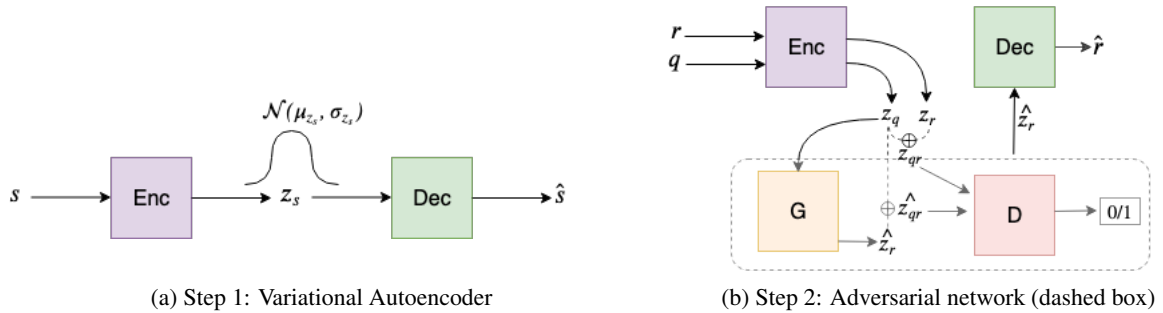


Figure 1: The framework of our proposed two-step training procedure.  $\oplus$  denotes concatenation.

which mitigates the mode-missing problem in GANs (Che et al., 2017), resulting in more relevant and diverse responses.

We evaluate our model on the deduplicated version (Bahuleyan et al., 2018) of the benchmark Daily-Dialog dataset (Li et al., 2017) and also the Switchboard dataset (Godfrey et al., 1992). Results indicate that responses generated by our model are more relevant to the input query/context, and are more diverse and fluent than the existing baselines.

The main contributions of our paper are as follows.

1. We propose a two-step framework of latent-space adversarial learning for generating diverse and relevant responses.
2. We propose a combination of adversarial loss and an auxiliary mean squared loss to help the GAN to converge faster and achieve better performance for dialog generation.

## 2 Approach

Figure 1 provides an overview of our proposed two-step approach.

**Step 1:** We first train an autoencoder, which takes an utterance  $s$  (either a query or a response) as input, gets its latent code  $z_s$  from the encoder, and then feeds it to a decoder for reconstructing. The autoencoder learns a real-valued vector representation of a generic sentence.

**Step 2:** We train an adversarial network on the latent  $z$  space for learning dialog generation. Given a query-response pair  $(q, r)$  in the training set, we use the trained encoder from Step 1 to obtain their latent variables  $z_q$  and  $z_r$ . The query latent variable  $z_q$  is fed to a generator  $G$ , that maps it to the corresponding response’s latent variable  $\hat{z}_r$ . When training the generator, we aim to match  $z_r$  and  $\hat{z}_r$  through the adversarial loss combined with a mean-squared error loss. In here, the adversarial loss further involves a discriminator that classifies the predicted response representation  $\hat{z}_r$  versus the encoded representation of the true response  $z_r$ , conditioned on the query  $z_q$ .

The details of our approach will be introduced in the rest of this section.

### 2.1 Step 1: Training an Autoencoder

In Step 1, our primary goal is to learn a continuous representation of all utterances in the dialog corpus. The mapping from a sentence to its continuous representation should ideally be invertible so that our adversarial loss (in Step 2) could be applied to the continuous space to generate dialog responses.

In particular, we adopt a variational autoencoder (Kingma and Welling, 2014, VAE) for our first step. A VAE encodes an input sentence  $s$  to a probabilistic, latent continuous representation  $z$ , from which the input sentence  $s$  is reconstructed.

We first impose a prior distribution on  $z$ , which is typically set to standard normal  $p(z) = \mathcal{N}(\mathbf{0}, \mathbf{I})$ . Given the sentence  $s$ , VAE encodes a posterior distribution  $q_E(z|s) = \mathcal{N}(\mu, \text{diag } \sigma^2)$ , where  $\mu$  and  $\sigma$  are predicted by the encoder of VAE. The training objective is to minimize the expected reconstruction

loss, penalized by a KL divergence term between the posterior and the prior. This is given by

$$J_{\text{AE}}(\boldsymbol{\theta}_{\text{Enc}}, \boldsymbol{\theta}_{\text{Dec}}) = -\mathbb{E}_{q_E(\mathbf{z}|s)}[\log p(s|\mathbf{z})] + \lambda_{\text{KL}}\text{KL}(q_E(\mathbf{z}|s)||p(\mathbf{z})) \quad (1)$$

where  $\lambda_{\text{KL}}$  balances the two terms.

Compared with a deterministic autoencoder, VAE learns a smoother latent space by its KL regularization. This is helpful during the second step, where a GAN is trained to predict the latent representation of a response for decoding.

## 2.2 Step 2: Predicting the Representation of the Response

In Step 2, the main objective is to predict the representation of the response given dialog context (such as the previous utterance). In this way, the predicted latent representation of the response can be fed to the trained decoder in Step 1 to generate the response utterance.

To predict the response representation, we re-use the encoder in Section 2.1 to capture the meaning of the context query as  $z_q$ . Then we have a two-layer perceptron (with a ReLU activation function in the hidden layer) to predict the representation of the utterance to be generated, denoted by  $\hat{z}_r = G(z_q)$ .

For adversarial training, we also encode the representation of the ground truth reply  $r$  as  $z_r$  by the encoder in Step 1. We train an adversarial discriminator  $D$  to classify whether the response representation is real or predicted. Such classification should be based on context, because we would learn not only if an utterance is appropriate as a reply, but also if the utterance is appropriate to a specific query. Therefore, we also feed the encoded context representation into the discriminator. The classification is denoted by  $D(z_r, z_q)$  or  $D(\hat{z}_r, z_q)$ , where we essentially concatenate the representations of the response and the query before feeding them to a logistic regression layer.

The adversarial loss for training the latent space is given by:

$$J_{\text{CGAN}} = \min_G \max_D V(D, G) \quad (2)$$

$$V(D, G) = \mathbb{E}_{(z_q, z_r) \sim \mathcal{D}_{\text{train}}} [\log D(z_r, z_q) + \log(1 - D(G(z_q), z_q))] \quad (3)$$

where  $\mathcal{D}_{\text{train}}$  is the training data.

In other words, the discriminator  $D$  is trained by maximizing  $V(D, G)$  so as to distinguish the true representation of a response and the predicted response representation given the query, whereas the generator  $G$  is trained to fool the discriminator by minimizing  $V(D, G)$ .

It should be emphasized that our model is different from adversarial autoencoders (Makhzani et al., 2015), because our discriminator takes the query into consideration. Our adversarial loss learns an implicit conditional distribution  $p(z_r|z_q)$ , instead of a marginal distribution  $p(z_r)$  as in Zhao et al. (2018).

Additionally, we introduce an auxiliary mean square error (MSE) loss to the objective function:

$$J_{\text{MSE}} = \|z_r - \hat{z}_r\|^2 \quad (4)$$

The MSE loss on the generator helps stabilize the GAN training and mitigate the mode-missing problem of GANs (Che et al., 2017). In summary, the overall training objective is given by

$$J = J_{\text{CGAN}} + \gamma J_{\text{MSE}} \quad (5)$$

where  $\gamma$  is a tunable hyperparameter that moderates the effect of the MSE loss.

For inference, our model first uses the pretrained VAE from Step 1 to encode an unseen query  $q^*$  as  $z_{q^*}$ . This encoded representation is then passed to the generator  $G$  to predict the response latent code  $G(z_{q^*})$ , which is finally fed to the decoder of the VAE from Step 1 to generate a response sentence.

In our experiments, we have two settings for dialog generation: single-turn and multi-turn. In the single-turn setting, we form query-response samples by extracting every pair of consecutive utterances of a conversation in the training data.

In the multi-turn setting, we have the query-response pairs by extracting every utterance with its preceding utterances in the *entire* conversation. The VAE in Step 1 remains the same, but we introduce

Model	BLEU			Diversity						Fluency
	Avg	Max	HM	Intra-1	Intra-2	Inter-1	Inter-2	ASL (14.43)	TTR	PPL
Single-turn results										
Seq2Seq	0.143	0.217	0.172	<b>0.99</b>	<b>0.99</b>	0.46	0.49	4.63	0.019	18.45
WED-S	0.215	0.357	0.268	0.94	<b>0.99</b>	0.48	0.74	10.42	0.034	33.91
DialogWAE	0.296	0.356	0.323	0.85	0.97	0.42	0.74	19.34	0.005	20
VAE-M (ours)	0.191	0.293	0.231	0.98	<b>0.99</b>	<b>0.5</b>	0.79	9.36	0.029	19.7
VAE-A (ours)	0.295	0.359	0.323	0.93	<b>0.99</b>	0.46	0.76	13.64	0.035	21.38
VAE-AM (ours)	<b>0.306</b>	<b>0.367</b>	<b>0.334</b>	0.91	<b>0.99</b>	0.46	<b>0.82</b>	16.90	<b>0.034</b>	<b>17.01</b>
Multi-turn results										
HRED*	0.232	0.232	0.232	0.94	0.97	0.09	0.10	10.1	-	-
CVAE*	0.222	0.265	0.242	0.94	0.97	0.18	0.22	10.0	-	-
CVAE-CO*	0.244	0.259	0.251	0.82	0.91	0.11	0.13	11.2	-	-
VHCR*	0.266	0.289	0.277	0.85	0.97	0.42	0.74	16.9	-	-
DialogWAE	0.279	0.365	0.316	0.79	0.92	0.35	0.68	19.84	0.007	161.86
VAE-AM (ours)	<b>0.314</b>	<b>0.371</b>	<b>0.340</b>	<b>0.847</b>	<b>0.98</b>	<b>0.41</b>	<b>0.73</b>	15.3	<b>0.036</b>	<b>119.39</b>

Table 1: Results on the DailyDialog dataset. BLEU scores are computed by the average/maximum of 10 randomly sampled replies. HM is the harmonic mean of average and maximum BLEU scores. Suffix A: adversarial loss; suffix M: MSE loss; suffix AM: both adversarial and MSE losses. \* denotes results taken from Gu et al. (2019), whose training and test sets contain duplicate samples. The Bold font shows the best performance on the de-duplicated dataset. The number in the bracket of the ASL column is the groundtruth average sentence length.

Model	BLEU			Diversity						Fluency
	Avg	Max	HM	Intra-1	Intra-2	Inter-1	Inter-2	ASL (8.49)	TTR	PPL
Single-turn results										
Seq2Seq	0.088	0.176	0.118	<b>0.989</b>	0.956	<b>0.816</b>	<b>0.927</b>	2.66	0.026	23.62
WED-S	0.193	<b>0.395</b>	0.259	0.941	0.989	0.404	0.525	10.41	0.032	35.63
DialogWAE	0.235	0.375	0.289	0.739	0.712	0.354	0.571	10.32	0.017	25.36
VAE-M (ours)	0.231	0.3	0.261	0.954	0.998	0.322	0.479	8.84	0.045	24.58
VAE-A (ours)	0.229	0.376	0.285	0.725	0.751	0.218	0.354	11.73	0.053	27.55
VAE-AM (ours)	<b>0.259</b>	0.364	<b>0.303</b>	<b>0.989</b>	<b>0.999</b>	0.436	0.569	7.08	<b>0.062</b>	<b>21.87</b>
Multi-turn results										
HRED*	0.262	0.262	0.262	0.813	0.452	0.081	0.045	12.1	-	-
CVAE*	0.258	0.295	0.275	0.803	0.415	0.112	0.102	12.4	-	-
CVAE-CO*	0.269	0.299	0.283	0.863	0.581	0.111	0.110	10.3	-	-
VHCR*	0.234	0.276	0.254	0.877	0.536	0.130	0.131	9.29	-	-
DialogWAE	0.267	<b>0.394</b>	<b>0.318</b>	0.779	0.844	0.325	0.513	14.7	0.019	183.82
VAE-AM (ours)	<b>0.271</b>	0.372	0.313	<b>0.954</b>	<b>0.966</b>	<b>0.412</b>	<b>0.559</b>	8.72	<b>0.065</b>	<b>136.81</b>

Table 2: Results on the Switchboard dataset. \* denotes the numbers taken from Gu et al. (2019).

another RNN to encode context. Specifically, it is built upon the VAE’s encoded representation of each utterance, and yields a fixed-length vector representation of the entire context. During the adversarial training, we concatenate the context vector with the query (immediate preceding utterance) representation before feeding them to the generator. In this way, our generator now also takes the context into account when predicting the response latent code. A similar adjustment is applied during inference as well.

### 3 Experiments

We conduct experiments on the DailyDialog dataset (Li et al., 2017), a manually labeled multi-turn dialog dataset, and the Switchboard dataset (Godfrey et al., 1992), a dialog dataset containing transcripts of telephonic conversations. For DailyDialog, we use the original splits after removing duplicates following Bahuleyan et al. (2019). We use the the AllenNLP framework (Gardner et al., 2018) to implement all our models. Appendix A presents more experimental details and hyper-parameters.

We use the following baseline models for comparison:

- **Seq2Seq.** The standard sequence to sequence model based on LSTM.

- **WED-S.** A stochastic Wasserstein encoder-decoder model (Bahuleyan et al., 2019).
- **DialogWAE.** A model based on adversarial regularization of autoencoders (Gu et al., 2019).
- **HRED.** A generalized Seq2Seq model that uses hierarchical RNN encoder (Serban et al., 2016).
- **CVAE.** A conditional VAE model with KL annealing (Zhao et al., 2017).
- **CVAE-CO.** A collaborative conditional VAE model (Shen et al., 2018).

### 3.1 Results and Analysis

The results for the Daily Dialog and the Switchboard datasets are shown in Tables 1 and 2, respectively. The generated responses are evaluated by the following criteria:

**Overall quality.** We measure the quality of the generated responses by BLEU scores (Papineni et al., 2002), for which we adopt the smoothing techniques in Gu et al. (2019). For each query, we generate 10 responses for a query, and compute the average and maximum BLEU scores. Then we also compute the harmonic mean of the average and the maximum BLEU scores.<sup>2</sup> Our model is either the best-performing model or highly competitive in terms of the BLEU scores. The DialogWAE model also achieves high BLEU scores, while the Seq2Seq model is the worst-performing model.

**Diversity.** We measure the diversity of dialog generation in two aspects:

- **Intra-diversity.** The Intra-diversity score measures the proportion of distinct unigrams and bigrams in each response. It is similar for most models.
- **Inter-diversity.** The Inter-diversity scores measure the proportion of distinct unigrams and bigrams across all 10 responses.

We note that our model performs the best across Inter-diversity metrics. We further use other diversity indicators such as the Average Sentence Length (ASL) of the responses. We see that diversity scores for the Seq2Seq model are very high on the Switchboard dataset; however, it has the lowest ASL score as well. This observation is within expectations, and the Seq2Seq model does not generate diverse responses overall. DialogWAE generates longer responses on average; however, our model is closer to the ground truth ASL (14.43 for DailyDialog and 8.49 for Switchboard). We also note that our model achieves good Type-Token Ratio (TTR) scores,<sup>3</sup> indicating diverse word choices when compared with other models.

**Fluency.** We compute the PPL scores of generated responses to measure fluency. We notice that our model achieves the best PPL scores, although DialogWAE is quite close. The Seq2Seq model also achieves low PPL, but this is mainly due to the short and generic responses. Interestingly, PPL scores are generally higher in the multi-turn setting, which may be attributed to the increased complexity of the output when more context is given.

**Analysis of Losses.** Combining the MSE and adversarial losses leads to significant improvements across all metrics, including the BLEU scores, response diversity (Inter-1 and Inter-2), and fluency (PPL). In our experiments, we also notice that the MSE term leads to quicker and more stable convergence of the GAN (within 6 epochs), making training easier.

We present human evaluation in Appendix B and a case study in Appendix C.

## 4 Conclusion

We propose an effective two-stage model for dialog generation. We make use of sentence representations learned by a VAE and train a adversarial network on VAE’s latent space to generate diverse responses given a query and context. We observe that our model outperforms existing state-of-the-art approaches by generating more diverse, fluent, and relevant sentences.

<sup>2</sup>The evaluation protocol follows Gu et al. (2019) with code at <https://github.com/guxd/DialogWAE/blob/29f206af05bfe5fe28fec4448e208310a7c9258d/experiments/metrics.py#L90>

<sup>3</sup>TTR is computed in the corpus level, whereas Inter- $n$  diversity is the average of per-sample distinct unigram ratio.

## Acknowledgments

This research was partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) under grant Nos. RGPIN-2019-04897 and RGPIN-2020-04465. Lili Mou is also supported by the Amii Fellow Program and the Canadian CIFAR AI Chair Program. This research was enabled in part by the support of Compute Canada ([www.compute-canada.ca](http://www.compute-canada.ca)).

## References

- Hareesh Bahuleyan, Lili Mou, Olga Vechtomova, and Pascal Poupart. 2018. Variational attention for sequence-to-sequence models. In *COLING*, pages 1672–1682.
- Hareesh Bahuleyan, Lili Mou, Hao Zhou, and Olga Vechtomova. 2019. Stochastic Wasserstein autoencoder for probabilistic sentence generation. In *NAACL-HLT, Volume 1*, pages 4068–4076.
- Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. Generating sentences from a continuous space. In *CoNLL*, pages 10–21.
- Tong Che, Yanran Li, Athul Paul Jacob, Yoshua Bengio, and Wenjie Li. 2017. Mode regularized generative adversarial networks. In *ICLR*.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke S. Zettlemoyer. 2018. AllenNLP: A deep semantic natural language processing platform. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6.
- John J. Godfrey, Edward C. Holliman, and Jane McDaniel. 1992. Switchboard: Telephone speech corpus for research and development. In *ICASSP, Volume 1*, page 517–520.
- Xiaodong Gu, Kyunghyun Cho, Jung-Woo Ha, and Sunghun Kim. 2019. Dialogwae: Multimodal response generation with conditional wasserstein auto-encoder. In *ICLR*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation, Volume 9*, pages 1735–1780.
- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Diederik P. Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *ICLR*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *NAACL-HLT*, pages 110–119.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. DailyDialog: A manually labelled multi-turn dialogue dataset. In *IJCNLP*, pages 986–995.
- Andrew L. Maas, Awni Y. Hannun, and Andrew Y. Ng. 2013. Rectifier nonlinearities improve neural network acoustic models. In *ICML Workshop on Deep Learning for Audio, Speech and Language Processing*.
- Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, and Ian J. Goodfellow. 2015. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*.
- Qi Mao, Hsin-Ying Lee, Hung-Yu Tseng, Siwei Ma, and Ming-Hsuan Yang. 2019. Mode seeking generative adversarial networks for diverse image synthesis. In *CVPR*, pages 1429–1437.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: A method for automatic evaluation of machine translation. In *ACL*, pages 311–318.
- Iulian V. Serban, Alessandro Sordani, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*, page 3776–3783.
- Iulian Vlad Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *AAAI*, page 3295–3301.

- Xiaoyu Shen, Hui Su, Shuzi Niu, and Vera Demberg. 2018. Improving variational encoder-decoders in dialogue generation. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *AAAI*, pages 5456–5463.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*, pages 3104–3112.
- Hoang Thanh-Tung, Truyen Tran, and Svetha Venkatesh. 2019. Improving generalization and stability of generative adversarial networks. *arXiv preprint arXiv:1902.03984*.
- Bolin Wei, Shuai Lu, Lili Mou, Hao Zhou, Pascal Poupart, Ge Li, and Zhi Jin. 2019. Why do neural dialog systems generate short and meaningless replies? a comparison between dialog and translation. In *ICASSP*, pages 7290–7294.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *ACL*, pages 654–664.
- Junbo Zhao, Yoon Kim, Kelly Zhang, Alexander Rush, and Yann LeCun. 2018. Adversarially regularized autoencoders. In *ICML*, pages 5902–5911.

## A Hyperparameter Settings and Training

**Single-turn.** In this setting, we first train a VAE on the entire corpus. We use a single-layer encoder with Bidirectional LSTMs (Hochreiter and Schmidhuber, 1997) and a unidirectional LSTM layer for the decoder of the VAE. Both use a hidden size of 512. The dimension of our latent vectors is 128, and that of the word embeddings is 300. Further, we adopt KL-annealing and word dropout from Bowman et al. (2016) to stabilize VAE’s training. We use a word dropout probability of 0.5 and a sigmoid annealing schedule to anneal the KL weight to 0.15 for 4500 iterations. The performance statistics of VAE in Step 1 are shown in 3.

Model	KL	BLEU	Dist-1	Dist-2
VAE	18.8	0.18	0.32	0.49

Table 3: Performance of the VAE in Step 1 on the DailyDialog dataset. BLEU is the reconstruction BLEU-4, Dist-1 and Dist-2 are distinct unigrams and bigrams in generated samples, KL is the validation KL in the best epoch (ELBO).

For the GAN, we use a 2-layer feed-forward network with a hidden layer of 256 units as the generator, along with batch normalization (Ioffe and Szegedy, 2015) and LeakyReLU activation (Maas et al., 2013). The discriminator shares a similar architecture. We use Adam (Kingma and Ba, 2015) to optimize all our networks.

**Multi-turn.** In this setting, the VAE’s architecture remains the same as the Single-turn setting. We introduce another BiLSTM encoder with hidden size of 512, which is fed with the VAE-encoded representations of the context sentences. Other hyperparameters are kept the same. For implementation, our generator predicts the response representation at each turn, but we use teacher-forcing, assuming the context is the actual previous utterances.

## B Human Evaluation

In addition to automatic metrics, we also evaluate our model and compare it with DialogWAE (Gu et al., 2019) using human evaluation. Five human judges rate the response of each model on a scale of 1–5, according to two criteria: 1) Relevance to the query, and 2) Fluency of the generated response. As can be seen from Table 4 our model is competitive and achieves better scores on both criteria.

Model	Relevance	Fluency
VAE-AM	2.3625	3.2225
DialogWAE	2.3125	3.1775

Table 4: Human evaluation results

## C Case Study

We conduct a case study on VAE-AM’s responses with DialogWAE’s in the single-turn and multi-turn settings. For both settings, we show the best response out of 10 generated responses for each input.

Table 5 shows that DialogWAE generates generic responses, although its sentences are longer. Our model has more concrete concepts, such as “leaflets” and “money.”

Table 6 shows generated responses in the multi-turn setting. We see that both models are more informative as we have more context compared with Table 5. However, DialogWAE is repeating the same concept, such as “a lot of” in Examples 2, 3, and 5, and “have to pay” in Examples 2 and 5. This observation is also consistent with the TTR results in Tables 1 and 2.



#	Generated Responses
1	<p><b>Query:</b> o . k . well , go to google.com , and type in battle of gettysburg .</p> <p><b>VAE-AM:</b> oh , there are a lot of things that i need to do .</p> <p><b>DialogWAE:</b> then ill have to check the details . what can i do for you today ?</p>
2	<p><b>Query:</b> yes . i m going to the untied states for travel . what places do you recommend ?</p> <p><b>VAE-AM:</b> here are some leaflets detailing everything . i have nt seen you in ages .</p> <p><b>DialogWAE:</b> it s very nice to see you again , but i think it s a good idea to do .</p>
3	<p><b>Query:</b> medium one will be perfect , thank you</p> <p><b>VAE-AM:</b> is that all ?</p> <p><b>DialogWAE:</b> in that case , i have to be able to get some of them .</p>
4	<p><b>Query:</b> the annual charge rate is NUM .</p> <p><b>VAE-AM:</b> that s a lot of money . can you make it a little cheaper ?</p> <p><b>DialogWAE:</b> then what about the &lt;UNK&gt; ? i mean , i am not sure .</p>
5	<p><b>Query:</b> but it is too loose.do you have the smaller size ?</p> <p><b>VAE-AM:</b> ok , we will . how about this one ?</p> <p><b>DialogWAE:</b> in that case , i have to be able to get some of them . i think i can get a new one .</p>

Table 5: Generated responses in the single-turn setting

#	Generated Responses
1	<p><b>Query:</b> allan, i am really disappointed by your lack of effort on this project .</p> <p><b>VAE-AM:</b> it seems you are not aware of the actual work split . i have done more than anyone else.</p> <p><b>DialogWAE:</b> ok , it's not that bad . but it's not that bad . i don't want to have any other questions about it .</p>
2	<p><b>Query:</b> when is he coming ?</p> <p><b>VAE-AM:</b> next week . he is enjoying his month long vacation in miami .</p> <p><b>DialogWAE:</b> sorry , he's not very serious . he's got a lot of money to do .</p>
3	<p><b>Query:</b> i ll try my best to make a good impression on the interviewer .</p> <p><b>VAE-AM:</b> do nt worry pal. only weak candidates should. just dress properly and try to make it alive there haha ...</p> <p><b>DialogWAE:</b> that is ok , we ' ll be able to get a raise in the next few days . we will have to pay you a lot of money .</p>
4	<p><b>Query:</b> can you tell me where the pots and pans are ?</p> <p><b>VAE-AM:</b> absolutely not . i m not letting you cook today . you messed it real bad yesterday .</p> <p><b>DialogWAE:</b> yes , of course .</p>
5	<p><b>Query:</b> what kind of personal information ?</p> <p><b>VAE-AM:</b> your name , height , sex , and contact . also tick this check box if you have faced this issue before .</p> <p><b>DialogWAE:</b> not at the moment . but there is a lot of money in the city , but we have to pay a lot of money for the money .</p>

Table 6: Generated responses in the multi-turn setting