

# Implicit Deep Latent Variable Models for Text Generation

Le Fang<sup>†</sup>, Chunyuan Li<sup>§</sup>, Jianfeng Gao<sup>§</sup>, Wen Dong<sup>†</sup>, Changyou Chen<sup>†</sup>

<sup>†</sup>University at Buffalo, <sup>§</sup>Microsoft Research, Redmond  
{lefang, wendong, changyou}@buffalo.edu  
{chunyl, jfgao}@microsoft.com

## Abstract

Deep latent variable models (LVM) such as variational auto-encoder (VAE) have recently played an important role in text generation. One key factor is the exploitation of smooth latent structures to guide the generation. However, the representation power of VAEs is limited due to two reasons: (1) the Gaussian assumption is often made on the variational posteriors; and meanwhile (2) a notorious “posterior collapse” issue occurs. In this paper, we advocate sample-based representations of variational distributions for natural language, leading to implicit latent features, which can provide flexible representation power compared with Gaussian-based posteriors. We further develop an LVM to directly match the aggregated posterior to the prior. It can be viewed as a natural extension of VAEs with a regularization of maximizing mutual information, mitigating the “posterior collapse” issue. We demonstrate the effectiveness and versatility of our models in various text generation scenarios, including language modeling, unaligned style transfer, and dialog response generation. The source code to reproduce our experimental results is available on GitHub<sup>1</sup>.

## 1 Introduction

Deep latent variable models (LVM) such as variational auto-encoder (VAE) (Kingma and Welling, 2013; Rezende et al., 2014) are successfully applied for many natural language processing tasks, including language modeling (Bowman et al., 2015; Miao et al., 2016), dialogue response generation (Zhao et al., 2017b), controllable text generation (Hu et al., 2017) and neural machine translation (Shah and Barber, 2018) etc. One advantage of VAEs is the flexible distribution-based latent representation. It captures holistic properties of input, such as style, topic, and high-level linguistics/semantic features, which further guide the generation of diverse and relevant sentences.

However, the representation capacity of VAEs is restrictive due to two reasons. The first reason is rooted in the assumption of variational posteriors, which usually follow spherical Gaussian distributions with diagonal co-variance matrices. It has been shown that an approximation gap generally exists between the true posterior and the best possible variational posterior in a restricted family (Cremer et al., 2018). Consequently, the gap may militate against learning an optimal generative model, as its parameters may be always updated based on sub-optimal posteriors (Kim et al., 2018). The second reason is the so-called *posterior collapse* issue, which occurs when learning VAEs with an auto-regressive decoder (Bowman et al., 2015). It produces undesirable outcomes: the encoder yields meaningless posteriors that are very close to the prior, while the decoder tends to ignore the latent codes in generation (Bowman et al., 2015). Several attempts have been made to alleviate this issue (Bowman et al., 2015; Higgins et al., 2017; Zhao et al., 2017a; Fu et al., 2019; He et al., 2019).

These two seemingly unrelated issues are studied independently. In this paper, we argue that the posterior collapse issue is partially due to the restrictive Gaussian assumption, as it limits the optimization space of the encoder/decoder in a given distribution family. (i) To break the assumption, we propose to use sample-based representations for natural language, thus leading to implicit latent features. Such a representation is much more expressive than Gaussian-based posteriors. (ii) This implicit representation allows us to extend VAE and develop new LVM that further mitigate the posterior collapse issue. It represents all the sentences in the dataset as posterior samples in the latent space, and matches the aggregated posterior samples to the prior distribution. Consequently,

<sup>1</sup><https://github.com/fangleai/Implicit-LVM>

latent features are encouraged to cooperate and behave diversely to capture meaningful information for each sentence.

However, learning with implicit representations faces one challenge: it is intractable to evaluate the KL divergence term in the objectives. We overcome the issue by introducing a conjugate-dual form of the KL divergence (Rockafellar et al., 1966; Dai et al., 2018). It facilitates learning via training an auxiliary dual function. The effectiveness of our models is validated by producing consistently state-of-the-art results on a broad range of generation tasks, including language modeling, unsupervised style transfer, and dialog response generation.

## 2 Preliminaries

When applied to text generation, VAEs (Bowman et al., 2015) consist of two parts, a generative network (decoder) and an inference network (encoder). Given a training dataset  $\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^{|\mathcal{D}|}$ , where  $\mathbf{x}_i = [x_{1i}, \dots, x_{Ti}]$  represents  $i$ th sentence of length  $T$ . Starting from a prior distribution  $p(\mathbf{z})$ , VAE generates a sentence  $\mathbf{x}$  using the deep generative network  $p_{\theta}(\mathbf{x}|\mathbf{z})$ , where  $\theta$  is the network parameter. Therefore, the joint distribution  $p_{\theta}(\mathbf{x}, \mathbf{z})$  is defined as  $p(\mathbf{z})p_{\theta}(\mathbf{x}|\mathbf{z})$ . The prior  $p(\mathbf{z})$  is typically assumed as a standard multivariate Gaussian. Due to the sequential nature of natural language, the decoder  $p_{\theta}(\mathbf{x}|\mathbf{z})$  takes an auto-regressive form  $p_{\theta}(\mathbf{x}|\mathbf{z}) = \prod_{t=1}^T p_{\theta}(x_t|x_{<t}, \mathbf{z})$ . The goal of model training is to maximize the marginal data log-likelihood  $\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \log p_{\theta}(\mathbf{x})$ .

However, it is intractable to perform posterior inference. An  $\phi$ -parameterized encoder is introduced to approximate  $p_{\theta}(\mathbf{z}|\mathbf{x}) \propto p_{\theta}(\mathbf{x}|\mathbf{z})p(\mathbf{z})$  with a variational distribution  $q_{\phi}(\mathbf{z}|\mathbf{x})$ . Variational inference is employed for VAE learning, yielding following evidence lower bound (ELBO):

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \log p_{\theta}(\mathbf{x}) \geq \mathcal{L}_1 = \mathcal{L}_2 \quad \text{with} \quad \mathcal{L}_1 = \mathcal{L}_E + \mathcal{L}_R \quad \text{where} \quad (1)$$

$$\mathcal{L}_E = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[ \mathbb{E}_{\mathbf{z} \sim q_{\phi}(\mathbf{z}|\mathbf{x})} \log p_{\theta}(\mathbf{x}|\mathbf{z}) \right] \quad (2)$$

$$\mathcal{L}_R = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [-\text{KL}(q_{\phi}(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z}))] \quad (3)$$

$$\mathcal{L}_2 = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \log p_{\theta}(\mathbf{x}) + \mathcal{L}_G \quad \text{where} \quad (4)$$

$$\mathcal{L}_G = -\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \text{KL}(q_{\phi}(\mathbf{z}|\mathbf{x}) \parallel p_{\theta}(\mathbf{z}|\mathbf{x})) \quad (5)$$

Note that  $\mathcal{L}_1$  and  $\mathcal{L}_2$  provide two different views for the VAE objective:

- $\mathcal{L}_1$  consists of a reconstruction error term  $\mathcal{L}_E$  and a KL divergence regularization term

$\mathcal{L}_R$ . With a strong auto-regressive decoder  $p_{\theta}(\mathbf{x}|\mathbf{z})$ , the objective tends to degenerate all encoding distribution  $q_{\phi}(\mathbf{z}|\mathbf{x})$  to the prior, causing  $\mathcal{L}_R \rightarrow 0$ , i.e., the posterior collapse issue.

- $\mathcal{L}_2$  indicates that VAE requests a flexible encoding distribution family to minimize the approximation gap  $\mathcal{L}_G$  between the true posterior and the best possible encoding distribution. This motivates us to perform more flexible posterior inference with implicit representations.

## 3 The Proposed Models

We introduce a sample-based latent representation for natural language, and develop two models that leverage its advantages. (1) When replacing the Gaussian variational distributions with sample-based distributions in VAEs, we derive *implicit* VAE (iVAE). (2) We further extend VAE to maximize mutual information between latent representations and observed sentences, leading to a variant termed as iVAE<sub>MI</sub>.

### 3.1 Implicit VAE

**Implicit Representations** Instead of assuming an explicit density form such as Gaussian, we define a sampling mechanism to represent  $q_{\phi}(\mathbf{z}|\mathbf{x})$  as a set of sample  $\{\mathbf{z}_{\mathbf{x},i}\}_{i=1}^M$ , through the encoder as

$$\mathbf{z}_{\mathbf{x},i} = G_{\phi}(\mathbf{x}, \epsilon_i), \epsilon_i \sim q(\epsilon) \quad (6)$$

where the  $i$ -th sample is drawn from a neural network  $G_{\phi}$  that takes  $(\mathbf{x}, \epsilon_i)$  as input;  $q(\epsilon)$  is a simple distribution such as standard Gaussian. It is difficult to naively combine the random noise  $\epsilon$  with the sentence  $\mathbf{x}$  (a sequence of discrete tokens) as the input of  $G_{\phi}$ . Our solution is to concatenate noise  $\epsilon_i$  with hidden representations  $\mathbf{h}$  of  $\mathbf{x}$ .  $\mathbf{h}$  is generated using a LSTM encoder, as illustrated in Figure 1.

**Dual form of KL-divergence** Though theoretically promising, the implicit representations in (6) render difficulties in optimizing the KL term  $\mathcal{L}_R$  in (3), as the functional form is no longer tractable with implicit  $q_{\phi}(\mathbf{z}|\mathbf{x})$ . We resort to evaluating its dual form based on Fenchel duality theorem (Rockafellar et al., 1966; Dai et al., 2018):

$$\begin{aligned} & \text{KL}(q_{\phi}(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z})) \\ &= \max_{\nu} \mathbb{E}_{\mathbf{z} \sim q_{\phi}(\mathbf{z}|\mathbf{x})} \nu_{\psi}(\mathbf{x}, \mathbf{z}) - \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} \exp(\nu_{\psi}(\mathbf{x}, \mathbf{z})), \end{aligned} \quad (7)$$




Figure 1: Illustration of the proposed implicit LVMs.  $\nu(x, z)$  is used in iVAE, and  $\nu(z)$  is used in iVAE<sub>MI</sub>. In this example, the prior is  $p(z) = \mathcal{N}(0, 1)$ , the sample-based aggregated posterior  $q(z) = \int q_\phi(z|x)q(x)dx$  for four observations is shown, where the posterior  $q_\phi(z|x)$  for each observation is visualized in a different color.

where  $\nu_\psi(x, z)$  is an auxiliary dual function, parameterized by a neural network with weights  $\psi$ . By replacing the KL term with this dual form, the implicit VAE has the following objective:

$$\begin{aligned} \mathcal{L}_{\text{iVAE}} = & \mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{z \sim q_\phi(z|x)} \log p_\theta(x|z) \\ & - \mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{z \sim q_\phi(z|x)} \nu_\psi(x, z) \\ & + \mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{z \sim p(z)} \exp(\nu_\psi(x, z)), \end{aligned} \quad (8)$$

**Training scheme** Implicit VAE inherits the end-to-end training scheme of VAEs with extra work on training the auxiliary network  $\nu_\psi(x, z)$ :

- Sample a mini-batch of  $x_i \sim \mathcal{D}$ ,  $\epsilon_i \sim q(\epsilon)$ , and generate  $z_{x_i, \epsilon_i} = G(x_i, \epsilon_i; \phi)$ ; Sample a mini-batch of  $z_i \sim p(z)$ .

- Update  $\psi$  in  $\nu_\psi(x, z)$  to maximize

$$\sum_i \nu_\psi(x_i, z_{x_i, \epsilon_i}) - \sum_i \exp(\nu_\psi(x_i, z_i)) \quad (9)$$

- Update parameters  $\{\phi, \theta\}$  to maximize

$$\sum_i \log p_\theta(x_i | z_{x_i, \epsilon_i}) - \sum_i \nu_\psi(x_i, z_{x_i, \epsilon_i}) \quad (10)$$

In practice, we implement  $\nu_\psi(x, z)$  with a multi-layer perceptron (MLP), which takes the concatenation of  $h$  and  $z$ . In another word, the auxiliary network distinguishes between  $(x, z_x)$  and  $(x, z)$ , where  $z_x$  is drawn from the posterior and  $z$  is drawn from the prior, respectively. We found the MLP-parameterized auxiliary network converges faster than LSTM encoder and decoder (Hochreiter and Schmidhuber, 1997). This means that the auxiliary network practically provides an accurate approximation to the KL regularization  $\mathcal{L}_R$ .

### 3.2 Mutual Information Regularized iVAE

It is noted the inherent deficiency of the original VAE objective in (3): the KL divergence regularization term matches each posterior distribution independently to the same prior. This is prone to posterior collapse in text generation, due to a strong auto-regressive decoder  $p_\theta(x|z)$ . When sequentially generating  $x_t$ , the model learns to solely rely on the ground-truth  $[x_1, \dots, x_{t-1}]$ , and ignore the dependency from  $z$  (Fu et al., 2019). It results in the learned variational posteriors  $q_\phi(z|x)$  to exactly match  $p(z)$ , without discriminating data  $x$ .

To better regularize the latent space, we propose to replace  $\mathcal{L}_R = \mathbb{E}_{x \sim \mathcal{D}} [-\text{KL}(q_\phi(z|x) \parallel p(z))]$  in (3), with the following KL divergence:

$$-\text{KL}(q_\phi(z) \parallel p(z)), \quad (11)$$

where  $q_\phi(z) = \int q(x)q_\phi(z|x)dx$  is the *aggregated posterior*,  $q(x)$  is the empirical data distribution for the training dataset  $\mathcal{D}$ . The integral is estimated by ancestral sampling in practice, *i.e.* we first sample  $x$  from dataset and then sample  $z \sim q_\phi(z|x)$ .

In (11), variational posterior is regularized as a whole  $q_\phi(z)$ , encouraging posterior samples from different sentences to cooperate to satisfy the objective. It implies a solution that each sentence is represented as a local region in the latent space, the aggregated representation of all sentences match the prior; This avoids the degenerated solution from (3) that the feature representation of individual sentence spans over the whole space.

**Connection to mutual information** The proposed latent variable model coincides with (Zhao et al., 2018, 2017a) where mutual information is introduced into the optimization, based on the following decomposition result (Please see detailed

proof in Appendix A):

$$\begin{aligned} -\text{KL}(q_\phi(z) \parallel p(z)) &= I(x, z) \\ &\quad - \mathbb{E}_x \text{KL}(q_\phi(z|x) \parallel p(z)), \end{aligned}$$

where  $I(x, z)$  is the mutual information between  $z$  and  $x$  under the joint distribution  $q_\phi(x, z) = q(x)q_\phi(z|x)$ . Therefore, the objective in (11) also maximizes the mutual information between individual sentences and their latent features. We term the new LVM objective as iVAE<sub>MI</sub>:

$$\begin{aligned} \mathcal{L}_{\text{iVAE}_{\text{MI}}} &= \mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{z \sim q_\phi(z|x)} \log p_\theta(x|z) \\ &\quad - \text{KL}(q_\phi(z) \parallel p(z)) \end{aligned} \quad (12)$$

**Training scheme** Note that the aggregated posterior  $q_\phi(z)$  is also a sample-based distribution. Similarly, we evaluate (12) through its dual form:

$$\begin{aligned} \text{KL}(q_\phi(z) \parallel p(z)) & \quad (13) \\ = \max_{\nu} \mathbb{E}_{z \sim q_\phi(z)} \nu_\psi(z) - \mathbb{E}_{z \sim p(z)} \exp(\nu_\psi(z)). \end{aligned}$$

Therefore, iVAE<sub>MI</sub> in (12) can be written as:

$$\begin{aligned} \mathcal{L}_{\text{iVAE}_{\text{MI}}} &= \mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{z \sim q_\phi(z|x)} \log p_\theta(x|z) \quad (14) \\ &\quad - \mathbb{E}_{z \sim q_\phi(z)} \nu_\psi(z) + \mathbb{E}_{z \sim p(z)} \exp(\nu_\psi(z)), \end{aligned}$$

where the auxiliary network  $\nu_\psi(z)$  is parameterized as a neural network. Different from iVAE,  $\nu_\psi(z)$  in iVAE<sub>MI</sub> only takes posterior samples as input. The training algorithm is similar to iVAE in Section 3.1, except a different auxiliary network  $\nu_\psi(z)$ . In Appendix B, we show the full algorithm of iVAE<sub>MI</sub>.

We illustrate the proposed methods in Figure 1. Note that both iVAE and iVAE<sub>MI</sub> share the same model architecture, except a different auxiliary network  $\nu$ .

## 4 Related Work

### 4.1 Solutions to posterior collapse

Several attempts have been made to alleviate the posterior collapse issue. The KL annealing scheme in language VAEs has been first used in (Bowman et al., 2015). An effective cyclical KL annealing schedule is used in (Fu et al., 2019), where the KL annealing process is repeated multiple times. KL term weighting scheme is also adopted in  $\beta$ -VAE (Higgins et al., 2017) for disentanglement. On model architecture side, dilated CNN was considered to replace auto-regressive LSTMs for decoding (Yang et al., 2017). The bag-of-word auxiliary

loss was proposed to improve the dependence on latent representations in generation (Zhao et al., 2017b). More recently, lagging inference proposes to aggressively update encoder multiple times before a single decoder update (He et al., 2019). Semi-amortized VAE refines variational parameters from an amortized encoder per instance with stochastic variational inference (Kim et al., 2018).

All these efforts utilize the Gaussian-based forms for posterior inference. Our paper is among the first ones to attribute posterior collapse issue to the restrictive Gaussian assumption, and advocate more flexible sample-based representations.

### 4.2 Implicit Feature Learning

Sample-based distributions, as well as implicit features, have been widely used in representation learning (Donahue et al., 2017; Li et al., 2017a). Vanilla autoencoders learn point masses of latent features rather than their distributions. Adversarial variational Bayes introduces an auxiliary discriminator network like GANs (Goodfellow et al., 2014; Makhzani et al., 2015) to learn almost arbitrarily distributed latent variables (Mescheder et al., 2017; Pu et al., 2017b). We explore the similar spirit in the natural language processing (NLP) domain. Amortized MCMC and particle based methods are introduced for LVM learning in (Li et al., 2017d; Pu et al., 2017a; Chen et al., 2018). Coupled variational Bayes (Dai et al., 2018) emphasizes an optimization embedding, *i.e.*, a flow of particles, in a general setting of non-parametric variational inference. It also utilizes similar dual form with auxiliary function  $\nu_\psi(x, z)$  to evaluate KL divergence. Adversarially regularized autoencoders (Makhzani et al., 2015; Kim et al., 2017) use similar objectives with iVAEs, in the form of a reconstruction error plus a specific regularization evaluated with implicit samples. Mutual information has also been considered into regularization in (Zhao et al., 2018, 2017a) to obtain more informative representations.

Most previous works focus on image domain. It is largely unexplored in NLP. Further, the auto-regressive decoder renders an additional challenge when applying implicit latent representations. Adversarial training with samples can be empirically unstable, and slow even applying recent stabilization techniques in GANs (Arjovsky et al., 2017; Gulrajani et al., 2017). To the best of our knowledge, this paper presents the first to effectively apply implicit feature representations, to NLP.



Dataset	PTB	Yahoo	Yelp
num. sents.	42,000	100,000	100,000
min. len.	2	21	21
max. len.	78	201	201
avg. len.	21.9	79.9	96.7

Table 1: Statistics of datasets for language modeling.

## 5 Experiments

In this section, the effectiveness of our methods is validated by largely producing state-of-the-art metrics on a broad range of text generation tasks under various scenarios.

### 5.1 Language Modeling

**Datasets.** We consider three public datasets, the Penn Treebank (PTB) (Marcus et al., 1993; Bowman et al., 2015), Yahoo, and Yelp corpora (Yang et al., 2017; He et al., 2019). PTB is a relatively small dataset with sentences of varying lengths, whereas Yahoo and Yelp contain larger amounts of with long sentences. Detailed statistics of these datasets are shown in Table 1.

**Settings** We implement both encoder and decoder as one-layer LSTM. As illustrated in Figure 1, concatenation of the last hidden state of the encoder LSTM and a Gaussian noise is fed into a MLP to draw a posterior sample. When decoding, a latent sample is fed into a MLP to produce the initial hidden state of the decoder LSTM. A KL-term annealing scheme is also used (Bowman et al., 2015). We list more detailed hyper-parameters and architectures in Appendix C.1.2.

**Baselines.** We compare with several state-of-the-art VAE language modeling methods, including (1) VAEs with a monotonically KL-annealing schedule (Bowman et al., 2015); (2)  $\beta$ -VAE (Higgins et al., 2017), VAEs with a small penalty on KL term scaled by  $\beta$ ; (3) SA-VAE (Kim et al., 2018), mixing instance-specific variational inference with amortized inference; (4) Cyclical VAE (Fu et al., 2019) that periodically anneals the KL term; (5) Lagging VAE, which aggressively updates encoder multiple times before a single decoder update.

**Evaluation metrics.** Two categories of metrics are used to study VAEs for language modeling:

- To characterize the modeling ability of the observed sentences, we use the negative ELBO

Methods	-ELBO↓	PPL↓	KL↑	MI↑	AU↑
Dataset: PTB					
VAE	102.6	108.26	1.08	0.8	2
$\beta(0.5)$ -VAE	104.5	117.92	<b>7.50</b>	3.1	5
SA-VAE	102.6	107.71	1.23	0.7	2
Cyc-VAE	103.1	110.50	3.48	1.8	5
iVAE	<b>87.6</b>	<b>54.46</b>	6.32	<b>3.5</b>	<b>32</b>
iVAE <sub>MI</sub>	<b>87.2</b>	<b>53.44</b>	<b>12.51</b>	<b>12.2</b>	<b>32</b>
Dataset: Yahoo					
VAE	328.6	61.21	0.0	0.0	0
$\beta(0.4)$ -VAE	328.7	61.29	6.3	2.8	8
SA-VAE	327.2	60.15	5.2	2.7	10
Lag-VAE	326.7	59.77	5.7	2.9	15
iVAE	<b>309.5</b>	<b>48.22</b>	<b>8.0</b>	<b>4.4</b>	<b>32</b>
iVAE <sub>MI</sub>	<b>309.1</b>	<b>47.93</b>	<b>11.4</b>	<b>10.7</b>	<b>32</b>
Dataset: Yelp					
VAE	357.9	40.56	0.0	0.0	0
$\beta(0.4)$ -VAE	358.2	40.69	4.2	2.0	4
SA-VAE	355.9	39.73	2.8	1.7	8
Lag-VAE	355.9	39.73	3.8	2.4	11
iVAE	<b>348.2</b>	<b>36.70</b>	<b>7.6</b>	<b>4.6</b>	<b>32</b>
iVAE <sub>MI</sub>	<b>348.7</b>	<b>36.88</b>	<b>11.6</b>	<b>11.0</b>	<b>32</b>

Table 2: Language modeling on three datasets.

as the sum of reconstruction loss and KL term, as well as perplexity (PPL).

- Compared with traditional neural language models, VAEs has its unique advantages in feature learning. To measure the quality of learned features, we consider (1) KL:  $\text{KL}(q_\phi(z|x) \parallel p(z))$ ; (2) Mutual information (MI)  $I(x, z)$  under the joint distribution  $q_\phi(x, z)$ ; (3) Number of active units (AU) of latent representation. The activity of a latent dimension  $z$  is measured as  $A_z = \text{Cov}_x(\mathbb{E}_{z \sim q(z|x)}[z])$ , which is defined as active if  $A_z > 0.01$ .

The evaluation of implicit LVMs is unexplored in language models, as there is no analytical forms for the KL term. We consider to evaluate both  $\text{KL}(q_\phi(z) \parallel p(z))$  and  $\text{KL}(q_\phi(z|x) \parallel p(z))$  by training a fully connected  $\nu$  network in Eq. (7) and (13). To avoid the inconsistency between  $\nu(x, z)$  and  $\nu(z)$  networks due to training, we train them using the same data and optimizer in every iteration. We evaluate each distribution  $q_\phi(z|x)$  with 128 code samples per  $x$ .

We report the results in Table 2. A better language model would pursue a lower negative ELBO (also lower reconstruction errors, lower PPL), and make sufficient use of the latent space (*i.e.*, maintain relatively high KL term, higher mutual infor-

Dataset	PTB		Yahoo		Yelp	
	Re.↓	Abs.↓	Re.↓	Abs.↓	Re.↓	Abs.↓
VAE/ $\beta$ -	1.0	1.3	1.0	5.3	1.0	5.7
SA-VAE	5.5	7.1	9.9	52.9	10.3	59.3
Lag-VAE	-	-	2.2	11.7	3.7	21.4
iVAEs	1.4	1.8	1.3	6.9	1.3	7.5

Table 3: Total training time in hours: absolute time and relative time versus VAE.

$t=0$	in new york the company declined comment
$t=0.1$	in new york the company declined comment
$t=0.2$	in new york the transaction was suspended
$t=0.3$	in the securities company said yesterday
$t=0.4$	in other board the transaction had disclosed
$t=0.5$	other of those has been available
$t=0.6$	both of companies have been unchanged
$t=0.7$	both men have received a plan to restructure
$t=0.8$	and to reduce that it owns
$t=0.9$	and to continue to make prices
$t=1$	and they plan to buy more today

Table 4: Interpolating latent representation.

mation and more active units). Under all these metrics, the proposed iVAEs achieve much better performance. The posterior collapse issue is largely alleviated as indicated by the improved KL and MI values, especially with iVAE<sub>MI</sub> which directly takes mutual information into account.

The comparison on training time is shown in Table 3. iVAE and iVAE<sub>MI</sub> requires updating an auxiliary network, it spends 30% more time than traditional VAEs. This is more efficient than SA-VAE and Lag-VAE.

**Latent space interpolation.** One favorable property of VAEs (Bowman et al., 2015; Zhao et al., 2018) is to provide smooth latent representations that capture sentence semantics. We demonstrate this by interpolating two latent feature, each of which represents a unique sentence. Table 4 shows the generated examples. We take two sentences  $x_1$  and  $x_2$ , and obtain their latent features as the sample-averaging results for  $z_1$  and  $z_2$ , respectively, from the implicit encoder, and then greedily decode conditioned on the interpolated feature  $z_t = z_1 \cdot (1 - t) + z_2 \cdot t$  with  $t$  increased from 0 to 1 by a step size of 0.1. It generates sentences with smooth semantic evolution.

**Improved Decoder.** One might wonder whether the improvements come from simply having a more flexible encoder during evaluation, rather than from

Model	Forward↓	Reverse↓
VAE	18494	10149
Cyc-VAE	3390	5587
AE	672	2589
$\beta(0)$ -VAE	625	1897
$\beta(0.5)$ -VAE	939	4078
SA-VAE	341	10651
iVAE	<b>116</b>	<b>1520</b>
iVAE <sub>MI</sub>	<b>134</b>	<b>1137</b>

Table 5: Forward and reverse PPL on PTB.

utilizing high-quality latent features, and learning a better decoder. We use PTB to confirm our findings. We draw samples from the prior  $p(z)$ , and greedily decode them using the trained decoder. The quality of the generated text is evaluated by an external library “KenLM Language Model Toolkit” (Heafield et al., 2013) with two metrics (Kim et al., 2017): (1) Forward PPL: the fluency of the generated text based on language models derived from the PTB training corpus; (2) Reverse PPL: the fluency of PTB corpus based on language model derived from the generated text, which measures the extent to which the generations are representative of the PTB underlying language model. For both the PPL numbers, the lower the better. We use  $n=5$  for  $n$ -gram language models in “KenLM”.

As shown in Table 5, implicit LVMs outperform others in both PPLs, which confirms that the implicit representation can lead to better decoders. The vanilla VAE model performs the worst. This is expected, as the posterior collapse issue results in poor utilization of a latent space. Besides, we can see that iVAE<sub>MI</sub> generates comparably fluent but more diverse text than pure iVAE, from the lower reverse PPL values. This is reasonable, due to the ability of iVAE<sub>MI</sub> to encourage diverse latent samples per sentence with the aggregated regularization in the latent space.

## 5.2 Unaligned style transfer

We next consider the task of unaligned style transfer, which represents a scenario to generate text with desired specifications. The goal is to control one style aspect of a sentence that is independent of its content. We consider non-parallel corpora of sentences, where the sentences in two corpora have the same content distribution but with different styles, and no paired sentence is provided.

**Model Extension.** The success of this task depends on the exploitation of the distributional equivalence of content to learn sentiment independent content code and decoding it to a different style. To ensure such independence, we extend iVAE<sub>MI</sub> by adding a sentiment classifier loss to its objective (14), similar to the previous style transfer methods (Shen et al., 2017; Kim et al., 2017). Let  $y$  be the style attribute,  $x_p$  and  $x_n$  (with corresponding features  $z_p, z_n$ ) be sentences with positive and negative sentiments, respectively. The style classifier loss  $\mathcal{L}_{\text{class}}(z_p, z_n)$  is the cross entropy loss of a binary classifier.

The classifier and encoder are trained adversarially: (1) the classifier is trained to distinguish latent features with different sentiments; (2) the encoder is trained to fool the classifier in order to remove distinctions of content features from sentiments. In practice, the classifier is implemented as a MLP. We implement two separate decoder LSTMs for clean sentiment decoding: one for positive  $p(x|z, y = 1)$ , and one for negative sentiment  $p(x|z, y = 0)$ . The prior  $p(z)$  is also implemented as an implicit distribution, via transforming noise from a standard Gaussian through a MLP. Appendix C.2.2 lists more details.

**Baseline.** We compare iVAE<sub>MI</sub> with a state-of-the-art unaligned sentiment transferer: the adversarially regularized autoencoder (ARAE) (Kim et al., 2017), which directly regularizes the latent space with Wasserstein distance measure.

**Datasets.** Following (Shen et al., 2017), the Yelp restaurant reviews dataset is processed from the original Yelp dataset in language modeling. Reviews with user rating above three are considered *positive*, and those below three are considered *negative*. The pre-processing allows sentiment analysis on sentence level with feasible sentiment, ending up with shorter sentences with each at most 15 words than those in language modeling. Finally, we get two sets of unaligned reviews: 250K negative sentences, and 350K positive ones. Other dataset details are shown in Appendix C.2.1.

**Evaluation Metrics.** (1) Acc: the accuracy of transferring sentences into another sentiment measured by an automatic classifier: the “fasttext” library (Joulin et al., 2017); (2) BLEU: the consistency between the transferred text and the original; (3) PPL: the reconstruction perplexity of original sentences without altering sentiment; (4) RPPL:

**Input:** it was super dry and had a **weird taste** to the entire **slice** .  
**ARAE:** it was **super nice** and **the owner** was **super sweet and helpful** .  
**iVAE<sub>MI</sub>:** it was **super tasty** and a good size with the best in the **burgh** .

**Input:** so i only had **half** of the regular **fries and my soda** .  
**ARAE:** it 's the **best to eat** and **had a great meal** .  
**iVAE<sub>MI</sub>:** so i had a **huge side** and the price was great .

**Input:** i am just **not a fan** of this kind of **pizza** .  
**ARAE:** i am very **pleased** and will definitely use **this place** .  
**iVAE<sub>MI</sub>:** i am just **a fan** of **the chicken and egg roll** .

**Input:** i have eaten the **lunch buffet** and it was **outstanding** !  
**ARAE:** once again , i was **told by the wait** and was **seated** .  
**iVAE<sub>MI</sub>:** we were **not impressed** with **the buffet** there last night .

**Input:** my favorite food is **kung pao beef** , it is **delicious** .  
**ARAE:** my **husband** was on the **phone** , which i tried it .  
**iVAE<sub>MI</sub>:** **my chicken** was n't warm , though it is n't **delicious** .

**Input:** overall , it was a very **positive dining experience** .  
**ARAE:** overall , it was very **rude and unprofessional** .  
**iVAE<sub>MI</sub>:** overall , it was a nightmare of **terrible experience** .

Table 6: Sentiment transfer on Yelp. (Up: From negative to positive, Down: From positive to negative.)

Model	Acc↑	BLEU↑	PPL↓	RPPL↓	Flu↑	Sim↑
ARAE	<b>95</b>	32.5	6.8	395	3.6	3.5
iVAE <sub>MI</sub>	92	<b>36.7</b>	<b>6.2</b>	<b>285</b>	<b>3.8</b>	<b>3.9</b>

Table 7: Sentiment Transfer on Yelp.

the reverse perplexity that evaluates the training corpus based on language model derived from the generated text, which measures the extent to which the generations are representative of the training corpus; (5) Flu: human evaluated index on the fluency of transferred sentences when read alone (1-5, 5 being most fluent as natural language); (6) Sim: the human evaluated similarity between the original and the transferred sentences in terms of their contents (1-5, 5 being most similar). Note that the similarity measure doesn’t care sentiment but only the topic covered by sentences. For human evaluation, we show 1000 randomly selected pairs of original and transferred sentences to crowdsourcing readers, and ask them to evaluate the “Flu” and “Sim” metrics stated above. Each measure is averaged among crowdsourcing readers.

As shown in Table 7, iVAE<sub>MI</sub> outperforms ARAE in metrics except Acc, showing that iVAE<sub>MI</sub> captures informative representations, generates consistently opposite sentences with similar grammatical structure and reserved semantic meaning. Both methods perform successful sentiment transfer as shown by Acc values. iVAE<sub>MI</sub> achieves a little lower Acc due to much more content reserving, even word reserving, of the source sentences.

Table 6 presents some examples. In each box, we show the source sentence, the transferred sentence by ARAE and iVAE<sub>MI</sub>, respectively. We observe that ARAE usually generates new sentences that miss the content of the source, while iVAE<sub>MI</sub> shows better content-preserving.

### 5.3 Dialog response generation

We consider the open-domain dialog response generation task, where we need to generate a natural language response given a dialog history. It is crucial to learn a meaningful latent feature representation of the dialog history in order to generate a consistent, relevant, and contentful response that is likely to drive the conversation (Gao et al., 2019).

**Datasets.** We consider two mainstream datasets in recent studies (Zhao et al., 2017b, 2018; Fu et al., 2019; Gu et al., 2018): Switchboard (Godfrey and Holliman, 1997) and Dailydialog (Li et al., 2017c). Switchboard contains 2,400 two-way telephone conversations under 70 specified topics. Dailydialog has 13,118 daily conversations for a English learner. We process each utterance as the response of previous 10 context utterances from both speakers. The datasets are separated into training, validation, and test sets as convention: 2316:60:62 for Switchboard and 10:1:1 for Dailydialog, respectively.

**Model Extension.** We adapt iVAE<sub>MI</sub> by integrating the context embedding  $c$  into all model components. The prior  $p(z|c)$  is defined as an implicit mapping between context embedding  $c$  and prior samples, which is not pre-fixed but learned together with the variational posterior for more modeling flexibility. The encoder  $q(z|x, c)$ , auxiliary dual function  $\nu_\psi(z, c)$  and decoder  $p(x|c, z)$  depend on context embedding  $c$  as well. Both encoder and decoder are implemented as GRUs. The utterance encoder is a bidirectional GRU with 300 hidden units in each direction. The context encoder and decoder are both GRUs with 300 hidden units. Appendix C.3.1 presents more training details.

**Baseline.** We benchmark representative baselines and state-of-the-art approaches, include: SeqGAN, a GAN based model for sequence generation (Li et al., 2017b); CVAE baseline (Zhao et al., 2017b); dialogue WAE, a conditional Wasserstein auto-encoder for response generation (Gu et al., 2018).

Metrics	SeqGAN	CVAE	WAE	iVAE <sub>MI</sub>
Dataset: Switchboard				
BLEU-R $\uparrow$	0.282	0.295	0.394	<b>0.427</b>
BLEU-P $\uparrow$	<b>0.282</b>	0.258	0.254	0.254
BLEU-F1 $\uparrow$	0.282	0.275	0.309	<b>0.319</b>
BOW-A $\uparrow$	0.817	0.836	0.897	<b>0.930</b>
BOW-E $\uparrow$	0.515	0.572	0.627	<b>0.670</b>
BOW-G $\uparrow$	0.748	0.846	0.887	<b>0.900</b>
Intra-dist1 $\uparrow$	0.705	0.803	0.713	<b>0.828</b>
Intra-dist2 $\uparrow$	0.521	0.415	0.651	<b>0.692</b>
Inter-dist1 $\uparrow$	0.070	0.112	0.245	<b>0.391</b>
Inter-dist2 $\uparrow$	0.052	0.102	0.413	<b>0.668</b>
Dataset: Dailydialog				
BLEU-R $\uparrow$	0.270	0.265	0.341	<b>0.355</b>
BLEU-P $\uparrow$	0.270	0.222	<b>0.278</b>	0.239
BLEU-F1 $\uparrow$	0.270	0.242	<b>0.306</b>	0.285
BOW-A $\uparrow$	0.907	0.923	0.948	<b>0.951</b>
BOW-E $\uparrow$	0.495	0.543	0.578	<b>0.609</b>
BOW-G $\uparrow$	0.774	0.811	0.846	<b>0.872</b>
Intra-dist1 $\uparrow$	0.747	<b>0.938</b>	0.830	0.897
Intra-dist2 $\uparrow$	0.806	0.973	0.940	<b>0.975</b>
Inter-dist1 $\uparrow$	0.075	0.177	0.327	<b>0.501</b>
Inter-dist2 $\uparrow$	0.081	0.222	0.583	<b>0.868</b>

Table 8: Dialog response generation on two datasets.

**Evaluation Metrics.** We adopt several widely used numerical metrics to systematically evaluate the response generation, including BLEU score (Papineni et al., 2002), BOW Embedding (Liu et al., 2016) and Distinct (Li et al., 2015), as used in (Gu et al., 2018). For each testing context, we sample 10 responses from each model.

- BLEU measures how much a generated response contains n-gram overlaps with the references. We compute 4-gram BLEU. For each test context, we sample 10 responses from the models and compute their BLEU scores. BLEU precision and recall is defined as the average and maximum scores, respectively (Zhao et al., 2017b).
- BOW embedding is the cosine similarity of bag-of-words embeddings between the generations and the reference. We compute three different BOW embedding: greedy, average, and extreme.
- Distinct evaluates the diversity of the generated responses: dist-n is defined as the ratio of unique n-grams ( $n=1,2$ ) over all n-grams in the generated responses. We evaluate diversities for both within each sampled response and among all responses as intra-dist and inter-dist, respectively.



Tables 8 show the performance comparison. iVAE<sub>MI</sub> achieves consistent improvement on a majority of the metrics. Especially, the BOW embeddings and Distinct get significantly improvement, which implies that iVAE<sub>MI</sub> produces both meaningful and diverse latent representations.

## 6 Conclusion

We present two types of implicit deep latent variable models, iVAE and iVAE<sub>MI</sub>. Core to these two models is the sample-based representation of the latent features in LVM, in replacement of traditional Gaussian-based distributions. Extensive experiments show that the proposed implicit LVM models consistently outperform the vanilla VAEs on three tasks, including language modeling, style transfer and dialog response generation.

## Acknowledgments

We acknowledge Yufan Zhou for helpful discussions and the reviewers for their comments.

## References

- Martín Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein gan. *CoRR*, abs/1701.07875.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. 2015. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349*.
- Changyou Chen, Chunyuan Li, Liqun Chen, Wenlin Wang, Yunchen Pu, and Lawrence Carin. 2018. Continuous-time flows for efficient inference and density estimation. *ICML*.
- Chris Cremer, Xuechen Li, and David Duvenaud. 2018. Inference suboptimality in variational autoencoders. *arXiv preprint arXiv:1801.03558*.
- Bo Dai, Hanjun Dai, Niao He, Weiyang Liu, Zhen Liu, Jianshu Chen, Lin Xiao, and Le Song. 2018. Coupled variational bayes via optimization embedding. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9713–9723. Curran Associates, Inc.
- Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. 2017. Adversarial feature learning. *ICLR*.
- Hao Fu, Chunyuan Li, Xiaodong Liu, Jianfeng Gao, Asli Celikyilmaz, and Lawrence Carin. 2019. Cyclical annealing schedule: A simple approach to mitigating kl vanishing. *NAACL*.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2019. Neural approaches to conversational ai. *Foundations and Trends® in Information Retrieval*, 13(2-3):127–298.
- John J Godfrey and Edward Holliman. 1997. Switchboard-1 release 2. *Linguistic Data Consortium, Philadelphia*, 926:927.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. **Generative adversarial nets**. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc.
- Xiaodong Gu, Kyunghyun Cho, Jungwoo Ha, and Sunghun Kim. 2018. Dialogwae: Multimodal response generation with conditional wasserstein auto-encoder. *arXiv preprint arXiv:1805.12352*.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. 2017. **Improved training of wasserstein gans**. In *Advances in Neural Information Processing Systems 30*, pages 5769–5779. Arxiv: 1704.00028.
- Junxian He, Daniel Spokoyny, Graham Neubig, and Taylor Berg-Kirkpatrick. 2019. Lagging inference networks and posterior collapse in variational autoencoders. *arXiv preprint arXiv:1901.05534*.
- Kenneth Heafield, Ivan Pouzyrevsky, Jonathan H. Clark, and Philipp Koehn. 2013. **Scalable modified Kneser-Ney language model estimation**. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 690–696, Sofia, Bulgaria.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. 2017. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, volume 3.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. 2017. Toward controlled generation of text. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1587–1596. JMLR. org.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431. Association for Computational Linguistics.

- Yoon Kim, Sam Wiseman, Andrew Miller, David Sonntag, and Alexander Rush. 2018. Semi-amortized variational autoencoders. In *International Conference on Machine Learning*, pages 2683–2692.
- Yoon Kim, Kelly Zhang, Alexander M Rush, Yann LeCun, et al. 2017. Adversarially regularized autoencoders. *arXiv preprint arXiv:1706.04223*.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Chunyu Li, Hao Liu, Changyou Chen, Yuchen Pu, Liqun Chen, Ricardo Henao, and Lawrence Carin. 2017a. Alice: Towards understanding adversarial learning for joint distribution matching. In *Advances in Neural Information Processing Systems*, pages 5495–5503.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017b. Adversarial learning for neural dialogue generation. *arXiv preprint arXiv:1701.06547*.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017c. Dailydialog: A manually labelled multi-turn dialogue dataset. *arXiv preprint arXiv:1710.03957*.
- Yingzhen Li, Richard E Turner, and Qiang Liu. 2017d. Approximate inference with amortised mcmc. *arXiv preprint arXiv:1702.08343*.
- Chia-Wei Liu, Ryan Lowe, Iulian V Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. *arXiv preprint arXiv:1603.08023*.
- Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. 2015. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*.
- Mitchell Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of english: The penn treebank. *Computational Linguistics*.
- Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. 2017. Adversarial variational bayes: Unifying variational autoencoders and generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2391–2400. JMLR. org.
- Yishu Miao, Lei Yu, and Phil Blunsom. 2016. Neural variational inference for text processing. In *International conference on machine learning*, pages 1727–1736.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Yuchen Pu, Zhe Gan, Ricardo Henao, Chunyu Li, Shaobo Han, and Lawrence Carin. 2017a. Vae learning via stein variational gradient descent. In *Advances in Neural Information Processing Systems*, pages 4236–4245.
- Yuchen Pu, Weiyao Wang, Ricardo Henao, Liqun Chen, Zhe Gan, Chunyu Li, and Lawrence Carin. 2017b. Adversarial symmetric variational autoencoder. In *Advances in Neural Information Processing Systems*, pages 4330–4339.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pages 1278–1286.
- R Tyrrell Rockafellar et al. 1966. Extension of fenchel’ duality theorem for convex functions. *Duke mathematical journal*, 33(1):81–89.
- Harshil Shah and David Barber. 2018. **Generative neural machine translation**. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 1346–1355. Curran Associates, Inc.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. In *Advances in neural information processing systems*, pages 6830–6841.
- Zichao Yang, Zhiting Hu, Ruslan Salakhutdinov, and Taylor Berg-Kirkpatrick. 2017. Improved variational autoencoders for text modeling using dilated convolutions. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3881–3890. JMLR. org.
- Shengjia Zhao, Jiaming Song, and Stefano Ermon. 2017a. InfoVAE: Information maximizing variational autoencoders. *arXiv preprint arXiv:1706.02262*.
- Tiancheng Zhao, Kyusong Lee, and Maxine Eskenazi. 2018. Unsupervised discrete sentence representation learning for interpretable neural dialog generation. *arXiv preprint arXiv:1804.08069*.

Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi.  
2017b. Learning discourse-level diversity for neural  
dialog models using conditional variational autoen-  
coders. *arXiv preprint arXiv:1703.10960*.

# Appendix

## A Decomposition of the regularization term

We have the following decomposition

$$\begin{aligned}
& -\text{KL}(q_\phi(z) \parallel p(z)) \\
& = -\mathbb{E}_{z \sim q_\phi(z)} [\log q_\phi(z) - \log p(z)] \\
& = H(q_\phi(z)) + \mathbb{E}_{z \sim q_\phi(z)} \log p(z) \\
& = H(q_\phi(z)) + \mathbb{E}_x \mathbb{E}_{z \sim q_\phi(z|x)} \log q_\phi(z|x) \\
& \quad - \mathbb{E}_x \mathbb{E}_{z \sim q_\phi(z|x)} \log q_\phi(z|x) + \mathbb{E}_{z \sim q_\phi(z)} \log p(z) \\
& = H(q_\phi(z)) - H(q_\phi(z|x)) \\
& \quad - \mathbb{E}_x \mathbb{E}_{z \sim q_\phi(z|x)} \log q_\phi(z|x) + \mathbb{E}_{z \sim q_\phi(z)} \log p(z) \\
& = I(x, z) - \mathbb{E}_x \mathbb{E}_{z \sim q_\phi(z|x)} \log q_\phi(z|x) \\
& \quad + \mathbb{E}_x \mathbb{E}_{z \sim q_\phi(z|x)} \log p(z) \\
& = I(x, z) - \mathbb{E}_x \text{KL}(q_\phi(z|x) \parallel p(z))
\end{aligned}$$

where  $q_\phi(z) = \int q(x)q_\phi(z|x)dx = \mathbb{E}_x [q_\phi(z|x)]$ ,  $q(x)$  is the empirical data distribution for the training dataset  $\mathcal{D}$ ;  $H(\cdot)$  is the Shannon entropy and  $I(x, z) = H(z) - H(z|x)$  is the mutual information between  $z$  and  $x$  under the joint distribution  $q_\phi(x, z) = q(x)q_\phi(z|x)$ .

## B Training scheme of iVAE<sub>MI</sub>

The full training scheme of iVAE<sub>MI</sub> is as following after substituting the auxiliary network of implicit VAE:

- Sample mini batch  $x_i \sim \mathcal{D}$ ,  $\epsilon_i \sim q(\epsilon)$ ,  $z_i \sim p(z)$  and generate  $z_{x_i, \epsilon_i} = G(x_i, \epsilon_i; \phi)$ ;
- Update  $\nu_\psi(z)$  net to maximize  $\sum_i \nu_\psi(z_{x_i, \epsilon_i}) - \sum_i \exp(\nu_\psi(z_i))$  according to Eq. (14);
- Update parameters  $\phi$ ,  $\theta$  to maximize  $\sum_i \log p_\theta(x_i | z_{x_i, \epsilon_i}) - \sum_i \nu_\psi(z_{x_i, \epsilon_i})$  according to Eq. (14).

## C Experimental details

### C.1 Language modeling

#### C.1.1 Data details

The details of the three datasets are shown below. Each datasets contain “train”, “validate” and “test” text files with each sentence in one line.

### C.1.2 More training details

We basically follow the architecture specifications in (Kim et al., 2018). They have an encoder predicting posterior mean and variance and we have an implicit encoder which predicts code samples directly from LSTM encoding and Gaussian noise.

We implement both encoder and decoder as one-layer LSTMs with 1024 hidden units and 512-dimensional word embeddings (in PTB, these numbers are 512 and 256). The  $\nu$  network is a fully connected network whose hidden units are 512-512-512-1. The vocabulary size is 20K for Yahoo and Yelp and 10K for PTB. The last hidden state of the encoder concatenated with a 32-dimensional Gaussian noise is used to sample 32-dimensional latent codes, which is then adopted to predict the initial hidden state of the decoder LSTM and additionally fed as input to each step at the LSTM decoder. A KL-cost annealing strategy is commonly used (Bowman et al., 2015), where the scalar weight on the KL term is increased linearly from 0.1 to 1.0 each batch over 10 epochs. There are dropout layers with probability 0.5 between the input-to-hidden layer and the hidden-to-output layer on the decoder LSTM only.

All models are trained with Adam optimizer, when encoder and decoder LSTM use learning rate 1e-3 or 8e-4, and  $\nu$  network uses learning rate 1e-5 and usually update 5 times versus a single LSTM end-to-end update. We train for 30~40 epochs, which is enough for convergence for all models. Model parameters are initialized over  $U(-0.01, 0.01)$  and gradients are clipped at 5.0.

SA-VAE runs with 10 refinement steps.

## C.2 Unsupervised style transfer

### C.2.1 Processing details

The pre-processing allows sentiment analysis on sentence level with feasible sentiment, ending up with shorter sentences with each at most 15 words than those in language modeling. Finally, we get two sets of unaligned reviews: 250K negative sentences, and 350K positive ones. They are then divided in ratio 7:1:2 to train, validate, and test data.

### C.2.2 More training details

The vocabulary size is 10K after replacing words occurring less than 5 times with the “<unk>” token. The encoder and decoder are all an one-layer LSTM with 128 hidden units. The word embed-



ding size is 128 with latent content code dimension 32. When LSTMs are trained with SGD optimizer with 1.0 learning rate, other fully connected neural nets are all trained with Adam optimizer with initial learning rate 1e-4. A grad clipping on the LSTMs is performed before each update, with max grad norm set to 1. The prior net,  $\nu$  auxiliary net, and style classifier are all implemented with 2 hidden layer (not include output layer) with 128-128 hidden units. The  $\nu$  network is trained 5 iterations versus other networks with single iteration with each mini-batch for performing accurate regularization. The “fasttext” python library (Joulin et al., 2017) is needed to evaluate accuracy, and “KenLM Language Model Toolkit” (Heafield et al., 2013) is needed to evaluate reverse PPL.

### C.3 Dialog response generation

#### C.3.1 More training details

Both encoder and decoder are implemented as GRUs. The utterance encoder is a bidirectional GRU with 300 hidden units in each direction. The context encoder and decoder are both GRUs with 300 hidden units. The prior and the recognition networks (implicit mappings) are both 2-layer feed-forward networks of size 200 with tanh non-linearity to sample Gaussian noise plus 3-layer feed-forward networks to implicitly generate samples with ReLU non-linearity and hidden sizes of 200, 200 and 400, respectively. The dimension of a latent variable  $z$  is set to 200. The initial weights for all fully connected layers are sampled from a uniform distribution  $[-0.02, 0.02]$ . We set the vocabulary size to 10,000 and define all the out-of-vocabulary words to a special token “<unk>”. The word embedding size is 200 and initialized with pre-trained Glove vectors (Pennington et al., 2014) on Twitter. The size of context window is set to 10 with a maximum utterance length of 40. We sample responses with greedy decoding so that the randomness entirely come from the latent variables. The GRUs are trained by SGD optimizer with an initial learning rate of 1.0, decay rate 0.4 and gradient clipping at 1. The fully connected neural nets are trained by RMSprop optimizer with fixed learning rates of 5e-5, except the  $\nu$  net is trained with learning rate 1e-5 and repeatedly 5 inner iterations.