

Stochastic Wasserstein Autoencoder for Probabilistic Sentence Generation

Hareesh Bahuleyan,[†] Lili Mou,[†] Hao Zhou,[‡] Olga Vechtomova[†]

[†]University of Waterloo [‡]ByteDance AI Lab

hpallika@uwaterloo.ca doublepower.mou@gmail.com
zhouhao.nlp@bytedance.com ovechtomova@uwaterloo.ca

Abstract

The variational autoencoder (VAE) imposes a probabilistic distribution (typically Gaussian) on the latent space and penalizes the Kullback–Leibler (KL) divergence between the posterior and prior. In NLP, VAEs are extremely difficult to train due to the problem of KL collapsing to zero. One has to implement various heuristics such as KL weight annealing and word dropout in a carefully engineered manner to successfully train a VAE for text. In this paper, we propose to use the Wasserstein autoencoder (WAE) for probabilistic sentence generation, where the encoder could be either stochastic or deterministic. We show theoretically and empirically that, in the original WAE, the stochastically encoded Gaussian distribution tends to become a Dirac-delta function, and we propose a variant of WAE that encourages the stochasticity of the encoder. Experimental results show that the latent space learned by WAE exhibits properties of continuity and smoothness as in VAEs, while simultaneously achieving much higher BLEU scores for sentence reconstruction.¹

1 Introduction

Natural language sentence generation in the deep learning regime typically uses a recurrent neural network (RNN) to predict the most probable next word given previous words (Mikolov et al., 2010). Such RNN architecture can be further conditioned on some source information, for example, an input sentence, resulting in a sequence-to-sequence (Seq2Seq) model.

Traditionally, sentence generation is accomplished in a *deterministic* fashion, i.e., the model uses a deterministic neural network to encode an

input sentence to some hidden representations, from which it then decodes an output sentence using another deterministic neural network.

Bowman et al. (2016) propose to use the variational autoencoder (VAE, Kingma and Welling, 2014) to map an input sentence to a probabilistic continuous latent space. VAE makes it possible to generate sentences from a distribution, which is desired in various applications. For example, in an open-domain dialog system, the information of an utterance and its response is not necessarily a one-to-one mapping, and multiple plausible responses could be suitable for a given input. Probabilistic sentence generation makes the dialog system more diversified and more meaningful (Serban et al., 2017; Bahuleyan et al., 2018). Besides, probabilistic modeling of the hidden representations serves as a way of posterior regularization (Zhang et al., 2016), facilitating interpolation (Bowman et al., 2016) and manipulation of the latent representation (Hu et al., 2017).

However, training VAEs in NLP is more difficult than the image domain (Kingma and Welling, 2014). The VAE training involves a reconstruction loss and a Kullback–Leibler (KL) divergence between the posterior and prior of the latent space. In NLP, the KL term tends to vanish to zero during training, leading to an ineffective latent space. Previous work has proposed various engineering tricks to alleviate this problem, including KL annealing and word dropout (Bowman et al., 2016).

In this paper, we address the difficulty of training VAE sentence generators by using a Wasserstein autoencoder (WAE, Tolstikhin et al., 2018). WAE modifies VAE in that it requires the integration of the posterior to be close to its prior, where the closeness is measured with empirical samples drawn from the distributions. In this way, the encoder could be either stochastic or deterministic, but the model still retains probabilistic properties.

¹Our code is available at https://github.com/HareeshBahuleyan/probabilistic_nlg
A preliminary version of this paper was preprinted at <https://arxiv.org/abs/1806.08462>




Figure 1: The latent space of the deterministic autoencoder (DAE), variational autoencoder (VAE), Wasserstein autoencoder (WAE), as well as WAE with by our KL penalty. Blue circles: Posterior or aggregated posterior distributions of data in the latent space. Red circles: Regularizations of the posterior.

Moreover, we show both theoretically and empirically that the stochastic Gaussian encoder in the original form tends to be a Dirac-delta function. We thus propose a WAE variant that encourages the encoder’s stochasticity by penalizing an auxiliary KL term.

Experiments show that the sentences generated by WAE exhibit properties of continuity and smoothness as in VAE, while achieving a much higher reconstruction performance. Our proposed variant further encourages the stochasticity of the encoder. More importantly, WAE is robust to hyperparameters and much easier to train, without the need for KL annealing or word dropout as in VAE. In a dialog system, we demonstrate that WAEs are capable of generating better quality and more diverse sentences than VAE.

2 Probabilistic Sentence Generation

Base Model: Deterministic Autoencoder (DAE). DAE encodes an input sentence with a recurrent neural network (RNN) and then decodes the same sentence through another RNN.

For the encoder, the hidden state of the last word is represented as the latent space of the input sentence \mathbf{x} . The latent representation is denoted as \mathbf{z} . We feed \mathbf{z} to the decoder RNN, which predicts one word at a time using a softmax layer, given by $p(x_t|\mathbf{z}, \mathbf{x}_{<t})$. The training objective for DAE is the sequence-aggregated cross-entropy loss, given by

$$J = - \sum_{n=1}^N \sum_{t=1}^{|\mathbf{x}^{(n)}|} \log p(x_t^{(n)}|\mathbf{z}^{(n)}, \mathbf{x}_{<t}^{(n)}) \quad (1)$$

where superscript (n) indicates the n th data point among $1, \dots, N$.

In DAE, the latent space is encoded and then decoded in a deterministic way, i.e., there is no probabilistic modeling of the hidden space. The hidden

representations of data may be located on an arbitrary manifold (Figure 1a), which is not suitable for probabilistic generation.

Variational Autoencoder (VAE). VAE extends DAE by imposing a prior distribution $p(\mathbf{z})$ on the latent variable \mathbf{z} , which is typically set to the standard normal $\mathcal{N}(\mathbf{0}, \mathbf{I})$ (Kingma and Welling, 2014). Given an input sentence \mathbf{x} , we would like to model the posterior of \mathbf{z} by another normal distribution, $q(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_{\text{post}}, \text{diag } \boldsymbol{\sigma}_{\text{post}}^2)$, where $\boldsymbol{\mu}_{\text{post}}$ and $\boldsymbol{\sigma}_{\text{post}}^2$ are the outputs of the encoder.

In the training of VAE, \mathbf{z} is sampled from $q(\mathbf{z}|\mathbf{x})$, and the training objective is to maximize a variational lower bound of the likelihood of data. This is equivalent to minimizing the (expected) reconstruction loss similar to (1), while being regularized by the KL divergence between $q(\mathbf{z}|\mathbf{x})$ and $p(\mathbf{z})$, given by

$$J = \sum_{n=1}^N \left[- \mathbb{E}_{\mathbf{z}^{(n)} \sim q} \sum_{t=1}^{|\mathbf{x}^{(n)}|} \log p(x_t^{(n)}|\mathbf{z}^{(n)}, \mathbf{x}_{<t}^{(n)}) + \lambda_{\text{VAE}} \cdot \text{KL}(q(\mathbf{z}^{(n)}|\mathbf{x}^{(n)})||p(\mathbf{z})) \right] \quad (2)$$

where in the expectation $\mathbf{z}^{(n)}$ is sampled from $q(\mathbf{z}|\mathbf{x}^{(n)})$ and λ_{VAE} is a hyperparameter balancing the two terms.

Since VAE penalizes the divergence of \mathbf{z} ’s posterior from its prior, it serves as a way of posterior regularization, making it possible to generate sentences from the continuous latent space.

However, the two objectives in (2) are contradictory to each other, as argued by Tolstikhin et al. (2018). VAE pushes the posterior of \mathbf{z} , given any input $\mathbf{x}^{(n)}$, to be close to its prior, i.e., every blue ellipse in Figure 1b should be close to the red one. This makes perfect reconstruction impossible.

Further, VAE is difficult to train in NLP due to the problem of KL collapse, where the KL term tends to be zero, meaning that the encoder captures no information and the decoder learns an unconditioned language model. This phenomenon is observed in variational auto-regressive decoders using RNN. To alleviate this problem, existing tricks include KL annealing and word dropout (Bowman et al., 2016), but both require extensive engineering.

Wasserstein Autoencoder (WAE). An alternative way of posterior regularization is to impose a constraint that the aggregated posterior of \mathbf{z} should be the same as its prior (Tolstikhin et al.,

2018), i.e., $q(z) \stackrel{\text{def}}{=} \sum_{\mathbf{x}} q(z|\mathbf{x})p_{\mathcal{D}}(\mathbf{x}) \stackrel{\text{set}}{=} p(z)$, where $p_{\mathcal{D}}$ is the data distribution. This is also demonstrated in Figure 1c. By contrast, VAE requires that $q(z|\mathbf{x})$ should be close to $p(z)$ for every input sentence \mathbf{x} .

For computational purposes, Tolstikhin et al. (2018) relax the above constraint by penalizing the Wasserstein distance between $q(z)$ and $p(z)$. In particular, it is computed by the Maximum Mean Discrepancy (MMD), defined as

$$\text{MMD} = \left\| \int k(z, \cdot) dP(z) - \int k(z, \cdot) dQ(z) \right\|_{\mathcal{H}_k}$$

where $P(z)$ and $Q(z)$ are cumulative density functions. \mathcal{H}_k refers to the reproducing kernel Hilbert space defined by the kernel k , which is often chosen as the inverse multiquadratic kernel $k(x, y) = \frac{C}{C + \|x - y\|_2^2}$ for high-dimensional Gaussians.

One advantage of the Wasserstein distance is that it can be estimated by empirical samples as

$$\begin{aligned} \widehat{\text{MMD}} &= \frac{1}{N(N-1)} \sum_{n \neq m} k(\mathbf{z}^{(n)}, \mathbf{z}^{(m)}) \\ &+ \frac{1}{N(N-1)} \sum_{n \neq m} k(\tilde{\mathbf{z}}^{(n)}, \tilde{\mathbf{z}}^{(m)}) - \frac{1}{N^2} \sum_{n, m} k(\mathbf{z}^{(n)}, \tilde{\mathbf{z}}^{(m)}) \end{aligned} \quad (3)$$

where $\tilde{\mathbf{z}}^{(n)}$ is a sample from the prior $p(z)$, and $\mathbf{z}^{(n)}$ is a sample from the aggregated posterior $q(z)$, which is obtained by sampling $\mathbf{x}^{(n)}$ from the data distribution and then sampling $\mathbf{z}^{(n)}$ from $q(z|\mathbf{x}^{(n)})$. In summary, the training objective of WAE is

$$J_{\text{WAE}} = - \sum_{n=1}^N \sum_{t=1}^{|\mathbf{x}^{(n)}|} \log p(\mathbf{x}_t^{(n)} | \mathbf{z}^{(n)}, \mathbf{x}_{<t}^{(n)}) + \lambda_{\text{WAE}} \cdot \widehat{\text{MMD}} \quad (4)$$

where λ_{WAE} balances the MMD penalty and the reconstruction loss.

Alternatively, the dual form (adversarial loss) can also be used for WAE (Zhao et al., 2018). In our preliminary experiments, we found MMD similar to but slightly better than the adversarial loss. The difference between our work and Zhao et al. (2018)—who extend the original WAE to sequence generation—is that we address the KL annealing problem of VAE and further analyze the stochasticity of WAE from a theoretical perspective, as follows.

WAE with Auxiliary Loss. In WAE, the aggregated posterior $q(z)$ involves an integration of data distribution, which allows using a deterministic function to encode \mathbf{z} as $\mathbf{z} = \mathbf{f}_{\text{encode}}(\mathbf{x})$ as suggested by Tolstikhin et al. (2018). This would largely alleviate the training difficulties as in VAE, because backpropagating gradient into the encoder no longer involves a stochastic layer.

The stochasticity of the encoder, however, is still a desired property in some applications, for example, generating diverse responses in a dialog system. We show both theoretically and empirically that a *dangling* Gaussian stochastic encoder could possibly degrade to a deterministic one.

Theorem 1. *Suppose we have a Gaussian family $\mathcal{N}(\boldsymbol{\mu}, \text{diag } \boldsymbol{\sigma}^2)$, where $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are parameters. The covariance is diagonal, meaning that the variables are independent. If the gradient of $\boldsymbol{\sigma}$ completely comes from sample gradient and $\boldsymbol{\sigma}$ is small at the beginning of training, then the Gaussian converges to a Dirac delta function with stochastic gradient descent, i.e., $\boldsymbol{\sigma} \rightarrow \mathbf{0}$. (See Appendix A for the proof.)* \square

To alleviate this problem, we propose a simple heuristic that encourages the stochasticity of the encoder. In particular, we penalize, for every data point, a KL term between the predicted posterior $q(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_{\text{post}}, \text{diag } \boldsymbol{\sigma}_{\text{post}}^2)$ and a Gaussian with covariance \mathbf{I} centered at the predicted mean, i.e., $\mathcal{N}(\boldsymbol{\mu}_{\text{post}}, \mathbf{I})$. This is shown in Figure 1d, where each posterior is encouraged to stretch with covariance \mathbf{I} . Formally, the loss is

$$\begin{aligned} J &= J_{\text{rec}} + \lambda_{\text{WAE}} \cdot \widehat{\text{MMD}} \\ &+ \lambda_{\text{KL}} \sum_n \text{KL} \left(\mathcal{N}(\boldsymbol{\mu}_{\text{post}}^{(n)}, \text{diag}(\boldsymbol{\sigma}_{\text{post}}^{(n)})^2) \parallel \mathcal{N}(\boldsymbol{\mu}_{\text{post}}^{(n)}, \mathbf{I}) \right) \end{aligned} \quad (5)$$

While our approach appears heuristic, the next theorem shows its theoretical justification.

Theorem 2. *Objective (5) is a relaxed optimization of the WAE loss (4) with a constraint on $\boldsymbol{\sigma}_{\text{post}}$. (See Appendix B for the proof.)* \square

We will show empirically that such auxiliary loss enables us to generate smoother and more diverse sentences in WAE. It, however, does not suffer from KL collapse as in VAEs. The auxiliary KL loss that we define for stochastic WAE is computed against a target distribution $\mathcal{N}(\boldsymbol{\mu}_{\text{post}}^{(n)}, \mathbf{I})$ for each data sample $\mathbf{x}^{(n)}$. Here, the predicted posterior mean itself is used in the target distribution.

	BLEU [†]	PPL [↓]	UniKL [↓]	Entropy	AvgLen
Corpus	-	-	-	→ 5.65	→ 9.6
DAE	86.35	146.2	0.178	6.23	11.0
VAE (KL-annealed)	43.18	79.4	0.081	5.04	8.8
WAE-D $\lambda_{\text{WAE}} = 3$	86.03	113.8	0.071	5.59	10.0
WAE-D $\lambda_{\text{WAE}} = 10$	84.29	104.9	0.073	5.57	9.9
WAE-S $\lambda_{\text{KL}} = 0.0$	75.66	115.2	0.069	5.61	9.9
WAE-S $\lambda_{\text{KL}} = 0.01$	82.01	84.9	0.058	5.26	9.4
WAE-S $\lambda_{\text{KL}} = 0.1$	47.63	62.5	0.150	4.65	8.7

Table 1: Results of SNLI-style sentence generation, where WAE is compared with DAE and VAE. **D** and **S** refer to the deterministic and stochastic encoders, respectively. ^{†/↓}The larger/lower, the better. For **Entropy** and **AvgLen**, the closer to corpus statistics, the better (indicated by the → arrow).

As a result, this KL term does not force the model to learn the same posterior for all data samples (as in VAE), and thus, the decoder does not degrade to an unconditioned language model.

3 Experiments

We evaluate WAE in sentence generation on the Stanford Natural Language Inference (SNLI) dataset (Bowman et al., 2015) as well as dialog response generation. All models use single-layer RNN with long short term memory (LSTM) units for both the encoder and decoder. Appendix C details our experimental settings.

VAE training. VAE is notoriously difficult to train in the RNN setting. While different researchers have their own practice of training VAE, we follow our previous experience (Bahuleyan et al., 2018) and adopt the following tricks to stabilize the training: (1) λ_{VAE} was annealed in a sigmoid manner. We monitored the value of $\lambda \cdot \text{KL}$ and stop annealing once it reached its peak value, known as *peaking annealing*. (2) For word dropout, we started with no dropout, and gradually increased the dropout rate by 0.05 every epoch until it reached a value of 0.5. The effect of KL annealing is further analyzed in Appendix D.

3.1 SNLI Generation

The SNLI sentences are written by crowd-sourcing human workers in an image captioning task. It is a massive corpus but with comparatively simple sentences (examples shown in Table 4). This task could be thought of as domain-specific sentence generation, analogous to hand written digit generation in computer vision.

In Table 1, we compare all methods in two aspects. (1) We evaluate by **BLEU** (Papineni et al.,

2002) how an autoencoder preserves input information in a reconstruction task. (2) We also evaluate the quality of probabilistic sentence generation from the latent space. Although there is no probabilistic modeling of the latent space in DAE, we nevertheless draw samples from $\mathcal{N}(\mathbf{0}, \mathbf{I})$, which could serve as a non-informative prior. Perplexity (**PPL**) evaluates how fluent the generated sentences are. This is given by a third-party n -gram language model trained on the Wikipedia dataset. The unigram-KL (**UniKL**) evaluates if the word distribution of the generated sentences is close to that of the training corpus. Other surface metrics (entropy of the word distribution and average sentence length) also measure the similarity of the latent space generated sentence set to that of the corpus.

We see that DAE achieves the best BLEU score, which is not surprising because DAE directly optimizes the maximum likelihood of data as a surrogate of word prediction accuracy. Consequently, DAE performs poorly for probabilistic sentence generation as indicated by the other metrics.

VAE and WAE have additional penalties that depart from the goal of reconstruction. However, we see that WAEs, when trained with appropriate hyperparameters ($\lambda_{\text{WAE}}, \lambda_{\text{KL}}$), achieve close performance to DAE, outperforming VAE by 40 BLEU points. This is because VAE encodes each input’s posterior to be close to the prior, from which it is impossible to perfectly reconstruct the data.

Comparing the deterministic and stochastic encoders in WAE, we observe the same trade-off between reconstruction and sampling. However, our proposed stochastic encoder, with $\lambda_{\text{KL}} = 0.1$ for WAE, consistently outperforms VAE in the contradictory metrics BLEU and PPL. The hyperparameters $\lambda_{\text{WAE}} = 10.0$ and $\lambda_{\text{KL}} = 0.01$ appear to have the best balance between reconstruction, sentence fluency, as well as similarity to the original corpus.

Moreover, all our WAEs are trained without annealing or word dropout. It is significantly simpler than training a VAE, whose KL annealing typically involves a number of engineering tricks, such as the time step when KL is included, the slope of annealing, and the stopping criterion for annealing.

	BLEU-2	BLEU-4	Entropy	Dist-1	Dist-2
Test Set	-	-	6.15	0.077	0.414
DED	3.96	0.85	5.55	0.044	0.275
VED	3.26	0.59	5.45	0.053	0.204
WED-D	4.05	0.98	5.53	0.042	0.272
WED-S	3.72	0.69	5.59	0.066	0.309

Table 2: Results on dialog generation, where VED/WED hyperparameters for each model were chosen by Table 1.

3.2 Dialog Generation

We extend WAE to an encoder-decoder framework (denoted by WED) and evaluate it on the DailyDialog corpus (Li et al., 2017).² We follow Bahuleyan et al. (2018), using the encoder to capture an utterance and the decoder to generate a reply.

Table 2 shows that WED with a deterministic encoder (WED-D) is better than the variational encoder-decoder (VED) in BLEU scores, but the generated sentences lack variety, which is measured by output entropy and the percentage of distinct unigrams and bigrams (Dist-1/Dist-2, Li et al., 2016), evaluated on the generated test set responses.

We then applied our stochastic encoder for WED and see that, equipped with our KL-penalized stochastic encoder, WED-S outperforms DED, VED, and WED-D in all diversity measures. WED-S also outperforms VED in generation quality, consistent with the results in Table 1.

4 Conclusion

In this paper, we address the difficulty of training VAE by using a Wasserstein autoencoder (WAE) for probabilistic sentence generation. WAE implementation can be carried out with either a deterministic encoder or a stochastic one. The deterministic version achieves high reconstruction performance, but lacks diversity for generation. The stochastic encoder in the original form may collapse to a Dirac delta function, shown by both a theorem and empirical results. We thus propose to encourage stochasticity by penalizing a heuristic

²In our pilot experiment, we obtained a BLEU-4 score of 6 by training a pure Seq2Seq model with LSTM units for 200 epochs, whereas Li et al. (2017) report 0.009 BLEU-4 and Luo et al. (2018) report 2.84 BLEU-4. Due to our unreasonably high performance, we investigated this in depth and found that the training and test sets of the DailyDialog corpus have overlaps. For the results reported in our paper, we have removed duplicate data in the test set, which is also available on our website (Footnote 1). To the best of our knowledge, we are the first to figure out the problem, which, unfortunately, makes comparison with previous work impossible.

KL loss for WAE, which turns out to be a relaxed optimization of the Wasserstein distance with a constraint on the posterior family.

We evaluated our model on both SNLI sentence generation and dialog systems. We see that WAE achieves high reconstruction performance as DAE, while retaining the probabilistic property as VAE. Our KL-penalty further improves the stochasticity of WAE, as we achieve the highest performance in all diversity measures.

Acknowledgments

We would like to acknowledge Yiping Song and Zhiliang Tian for their independent investigation on the DailyDialog corpus. We also thank Yanran Li, one of the authors who released DailyDialog, for discussion on this issue. This work was supported in part by the NSERC grant RGPIN-261439-2013 and an Amazon Research Award.

References

- Hareesh Bahuleyan, Lili Mou, Olga Vechtomova, and Pascal Poupart. 2018. [Variational attention for sequence-to-sequence models](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1672–1682.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. 2016. [Generating sentences from a continuous space](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21.
- Christopher P Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. 2017. [Understanding disentangling in beta-VAE](#). In *Proceedings of the Workshop on Learning Disentangled Representations: From Perception to Control*.
- 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 *Proceedings of the International Conference on Learning Representations*.
- 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, pages 1587–1596.
- Diederik P Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *Proceedings of the International Conference on Learning Representations*.
- Diederik P Kingma and Max Welling. 2014. [Auto-encoding variational Bayes](#). In *Proceedings of International Conference on Learning Representations*, pages 10–21.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [DailyDialog: a manually labelled multi-turn dialogue dataset](#). In *Proceedings of the 8th International Joint Conference on Natural Language Processing*, pages 986–995.
- Liangchen Luo, Jingjing Xu, Junyang Lin, Qi Zeng, and Xu Sun. 2018. [An auto-encoder matching model for learning utterance-level semantic dependency in dialogue generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 702–707.
- Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. 2010. [Recurrent neural network based language model](#). In *Eleventh Annual Conference of the International Speech Communication Association*, pages 1045–1048.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318.
- Iulian Vlad Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron C Courville, and Yoshua Bengio. 2017. [A hierarchical latent variable encoder-decoder model for generating dialogues](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 3295–3301.
- Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Scholkopf. 2018. [Wasserstein auto-encoders](#). In *Proceedings of International Conference on Learning Representations*.
- Jiacheng Xu and Greg Durrett. 2018. [Spherical latent spaces for stable variational autoencoders](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4503–4513.
- Zichao Yang, Zhiting Hu, Ruslan Salakhutdinov, and Taylor Berg-Kirkpatrick. 2017. [Improved variational autoencoders for text modeling using dilated convolutions](#). In *Proceedings of the International Conference on Machine Learning*, pages 3881–3890.
- Biao Zhang, Deyi Xiong, jinsong su, Hong Duan, and Min Zhang. 2016. [Variational neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 521–530.
- Junbo Zhao, Yoon Kim, Kelly Zhang, Alexander Rush, and Yann LeCun. 2018. [Adversarially regularized autoencoders](#). In *Proceedings of the 35th International Conference on Machine Learning*, pages 5902–5911.

A Proof of Theorem 1

Theorem 1. *Suppose we have a Gaussian family $\mathcal{N}(\boldsymbol{\mu}, \text{diag } \boldsymbol{\sigma}^2)$, where $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are parameters. The covariance is diagonal, meaning that the variables are independent. If the gradient of $\boldsymbol{\sigma}$ completely comes from sample gradient and $\boldsymbol{\sigma}$ is small at the beginning of training, then the Gaussian converges to a Dirac delta function with stochastic gradient descent, i.e., $\boldsymbol{\sigma} \rightarrow \mathbf{0}$.*

Proof. For the predicted posterior $\mathcal{N}(\boldsymbol{\mu}, \text{diag } \boldsymbol{\sigma}^2)$ where all dimensions are independent, we consider a certain dimension, where the sample is $z_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$.

We denote the gradient of J wrt to z_i at $z_i = \mu_i$ by $g_i \triangleq \frac{\partial J}{\partial z_i} \Big|_{z_i = \mu_i}$. At a particular sample $z_i^{(j)}$ around μ_i , the gradient $g_i^{(j)}$ is

$$g_i^{(j)} \triangleq \frac{\partial J}{\partial z_i} \Big|_{z_i = z_i^{(j)}} \quad (6)$$

$$\approx g_i + \frac{\partial^2 J}{\partial \mu_i^2} (z_i^{(j)} - \mu_i) \quad (7)$$

$$\triangleq g_i + k(z_i^{(j)} - \mu_i) \quad (8)$$

$$= g_i + k(\mu_i + \epsilon^{(j)} \sigma_i - \mu_i) \quad (9)$$

$$= g_i + k\sigma_i \epsilon^{(j)} \quad (10)$$

where (7) is due to Taylor series approximation, if we assume σ_i^2 is small and thus $z_i^{(j)}$ is near μ_i . k denotes $\frac{\partial^2 J}{\partial \mu_i^2}$.

We compute the expected gradient wrt to σ_i for $\epsilon \sim \mathcal{N}(0, 1)$. The assumption of this theorem is that the gradient of μ_i and σ_i completely comes from the sample z_i . By the chain rule, we have

$$\mathbb{E}_{\epsilon^{(j)} \sim \mathcal{N}(0,1)} \left[\frac{\partial J}{\partial \sigma_i} \right] \quad (11)$$

$$= \mathbb{E}_{\epsilon^{(j)} \sim \mathcal{N}(0,1)} \left[\frac{\partial J}{\partial z_i^{(j)}} \cdot \frac{\partial z_i^{(j)}}{\partial \sigma_i} \right] \quad (12)$$

$$= \mathbb{E}_{\epsilon^{(j)} \sim \mathcal{N}(0,1)} \left[(g_i + k\sigma_i \epsilon^{(j)}) \cdot \epsilon^{(j)} \right] \quad (13)$$

$$= \mathbb{E}_{\epsilon^{(j)} \sim \mathcal{N}(0,1)} \left[g_i \epsilon_i^{(j)} + k\sigma_i (\epsilon^{(j)})^2 \right] \quad (14)$$

$$= k\sigma_i \quad (15)$$

Notice that $k > 0$ if we are near a local optimum (locally convex).

In other words, the expected gradient of σ_i is proportional to σ_i . According to stochastic gradient descent (SGD), σ_i will converge to zero. \square

The theorem assumes σ^2 is small, compared with how J changes in the latent space. In practice, the encoded vectors of different samples may vary a lot, whereas if we sample different vectors from a certain predicted multi-variate Gaussian, we would generally obtain the same sentence. Therefore, J is kind of smooth in the latent space. The phenomenon can also be verified empirically by plotting the histogram of σ in WAE with a stochastic Gaussian encoder (Figure 2). We see that if the KL coefficient λ_{KL} is 0, meaning that the gradient of σ comes only from the samples, then most σ 's collapse to 0.

Notice, however, that the theorem does not suggest a stochastic WAE and a deterministic WAE will yield exactly the same result, as their trajectories may be different.


B Proof of Theorem 2

Theorem 2. *Objective (5) is a relaxed optimization of the WAE loss (4) with a constraint on σ_{post} .*

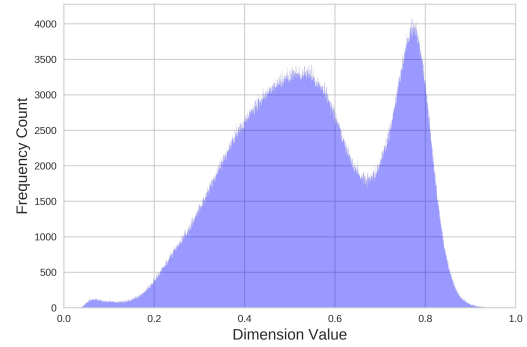
Proof. Objective (5) optimizes

$$J = J_{\text{rec}} + \lambda_{\text{WAE}} \cdot \widehat{\text{MMD}} + \lambda_{\text{KL}} \sum_n \text{KL} \left(\mathcal{N}(\boldsymbol{\mu}_{\text{post}}^{(n)}, \text{diag}(\boldsymbol{\sigma}_{\text{post}}^{(n)2}) \parallel \mathcal{N}(\boldsymbol{\mu}_{\text{post}}^{(n)}, \mathbf{I}) \right)$$


The first two terms are the WAE loss, whereas the last penalty relaxes the following optimization problem



(a) $\lambda_{\text{KL}} = 0$



(b) $\lambda_{\text{KL}} = 0.01$



(c) $\lambda_{\text{KL}} = 0.1$

Figure 2: The histograms of σ in the posterior of the WAE in the SNLI experiment. In the plot, there are 200 buckets in the range of (0, 1).

$$\text{minimize } J_{\text{rec}} + \lambda_{\text{WAE}} \cdot \widehat{\text{MMD}}$$

subject to

$$\sum_n \text{KL} \left(\mathcal{N}(\boldsymbol{\mu}_{\text{post}}^{(n)}, \text{diag}(\boldsymbol{\sigma}_{\text{post}}^{(n)2}) \parallel \mathcal{N}(\boldsymbol{\mu}_{\text{post}}^{(n)}, \mathbf{I}) \right) < C \quad (16)$$

for some constant C .

As known, the KL divergence between two

SNLI Experiment

LSTM Hidden Dimension	100d, single layer
Word Embeddings	300d, pretrained on SNLI Corpus
Latent Dimension	100d
Epochs	20
Learning Rate	Fixed rate of 0.001
Batch Size	128
Max Sequence Length	20
Vocab Size	30000

Dialog Experiment

LSTM Hidden Dimension	500d, single layer
Word Embeddings	300d, pretrained on Dialog Corpus
Latent Dimension	300d
Epochs	200
Learning Rate	Initial rate of 0.001, multiplicative decay of 0.98 until a minimum of 0.00001
Batch Size	128
Max Sequence Length	20
Vocab Size	20000

Table 3: Experimental settings.

(univariate) Gaussian distributions is

$$\begin{aligned} & \text{KL}(\mathcal{N}(\mu_1, \sigma_1) \parallel \mathcal{N}(\mu_2, \sigma_2)) \\ &= \log \frac{\sigma_2}{\sigma_1} + \frac{\sigma_1^2 + (\mu_1 - \mu_2)^2}{2\sigma_2^2} - \frac{1}{2} \end{aligned} \quad (17)$$

The constraint in (16) is equivalent to

$$\sum_n \sum_i \left[-\log \sigma_i^{(n)} + \frac{1}{2} (\sigma_i^{(n)})^2 \right] < C \quad (18)$$

In other words, our seemingly heuristic KL penalty optimizes the Wasserstein loss, while restricting the posterior family. \square

C Implementation Details

All models were trained with the Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. In all our experiments, we feed the sampled latent vector z to each time step of the decoder. Task-specific settings are listed in Table 3.

D VAE Training Difficulties

It is a common practice that training VAEs involves KL annealing and word dropout, which further consists of hacks for tuning hyperparameters. We conducted an experiment of training VAE without KL annealing. In Figure 3, we present the KL loss (weighted by λ_{VAE}) during the training process for different values of λ_{VAE} . The KL loss is believed to be an important diagnostic measure to indicate if the latent space is “variational” (Yang et al., 2017; Higgins et al., 2017; Burgess et al., 2017). We see that if the penalty is too large, KL




Figure 3: Learning curves of the KL term in the VAE loss function ($\lambda \cdot \text{KL}$) for different values of λ , and the variant where λ is annealed.

simply collapses to zero ignoring the entire input, in the case of which, the decoder becomes an unconditioned language model. On the other hand, if the KL penalty is too small, the model tends to become more *deterministic* and the KL term does not play a role in the training. This is expected since in the limit of λ_{VAE} to 0, the model would probably ignore the KL term and becomes a deterministic autoencoder (shown also by Theorem 1).

The VAE with collapsed KL does not exhibit interesting properties such as random sampling for probabilistic sequence generation (Bowman et al., 2016). As seen in Table 4, the generated sentences by VAE without annealing are very close to each other. This is because VAE’s encoder does not capture useful information in the latent space, which is simply ignored during the decoding phase. By sampling the latent space, we do not obtain varying sentences. The empirical evidence verifies our intuition.

A recent study (Xu and Durrett, 2018) propose to get rid of KL annealing by using the von Mises–Fisher (vMF) family of posterior and prior. In particular, they set the prior to the uniform distribution on a unit hypersphere, whereas the posterior family is normal distribution on the surface of the same sphere. They fix the standard deviation (parametrized by κ) of the posterior, so that their KL is a constant and annealing is not required. This, unfortunately, loses the privilege of learning uncertainty in the probabilistic modeling. Examples in Table 4 show that, while we have reproduced the reconstruction negative log-likelihood with vMF-VAE (the metric used in their paper), the generated sentences are of poor quality. As also suggested by Xu and Durrett (2018), if the posterior uncertainty in vMF is made learnable, it re-introduces the KL collapse problem, in which

Training Samples
<i>a mother and her child are outdoors. the people are opening presents. the girls are looking toward the water. a small boy walks down a wooden path in the woods. a person in a green jacket is surfing while holding on to a line.</i>
DAE
<i>two families walking in a towel down alaska sands a cot . a blade is rolling its nose furiously paper . a woman in blue shirts is passing by a some beach transporting his child are wearing overalls . a guys are blowing on professional thinks the horse .</i>
VAE without Annealing
<i>a man is playing a guitar . a man is playing with a dog . a man is playing with a dog . a man is playing a guitar . a man is playing with a dog .</i>
VAE with Annealing
<i>the band is sitting on the main street . couple dance on stage in a crowded room . two people run alone in an empty field . the group of people have gathered in a picture . a cruise ship is docking a boat ship .</i>
VAE vMF (κ fixed)
<i>a car is a and and a blue shirt top is . two children are playing on the group in are the the . the a child and a adult and the young is playing for a picture a are playing to a little is playing a background . .</i>
WAE-D ($\lambda_{\text{WAE}} = 10$)
<i>the lone man is working . the group of men is using ice at the sunset . a family is outside in the background . two women are standing on a busy street outside a fair a tourists is having fun on a sunny day</i>
WAE-S ($\lambda_{\text{WAE}} = 10, \lambda_{\text{KL}} = 0.01$)
<i>an asian man is dancing in a highland house . a person wearing a purple snowsuit jumps over the tree . the vocalist is at the music and dancing with a microphone . a young man is dressed in a white shirt cleaning clothes . three children lie together and a woman falls in a plane .</i>

Table 4: Sentences generated by randomly sampling from the prior for different models.

case, the KL annealing is still needed.

By contrast, WAEs for sequence-to-sequence models are trained without any additional optimization strategies such as annealing. Even in our stochastic encoder, the KL penalty does not make WAE an unconditioned language model, because it does not force the encoded posterior to be the same for different input sentences.

E Qualitative Samples

Table 4 shows sentences generated by randomly sampling points in the latent space for different models, along with sample sentences from the training set. They provide a qualitative understanding of each model’s performance.