

A Surprisingly Effective Fix for Deep Latent Variable Modeling of Text

Bohan Li^{*1}, Junxian He^{*1}, Graham Neubig¹, Taylor Berg-Kirkpatrick², Yiming Yang¹

¹Language Technologies Institute, Carnegie Mellon University

²Department of Computer Science and Engineering, University of California San Diego

{bohanl1, junxianh, gneubig, yiming}@cs.cmu.edu, tberg@eng.ucsd.edu

Abstract

When trained effectively, the Variational Autoencoder (VAE) is both a powerful language model and an effective representation learning framework. In practice, however, VAEs are trained with the *evidence lower bound* (ELBO) as a surrogate objective to the intractable marginal data likelihood. This approach to training yields unstable results, frequently leading to a disastrous local optimum known as *posterior collapse*. In this paper, we investigate a simple fix for posterior collapse which yields surprisingly effective results. The combination of two known heuristics, previously considered only in isolation, substantially improves held-out likelihood, reconstruction, and latent representation learning when compared with previous state-of-the-art methods. More interestingly, while our experiments demonstrate superiority on these principle evaluations, our method obtains a *worse ELBO*. We use these results to argue that the typical surrogate objective for VAEs may not be sufficient or necessarily appropriate for balancing the goals of representation learning and data distribution modeling.¹

1 Introduction

Latent variable models attempt to model observed data \mathbf{x} given latent variables \mathbf{z} , both for purposes of modeling data distributions $p(\mathbf{x})$ (e.g. language modeling) and learning representations \mathbf{z} for a particular \mathbf{x} (e.g. sentence embedding). Variational Autoencoders (VAEs) (Kingma and Welling, 2014) are a powerful framework for learning latent variable models using neural networks. The generative model of VAEs first samples a latent vector \mathbf{z} from a prior $p(\mathbf{z})$, then applies a neural decoder $p(\mathbf{x}|\mathbf{z})$ to produce \mathbf{x} conditioned on the latent code \mathbf{z} . VAEs are trained

to maximize the evidence lower bound (ELBO) of the intractable log marginal likelihood:

$$\mathbb{E}_{\mathbf{z} \sim q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\theta}(\mathbf{x}|\mathbf{z})] - D_{\text{KL}}(q_{\phi}(\mathbf{z}|\mathbf{x})||p(\mathbf{z})),$$

where $q_{\phi}(\mathbf{z}|\mathbf{x})$ represents an approximate posterior distribution (i.e. the *encoder* or *inference network*) and $p_{\theta}(\mathbf{x}|\mathbf{z})$ is the generative distribution (i.e. the *decoder*).

However, modeling text with VAEs has proven to be challenging, and is an open research problem (Yang et al., 2017; Xu and Durrett, 2018; Kim et al., 2018; Dieng et al., 2018; He et al., 2019; Pelsmaecker and Aziz, 2019). When a strong decoder (e.g. the LSTM (Hochreiter and Schmidhuber, 1997)) is employed, training often falls into a trivial local optimum where the decoder learns to ignore the latent variable and the encoder fails to encode any information. This phenomenon is referred to as “posterior collapse” (Bowman et al., 2016). Existing efforts tackling this problem include re-weighting the KL loss (Bowman et al., 2016; Kingma et al., 2016; Liu et al., 2019), changing the model (Yang et al., 2017; Semeniuta et al., 2017; Xu and Durrett, 2018), and modifying the training procedure (He et al., 2019).

After conducting an empirical examination of the state-of-the-art methods (Section 2), we find that they have difficulty striking a good balance between language modeling and representation learning. In this paper, we present a practically effective combination of two simple heuristic techniques for improving VAE learning: (1) pretraining the inference network using an autoencoder objective and (2) thresholding the KL term in the ELBO objective (also known as “free bits” (Kingma et al., 2016)). The former technique initializes VAE training with an inference network that encodes useful information about \mathbf{x} , biasing learning away from local optima where \mathbf{x} is ignored (i.e. posterior collapse). The latter tech-

^{*}Equal contribution.

¹Code is available at <https://github.com/bohanli/vae-pretraining-encoder>.

nique modifies the ELBO objective to prevent the KL term from dominating the encoder’s role in reconstruction (Alemi et al., 2018), again biasing learning to avoid posterior collapse (Chen et al., 2017; Zhao et al., 2019).

In experiments we find that these two techniques do not perform well in isolation. However, when combined, they substantially outperform all baselines across various metrics that evaluate both language modeling and representation learning capabilities. Finally, we find that our method tends to achieve a superior language modeling results in terms of perplexity but an inferior ELBO value, and we use these results to argue that ELBO is sub-optimal for language modeling even though it provides a formal lower bound on log marginal likelihood. Thus, we suggest that future research in this direction should be careful to monitor the gap between ELBO and log marginal likelihood, and reconsider using ELBO as the surrogate, especially for evaluation.

2 Analysis of Existing Methods

In this section, we first analyze and compare state-of-the-art solutions to posterior collapse² to obtain a holistic view of current progress and remaining challenges. Based on these observations, we then propose and evaluate a method that demonstrates substantially improved performance.

2.1 Evaluation

In this paper we focus on the standard VAE setting where both the prior $p(\mathbf{z})$ and the posterior $p_{\theta}(\mathbf{z}|\mathbf{x})$ are factorized Gaussians.³ We conduct preliminary experiments on the English Penn Treebank (PTB) (Marcus and Marcinkiewicz), a standard dataset for benchmarking VAE that has been used extensively in previous work (Bowman et al., 2016; Xu and Durrett, 2018; Liu et al., 2019). We use 32-dimension latent codes (full setup details can be found in Appendix A). We evaluate using the following metrics:

Perplexity (PPL). To compute perplexity, we first estimate the log marginal likelihood with 1000 importance weighted samples (Burda et al., 2016). Note that it is inappropriate to estimate

²Here we focus on different training techniques for existing models, not alternatives which require changing the underlying model (Yang et al., 2017; Xu and Durrett, 2018).

³Pelsmaeker and Aziz (2019) thoroughly investigate using more complicated priors/posteriors (Rezende and Mohamed, 2015) but find only marginal improvements.

PPL with ELBO directly since the gap between ELBO and log marginal likelihood might be large, particularly when the posterior does not collapse.

Reconstruction loss (Recon). Reconstruction loss is equivalent to the negative reconstruction term in ELBO: $-\mathbb{E}_{\mathbf{z}\sim q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\theta}(\mathbf{x}|\mathbf{z})]$. It characterizes how well the latent code can be used to recover the input.

Number of active units (AU, Burda et al. (2016)). Active units correspond to the dimensions of \mathbf{z} that covary with observations after the model is trained. More active units usually indicates richer latent representations (Burda et al., 2016). Specifically, a dimension is “active” when it is sensitive to the change in observations \mathbf{x} . Here we follow (Burda et al., 2016) and classify a latent dimension z as active if $\text{Cov}(\mathbf{x}, E_{z\sim q(\mathbf{z}|\mathbf{x})}[z]) > 0.01$.

In addition to the metrics above, we include KL between prior and posterior approximation, as well as the negative ELBO, for reference – though we find that these quantities are only partially descriptive of model quality. In Section 3.2 we also evaluate the latent space of learned models with specialized metrics such as reconstruction BLEU and classification accuracy.

2.2 Baselines

We experiment with several state-of-the-art techniques to mitigate posterior collapse, including several KL reweighting methods and the recently proposed aggressive training (He et al., 2019).

KL annealing (Bowman et al., 2016). KL annealing might be the most common method for reweighting. During annealing, the weight of the KL term is increased from a small value to 1.0 in the beginning of training.

Cyclic annealing (Liu et al., 2019). Cyclic annealing is another reweighting scheme proposed recently. It changes the weight of the KL term in a cyclic fashion, rather than monotonically increasing the weight.⁴

KL Thresholding / Free Bits (FB) (Kingma et al., 2016). FB replaces the KL term in ELBO with a hinge loss term that maxes each component

⁴We use the default cyclic annealing schedule from https://github.com/haofumml/cyclic_annealing in our codebase for a fair comparison.

of the original KL with a constant:

$$\sum_i \max[\lambda, D_{\text{KL}}(q_\phi(z_i|\mathbf{x})||p(z_i))] \quad (1)$$

Here, λ denotes the *target rate*, and z_i denotes the i^{th} dimension in the latent variable \mathbf{z} . Using the FB objective causes learning to give up trying to drive down KL for dimensions of \mathbf{z} that are already beneath the target rate. Pelsmaeker and Aziz (2019) conduct a comprehensive experimental evaluation of related methods and conclude that the FB objective is able to achieve comparable or superior performance (in terms of both language modeling and reconstruction) in comparison with other top-performing methods, many of which are substantially more complex. We notice that Pelsmaeker and Aziz (2019) experiment with a slightly different version of FB where the threshold is applied to the entire KL term directly, rather than on each dimension’s KL separately. We examine both versions here and refer to the single threshold version as “FBP” and the multiple threshold version (Eq. 1) as “FB”. For both FB and FBP, we vary the target rate λ and report the setting with the best validation PPL and the setting with the best balance between PPL and reconstruction loss.⁵

Aggressive training (He et al., 2019). He et al. (2019) observe that when posterior collapse occurs, the inference network often lags behind the generator during training. In contrast with the KL reweighting methods described above, He et al. (2019) propose an aggressive training schedule which iterates between multiple encoder update steps and one decoder update step to mitigate posterior collapse.⁶

Autoencoder (AE). We also include an autoencoder⁷ as a reference for reconstruction loss.

We show the results of these baselines trained on PTB in Table 1, where we find that it is difficult to balance language modeling (PPL) and representation learning (Recon and AU) – the systems with relatively good reconstruction (FBP, $\lambda = 7$) or higher AU (FB baselines) have suboptimal PPL, and the best PPL is achieved with sacrifice of Recon and AU. Note that good PPL indicates good LM, but good recon and AU indicates good representation learning. Without both, we do not really

⁵It is subjective to judge “balance”, thus we also report complete results for different target rates in Appendix C.

⁶We use the public code at <https://github.com/jxhe/vae-lagging-encoder>.

⁷Here, AE denotes the VAE trained without the KL term.

Table 1: Results on PTB test set for various baselines.

Method	PPL↓	Recon↓	AU↑	KL	-ELBO
AE	-	70.36	32	-	-
VAE	101.39	101.27	0	0.00	101.27
+ anneal	101.40	101.28	0	0.00	101.28
+ cyclic	108.97	101.85	5	1.37	103.22
+ aggressive	99.83	100.26	4	0.93	101.19
+ FBP ($\lambda = 7$)	102.82	95.63	4	7.05	102.67
+ FBP ($\lambda = 3$)	99.62	98.52	3	2.95	101.48
+ FB ($\lambda = 7$)	104.06	98.97	32	6.74	105.72
+ FB ($\lambda = 3$)	100.50	99.94	32	2.96	102.90

Table 2: Results on PTB test with encoder pretraining.

Method	PPL↓	Recon↓	AU↑	KL	-ELBO
AE	-	70.36	32	-	-
VAE	101.39	101.27	0	0.00	101.27
+ pretrain	102.26	101.46	0	0.00	101.46
+ pretrain + anneal	97.74	99.67	2	1.01	100.68

have a strong probabilistic model of language that captures latent factors.

We make two observations from the results in Table 1. First, reconstruction loss for an AE is substantially better than all VAE methods, which is intuitive since reconstruction is the only goal of training an AE. Second, models with high ELBO do not necessarily have good PPL (e.g. VAE+anneal); ELBO is not an ideal surrogate for evaluating language modeling performance.

2.3 Autoencoder-based Initialization

Based on the observations above we hypothesize that VAEs might benefit from initialization with an non-collapsed encoder, trained via an AE objective. Intuitively, if the encoder is providing useful information from the beginning of training, the decoder is more likely to make use of the latent code. In Table 2 we show the results of exploring this hypothesis on PTB. Even with encoder pretraining, we see that posterior collapse occurs immediately after beginning to update both encoder and decoder using the full ELBO objective. This indicates that the gradients of ELBO point towards a collapsed local optimum, even with biased initialization. When pretraining is combined with annealing, PPL improves substantially. However, the pretraining and anneal combination only has 2 active units and has small KL value – the latent representation is likely unsatisfactory. We speculate that this is because the annealing schedule eventually returns to the full ELBO objective which guides learning towards a (nearly) collapsed latent space. In the next section, we present an alternate approach using the KL thresholding / free bits method.

2.4 Our Method

In our proposed method, we initialize the inference network with an encoder that is pretrained using an autoencoder objective, as described above. Then, we train the VAE using the FB objective, also described above. We use the original FB which thresholds along each dimension.⁸ Thus, we combine two approaches so far considered independently: pretraining and KL thresholding. In this way, however, the VAE would start with a large KL and is thus trained with the full ELBO objective in the initial stage, which is bad as observed in Section 2.3. To remedy this, we apply an annealing weight to Eq. 1. We use the simplest linear annealing schedule that increases the weight from 0 to 1 in the first 10 epochs for all of our experiments. This approach can be viewed in connection with KL annealing: we train with zero KL weight until convergence, then reset the decoder and start increasing the KL weight with the free bits objective. Next we conduct comprehensive experiments across different datasets to validate our method.

3 Experiments

In this section we work with three text datasets: PTB (Marcus and Marcinkiewicz), Yahoo (Yang et al., 2017), and a downsampled version of SNLI (Bowman et al., 2015). We demonstrate the effectiveness of our method through language modeling, text reconstruction, and quality of the learned latent space. Complete experimental setup details can be found in Appendix A.

3.1 Language Modeling

For language modeling, we only report the results for PTB and Yahoo due to the space limit, but include the SNLI results in Appendix C. Since we have already shown that FBP outperforms FB on PPL in Section 2 (without pretraining), here we only include FBP as a baseline. For both FBP and our method we vary the target rate and report the settings that achieve competitive validation PPL.⁹

As shown in Table 3, our method with different target rates is able to consistently outperform all the baselines in terms of PPL. Meanwhile, for representation learning, our method beats all the

⁸Note that we do not pretrain the decoder and rather initialize it randomly. In our preliminary experiments, pretraining the decoder produced worse PPL compared with encoder-only pretraining.

⁹Results of all target rates can be found in Appendix C.

Table 3: Language modeling results on PTB and Yahoo test set. We bold the lines that represent the best average of language modeling and reconstruction loss. Cyclic* is from (Liu et al., 2019).

Method	PPL↓	Recon↓	AU↑	KL	-ELBO
PTB					
LSTM-LM	100.47	-	-	-	-
VAE	101.39	101.27	0	0.00	101.27
+ anneal	101.40	101.28	0	0.00	101.28
+ cyclic*	-	100.51	-	1.96	102.46
+ cyclic	108.97	101.85	5	1.37	103.22
+ aggressive	99.83	100.26	4	0.93	101.19
+ FBP ($\lambda = 3$)	99.62	98.52	3	2.95	101.48
+ FBP ($\lambda = 2$)	100.96	99.37	2	1.99	101.36
Ours ($\lambda = 8$)	98.07	92.60	32	10.95	103.56
Ours ($\lambda = 6$)	96.35	94.52	32	8.15	102.67
Ours ($\lambda = 4$)	96.17	96.91	32	4.99	101.90
Yahoo					
LSTM-LM	60.75	-	-	-	-
VAE	61.52	329.10	0	0.00	329.10
+ anneal	61.21	328.80	0	0.00	328.80
+ cyclic	66.93	333.80	4	2.83	336.63
+ aggressive	59.77	322.70	15	5.70	328.40
+ FBP ($\lambda = 9$)	62.59	322.91	6	9.08	331.99
+ FBP ($\lambda = 7$)	62.76	324.66	5	7.03	331.69
+ FBP ($\lambda = 5$)	62.78	326.26	3	5.07	331.32
+ FBP ($\lambda = 3$)	62.88	328.13	2	3.06	331.19
Ours ($\lambda = 6$)	59.23	317.39	32	12.09	329.48
Ours ($\lambda = 8$)	59.51	315.31	32	15.02	330.33
Ours ($\lambda = 9$)	59.60	315.09	32	15.49	330.58

baselines by a large margin on reconstruction loss with all latent units active. We also note that our method is not very sensitive to the target rate λ .

It is worth noticing that in some cases our method (e.g. $\lambda = 8$ in PTB and $\lambda = 8, 9$ in Yahoo) is able to outperform all the baselines but produce a bad ELBO (actually the worst on PTB). This suggests that ELBO might be a suboptimal surrogate for the log marginal likelihood sometimes, especially when KL is large, where the gap between ELBO and the marginal tends to be large as well.

3.2 Probing the Latent Space

We assess quality of learned latent space with SNLI through several metrics. For our method and FBP, we use the target rate where the best reconstruction loss is achieved while maintaining comparable PPL with aggressive training.¹⁰

Text Reconstruction We use greedy decoding and compute the BLEU score of the reconstructed sentence with the original one as the reference. The result is shown in Table 4. Unsurprisingly,

¹⁰Basically we try to tie PPL of different models and compare their latent space. Specifically, aggressive training has PPL 32.83, and PPL of the selected model for FBP and our method are 33.07 and 32.88, respectively.

the autoencoder achieves the highest BLEU score, meanwhile our method beats other VAE baselines.

Table 4: Reconstruction

Method	BLEU
AE	60.80
VAE	1.82
+ anneal	2.51
+ cyclic	4.39
+ aggressive	2.95
+ FBP ($\lambda = 7$)	8.07
Ours ($\lambda = 4$)	8.62

Table 5: Smoothness

Method	PCC
AE	0.620
VAE	0.039
+ anneal	0.009
+ cyclic	0.482
+ aggressive	0.209
+ FBP ($\lambda = 7$)	0.242
Ours ($\lambda = 4$)	0.683

Smoothness of Latent Space A major difference between VAE and AE is that VAE can learn a smooth latent space through the regularization from the Gaussian prior. In a smooth latent space, latent codes of similar sentences should be close to each other and vice versa. Therefore, we randomly sample 100k sentence pairs and evaluate the Pearson Correlation Coefficient (PCC) between the ℓ_2 distances of latent codes and edit word distances. As shown in Table 5, our method achieves a much higher PCC compared to the baselines.

Zhao et al. (2018) argue that a smooth latent space is beneficial for reconstructing noisy inputs. We follow their experiments and introduce noise to the input by randomly swapping words k times. As shown in Table 6, while AE achieves the best reconstruction when the noise is small ($k = 1$), its reconstruction deteriorates dramatically when $k > 1$, which suggests AE fails to learn a smooth latent space. In contrast, our method outperforms all the baselines by a large margin when $k > 1$.

Interpolation As illustrated in (Bowman et al., 2016), linear interpolation between latent variables is an intuitive way to qualitatively evaluate the smoothness of the latent space. We sample two latent codes \mathbf{z}_0 and \mathbf{z}_1 from the prior $p(\mathbf{z})$ (Table 7) and do linear interpolation between the two with evenly divided intervals.¹¹ For each interpolated point, we decode it greedily. Our method is able to generate grammatically plausible and semantically consistent interpolation in both cases.

3.3 Classification

To further evaluate the quality of the latent codes, we train a Gaussian mixture model (for unsupervised clustering) or a one-layer linear classifier (for supervised classification) on the pretrained latent codes. We work with a downsampled version

¹¹Interpolation examples from posterior samples $q_\phi(\mathbf{z}|\mathbf{x})$ are provided in Appendix 13.

Table 6: Noisy reconstruction loss (\downarrow) on SNLI. #swap denotes the number of word swaps.

#swap	1	2	3	4
AE	26.05	40.46	52.77	63.07
VAE	33.10	33.11	33.11	33.12
+ anneal	32.20	32.65	33.12	33.39
+ cyclic	31.83	32.87	33.73	34.38
+ aggressive	31.78	31.99	32.21	32.32
+ FBP ($\lambda = 7$)	29.93	32.59	34.90	36.77
Ours ($\lambda = 4$)	27.92	29.12	30.03	30.85

Table 7: Interpolation between prior samples on SNLI.

AE
people on their ground and they sit towards each other .
girls riding their cellphones and other people sit near papers .
girls riding in an area while not talk to dishes .
person riding in an area while carrying bags and papers .
someone riding in an office , selling a button .
three kid riding in <unk> signs a brick advertisement area .
Ours
a man with a cane is walking down the street .
a man with a cane is walking down the street .
a man in a blue shirt is eating food .
people are eating food .
people walk in a city .
people are outside in a city .

Table 8: Accuracy on Yelp of unsupervised and supervised classification. Evaluated via accuracy. #labeled denotes the number of labeled example during training.

#labeled	0	100	500	1k	2k	10k
AE	52.0	78.4	81.1	83.5	83.8	83.8
VAE	56.5	58.9	62.3	62.5	62.9	64.0
+ anneal	56.1	58.7	60.6	61.5	61.4	64.1
+ cyclic	59.3	78.1	79.8	81.1	81.7	83.1
+ aggressive	63.7	65.6	68.6	72.1	76.7	79.4
+ FBP ($\lambda = 9$)	60.7	73.3	75.0	76.1	77.6	79.5
Ours ($\lambda = 6$)	67.2	83.8	88.3	89.1	89.5	89.5

of Yelp sentiment dataset collected by Shen et al. (2017). We vary the number of labeled data¹² and the results are shown in Table 8. For FBP and our method, the target rate is selected in terms of the average validation accuracy. Our method consistently yields the best results on all settings – remarkably, its performance with only 100 labeled samples already surpasses others with 10k labels.

4 Conclusion

In this paper, we propose a simple training fix to tackle posterior collapse in VAEs. Extensive experiments demonstrate the effectiveness of our method on both representation learning and language modeling.

¹²Specifically, we train a VAE model on the Yelp dataset to obtain the latent codes, then we use the latent codes of labeled data to train the classifier.

References

- Alexander Alemi, Ben Poole, Ian Fischer, Joshua Dillon, Rif A Saurous, and Kevin Murphy. 2018. Fixing a broken ELBO. In *Proceedings of ICML*.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of EMNLP*.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. Generating sentences from a continuous space. In *Proceedings of CoNLL*.
- Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. 2016. Importance weighted autoencoders. In *Proceedings of ICLR*.
- Xi Chen, Diederik P Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2017. Variational lossy autoencoder. In *Proceedings of ICLR*.
- Adji B Dieng, Yoon Kim, Alexander M Rush, and David M Blei. 2018. Avoiding latent variable collapse with generative skip models. In *Proceedings of ICML workshop on Theoretical Foundations and Applications of Deep Generative Models*.
- Junxian He, Daniel Spokoyny, Graham Neubig, and Taylor Berg-Kirkpatrick. 2019. Lagging inference networks and posterior collapse in variational autoencoders. In *Proceedings of ICLR*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Yoon Kim, Sam Wiseman, Andrew Miller, David Sontag, and Alexander Rush. 2018. Semi-amortized variational autoencoders. In *Proceedings of ICML*.
- Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *Proceedings of ICLR*.
- Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. 2016. Improved variational inference with inverse autoregressive flow. In *Proceedings of NeurIPS*.
- Xiaodong Liu, Jianfeng Gao, Asli Celikyilmaz, Lawrence Carin, et al. 2019. Cyclical annealing schedule: A simple approach to mitigating KL vanishing. In *Proceedings of NAACL*.
- Mitchell P Marcus and Mary Ann Marcinkiewicz. Building a large annotated corpus of english: The penn treebank. *Computational Linguistics*, 19(2).
- Tom Pelsmaecker and Wilker Aziz. 2019. Effective estimation of deep generative language models. *arXiv preprint arXiv:1904.08194*.
- Daniilo Rezende and Shakir Mohamed. 2015. Variational inference with normalizing flows. In *Proceedings of ICML*.
- Stanislau Semeniuta, Aliaksei Severyn, and Erhardt Barth. 2017. A hybrid convolutional variational autoencoder for text generation. In *Proceedings of EMNLP*.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. In *Proceedings of NeurIPS*.
- Jiacheng Xu and Greg Durrett. 2018. Spherical latent spaces for stable variational autoencoders. In *Proceedings of EMNLP*.
- Zichao Yang, Zhiting Hu, Ruslan Salakhutdinov, and Taylor Berg-Kirkpatrick. 2017. Improved variational autoencoders for text modeling using dilated convolutions. In *Proceedings of ICML*.
- Junbo Zhao, Yoon Kim, Kelly Zhang, Alexander Rush, and Yann LeCun. 2018. Adversarially regularized autoencoders. In *Proceedings of ICML*.
- Shengjia Zhao, Jiaming Song, and Stefano Ermon. 2019. InfoVAE: Information maximizing variational autoencoders. In *Proceedings of AAAI*.

A Details of Experimental Setup

For SNLI, we randomly downsample a subset of it, which contains 100K/10K/10K sentences as training/validation/test. For the Yelp sentiment dataset, we also randomly downsample 100K/10K/10K sentences for training/validation/test, respectively. We use an one-layer LSTM for both the encoder and decoder and a latent vector size of 32.

We follow Kim et al. (2018); He et al. (2019) and use a single-layer LSTM the encoder and decoder in all of our experiments. The sizes of word embeddings and hidden states for different datasets are given in Table 9. We initialize the LSTM parameters with a uniform distribution $\mathcal{U}(-0.01, 0.01)$, and embeddings with another uniform distribution $\mathcal{U}(-0.1, 0.1)$.

Just as in Kim et al. (2018); He et al. (2019), for the decoder, we use dropout of 0.5 on both the input word embedding and the last dense before logits. During training, we use the SGD optimizer without momentum. Initialized with 0.5, the learning rate is decayed by $\times 0.5$ with a patience of 2 if the validation loss has not improved in the past 2 epochs. The maximum number of epochs is 100, but the training will stop early after 5 learning rate decays. For our VAE + anneal baseline, we use the simplest linear annealing schedule that increases the weight from 0 to 1 in the first 10 epochs, just the same as in our method (stated in Section 2.4).

Table 9: The sizes of word embeddings and hidden states for PTB, SNLI and Yahoo.

	PTB	SNLI	Yahoo
Word Embedding Size	256	128	512
Hidden Size of Encoder	256	512	1024
Hidden Size of Decoder	256	512	1024

B Copying Behaviour

We check to make sure that our model is not simply learning to copy sentences from the training set. We test for copying behaviour on the PTB dataset. Specifically, we sampled 300 sentences from the prior and retrieved their nearest neighbors in the training set. The average edit distance between the samples and their nearest neighbors from our method is 14.11 (the average training sentence length is 22.10), versus 13.68 from the collapsed VAE. This means that there is no obvious copying behaviour when KL grows in our method.

C Additional Results of Language Modeling

In Table 10, we provide a more detailed version of the language modeling results on PTB, SNLI and Yahoo. The full results of all the target rate trial for both FBP and our method are included. In addition to the metrics in Table 3, we also report NLL, Mutual Information I_q between \mathbf{z} and \mathbf{x} under $q_\phi(\mathbf{z}|\mathbf{x})$ (MI)¹³, and the perplexity computed by ELBO (ELBO PPL).

¹³For the estimation of I_q , we follow the same method as in He et al. (2019)

Table 10: Additional results of language modeling on PTB and SNLI.

Dataset	Method	NLL↓	PPL↓	Recon↓	MI↑	AU↑	-ELBO	ELBO PPL	KL
PTB	LSTM-LM	101.04	100.47	-	-	-	-	-	-
	AE	-	-	70.36	8.22	32	-	-	-
	VAE	101.23	101.39	101.27	0.01	0	101.27	101.58	0.00
	+ anneal	101.24	101.40	101.28	0.00	0	101.28	101.62	0.00
	+ cyclic (reported)	-	-	100.51	-	-	102.46	107.25	1.96
	+ cyclic	102.81	108.97	101.85	1.27	5	103.22	111.03	1.37
	+ cyclic (SGD)	102.06	105.28	102.14	0.00	0	102.14	105.67	0.00
	+ aggressive	100.89	99.83	100.26	0.83	4	101.19	101.17	0.93
	+ pretrained enc	101.42	102.26	101.46	0.0	0	101.46	102.45	0.00
	+ anneal enc	100.43	97.74	99.67	0.97	2	100.68	98.88	1.01
	+ FBP ($\lambda = 9$)	101.95	104.73	94.66	7.22	6	103.59	112.88	8.93
	+ FBP ($\lambda = 8$)	101.64	103.30	95.39	7.11	6	103.43	112.09	8.04
	+ FBP ($\lambda = 7$)	101.54	102.82	95.63	6.52	4	102.67	108.26	7.05
	+ FBP ($\lambda = 6$)	101.68	103.47	96.66	5.76	10	102.67	108.27	6.01
	+ FBP ($\lambda = 5$)	101.24	101.42	97.12	4.80	4	102.21	106.01	5.10
	+ FBP ($\lambda = 4$)	101.49	102.58	97.84	3.86	4	101.86	104.31	4.01
	+ FBP ($\lambda = 3$)	100.85	99.62	98.52	2.86	3	101.48	102.52	2.95
	+ FBP ($\lambda = 2$)	101.14	100.96	99.37	1.91	2	101.36	101.98	1.99
	IWAE ($k = 10$)	100.86	99.69	100.89	0.05	0	100.89	99.81	0.00
	+ pretrained enc	100.92	99.95	100.96	0.01	0	100.96	100.15	0.00
	Ours ($\lambda = 9$)	101.09	100.71	92.00	7.78	32	104.43	117.32	12.44
	Ours ($\lambda = 8$)	100.51	98.07	92.60	7.49	32	103.56	112.72	10.95
	Ours ($\lambda = 7$)	101.06	100.60	93.25	7.46	32	103.74	113.68	10.49
Ours ($\lambda = 6$)	100.12	96.35	94.52	6.30	32	102.67	108.25	8.15	
Ours ($\lambda = 5$)	100.23	96.86	95.87	5.31	32	102.41	106.97	6.54	
Ours ($\lambda = 4$)	100.08	96.17	96.91	4.08	32	101.90	104.51	4.99	
Ours ($\lambda = 3$)	100.21	96.75	97.71	3.19	32	101.56	102.90	3.85	
Ours ($\lambda = 2$)	100.41	97.65	98.73	2.21	32	101.38	102.07	2.65	
SNLI	LSTM-LM	32.97	21.44	-	-	-	-	-	-
	AE	-	-	8.68	9.18	32	-	-	-
	VAE	33.09	21.67	33.08	0.03	1	33.12	21.73	0.04
	+ anneal	33.01	21.50	31.66	1.45	2	33.07	21.63	1.42
	+ cyclic	34.04	23.67	30.69	3.60	5	34.32	24.29	3.63
	+ cyclic (SGD)	33.07	21.62	30.89	2.33	4	33.25	21.99	2.36
	+ aggressive	32.83	21.16	31.53	1.38	5	32.95	21.39	1.42
	+ FBP ($\lambda = 9$)	33.28	22.05	25.26	8.06	6	34.25	24.13	8.99
	+ FBP ($\lambda = 8$)	33.26	22.02	26.07	7.35	7	34.08	23.75	8.01
	+ FBP ($\lambda = 7$)	33.07	21.62	26.65	6.76	6	33.78	23.11	7.14
	+ FBP ($\lambda = 6$)	33.09	21.68	27.54	5.95	5	33.59	22.71	6.06
	+ FBP ($\lambda = 5$)	33.04	21.58	28.38	4.95	6	33.49	22.48	5.10
	+ FBP ($\lambda = 4$)	33.04	21.57	29.25	4.06	4	33.36	22.22	4.11
	+ FBP ($\lambda = 3$)	33.04	21.56	30.19	3.00	4	33.31	22.11	3.11
	+ FBP ($\lambda = 2$)	32.99	21.46	31.04	2.11	3	33.16	21.80	2.12
	Ours ($\lambda = 9$)	33.42	22.33	22.30	8.80	32	35.70	27.61	13.40
	Ours ($\lambda = 8$)	33.47	22.45	22.65	8.76	32	35.62	27.40	12.96
	Ours ($\lambda = 7$)	33.25	22.00	23.36	8.48	32	35.11	26.14	11.75
	Ours ($\lambda = 6$)	33.17	21.84	24.06	8.24	32	34.83	25.47	10.77
	Ours ($\lambda = 5$)	33.07	21.64	24.94	7.71	32	34.47	24.63	9.53
	Ours ($\lambda = 4$)	32.88	21.24	26.52	6.65	32	34.11	23.83	7.60
	Ours ($\lambda = 3$)	32.87	21.23	28.02	5.25	32	33.87	23.31	5.86
	Ours ($\lambda = 2$)	32.79	21.07	29.75	3.32	32	33.42	22.35	3.67

Table 11: Additional results of language modeling on Yahoo.

Dataset	Method	NLL↓	PPL↓	Recon↓	MI↑	AU↑	-ELBO	ELBO PPL	KL
Yahoo	LSTM-LM	328.00	60.75	-	-	-	-	-	-
	AE	-	-	278.76	9.26	32	-	-	-
	VAE	329.00	61.52	329.10	0.00	0	329.10	61.59	0.0
	+ anneal	328.60	61.21	328.80	0.00	0	328.80	61.36	0.0
	+ cyclic	335.74	66.93	333.80	2.77	4	336.63	67.69	2.83
	+ cyclic (SGD)	332.48	64.26	332.65	0.00	1	332.68	64.42	0.03
	+ aggressive	326.70	59.77	322.70	2.9	15	328.40	61.06	5.70
	+ FBP ($\lambda = 9$)	330.38	62.59	322.91	8.21	6	331.99	63.86	9.08
	+ FBP ($\lambda = 8$)	330.80	62.92	324.03	7.54	6	332.09	63.94	8.05
	+ FBP ($\lambda = 7$)	330.60	62.76	324.66	6.76	5	331.69	63.62	7.03
	+ FBP ($\lambda = 6$)	331.05	63.12	325.87	5.94	5	332.01	63.88	6.13
	+ FBP ($\lambda = 5$)	330.62	62.78	326.26	5.00	3	331.32	63.33	5.07
	+ FBP ($\lambda = 4$)	331.06	63.13	327.55	4.00	3	331.66	63.60	4.11
	+ FBP ($\lambda = 3$)	330.75	62.88	328.13	2.99	2	331.19	63.23	3.06
	+ FBP ($\lambda = 2$)	331.30	63.32	329.60	1.98	1	331.63	63.58	2.04
	Ours ($\lambda = 2$)	326.34	59.50	322.55	5.35	32	328.51	61.14	5.96
	Ours ($\lambda = 3$)	326.12	59.34	321.29	6.41	32	328.73	61.31	7.44
	Ours ($\lambda = 4$)	326.01	59.26	319.49	7.58	32	329.03	61.54	9.54
	Ours ($\lambda = 5$)	326.04	59.28	318.55	8.08	32	329.31	61.76	10.76
	Ours ($\lambda = 6$)	325.97	59.23	317.39	8.51	32	329.48	61.89	12.09
Ours ($\lambda = 7$)	326.08	59.31	316.42	8.78	32	329.76	62.10	13.34	
Ours ($\lambda = 8$)	326.35	59.51	315.31	8.99	32	330.33	62.55	15.02	
Ours ($\lambda = 9$)	326.47	59.60	315.09	9.03	32	330.58	62.75	15.49	

D Additional Qualitative Examples

D.1 Interpolation between Prior Samples

We randomly sample 10 pair of source and target latent code from the standard Gaussian prior and do linear interpolation. For the sampled and interpolated latent codes, we do greedy decoding. The results are shown in Table 12.

D.2 Interpolation between Posterior Samples

We randomly sample 10 pairs of source and target input sentences from the test set of SNLI. For each input sentence, we randomly sample a latent code from the approximated posterior $q(\mathbf{z}|\mathbf{x})$. Then we linearly interpolate between each pair of sampled source/target latent code. For the sampled and interpolated latent codes, we do greedy decoding. The results are shown in Table 13.

Table 12: Interpolation between prior samples on SNLI

EXAMPLE 1	EXAMPLE 6
<p>the kids are playing hide and seek in the classroom the girl is about to play the drums the girl is about to play in the sandbox the girl is about to play in the sandbox the girls are watching tv in the classroom . the girl is eating cake in the kitchen . the women are watching tv in the bar . the women are eating lunch . the women are eating lunch . the women are eating dinner . a woman is eating in a restaurant .</p>	<p>a man is jumping off a rock into the air . a man is sitting on a bench with a red umbrella . a man is sitting on a bench with a red umbrella . a man is sitting on a bench with a red umbrella . a woman is sitting on a bench in front of a building . a young man is sitting on a bench outside . a young man is taking pictures of a building . a large group of people are taking pictures of a building . a large group of people are taking pictures in the street . a large group of people are taking pictures in the street . a large group of people are taking pictures in the street .</p>
EXAMPLE 2	EXAMPLE 7
<p>a man with a cane is walking down the street . a man with a cane is walking down the street . a man with a cane is walking down the street . a man with a cane is walking down a sidewalk . a man in a blue shirt is eating food . man in a hat and jeans is walking down a sidewalk . people are eating food . people walk through a city street . people walk in a city . people walk in a city . people are outside in a city .</p>	<p>two men are swimming in the ocean . two men are swimming in the ocean . two men are on the beach . two people are at the beach . two people are at the beach . the man is at the park . a man is at the beach . a man is at the beach . a man is at the beach . a man is at the beach . a man is taking pictures of the ocean .</p>
EXAMPLE 3	EXAMPLE 8
<p>the man is going to the bathroom . the man is going to the bathroom . the man is going to the bathroom . the man is going to the bathroom . the man is playing music in the living room . the man is playing with his dog . the man is playing with a cat . the man is playing with a ball . the man is playing with a ball . the man is playing with a ball . the man is playing with a ball .</p>	<p>two people in bathing suits are in a park . two people in bathing suits are in a park . two people in blue shirts are in a field . two people in blue shirts are in a field . two people in a field are playing soccer . two people in a field of flowers . two people are at a beach . a man in a blue shirt is looking at a camera . a man in a blue shirt is playing basketball . a man in a blue shirt is playing basketball . a man sits at a carnival .</p>
EXAMPLE 4	EXAMPLE 9
<p>a person is about to get a picture taken a person is about to get a picture taken a person is about to get a picture taken a person is waiting for a friend to come to work a person is waiting for a friend to come a person is trying to find a speech the people are playing monopoly there are people playing monopoly there are people performing there are people performing surgery there are no people in this picture .</p>	<p>two women are outside . two women are outside . two women are outside . two women are sitting outside . a couple is sitting outside . a couple is sitting outside . a couple is sitting in a car . a man is reading a newspaper in a park . a man is reading a book in a park . a man is carrying a bag of food . a man is carrying a bag of food .</p>
EXAMPLE 5	EXAMPLE 10
<p>a girl sits on a bench in front of a large crowd . a girl sits on a bench in front of a large crowd . a man sits on a bench in front of a large crowd . a man sits on a bench in front of a large crowd . a woman wearing a blue shirt is walking on the beach . a woman wearing a blue shirt is walking on the beach . the woman is wearing a blue shirt . the woman is wearing a blue shirt . the woman is wearing a blue shirt . the woman is wearing a blue shirt . the woman is wearing a blue shirt . the woman is moving her legs .</p>	<p>the man is climbing the mountain . the man is making a noise . the man is making a noise . the man is making a noise . the man is making a noise . the people are enjoying the sunshine . the people are watching a movie . the women are eating a meal . the men are watching a movie . the men are watching a movie . two women sit in a circle together . two women sit in a circle together .</p>

Table 13: Interpolation between posterior samples on SNLI.

	EXAMPLE 1	EXAMPLE 6
SOURCE INPUT	a child is eating with utensils .	a middle-aged man with long , curly red-hair wearing a dark vest , shirt and pants is holding a microphone in front of a black backdrop .
TARGET INPUT	a youth wearing a blue and red jersey and yellow helmet is crouching in a football position	people are doing <unk>
POSTERIOR-SAMPLED SOURCE	a little girl is eating at a restaurant .	the two men are wearing jeans and a blue shirt , and a woman are holding a rope .
INTERPOLATION	a little girl is eating at a restaurant . a little girl is eating at a table . a little girl is playing with a toy . a little girl is playing with a ball . a little girl is wearing a pink shirt . a little girl is wearing a pink shirt and holding a popsicle . a man is wearing a blue shirt and a hat . a man is wearing a blue shirt and a hat . a man is wearing a black shirt and black pants	the two men are wearing jeans and a blue shirt . the two men are wearing jeans and a blue shirt . the two men are wearing white shirts and are playing a sport . the people are trying to find a cure for cancer . the people are playing soccer two people are playing baseball two people are playing baseball two people are playing baseball people are playing baseball
POSTERIOR-SAMPLED TARGET	a man wearing a black shirt and black pants is standing on a sidewalk	people are playing baseball
	EXAMPLE 2	EXAMPLE 7
SOURCE INPUT	the men are feeling competitive .	both men are wearing similar colors .
TARGET INPUT	a young woman is sitting in a field .	a huge animal surrounded
POSTERIOR-SAMPLED SOURCE	a man is climbing a tree .	a man is cutting a cake .
INTERPOLATION	a man is climbing a tree . a man is eating a pizza . a man is wearing a blue shirt . a man is wearing a blue shirt . a man is wearing a blue shirt . a man is standing in front of a large crowd . a woman is standing in front of a large crowd . a woman is standing in front of a large crowd . a woman is standing in front of a large crowd .	a man is cutting a cake . a man is painting a portrait of a woman . a man is painting a portrait of a woman . a man with a beard is playing guitar . a young boy with a blue shirt . a little girl with a pink shirt . a tall human with a shirt a tall human with a shirt a tall human with a shirt a tall human looking
POSTERIOR-SAMPLED TARGET	a woman is standing in front of a large crowd .	
	EXAMPLE 3	EXAMPLE 8
SOURCE INPUT	the animals are near the water .	a crowded city street in asia .
TARGET INPUT	a truck is going to tow an illegally parked white volkswagon .	a guy is in front of a business . a tree with flowers is in front .
POSTERIOR-SAMPLED SOURCE	two boys are at a beach .	a bunch of people are in a store .
INTERPOLATION	two boys are at a beach . two men are looking at a man in a wheelchair . the children are at the beach . the children are looking at the sky . a woman is looking at a man in a wheelchair . a woman is looking at a man in a wheelchair . a woman is looking at a map . a woman is waiting for a bus to come out of the road . a woman is waiting for a bus to come out of the city .	a bunch of people are in a park . a bunch of people are in a park . a few people are in a park . a few people are in a park . a few people are in a park . a few people are in front of a building . a lady is in a store with a man . a lady in a blue shirt is looking at a man in a blue shirt . a lady in a blue shirt is looking at a man in a blue shirt . a lady in a blue shirt is looking at a man in a blue shirt . a lady in a blue shirt and black pants is playing in a fountain with a small child in the background .
POSTERIOR-SAMPLED TARGET	a woman is waiting for a bus .	
	EXAMPLE 4	EXAMPLE 9
SOURCE INPUT	one young child in a swimsuit jumping off a blue inflatable slide with water .	a woman with 5 small children .
TARGET INPUT	a girl swings from a rope swing in front .	a man works on <unk> a circuit as he monitors the progress on a tablet device .
POSTERIOR-SAMPLED SOURCE	a young man in a blue shirt and black pants is standing by a large rock formation .	man with blue shirt and blue shirt is playing basketball .
INTERPOLATION	a young girl in a pink shirt and blue shorts is jumping into a pool . a woman in a pink shirt and black shorts is jumping into a pool . a woman in a pink shirt and black shorts is playing a game of soccer . a woman in a pink shirt and black shorts is playing a game of soccer . a woman with a red shirt and a black shirt is sitting in a chair . a woman with a red shirt and a black shirt is sitting in a chair . a woman with a red shirt and a black shirt is sitting in a chair . a woman with a red shirt and a black shirt is looking at a camera . a woman watches a man play a game of soccer .	man with blue shirt and blue shirt is playing basketball . man with blue shirt and blue shirt is playing basketball . man in blue shirt with a blue shirt on his head . a man with a hat is playing with a ball . a man in a blue shirt is looking at a plant . a man in a blue shirt is sitting on a bench with a shovel . a man in a blue shirt is sitting on a bench with a shovel . a man is on a skateboard in front of a building . a man is on a skateboard in front of a building with graffiti on it .
POSTERIOR-SAMPLED TARGET	a woman watches a man play a guitar in front of a crowd .	a man works on a project on a stove .
	EXAMPLE 5	EXAMPLE 10
SOURCE INPUT	a girl with glasses next red white and blue flags .	a man walking across a bridge near a steak restaurant .
TARGET INPUT	three greyhounds are taking a walk with their owner .	a woman in a white shirt and shorts is playing a red guitar .
POSTERIOR-SAMPLED SOURCE	the girl is drinking milk with the camera .	people in a park
INTERPOLATION	the girl is drinking milk with the camera . the girl is drinking milk with her hands . the girl is drinking water with a bucket . the girl is using a camera . two girls are outside with a blue umbrella . two girls are outside with a blue umbrella . two girls are outside with a dog . two girls are taking a picture of a tree . two guys are on a bench .	a woman in a blue shirt is eating a sandwich . a woman in a blue shirt is playing a game of tennis . a woman is playing a game of tennis . a woman is playing a game of tennis . a man is playing a guitar . a man is playing a guitar . a man is playing a guitar . a man is playing a guitar . a man is playing a guitar . a man is playing a guitar .
POSTERIOR-SAMPLED TARGET	two guys are on a boat .	a man is playing a guitar .