

# Variational Neural Machine Translation with Normalizing Flows

Hendra Setiawan    Matthias Sperber    Udhay Nallasamy    Matthias Paulik

Apple

{ hendra, sperber, udhay, mpaulik }@apple.com

## Abstract

Variational Neural Machine Translation (VNMT) is an attractive framework for modeling the generation of target translations, conditioned not only on the source sentence but also on some latent random variables. The latent variable modeling may introduce useful statistical dependencies that can improve translation accuracy. Unfortunately, learning informative latent variables is non-trivial, as the latent space can be prohibitively large, and the latent codes are prone to be ignored by many translation models at training time. Previous works impose strong assumptions on the distribution of the latent code and limit the choice of the NMT architecture. In this paper, we propose to apply the VNMT framework to the state-of-the-art Transformer and introduce a more flexible approximate posterior based on normalizing flows. We demonstrate the efficacy of our proposal under both in-domain and out-of-domain conditions, significantly outperforming strong baselines.

## 1 Introduction

Translation is inherently ambiguous. For a given source sentence, there can be multiple plausible translations due to the author’s stylistic preference, domain, and other factors. On the one hand, the introduction of neural machine translation (NMT) has significantly advanced the field (Bahdanau et al., 2015), continually producing state-of-the-art translation accuracy. On the other hand, the existing framework provides no explicit mechanisms to account for translation ambiguity.

Recently, there has been a growing interest in latent-variable NMT (LV-NMT) that seeks to incorporate latent random variables into NMT to account for the ambiguities mentioned above. For instance, Zhang et al. (2016) incorporated latent codes to capture underlying global semantics of source sentences into NMT, while Su et al. (2018)

proposed fine-grained latent codes at the word level. The learned codes, while not straightforward to analyze linguistically, are shown empirically to improve accuracy. Nevertheless, the introduction of latent random variables complicates the parameter estimation of these models, as it now involves intractable inference. In practice, prior work resorted to imposing strong assumptions on the latent code distribution, potentially compromising accuracy.

In this paper, we focus on improving Variational NMT (VNMT) (Zhang et al., 2016): a family of LV-NMT models that relies on the amortized variational method (Kingma and Welling, 2014) for inference. Our contributions are twofold. (1) We employ variational distributions based on normalizing flows (Rezende and Mohamed, 2015), instead of uni-modal Gaussian. Normalizing flows can yield complex distributions that may better match the latent code’s true posterior. (2) We employ the Transformer architecture (Vaswani et al., 2017), including *Transformer-Big*, as our VNMT’s generator network. We observed that the generator networks of most VNMT models belong to the RNN family that are relatively less powerful as a translation model than the Transformer.

We demonstrate the efficacy of our proposal on the German-English IWSLT’14 and English-German WMT’18 tasks, giving considerable improvements over strong non-latent Transformer baselines, and moderate improvements over Gaussian models. We further show that gains generalize to an out-of-domain condition and a simulated bimodal data condition.

## 2 VNMT with Normalizing Flows

**Background** Let  $x$  and  $y$  be a source sentence and its translation, drawn from a corpus  $\mathcal{D}$ . Our model seeks to find parameters  $\theta$  that maximize the marginal of a latent-variable model  $p_\theta(y, Z | x)$

where  $Z \in \mathbb{R}^D$  is a sentence-level latent code similar to (Zhang et al., 2016). VNMT models sidestep the marginalization by introducing variational distributions and seek to minimize this function (i.e., the Evidence Lower Bound or ELBO):

$$\sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \mathbb{E}_{q(Z|\mathbf{x}, \mathbf{y})} [\log p_\theta(\mathbf{y} | \mathbf{x}, Z)] - \text{KL}(q(Z | \mathbf{x}, \mathbf{y}) || p(Z | \mathbf{x})), \quad (1)$$

where  $q(Z | \mathbf{x}, \mathbf{y})$ ,  $p(Z | \mathbf{x})$  are the variational posterior and prior distribution of the latent codes, while  $p(\mathbf{y} | \mathbf{x}, Z)$  is a *generator* that models the generation of the translation conditioned on the latent code<sup>1</sup>. The ELBO is improved when the model learns a posterior distribution of latent codes that minimizes the reconstruction loss (the first term) while incurring a smaller amount of KL divergence penalty between the variational posterior and the prior (the second term).

The majority of VNMT models design their variational distributions to model unimodal distribution via isotropic Gaussians with diagonal covariance, which is the simplest form of prior and approximate posterior distribution. This assumption is computationally convenient because it permits a closed-form solution for computing the KL term and facilitates end-to-end gradient-based optimization via the re-parametrization trick (Rezende and Mohamed, 2015). However, such a simple distribution may not be expressive enough to approximate the true posterior distribution, which could be non-Gaussian, resulting in a loose gap between the ELBO and the true marginal likelihood. Therefore, we propose to employ more flexible posterior distributions in our VNMT model, while keeping the prior a Gaussian.

**Normalizing Flows-based Posterior** Rezende and Mohamed (2015) proposed Normalizing Flows (NF) as a way to introduce a more flexible posterior to Variational Autoencoder (VAE). The basic idea is to draw a sample,  $Z_0$ , from a simple (e.g., Gaussian) probability distribution and to apply  $K$  invertible parametric transformation functions ( $f_k$ ) called *flows* to transform the sample. The final latent code is given by  $Z_K = f_K(\dots f_2(f_1(Z_0))\dots)$  whose probability density function,  $q_\lambda(Z_K | \mathbf{x}, \mathbf{y})$ ,

<sup>1</sup>In VAE terms, the posterior and prior distributions are referred to as the encoders, while the generator is referred to as the decoder. As these terms have other specific meaning in NMT, we avoid to use them in this paper.

is defined via the change of variable theorem as follows:

$$q_0(Z_0 | \mathbf{x}, \mathbf{y}) \prod_{k=1}^K \left| \det \frac{\partial f_k(Z_{k-1}; \lambda_k(\mathbf{x}, \mathbf{y}))}{\partial Z_{k-1}} \right|^{-1},$$

where  $\lambda_k$  refers to the parameters of the  $k$ -th flow with  $\lambda_0$  corresponds to the parameters of a base distribution. In practice, we can only consider transformations, whose determinants of Jacobians (the second term) are invertible and computationally tractable.

For our model, we consider several NFs, namely *planar flows* (Rezende and Mohamed, 2015), *Sylvester flows* (van den Berg et al., 2018) and *affine coupling layer* (Dinh et al., 2017), which have been successfully applied in computer vision tasks.

Planar flows (PF) applies this function:

$$f_k(Z; \lambda_k(\mathbf{x}, \mathbf{y})) = Z + \mathbf{u} \cdot \tanh(\mathbf{w}^T Z + \mathbf{b}),$$

where  $\lambda_k = \{\mathbf{u}, \mathbf{w} \in \mathbb{R}^D, \mathbf{b} \in \mathbb{R}\}$ . Planar flows perform contraction or expansion to the direction perpendicular to the  $(\mathbf{w}^T Z + \mathbf{b})$  hyperplane.

Sylvester flows (SF) applies this function:

$$f_k(Z; \lambda_k(\mathbf{x}, \mathbf{y})) = Z + \mathbf{A} \cdot \tanh(\mathbf{B}Z + \mathbf{b}),$$

where  $\lambda_k = \{\mathbf{A}, \mathbf{B} \in \mathbb{R}^{M \times D}, \mathbf{b} \in \mathbb{R}^M\}$  and  $M$  is the number of hidden units. Planar flows are a special case of Sylvester flows where  $M = 1$ . In our experiments, we consider the orthogonal Sylvester flows (van den Berg et al., 2018), whose parameters are matrices with  $M$  orthogonal columns.

Meanwhile, the affine coupling layer (CL) first splits  $Z$  into  $Z^{d_1}, Z^{d_2} \in \mathbb{R}^{D/2}$  and applies the following function:

$$f_k(Z^{d_1}; \lambda_k(\mathbf{x}, \mathbf{y})) = Z^{d_1}, \\ f_k(Z^{d_2}; \lambda_k(\mathbf{x}, \mathbf{y}, Z^{d_1})) = Z^{d_2} \odot \exp(\mathbf{s}_k) + \mathbf{t}_k,$$

where it applies identity transform to  $Z^{d_1}$  and applies a scale-shift transform to  $Z^{d_2}$  according to  $\lambda_k = \{\mathbf{s}_k, \mathbf{t}_k\}$ , which are conditioned on  $Z^{d_1}$ ,  $\mathbf{x}$  and  $\mathbf{y}$ . CL is less expressive than PF and SF, but both sampling and computing the probability of arbitrary samples are easier. In practice, we follow (Dinh et al., 2017) to switch  $Z^{d_1}$  and  $Z^{d_2}$  alternately for subsequent flows.

As we adopt the amortized inference strategy, the parameters of these NFs are data-dependent. In our model, they are the output of 1-layer linear map

with inputs that depend on  $\mathbf{x}$  and  $\mathbf{y}$ . Also, as the introduction of normalizing flows no longer offers a simple closed-form solution, we modify the KL term in Eq. 1 into:

$$\mathbb{E}_{q_{\lambda}(Z|\mathbf{x},\mathbf{y})} [\log q_{\lambda}(Z | \mathbf{x}, \mathbf{y}) - \log p_{\psi}(Z | \mathbf{x})]$$

where we estimate the expectation w.r.t.  $q(Z_K|\mathbf{x}; \lambda)$  via  $L$  Monte-Carlo samples. We found that  $L = 1$  is sufficient, similar to (Zhang et al., 2016). To address variable-length inputs, we use the average of the embeddings of the source and target tokens via a mean-pooling layer, i.e.,  $\text{meanpool}(\mathbf{x})$  and  $\text{meanpool}(\mathbf{y})$  respectively.

**Transformer-based Generator** We incorporate the latent code to the Transformer model by mixing the code into the output of the Transformer decoder’s last layer ( $h_j$ ) as follows:

$$g_j = \delta([h_j; Z]), \quad h_j = (1 - g_j) * h_j + g_j * Z$$

where  $g_j$  controls the latent code’s contribution, and  $\delta(\cdot)$  is the sigmoid function. In the case of the dimension of the latent code ( $D$ ) doesn’t match the dimension of  $h_j$ , we apply a linear projection layer. Our preliminary experiments suggest that Transformer is less likely to ignore the latent code in this approach compared to other approaches we explored, e.g., incorporating the latent code as the first generated token as used in (Zhang et al., 2016).

**Prediction** Ultimately, we search for the most probable translation ( $\hat{\mathbf{y}}$ ) given a source sentence ( $\mathbf{x}$ ) through the evidence lower bound. However, sampling latent codes from the posterior distribution is not straightforward, since the posterior is conditioned on the sentence being predicted. Zhang et al. (2016) suggests taking the prior’s mean as the latent code. Unfortunately, as our prior is a Gaussian distribution, this strategy can diminish the benefit of employing normalizing flows posterior.

Eikema and Aziz (2018) explore two strategies, namely restricting the conditioning of the posterior to  $\mathbf{x}$  alone (dropping  $\mathbf{y}$ ) and introducing an auxiliary distribution,  $r(Z|\mathbf{x})$ , from which the latent codes are drawn. They found that the former is more accurate with the benefit of being simpler. This is confirmed by our preliminary experiments. We opt to adopt this strategy and use the mean of the posterior as the latent code at prediction time.

**Mitigating Posterior Collapse** As reported by previous work, VNMT models are prone to posterior collapse, where the training fails to learn informative latent code as indicated by the value of KL term that vanishes to 0. This phenomenon is often attributed to the strong generator (Alemi et al., 2018) employed by the models, in which case, the generator’s internal cells carry sufficient information to generate the translation. Significant research effort has been spent to weaken the generator network. Mitigating posterior collapse is crucial for our VNMT model as we employ the Transformer, an even stronger generator that comes with more direct connections between source and target sentences (Bahuleyan et al., 2018).

To remedy these issues, we adopt the  $\beta_C$ -VAE (Prokhorov et al., 2019) and compute the following modified KL term:  $\beta |KL - C|$  where  $\beta$  is the scaling factor while  $C$  is a rate to control the KL magnitude. When  $C > 0$ , the models are discouraged from ignoring the latent code. In our experiments, we set  $C = 0.1$  and  $\beta = 1$ . Additionally, we apply the standard practice of word dropping in our experiments.

**Related Work** VNMT comes in two flavors. The first variant models the conditional probability akin to a translation model, while the second one models the joint probability of the source and target sentences. Our model adopts the first variant similar to (Zhang et al., 2016; Su et al., 2018; Pagnoni et al., 2018), while (Eikema and Aziz, 2018; Shah and Barber, 2018) adopt the second variant. The majority of VNMT models employ RNN-based generators and assume isotropic Gaussian distribution, except for (McCarthy et al., 2019) and (Przystupa et al., 2019). The former employs the Transformer architecture but assumes a Gaussian posterior, while the latter employs the normalizing flows posterior (particularly planar flows) but uses an RNN-based generator. We combine more sophisticated normalizing flows and the more powerful Transformer architecture to produce state-of-the-art results.

### 3 Experimental Results

**Experimental Setup** We integrate our proposal into the Fairseq toolkit (Ott et al., 2019; Gehring et al., 2017a,b). We report results on the IWSLT’14 German-English (De-En) and the

WMT’18 English-German (En-De) tasks. For IWSLT’14, we replicate [Wu et al. \(2019\)](#); [Edunov et al. \(2018\)](#)’s setup with 160K training sentences and a 10K joint BPE vocabulary, while for WMT’18, we replicate [Edunov et al. \(2018\)](#)’s setup with 5.2M training sentences and a 32K joint BPE vocabulary. For WMT experiments, we report the accuracy using detokenized SacreBLEU ([Post, 2018](#)) to facilitate fair comparison with other published results. Note that tokenized BLEU score is often higher depending on the tokenizer, thus not comparable. We apply KL annealing schedule and token dropout similar to ([Bowman et al., 2016](#)), where we set the KL annealing to 80K updates and drop out 20% target tokens in the IWSLT and 10% in the WMT experiments.

The encoder and decoder of our Transformer generator have 6 blocks each. The number of attention heads, embedding dimension, and inner-layer dimensions are 4, 512, 1024 for IWSLT; and 16, 1024, 4096 for WMT. The WMT setup is often referred to as the *Transformer Big*. To our knowledge, these architectures represent the best configurations for our tasks. We set the latent dimension to  $D = 128$ , which is projected using a 1-layer linear map to the embedding space. We report decoding results with beam=5. For WMT experiments, we set the length penalty to 0.6. For all experiments with NF-based posterior, we employ flows of length 4, following the results of our pilot study.

**In-Domain Results** We present our IWSLT results in rows 1 to 6 of Table 1. The accuracy of the baseline Transformer model is reported in row (1), which matches the number reported by [Wu et al. \(2019\)](#). In row (2), we report a static  $Z$  experiment, where  $Z = \text{meanpool}(x)$ . We design this experiment to isolate the benefits of token dropping and utilizing average source embedding as context. As shown, the static  $Z$  provides +0.8 BLEU point gain. In row (3), we report the accuracy of our VNMT baseline when the approximate posterior is a Gaussian, which is +1.3 BLEU point from baseline or +0.5 point from the static  $Z$ , suggesting the efficacy of latent-variable modeling. We then report the accuracy of different variants of our model in rows (4) to (6), where we replace the Gaussian posterior with a cascade of 4 PF, SF and CL, respectively. For SF, we report the result with  $M = 8$  orthogonal columns in row (5).

As shown, these flows modestly add +0.2 to +0.3 points. It is worth noticing that the improvement introduces only around 5% additional parameters.

	System	#params	BLEU
1	Transformer IWSLT	42.9M	34.5
2	+ static $Z$	42.9M	35.3
3	+ $Z \sim$ Gaussian	43.6M	35.8
4	+ $Z \sim$ 4 x PF	44.2M	36.1
5	+ $Z \sim$ 4 x SF (M=8)	45.9M	36.0
6	+ $Z \sim$ 4 x CL	44.3M	36.1
7	(1) + distilled	42.9M	34.9
8	(6) + distilled	44.3M	<b>36.6</b>
9	( <a href="#">Edunov et al., 2018</a> )		29.0
10	Transformer Big	209.1M	28.9
11	+ static $Z$	209.1M	29.0
12	+ $Z \sim$ Gaussian	210.5M	29.1
13	+ $Z \sim$ 4 x PF	211.6M	29.3
14	+ $Z \sim$ 4 x SF (M=8)	215.3M	29.5
15	+ $Z \sim$ 4 x CL	210.6M	29.2
16	(10) + distilled	209.1M	29.2
17	(14) + distilled	215.3M	<b>29.9</b>

Table 1: The translation accuracy on the De-En IWSLT’14 task (rows 1-8), the En-De WMT’18 task (rows 10-17). Each task’s best results in the in-domain setting are *italicized*, while the results with added distilled data are in **bold**.

We report our WMT results that use the Transformer Big architecture in rows (10) to (15). For comparison, we quote the state-of-the-art result for this dataset from [Edunov et al. \(2018\)](#) in row (9), where the SacreBLEU score is obtained from [Edunov \(2019\)](#). As shown, our baseline result (row 10) is on par with the state-of-the-art result. The WMT results are consistent with the IWSLT experiments, where our models (rows 13-15) significantly outperform the baseline, even though they differ in terms of which normalizing flows perform the best. The gain over the VNMT baseline is slightly higher, perhaps because NF is more effective in larger datasets. In particular, we found that SF and PF perform better than CL, perhaps due to their simpler architecture, i.e., their posteriors are conditioned only on the source sentence, and their priors are uninformed Gaussian. Row (11) shows that the static  $Z$ ’s gain is minimal. In row (14), our best VNMT outperforms the state-of-the-art Transformer Big model by +0.6 BLEU while adding only 3% additional parameters.

**Simulated Bimodal Data** We conjecture that the gain partly comes from NF’s ability to capture non-Gaussian distribution. To investigate this, we artificially increase the modality of our training data, i.e., forcing all source sentences to have multiple translations. We perform the sequence-level knowledge distillation (Kim and Rush, 2016) with baseline systems as the teachers, creating additional data referred to as *distilled* data. We then train systems on this augmented training data, i.e., original + distilled data. Rows (7) and (16) show that the baseline systems benefit from the distilled data. Rows (8) and (17) show that our VNMT models gain more benefit, resulting in +2.1 and +0.9 BLEU points over non-latent baselines on IWSLT and WMT tasks respectively.

**Simulated Out-of-Domain Condition** We investigate whether the in-domain improvement carries to out-of-domain test sets. To simulate an out-of-domain condition, we utilize our existing setup where the domain of the De-En IWSLT task is TED talks while the domain of the En-De WMT task is news articles. In particular, we invert the IWSLT De-En test set, and decode the English sentences using our baseline and best WMT En-De systems of rows (10) and (14). For this inverted set, the accuracy of our baseline system is 27.9, while the accuracy of our best system is 28.8, which is +0.9 points better. For reference, the accuracy of the Gaussian system in row (11) is 28.2 BLEU. While more rigorous out-of-domain experiments are needed, this result gives a strong indication that our model is relatively robust for this out-of-domain test set.

**Translation Analysis** To better understand the effect of normalizing flows, we manually inspect our WMT outputs and showcase a few examples in Table 2. We compare the outputs of our best model that employs normalizing flows (VNMT-NF, row 14) with the baseline non-latent Transformer (row 10) and the baseline VNMT that employs Gaussian posterior (VNMT-G, row 12).

As shown, our VNMT model consistently improves upon gender consistency. In example 1, the translation of *the interior decorator* depends on the gender of its cataphora (*her*), which is feminine. While all systems translate the cataphora correctly to *ihrem*, the baseline and VNMT-G translate the

Example 1 Source	In <i>her</i> book , <i>the interior decorator</i> presents 17 housing models for independent living in old age .
Reference	In <i>ihrem</i> Buch stellt <i>die Innenarchitektin</i> 17 Wohnmodelle für ein selbstbestimmtes Wohnen im Alter vor .
Non-latent Baseline	In <i>ihrem</i> Buch präsentiert <i>der Innenarchitekt</i> 17 Wohnmodelle für ein unabhängiges Leben im Alter .
VNMT-G	In <i>ihrem</i> Buch stellt die <i>der Innenarchitekt</i> 17 Wohnmodelle für ein selbstbestimmtes Wohnen im Alter vor .
VNMT-NF	In <i>ihrem</i> Buch präsentiert <i>die Innen-dekuratorin</i> 17 Wohnmodelle für ein unabhängiges Leben im Alter .
Example 2 Source	Even though <i>she</i> earns S 3,000( 2,400 ) a month <i>as an administrator</i> and her husband works as well , the monthly family income is insufficient , she says .
Reference	Obwohl <i>sie</i> jeden Monat 3.000 Singapur-Dollar ( ca 1.730 Euro ) <i>als Verwaltungsmitarbeiterin</i> verdiene –truncated–
Non-latent Baseline	Obwohl <i>sie</i> pro Monat 3.000 S \$ ( 2.400 \$ ) <i>als Verwalter</i> verdient und auch ihr Mann arbeitet , ist das –truncated–
VNMT-G	Obwohl <i>sie</i> jeden Monat 3.000 Singapur -Dollar ( ca 1.730 Euro ) <i>als Verwaltungsmitarbeiterin</i> –truncated–
VNMT-NF	Obwohl <i>sie</i> S \$ 3.000 ( \$ 2.400 ) pro Monat <i>als Administratorin</i> verdient und ihr Mann auch –truncated–

Table 2: Translation examples with different gender consistency. Inconsistent, consistent translations and source words are in *red*, *orange*, *blue* respectively.

phrase to its masculine form. In contrast, the translation of our VNMT-NF produces the feminine translation, respecting the gender agreement. In example 2, only VNMT-NF and VNMT-G produce gender consistent translations.

## 4 Discussions and Conclusions

We present a Variational NMT model that outperforms a strong state-of-the-art non-latent NMT model. We show that the gain modestly comes from the introduction of a family of flexible distribution based on normalizing flows. We also demonstrate the robustness of our proposed model in an increased multimodality condition and on a simulated out-of-domain test set.

We plan to conduct a more in-depth investigation into actual multimodality condition with high-coverage sets of plausible translations. We conjecture that conditioning the posterior on the target sentences would be more beneficial. Also, we plan to consider more structured latent variables beyond modeling the sentence-level variation as well as to apply our VNMT model to more language pairs.

## References

- Alexander A. Alemi, Ben Poole, Ian Fischer, Joshua V. Dillon, Rif A. Saurous, and Kevin Murphy. 2018. **Fixing a broken ELBO**. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 159–168.
- Dzmitry Bahdanau, KyungHyun Cho, and Yoshua Bengio. 2015. **Neural machine translation by jointly learning to align and translate**. In *International Conference on Representation Learning (ICLR)*, San Diego, USA.
- 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, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Rianne van den Berg, Leonard Hasenclever, Jakub M. Tomczak, and Max Welling. 2018. **Sylvester normalizing flows for variational inference**. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, Tel Aviv-Yafo, Israel.
- Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew 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, Berlin, Germany. Association for Computational Linguistics.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. 2017. **Density estimation using Real NVP**.
- Sergey Edunov. 2019. <https://github.com/pytorch/fairseq/issues/506#issuecomment-464411433>.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. **Understanding back-translation at scale**. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 489–500, Brussels, Belgium.
- Bryan Eikema and Wilker Aziz. 2018. **Auto-encoding variational neural machine translation**. *arXiv:1807.10564v1*.
- Jonas Gehring, Michael Auli, David Grangier, and Yann N. Dauphin. 2017a. **A convolutional encoder model for neural machine translation**. In *Conference of the Association for Computational Linguistics (ACL)*, Vancouver, Canada.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017b. **Convolutional sequence to sequence learning**. In *International Conference on Machine Learning (ICML)*, Sydney, Australia.
- Yoon Kim and Alexander M. Rush. 2016. **Sequence-level knowledge distillation**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- Diederik P. Kingma and Max Welling. 2014. **Auto-encoding variational bayes**. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- Arya D. McCarthy, Xian Li, Jiatao Gu, and Ning Dong. 2019. **Improved variational neural machine translation by promoting mutual information**. <https://arxiv.org/abs/1909.09237>.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. **fairseq: A fast, extensible toolkit for sequence modeling**. In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Artidoro Pagnoni, Kevin Liu, and Shangyan Li. 2018. **Conditional variational autoencoder for neural machine translation**. *CoRR*, abs/1812.04405.
- Matt Post. 2018. **A call for clarity in reporting BLEU scores**. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Victor Prokhorov, Ehsan Shareghi, Yingzhen Li, Mohammad Taher Pilehvar, and Nigel Collier. 2019. **On the importance of the Kullback-Leibler divergence term in variational autoencoders for text generation**. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 118–127, Hong Kong. Association for Computational Linguistics.
- Michael Przystupa, Mark Schmidt, and Muhammad Abdul-Mageed. 2019. **Investigating the impact of normalizing flows on latent variable machine translation**. In *Proceedings of The Workshop on Invertible Neural Nets and Normalizing Flows. International Conference of Machine Learning*.
- Danilo Rezende and Shakir Mohamed. 2015. **Variational inference with normalizing flows**. In *International Conference on Machine Learning (ICML)*, pages 1530–1538, Lille, France.
- 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.
- Jinsong Su, Shan Wu, Deyi Xiong, Yaojie Lu, Xianpei Han, and Biao Zhang. 2018. **Variational recurrent neural machine translation**. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications*

of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018, pages 5488–5495.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Neural Information Processing Systems Conference (NIPS)*, pages 5998–6008, Long Beach, USA.

Felix Wu, Angela Fan, Alexei Baevski, Yann N. Dauphin, and Michael Auli. 2019. [Pay less attention with lightweight and dynamic convolutions](#). In *International Conference on Learning Representations (ICLR)*, New Orleans, USA.

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, Austin, Texas. Association for Computational Linguistics.

## A Word dropout

We investigate the effect of different dropout rate and summarize the results in Table 3. In particular, we take the VNMT baseline with Gaussian latent variable for IWSLT (row 3 in Table 1) and for WMT (row 12 in Table 1). As shown, word dropout is important for both setup but it is more so for IWSLT. It seems that tasks with low resources benefit more from word dropout. We also observe that above certain rate, word dropout hurts the performance.

Dropout rate	0.0	0.1	0.2	0.3
IWSLT	34.4	35.7	<b>35.8</b>	35.6
WMT	29.0	<b>29.1</b>	28.8	28.7

Table 3: Results of different dropout rate for IWSLT and WMT setup. The best results are in **bold**.

## B Latent Dimension

We report the results of varying the dimension of latent variable ( $D$ ) in Table 4. For this study, we use the VNMT baseline with Gaussian latent variable in IWSLT condition (row 3 in Table 1). Our experiments suggest that the latent dimension between 64 and 128 is optimal. The same conclusion holds for the WMT condition.

$D$	8	16	32	64	128	256
BLEU	35.6	35.5	35.4	35.7	<b>35.8</b>	35.4

Table 4: Results of different dropout rate for IWSLT. The best results are in **bold**.

## C Normalizing Flow Configuration

In the Experimental Results section, we report the accuracy for our models with 4 flows. In Table 5, we conduct experiments varying the number of flows for the IWSLT condition. Our baseline (num flows=0) is an NMT model with word dropout, which performs on par with the static  $Z$  experiment reported in Table 1’s row 3. These results suggest that increasing the number of flows improves accuracy, but the gain diminishes after 4 flows. The results are consistent for all normalizing flows that we considered. We also conduct experiments with employing more flows, but unfortunately, we observe either unstable training or lower accuracy.

Num Flows	PF	SF ( $M=8$ )	CL
0	35.3		
1	35.8	35.6	35.8
2	35.7	35.5	35.8
3	36.0	35.9	35.7
4	<b>36.1</b>	36.0	<b>36.1</b>
5	35.9	<b>36.1</b>	35.9
6	35.8	36.0	35.9

Table 5: Translation accuracy of VNMT models employing various number of flows in the IWSLT condition. The best results are in **bold**.

In Table 6, we conduct experiments varying the number of orthogonal columns ( $M$ ) in our Sylvester normalizing flows (SF) experiments. As shown, increasing  $M$  improves the accuracy up to  $M = 24$ . We see no additional gain from employing more additional orthogonal columns beyond 24. In Table 1, we report  $M = 8$ , because it introduces the least number of additional parameters.

$M$	2	4	8	16	24	32
BLEU	35.7	35.5	36.0	36.0	<b>36.2</b>	35.9

Table 6: Results of different number of orthogonal columns for SF. The best results are in **bold**.