

# Do sequence-to-sequence VAEs learn global features of sentences?

**Tom Bosc**

Mila, Université de Montréal  
bosct@mila.quebec

**Pascal Vincent**

Mila, Université de Montréal, CIFAR  
vincentp@iro.umontreal.ca

## Abstract

Autoregressive language models are powerful and relatively easy to train. However, these models are usually trained without explicit conditioning labels and do not offer easy ways to control global aspects such as sentiment or topic during generation. Bowman et al. (2016) adapted the Variational Autoencoder (VAE) for natural language with the sequence-to-sequence architecture and claimed that the latent vector was able to capture such global features in an unsupervised manner. We question this claim. We measure which words benefit most from the latent information by decomposing the reconstruction loss per position in the sentence. Using this method, we find that VAEs are prone to memorizing the first words and the sentence length, producing local features of limited usefulness. To alleviate this, we investigate alternative architectures based on bag-of-words assumptions and language model pretraining. These variants learn latent variables that are more global, i.e., more predictive of topic or sentiment labels. Moreover, using reconstructions, we observe that they decrease memorization: the first word and the sentence length are not recovered as accurately than with the baselines, consequently yielding more diverse reconstructions.

## 1 Introduction

The problem of generating natural language underlies many classical NLP tasks such as translation, summarization, paraphrasing, etc. The problem is often formulated as learning a probabilistic model of sentences, then searching for probable sentences under this model. Expressive language models are typically built using neural networks (Bengio et al., 2003; Mikolov et al., 2010).

Whether based on LSTMs (Hochreiter and Schmidhuber, 1997; Sundermeyer et al., 2012) or Transformers (Vaswani et al., 2017; Radford et al.,

2019), language models are mostly autoregressive: the probability of a sentence is the product of the probability of each word given the previous words. By contrast, Bowman et al. (2016) built a Variational Autoencoder (VAE) (Kingma and Welling, 2013; Rezende et al., 2014) out of a sequence-to-sequence architecture (*seq2seq*) (Sutskever et al., 2014). It generates text in a two-step process: first, a latent vector is sampled from a prior distribution; then, words are sampled from the probability distribution produced by the autoregressive decoder, conditionally on the latent vector. The goal was to encourage a useful information decomposition, where latent vectors would “explicitly model holistic properties of sentences such as style, topic, and high-level syntactic features” (Bowman et al., 2016), while the more local correlations would be handled by the recurrent decoder.

In principle, such a decomposition can be the basis for many applications. For example, using a single, unannotated corpus, it could enable paraphrasing (Roy and Grangier, 2019) or style transfer (Xu et al., 2019). For tasks requiring conditional generation such as machine translation or dialogue modeling, we could enforce a level of formality or impose a certain tone by clamping the latent vector. Moreover, latent-variable models can represent multimodal distributions. Thus, for these conditional tasks, the latent variable can be used as a source of stochasticity to ensure more diverse translations (Pagnoni et al., 2018) or answers in a dialogue (Serban et al., 2017).

Despite its conceptual appeal, Bowman et al. (2016)’s VAE suffers from the *posterior collapse* problem: early on during training, the KL term in the VAE optimization objective goes to 0, such that the approximate posterior becomes the prior and no information is encoded in the latent variable. Free bits are a popular workaround (Kingma et al., 2016) to ensure that the KL term is above a cer-

tain level, thereby enforcing that *some* information about the input is encoded. But this information is not necessarily global. After all, posterior collapse can be solved trivially, without any learning, using encoders that copy parts of the inputs in the latent variable, yielding very local and useless features.

In Section 3, we show that encoders learn to partially memorize the first few words and the document lengths, as was first discovered by Kim et al. (2018). To do so, we compare the average values of the reconstruction loss at different positions in the sentence to that of an unconditional language model. We elaborate on the negative consequences of this finding for generative models of texts. In Section 4, we propose three simple variants of the model and the training procedure, in order to alleviate memorization and to yield more useful global features. In Section 5, we empirically confirm that our variants produce more global features, i.e., features more predictive of global aspects of documents such as topic and sentiment. They do so while memorizing the first word and the sentence length less often, as shown in Section 6.

## 2 Model and datasets

Firstly, we describe the VAE based on the seq2seq architecture of Bowman et al. (2016). A document, sentence or paragraph, of  $L$  words  $x = (x_1, \dots, x_L)$  is embedded in  $L$  vectors  $(e_1, \dots, e_L)$ . An LSTM encoder processes these embeddings to produce hidden states:

$$h_1, \dots, h_L = \text{LSTM}(e_1, \dots, e_L)$$

In general, the encoder produces a vector  $r$  that represents the entire document. In the original model, this vector is the hidden state of the last word  $r = h_L$ , but we introduce variants later on. This representation is transformed by linear functions  $L_1$  and  $L_2$ , yielding the variational parameters that are specific to each input document:

$$\begin{aligned} \mu &= L_1 r \\ \sigma^2 &= \exp(L_2 r) \end{aligned}$$

These two vectors of dimension  $d$  fully determine the approximate posterior, a multivariate normal with a diagonal covariance matrix,  $q_\phi(z|x) = \mathcal{N}(z|\mu, \text{diag}(\sigma^2))$ , where  $\phi$  is the set of all encoder parameters (the parameters of the LSTM,  $L_1$  and  $L_2$ ). Then, a sample  $z$  is drawn from the approximate posterior, and the decoder, another LSTM,

produces a sequence of hidden states:

$$h'_1, \dots, h'_L = \text{LSTM}([e_{\text{BOS}}; z], [e_1; z], \dots, [e_L; z])$$

where BOS is a special token indicating the beginning of the sentence and  $[\cdot; \cdot]$  denotes the concatenation of vectors. Finally, each hidden state at position  $i$  is transformed to produce a probability distribution of the word at position  $i + 1$ :

$$p_\theta(x_{i+1}|x_{1,\dots,i}, z) = \text{softmax}(Wh'_i + b)$$

where  $\text{softmax}(v_i) = e^{v_i} / \sum_j e^{v_j}$  and  $\theta$  is the set of parameters of the decoder (the parameters of the LSTM decoder,  $W$  and  $b$ ). An EOS token indicating the end of the sentence is appended to every document.

For each document  $x$ , the lower-bound on the marginal log-likelihood (*ELBo*) is:

$$\begin{aligned} \text{ELBo}(x, \phi, \theta) &= -D_{\text{KL}}(q_\phi(z|x)||p(z)) + \\ &\quad \mathbb{E}_{q_\phi}[\log p_\theta(x|z)] \\ &\leq \log p(x) \end{aligned}$$

On the entire training set  $\{x^{(1)}, \dots, x^{(N)}\}$ , the objective is:

$$\arg \max_{\phi, \theta} \sum_{j=1}^N \text{ELBo}(x^{(j)}, \phi, \theta)$$

### 2.1 Dealing with posterior collapse

Following Alemi et al. (2018), we call the average value of the KL term the *rate*. It measures how much information is encoded on average about the datapoint  $x$  by the approximate posterior  $q_\phi(z|x)$ . When the rate goes to 0, the posterior is said to *collapse*, meaning that  $q_\phi(z|x) \approx p(z)$  and that the latent variable  $z$  sampled to train the decoder does not contain any information about the input  $x$ .

To prevent this, we can modify the KL term to make sure it is above a target rate using a variety of techniques (see Appendix A.1 for a small survey). We use the free bits formulation of the  $\delta$ -VAE (Razavi et al., 2019). For a desired rate  $\lambda$ , the modified negative ELBo is:

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

Seq2seq VAEs are prone to posterior collapse, so in practice, the rates obtained are very close to the target rates  $\lambda$ .

As observed by Alemi et al. (2018), different models or sets of hyperparameters for a given

model can yield very similar values of ELBos despite reaching very different rates. Thus, for our purposes, the free bits modification is also useful to compare models with similar capacity.

## 2.2 Variants

Throughout the paper, we use variants of the original architecture and training procedure. In general, these variants use free bits objectives, but reach lower perplexities than what free bits alone allow.

Li et al. (2019)’s method is the following: pre-train an *AE*, reinitialize the weights of the decoder, train the entire model again end-to-end with the VAE objective. The sentence representation  $r$  is also the last hidden state of the LSTM encoder, so we call this method *last-PreAE*.

In the second variant, proposed by Long et al. (2019), the representation of the document  $r$  is the component-wise maximum over hidden states  $h_i$ , i.e.,  $r^j = \max_i h_i^j$ . We call this model *max*. In later experiments, we also consider a hybrid of the two techniques, *max-PreAE*.

We chose these two baselines because they are relatively recent and outperform or perform on par with other recent methods such that cyclical learning rates (Fu et al., 2019) or aggressive training (He et al., 2019). Moreover, the pooling encoder of Long et al. (2019) is particularly interesting: since pooling operators aggregate information over sets of vectors, they might prevent the copying of local information in the latent variable.

We make slight, beneficial modifications to these two methods. We remove KL annealing, which is not only redundant with the free bits technique but also increases the rate erratically (Pelsmaeker and Aziz, 2019). Moreover, for Li et al. (2019)’s method, we use  $\delta$ -VAE-style free bits instead of the original free bits to get rid of the unnecessary constraint that the free bits be balanced across components. For more details, see Appendix A. In summary, all of our experiments use  $\delta$ -VAE-style free bits without KL annealing.

Finally, *AE* denotes the deterministic autoencoder trained only with the reconstruction loss.

## 2.3 Datasets

We train VAEs on four small versions of the AG-News, Amazon, Yahoo, and Yelp datasets created by Zhang et al. (2015). Each document is written in English and consists of one or several sentences. Each document is labeled manually according to its main topic or the sentiment it expresses, and

the labels are close to uniformly balanced over all the datasets. For faster training, we use smaller datasets. The characteristics of these datasets are detailed in Table 5 in the Appendix.

## 3 Encoders partially memorize the first words and sentence length

The ELBo objective trades off the KL term against the reconstruction term. To minimize the objective, it is worth increasing the KL term only if the reconstruction term is decreased by the same amount or more. With free bits, the encoder is allowed to store a fixed amount of information for free. The objective becomes to minimize the reconstruction cost using the “free storage” as efficiently as possible.

There are many solutions to this objective that are undesirable. For instance, we could program an encoder that would encode the words into the latent variable losslessly, until all the free bits are used. However, this model would not be more useful than a standard, left-to-right autoregressive models. Therefore, it is necessary to check that such useless, purely local features are not learned.

In order to visualize *what* information is stored in the latents, our method is to look at *where* gains are seen in the reconstruction loss. Since the loss is a sum over documents and positions in these documents, these gains could be concentrated: i) on certain documents, for example, on large documents or documents containing rarer words; ii) at certain positions in the sentence, for example, in the beginning or in the middle of the sentence. We investigate the latter possibility.<sup>1</sup>

### 3.1 Visualizing the reconstruction loss

Concretely, we compare the reconstruction loss of different models at different positions in the sentence. The baseline is a LSTM trained with a language model objective (*LSTM-LM*). It has the same size as the decoders of the autoencoder models.<sup>2</sup> Since the posterior collapse makes VAEs behave exactly like the *LSTM-LM*, the reconstruction losses between the VAEs and the *LSTM-LM* are directly comparable. Additionally, the deterministic *AE* gives us the reconstruction error that is reachable with a latent space constrained only by its dimension  $d$ , but not by any target rate  $\lambda$  (equivalent to an infinite target rate).

<sup>1</sup>PyTorch (Paszke et al., 2019) implementation available at <https://github.com/tombosc/exps-s2svae>.

<sup>2</sup>Only the input dimensions slightly change because in VAEs, the inputs of the decoder also include the latent vector.




Figure 1: *Left*: Reconstruction loss on Yahoo dataset per each position in the sentence, averaged over sentences of 15 words (error bars: min, max on 3 runs); *Right*: Relative improvement compared to baseline LSTM. Seq2seq autoencoders consistently store information about the first couple of words as well as the sentence length in priority.

In Figure 1, the left-hand side plot shows the reconstruction losses of different models and different target rates  $\lambda$  on the Yahoo dataset. As expected, for all models, raising the target rate lowers the reconstruction cost. Remarkably, these gains are very focused around the beginning and the end of documents. For a clearer picture of the gains at the end of the sentence, we plot the *relative improvement* in reconstruction with respect to the baseline (right-hand side of Figure 1) using:

$$\tilde{r}(i) = \frac{\max(r_{\text{LSTM}}(i) - r(i), 0)}{r_{\text{LSTM}}(i)}$$

where  $r_{\text{LSTM}}(i)$  is the loss of the LSTM.

All the models reconstruct the first couple of words and the penultimate token better than the *LSTM-LM*. On the three other datasets, there are similar peaks on relative improvements in the beginning and the end of sentences (Appendix B.1).

It is not obvious that a lower reconstruction at a given position corresponds to information stored about the word in that position in the latent vector. Indeed, words are not independently modeled. However, we argue that it is roughly the case because the decoder is factorised from left-to-right and because correlations between words decrease with their distance in the sentence. The argument is detailed in the Appendix B.2.

How much do these gains on the reconstruction loss translate to decoding the first words and the document lengths more accurately? To find out, we compare regular VAEs to fixed-encoder, ideal VAEs that encode the true label perfectly and exclusively (in other words, VAEs whose latent variable is the ground-truth label). On sentence reconstruction, we found that regular VAEs decoded the first word 2 to 5 times more often than the baselines,

indicating memorization of the first word. We also found similar but less dramatic results for sentence length (see Appendix B.3 for details).

This phenomenon was already noticed by Kim et al. (2018), using a different method (saliency measures, see Appendix D.2 for details).

To sum up, compared to an unconditional *LSTM-LM*, the seq2seq VAEs incur a much lower reconstruction loss on the first tokens and towards the end of the sentence (around 50% less with  $\lambda = 8$ ). Moreover, if the latent variable of the VAEs did encode the label perfectly and exclusively, they would reconstruct the first words or recover sentence length with much lower accuracy than what is observed. Therefore, we conclude that seq2seq VAEs are biased towards memorizing the first few words and the sentence length.

### 3.2 The problem with memorization

One could argue that this is a superficial problem, as we can always give the model more free bits and decrease the loss in intermediary positions. However, this is not so simple because increasing capacity leads to a worse model fit, as was noted by Alemi et al. (2018). More specifically, on text data, Prokhorov et al. (2019) noted that the coherence of samples decreases as the target rate increases. Pelsmaeker and Aziz (2019) reported similar findings, and also, that more complex priors or posteriors do not help. Therefore, given current techniques, higher rates come at the cost of worse modeling of the data and therefore, we should strive for latent-variable models that store less information, but more global information.

Secondly, for controllable generation, conditioning on memorized information is useless. When the first words are encoded in the latent variable,

the factorization of the VAE becomes the same as that of the usual autoregressive models, which are naturally able to continue a given beginning of the sentence (a “prompt”). Similarly, document length is easily controlled by stopping the sampling after producing the desired number of words.<sup>3</sup> Finally, even for semi-supervised learning, a classifier that would only use the first few words and the sentence length would be suboptimal.

If these arguments are correct, it is doubtful that common seq2seq VAE architectures and training procedures in the low-capacity regime would learn useful representations. This is precisely the third problem: most of the KL values reported in the literature are low.<sup>4</sup> Therefore, it is not clear whether the reported gains in performance (however measured) are significant, and if they are, what exactly cause these gains.

## 4 Improving existing models

What architectures could avoid learning to memorize? We investigate simple variants and for a more thorough comparison with existing models, we refer to Appendix D.1.

Our first variant uses a simple bag-of-words (*BoW*) encoder in place of the LSTM encoder. The sentence representation is  $r^j = \max_i e_i^j$ , where the exponents denote components, and the indices denote positions in the sentence. We call it *BoW-max-LSTM*. It is similar to the max-pooling model of Long et al. (2019) except that the maximum is taken over embeddings rather than LSTM hidden states. As Long et al. (2019) reported, the max-pooling operator is better than the average operator, both when the encoder is a LSTM and *BoW* (possibly because the maximum introduces a non-linearity). Therefore, we use the maximum operator. A priori, we think that since word order is not provided to the encoder, the encoder should be unable to memorize the first words.

For our second variant we use a unigram decoder (*Uni*) in place of an LSTM decoder. It produces a single output probability distribution

for all positions in the sentence  $i$ , conditioned only on the latent variable  $z$ . This distribution is obtained by applying a one-hidden layer MLP followed by softmax to the latent vector:  $p_\theta(x_i|z) = \text{softmax}(W_2 \text{ReLU}(W_1 z) + b)$ , where  $\text{ReLU}(x) = \max(x, 0)$  (Nair and Hinton, 2010). We hope that the encoder will learn representations that do not focus on the first words, because the decoder should not need this particular information. We can use any encoder in combination of this decoder and if we use a *BoW* encoder, we obtain the NVDM model of Miao et al. (2016).

Both the *BoW* encoders and *Uni* decoders variants might benefit from the *PreAE* pretraining technique, but we leave this for future work.

Lastly, the pretrained LM (*PreLM*) variant is obtained in two training steps. First, we pretrain a *LSTM-LM*. Then, it is used as an encoder with fixed weights. We use average pooling over the hidden states to get a sentence representation, i.e.,  $r = \frac{1}{L} \sum_{i=1}^L h_i$ , and learn the transformations  $L_1$  and  $L_2$  that compute the variational parameters. Initially, we tried to use max-pooling but the training was extremely unstable. The LM objective requires the hidden state to capture both close correlations between words but also more global information to predict long-distance correlations. The hope is that this global information can be retrieved via pooling and encoded in the variational parameters. The *PreLM* variant is nothing more than the use of a pretrained LM as a feature extractor (Peters et al., 2018). While Yang et al. (2017) and Kim et al. (2018) both consider the use of pretrained LMs as encoders, the weights are not frozen such that it is hard to disentangle the impact of pretraining from subsequent training. In contrast, we freeze the weights so that the effect of pretraining can not be overridden. To isolate the effect of this training procedure independently of the architecture, we keep the same LSTM instead of using more powerful architectures such as Transformers.

## 5 Semi-supervised learning evaluation

We turn to the semi-supervised learning (*SSL*) setting to compare the learned representations of our variants. For the purpose of controllable text generation, we assume that the global information that is desirable to capture is the topic or sentiment. There are two training phases: first, an unsupervised pretraining phase where VAEs are trained; second, a supervised learning phase where classi-

<sup>3</sup>Or by explicitly conditioning on the sentence length. It can be useful for unsupervised summarization (Schumann, 2018), in flow-based approaches (Ziegler and Rush, 2019), or more broadly for the decoder to plan sentence construction.

<sup>4</sup>Most papers do not report the log base (1 bit is  $\ln(2) \approx 0.693$  nats). Here are some reported rates of the best models: Bowman et al. (2015): 2.0 (PTB) ; Long et al. (2019): 3.7 (Yahoo), 3.1 (Yelp); Li et al. (2019): 15.02 (Yahoo), 8.15 (PTB); He et al. (2019): 5.6 (Yahoo), 3.4 (Yelp); Fu et al. (2019): 1.955 (PTB), ...

fiers are trained to predict ground-truth labels given the latent vectors encoded with the encoders of the VAEs. This is essentially the same setup as *MI* from Kingma et al. (2014).<sup>5</sup> The small and large data-regimes give us complementary information: with many labels and complex classifiers, we quantify *how much* of the information pertaining to the labels is encoded; with few labels and simple classifiers, *how accessible* the information is.

For each dataset, we subsample  $g = 5$  balanced labeled datasets for each different data-regimes, containing 5, 50, 500, and 5000 examples per class. These labeled datasets are used for training and validating during the supervised learning phase.<sup>6</sup> Each model is trained with  $s = 3$  seeds. The performance of the classifiers are measured by the macro F1-score on the entire test sets.

To select hyperparameters on each subsample, we use repeated stratified K-fold cross-validation (Moss et al., 2018) as detailed in the Appendix E.1. We obtain the test set F1-scores  $F_{ij}$ , where  $i$  is the subsample seed and  $j$  is the parameter initialisation seed, and report  $\bar{F}_{..}$ , the average F1-score over  $i$  and  $j$ . We note  $\bar{F}_{.j}$  the empirical average F1-score for a given parameter initialisation  $j$  and decompose the variance into:

- $\sigma_{\text{init}} = \left(\frac{1}{s-1} \sum_{j=1}^s g(\bar{F}_{.j} - \bar{F}_{..})^2\right)^{\frac{1}{2}}$ , which quantifies the variability due to the initialisation of the model,
- $\sigma = \left(\frac{1}{g} \sum_{i=1}^g \frac{1}{s-1} \sum_{j=1}^s (F_{ij} - \bar{F}_{.j})^2\right)^{\frac{1}{2}}$ , which quantifies the remaining variability.

In the context of ANOVA with a linear model and a single factor, these quantities are the square roots of  $MS_T$  and  $MS_E$  (see Appendix E.2).

Finally, we also add a data-regime where the entire labeled training set is used in the supervised learning phase. In this setting, we use more expressive one-hidden-layer MLP classifiers, with early stopping on the validation set. Thus, we can check that our conclusions in the large data-regime do not depend on the model selection procedure and the choice of the classifier.

<sup>5</sup>We could integrate the labels into the generative model as a random variable that is either observed or missing to obtain better results (Kingma et al., 2014). Still, our goal is to study the inductive bias of the seq2seq VAE as an unsupervised learning method, so we do not train the encoder using the labels.

<sup>6</sup>It is especially important to use several subsamples in the low data-regimes where subsamples containing unrepresentative texts or noisy labels are not unlikely.

For each class of model, we perform a grid search over target rates and latent vector sizes. We search for target rates  $\lambda$  in  $\{2, 8\}$ : large enough to capture label information but small enough to avoid underfitting, as explained above. The size of latent vectors  $d$  are chosen in  $\{4, 16\}$ . They should be small enough for extremely low-data regimes. For instance, on Yelp, the smallest data regime (5 per class) uses only 8 examples to train the classifier and 2 to do cross-validation. A thorough explanation is presented in Appendix C, along with the values of hyperparameters held constant.

What representation should be used as inputs to the classifiers? Kingma et al. (2016) use samples from the approximate posterior  $q_\phi(z|x) = \mathcal{N}(z|\mu, \text{diag}(\sigma^2))$ , but in the NLP literature, most evaluations focus on  $\mu$  without mention or justification. To evaluate the VAE as a generative model, we claim that only noisy samples  $z$  should be used. In fact, using a model with a rate close to 0 on Yelp, we can recover the label with a high F1-score of 81.5% by using  $\mu$ , whereas, as expected, noisy samples  $z$  do not do better than random (50%). The information contained in  $\mu$  is misleading because it is not transmitted to the decoder and not used directly during generation. Therefore, we use samples  $z$  (cf. Appendix E.3 for details).

## 5.1 Results

Table 1 contains the results of the SSL experiments in the smallest and largest data-regimes. The results for intermediary data-regimes as well as for baseline models without pretraining, which underperform, are presented in the Appendix, Table 8. The proposed variants are either on par or improve significantly over the baselines. In the large data-regime, *BoW-max-LSTM* and *LSTM-avg-LSTM-PreLM* perform best on average while *LSTM-last-Uni* performs the worst and suffers from unstable training on AGNews. In the small data-regime, the picture is less clear because there is more variance.

On AGNews and Yelp, in the large data-regime, our variants do not seem to improve over the baselines. However, on Amazon and Yahoo, in the large data-regime, the variants seem to improve by 5 in F1-score. Why do the gains vary so widely depending on the datasets? We posit that, on some datasets, the first words are enough to predict the labels correctly. We train bag-of-words classifiers<sup>7</sup> using either i) only the first three words or ii) all

<sup>7</sup>fastText classifiers (Joulin et al., 2017) with embedding

	Enc.	$r$	Dec.	Pre.	5		All	
					F1	$\sigma_{init}$	F1	$\sigma_{init}$
AGNews	LSTM	last	LSTM	AE	65.8	$\frac{3.3}{3.3}$	83.4	$\frac{-}{0.3}$
	LSTM	max	LSTM	AE	55.7	$\frac{4.5}{18.7}$	83.3	$\frac{-}{0.4}$
	BoW	max	LSTM	-	72.7	$\frac{2.0}{5.9}$	83.1	$\frac{-}{0.3}$
	LSTM	max	Uni	-	71.6	$\frac{0.1}{5.5}$	83.9	$\frac{-}{0.3}$
	LSTM	last	Uni	-	54.8	$\frac{5.2}{57.1}$	59.3	$\frac{-}{40.9}$
	BoW	max	Uni	-	71.8	$\frac{4.5}{1.8}$	83.1	$\frac{-}{0.5}$
	LSTM	avg	LSTM	LM	70.8	$\frac{4.8}{4.3}$	83.5	$\frac{-}{0.1}$
Amazon	LSTM	last	LSTM	AE	20.0	$\frac{2.2}{0.9}$	28.1	$\frac{-}{1.0}$
	LSTM	max	LSTM	AE	22.3	$\frac{2.6}{0.7}$	34.0	$\frac{-}{1.6}$
	BoW	max	LSTM	-	21.0	$\frac{2.6}{1.1}$	38.9	$\frac{-}{0.7}$
	LSTM	max	Uni	-	21.8	$\frac{3.1}{1.6}$	38.2	$\frac{-}{0.5}$
	LSTM	last	Uni	-	24.0	$\frac{3.0}{1.0}$	36.8	$\frac{-}{0.9}$
	BoW	max	Uni	-	25.4	$\frac{3.2}{0.2}$	37.9	$\frac{-}{0.2}$
	LSTM	avg	LSTM	LM	21.8	$\frac{3.8}{0.6}$	40.0	$\frac{-}{0.4}$
Yahoo	LSTM	last	LSTM	AE	20.7	$\frac{0.7}{0.5}$	37.2	$\frac{-}{0.7}$
	LSTM	max	LSTM	AE	20.8	$\frac{1.3}{2.3}$	36.6	$\frac{-}{0.7}$
	BoW	max	LSTM	-	23.4	$\frac{2.1}{2.9}$	42.6	$\frac{-}{0.2}$
	LSTM	max	Uni	-	24.9	$\frac{1.3}{2.2}$	38.9	$\frac{-}{1.7}$
	LSTM	last	Uni	-	24.5	$\frac{3.8}{1.7}$	37.1	$\frac{-}{2.3}$
	BoW	max	Uni	-	24.1	$\frac{2.9}{2.7}$	40.1	$\frac{-}{0.7}$
	LSTM	avg	LSTM	LM	21.9	$\frac{2.3}{1.3}$	41.7	$\frac{-}{0.3}$
Yelp	LSTM	last	LSTM	AE	59.3	$\frac{5.4}{2.9}$	67.9	$\frac{-}{0.1}$
	LSTM	max	LSTM	AE	59.9	$\frac{10.4}{7.9}$	84.1	$\frac{-}{0.7}$
	BoW	max	LSTM	-	67.1	$\frac{10.1}{15.7}$	85.0	$\frac{-}{0.2}$
	LSTM	max	Uni	-	62.3	$\frac{4.6}{3.8}$	83.1	$\frac{-}{0.5}$
	LSTM	last	Uni	-	65.0	$\frac{8.0}{4.4}$	81.6	$\frac{-}{0.5}$
	BoW	max	Uni	-	59.9	$\frac{7.2}{3.7}$	83.3	$\frac{-}{0.4}$
	LSTM	avg	LSTM	LM	63.6	$\frac{7.4}{5.4}$	84.4	$\frac{-}{0.5}$

Table 1: Using *BoW* encoders, *Uni* decoders or *PreLM* pretraining, the learned representations are more predictive of the labels (sentiment or topic).

the words as features on the entire datasets. If the *three-words* classifiers are as good as the *all-words* classifiers, we expect that the original VAE variants will perform well: in that case, encoding information about the first words is not harmful, it could be a rather useful inductive bias. Conversely, if the first three words are not predictive of the label, the original VAEs will perform badly.

As reported in the Appendix, Table 9, on AGNews and Yelp, classifiers trained on the first three words have a performance somewhat close to the classifier trained on all the words, reaching 80.8% and 85.4% of its scores respectively. For instance, on AGNews, the first words are often nouns that directly gives the topic of the news item: country names for the politics category, firm names for the technology category, athlete or team names for the sports category, etc. On the two other datasets, the performance decays a lot if we only use the

dimension of 200 and the default parameters.

Enc.	$r$	Pre.	Agree.	1st (%)	Len (%)	$\approx$ PPL
LSTM	last	AE	80.2 $\pm$ 1.0	29.6 $\pm$ 1.1	3.6 $\pm$ 0.1	<b>34.8<math>\pm</math>0.4</b>
LSTM	max	AE	79.5 $\pm$ 0.9	31.7 $\pm$ 1.1	3.7 $\pm$ 0.5	<b>34.7<math>\pm</math>0.4</b>
BoW	max	-	78.0 $\pm$ 1.3	18.9 $\pm$ 1.2	<b>2.7<math>\pm</math>0.3</b>	36.1 $\pm$ 0.6
BoW	max	Uni	81.3 $\pm$ 0.1	<b>13.9<math>\pm</math>0.3</b>	3.1 $\pm$ 0.1	36.3 $\pm$ 0.7
LSTM	max	Uni	<b>82.0<math>\pm</math>0.4</b>	<b>13.9<math>\pm</math>0.2</b>	3.3 $\pm$ 0.4	36.0 $\pm$ 0.4
LSTM	avg	LM	79.2 $\pm$ 0.4	22.2 $\pm$ 0.8	3.2 $\pm$ 0.2	<b>35.0<math>\pm</math>0.3</b>
LSTM	last	AE	24.5 $\pm$ 0.4	42.4 $\pm$ 2.3	13.0 $\pm$ 1.6	<b>44.5<math>\pm</math>0.2</b>
LSTM	max	AE	30.8 $\pm$ 1.1	41.7 $\pm$ 0.8	11.5 $\pm$ 1.0	<b>44.4<math>\pm</math>0.3</b>
BoW	max	-	34.2 $\pm$ 0.5	33.3 $\pm$ 0.7	<b>9.9<math>\pm</math>0.7</b>	45.3 $\pm$ 0.5
BoW	max	Uni	33.3 $\pm$ 0.4	<b>21.5<math>\pm</math>0.3</b>	11.8 $\pm$ 0.5	45.3 $\pm$ 0.4
LSTM	max	Uni	34.1 $\pm$ 0.5	22.1 $\pm$ 0.1	11.7 $\pm$ 0.6	45.4 $\pm$ 0.6
LSTM	avg	LM	<b>35.8<math>\pm</math>0.4</b>	38.3 $\pm$ 0.9	11.5 $\pm$ 1.0	<b>44.2<math>\pm</math>0.4</b>
LSTM	last	AE	23.8 $\pm$ 0.2	56.6 $\pm$ 1.0	17.1 $\pm$ 1.1	<b>48.8<math>\pm</math>0.2</b>
LSTM	max	AE	22.9 $\pm$ 0.8	58.7 $\pm$ 1.7	18.4 $\pm$ 0.8	<b>48.6<math>\pm</math>0.1</b>
BoW	max	-	<b>26.9<math>\pm</math>0.5</b>	49.3 $\pm$ 1.2	11.8 $\pm$ 0.3	49.7 $\pm$ 0.4
BoW	max	Uni	<b>26.8<math>\pm</math>0.6</b>	<b>37.6<math>\pm</math>0.9</b>	<b>10.6<math>\pm</math>0.4</b>	49.8 $\pm$ 0.1
LSTM	max	Uni	<b>27.1<math>\pm</math>1.0</b>	<b>37.7<math>\pm</math>1.6</b>	<b>11.0<math>\pm</math>0.3</b>	50.0 $\pm$ 0.4
LSTM	avg	LM	<b>26.7<math>\pm</math>0.2</b>	51.9 $\pm$ 0.5	16.7 $\pm$ 1.8	<b>48.5<math>\pm</math>0.1</b>
LSTM	last	AE	<b>81.7<math>\pm</math>1.3</b>	53.0 $\pm$ 0.5	33.7 $\pm$ 1.7	<b>31.7<math>\pm</math>0.3</b>
LSTM	max	AE	81.3 $\pm$ 0.7	52.4 $\pm$ 0.5	29.5 $\pm$ 2.5	<b>31.8<math>\pm</math>0.1</b>
BoW	max	-	<b>82.2<math>\pm</math>0.5</b>	36.4 $\pm$ 0.3	22.4 $\pm$ 0.5	32.3 $\pm$ 0.4
BoW	max	Uni	80.4 $\pm$ 0.4	<b>30.6<math>\pm</math>0.5</b>	<b>15.4<math>\pm</math>0.4</b>	32.8 $\pm$ 0.1
LSTM	max	Uni	80.9 $\pm$ 0.4	32.0 $\pm$ 0.4	17.2 $\pm$ 0.7	33.1 $\pm$ 0.3
LSTM	avg	LM	<b>82.3<math>\pm</math>0.7</b>	47.7 $\pm$ 0.4	24.1 $\pm$ 0.4	<b>31.9<math>\pm</math>0.2</b>

Table 2: Our variants reconstruct inputs with higher agreement, less memorization of the 1st words and lengths and a negligible loss in likelihood. Best score and scores within one standard deviation are bolded.

first three words: *three-words* F1-scores make up for 60.7% and 30.3% of *all-words* F1-scores on Amazon and Yahoo. This explains why the original VAE can perform on par or slightly better than our variants on certain datasets for which the first words are very predictive of the labels. This also proves that using several datasets is necessary to draw robust conclusions.

Despite similar asymptotic performance on AGNews and Yelp, our variants clearly improve over the baselines in the small data-regime, which suggests that the encoded information is quantitatively different. This is confirmed in the next section.

It might be surprising that *LSTM-max-LSTM* models are inferior to *BoW-max-LSTM* models. In Appendix E.4, we show that with recurrent encoders, some components of the hidden states are consistently maximized at certain early positions in the sentence. This explains why the power of LSTMs can be undesirable, and why the simpler *BoW* encoders perform better.

## 6 Text generation evaluation

How do these different variants perform during generation? We expect that the SSL classification

performances would correlate with the abilities of the decoders to reconstruct documents that exhibit a similar global aspect than the encoded documents.

To measure the agreement in label between the source document and its reconstruction, we adapt the evaluation procedure used by Ficler and Goldberg (2017) so that no human annotators or heuristics are required (see Appendix D.2). First, a classifier is trained to predict the label on the source dataset. Then, for each model, we encode the documents, reconstruct them, and classify these reconstructions using the classifier. The *agreement* is the F1-scores between the original labels and the labels given by the classifiers on the generated samples.

To quantify memorization, we measure the reconstruction accuracy of the first word and the ratio of identical sentence length between sources and reconstructions. Finally, to verify that our bag-of-words assumptions do not hurt the overall fit to the data, we estimate the negative log-likelihood via the importance-weighted lower bound (Burda et al., 2015) (500 samples) to compute an approximate perplexity per word ( $\approx PPL$ ).

We use two decoding schemes: beam search with a beam of size 5 and greedy decoding. We fix  $\lambda = 8$ ,  $d = 16$  on all models, with three seeds. For the *Uni* decoder, we drop *LSTM-last-Uni* which underperformed by a large margin in the SSL setting, and for the other *Uni* models, we freeze the encoder,  $L_1$  and  $L_2$  and train a new recurrent decoder using the reconstruction loss only. Essentially, the *Uni* decoder is an *auxiliary decoder*, as described by De Fauw et al. (2019) (see Appendix D.1 for details) and we denote this technique by *PreUni*.

Table 2 show the results for beam search decoding.<sup>8</sup> There is a close correspondence between agreement and performance on the SSL tasks in the large data-regime. Our variants have a higher agreement than the baselines, especially on Amazon and Yahoo datasets for which the memorization of the first words is especially harmful.

The baselines reconstruct the first words with very high accuracy (more than 50% of the time on Yahoo and Yelp) while our variants mitigate this memorization. For instance, *PreUni* models recover the first word around 2 or 1.5 times less often.

Let us focus on AGNews and Yelp, where the first words are very predictive of the labels. Both

---

<sup>8</sup>Similar results were obtained using greedy decoding, albeit sometimes consistently shifted.

baselines and variants have roughly similarly high agreement. However, our variants produce more diverse beginnings, while still managing to reproduce the topic or sentiment of the original document. On the other hand, the reconstructions of the baselines exhibit the same labels as the sources mostly as a side-effect of starting with the same words. This also explains that in the SSL setting, despite similar performances asymptotically, our variants were much more efficient using five examples per class. Memorization of the first words does not abstract away from the particular words and therefore, the amount of data required to learn a good classifier will be high, compared to a model which truly *infer* unobserved characteristics of documents.

Both *BoW* encoders and *Uni* decoders lower memorization, so bag-of-words assumptions are efficient for dealing with the memorization problem. Still, *BoW-Max* and *LSTM-Max* with *PreUni* pretraining yield very close performance despite having a different encoder, showing that the decoder has a far greater influence than the encoder. This is consistent with McCoy et al. (2019)’s findings (see Section D.2 in Appendix for details).

Finally, there seems to be a trade-off between the global character of the latent information and the fit of the model to the data. *BoW* and *Uni* variants have perplexity roughly one unit above the baselines, a significant but small difference.

In Appendix F, we perform a qualitative analysis of reconstruction samples to illustrate these conclusions. It also sheds light on the inherent difficulty of the Yahoo dataset.

To recapitulate, the bag-of-words assumptions decrease the memorization of the first word and of the sentence length in the latent variable while increasing the agreement between the labels of the source and of the reconstruction. This is achieved at the cost of a small increase in perplexity.

## 7 Conclusion

Eliminating posterior collapse is necessary to get useful VAE models, but not sufficient. Although recent incarnations of the seq2seq VAE fix the posterior collapse, they partially memorize the first few words and the document lengths. Depending on the data, these local features are sometimes not very correlated with global aspects like topic or sentiment. Therefore, they are of limited use for controllable and diverse text generation.

To learn to infer more global features, we ex-



explored alternative architectures based on bag-of-word assumptions on the encoder or decoder side, as well as a pretraining procedure. These variants are all effective, in particular, the unigram decoder used as an auxiliary decoder (De Fauw et al., 2019). The latent variable is more predictive of global features and memorisation of the first words and sentence length is decreased. Thus, these models are more suitable for diverse and controllable generation.

Methodologically, we introduced a simple way to examine the content of latent variables by looking at the reconstruction loss per position. We also presented a reliable way to perform semi-supervised learning experiments to analyze the content of the variable, free of the problems that one can find in past work (incorrect model selection for small data-regimes, use of samples instead of variational parameters as inputs). We showed that there are particularly difficult datasets for which the first words are not very predictive of their labels, and therefore, these datasets should be systematically used in evaluations. Moreover, the agreement metric is another complementary evaluation that is automatic and focused on generation. We hope that these methods will see widespread adoption for measuring progress more reliably.

A promising research direction is to investigate the root cause behind memorization. A simple reason for the memorization of the first few words could be that, in the beginning of training, the reconstruction loss is higher on these words (see *LSTM-LM* in Figures 1, 2, 3, 4). These early errors should therefore account for a proportionally large part of the gradients and pressure the encoder to store information about the first words. If that is correct, the left-to-right factorization of the decoder could be at fault, which would explain the successes of the unigram decoders. More powerful decoders with alternative factorizations could avoid this issue, for example, non-autoregressive Transformers (Gu et al., 2017) or Transformers with flexible word orders (Gu et al., 2019).

VAEs operate on uncorrupted inputs and learn a corruption process in the latent space. In contrast, models in the BERT family (Devlin et al., 2018) are given corrupted inputs and are penalized only on these corrupted inputs, thereby avoid memorization altogether. Therefore, another research avenue would be to blend the two frameworks (Im et al., 2017).

## Acknowledgements

We thank NSERC for financial support, Calcul Canada for computational resources and Siva Reddy, Loren Lugosch, Makesh Narasimhan and Arjun Akula for their comments on the draft, as well as the reviewers and the meta-reviewer.

## References

- Alexander Alemi, Ben Poole, Ian Fischer, Joshua Dillon, Rif A Saurous, and Kevin Murphy. 2018. Fixing a Broken ELBO. In *International Conference on Machine Learning*, pages 159–168.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155.
- 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 (EMNLP)*. Association for Computational Linguistics.
- 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.
- Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. 2015. Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*.
- Gavin C Cawley and Nicola LC Talbot. 2010. On overfitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11(Jul):2079–2107.
- Xi Chen, Diederik P Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*.
- Jeffrey De Fauw, Sander Dieleman, and Karen Simonyan. 2019. Hierarchical autoregressive image models with auxiliary decoders. *arXiv preprint arXiv:1903.04933*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Jessica Fidler and Yoav Goldberg. 2017. Controlling Linguistic Style Aspects in Neural Language Generation. In *Proceedings of the Workshop on Stylistic Variation*, pages 94–104.

- Hao Fu, Chunyuan Li, Xiaodong Liu, Jianfeng Gao, Asli Celikyilmaz, and Lawrence Carin. 2019. [Cyclical Annealing Schedule: A Simple Approach to Mitigating KL Vanishing](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 240–250, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. 2017. Non-autoregressive neural machine translation. *arXiv preprint arXiv:1711.02281*.
- Jiatao Gu, Qi Liu, and Kyunghyun Cho. 2019. Insertion-based decoding with automatically inferred generation order. *Transactions of the Association for Computational Linguistics*, 7:661–676.
- Junxian He, Daniel Spokoyny, Graham Neubig, and Taylor Berg-Kirkpatrick. 2019. [Lagging Inference Networks and Posterior Collapse in Variational Autoencoders](#). In *International Conference on Learning Representations*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long Short-Term Memory](#). *Neural Comput.*, 9(8):1735–1780.
- Daniel Im Jiwoong Im, Sungjin Ahn, Roland Memisevic, and Yoshua Bengio. 2017. Denoising criterion for variational auto-encoding framework. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. [Bag of Tricks for Efficient Text Classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Yoon Kim, Sam Wiseman, Andrew Miller, David Sonntag, and Alexander Rush. 2018. Semi-Amortized Variational Autoencoders. In *International Conference on Machine Learning*, pages 2683–2692.
- Diederik P. Kingma, Danilo Jimenez Rezende, Shakir Mohamed, and Max Welling. 2014. [Semi-Supervised Learning with Deep Generative Models](#). *CoRR*, abs/1406.5298.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. 2016. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pages 4743–4751.
- Bohan Li, Junxian He, Graham Neubig, Taylor Berg-Kirkpatrick, and Yiming Yang. 2019. [A Surprisingly Effective Fix for Deep Latent Variable Modeling of Text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3601–3612, Hong Kong, China. Association for Computational Linguistics.
- Teng Long, Yanshuai Cao, and Jackie Chi Kit Cheung. 2019. [Preventing Posterior Collapse in Sequence VAEs with Pooling](#). *arXiv:1911.03976 [cs, stat]*. ArXiv: 1911.03976.
- R. Thomas McCoy, Tal Linzen, Ewan Dunbar, and Paul Smolensky. 2019. [RNNs implicitly implement tensor-product representations](#). In *International Conference on Learning Representations*.
- Yishu Miao, Lei Yu, and Phil Blunsom. 2016. Neural Variational Inference for Text Processing. In *International Conference on Machine Learning*, pages 1727–1736.
- 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*.
- Henry B Moss, David S Leslie, and Paul Rayson. 2018. Using JK fold Cross Validation to Reduce Variance When Tuning NLP Models. *arXiv preprint arXiv:1806.07139*.
- Vinod Nair and Geoffrey E Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814.
- Gary W Oehlert. 2010. *A first course in design and analysis of experiments*.
- Artidoro Pagnoni, Kevin Liu, and Shangyan Li. 2018. Conditional variational autoencoder for neural machine translation. *arXiv preprint arXiv:1812.04405*.
- Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up?: sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10*, pages 79–86. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [PyTorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle,

- A. Beygelzimer, F. d'Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Tom Pelsmaeker and Wilker Aziz. 2019. Effective Estimation of Deep Generative Language Models. *arXiv preprint arXiv:1904.08194*.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- 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.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. Technical report.
- Ali Razavi, Aaron van den Oord, Ben Poole, and Oriol Vinyals. 2019. [Preventing Posterior Collapse with delta-VAEs](#). In *International Conference on Learning Representations*.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *International Conference on Machine Learning*, pages 1278–1286.
- Aurko Roy and David Grangier. 2019. Unsupervised paraphrasing without translation. *arXiv preprint arXiv:1905.12752*.
- Raphael Schumann. 2018. Unsupervised abstractive sentence summarization using length controlled variational autoencoder. *arXiv preprint arXiv:1809.05233*.
- Iulian Vlad Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Xing Shi, Kevin Knight, and Deniz Yuret. 2016. [Why Neural Translations are the Right Length](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2278–2282, Austin, Texas. Association for Computational Linguistics.
- Paul Smolensky. 1990. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial intelligence*, 46(1-2):159–216.
- Martin Sundermeyer, Ralf Schlüter, and Hermann Ney. 2012. Lstm neural networks for language modeling. In *Thirteenth annual conference of the international speech communication association*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- Mirac Suzgun, Yonatan Belinkov, Stuart Shieber, and Sebastian Gehrmann. 2019. [LSTM Networks Can Perform Dynamic Counting](#). In *Proceedings of the Workshop on Deep Learning and Formal Languages: Building Bridges*, pages 44–54, Florence. Association for Computational Linguistics.
- 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Paul J Werbos. 1990. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560.
- Peng Xu, Yanshuai Cao, and Jackie Chi Kit Cheung. 2019. Unsupervised controllable text generation with global variation discovery and disentanglement. *arXiv preprint arXiv:1905.11975*.
- Zichao Yang, Zhiting Hu, Ruslan Salakhutdinov, and Taylor Berg-Kirkpatrick. 2017. Improved variational autoencoders for text modeling using dilated convolutions. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3881–3890. JMLR. org.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level Convolutional Networks for Text Classification](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 649–657. Curran Associates, Inc.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. [Learning Discourse-level Diversity for Neural Dialog Models using Conditional Variational Autoencoders](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664, Vancouver, Canada. Association for Computational Linguistics.
- Zachary Ziegler and Alexander Rush. 2019. Latent Normalizing Flows for Discrete Sequences. In *International Conference on Machine Learning*, pages 7673–7682.

## A On the use of KL annealing, the choice of the free bits flavor and resetting the decoder

Li et al. (2019) evaluated their models in the SSL setting (Section 3.3 of their paper). However, their experimental setting is not very rigorous. In the case of the 100 labeled examples, hyperparameter selection is done on a very large validation set of 10000 examples. However, the validation set here should be seen as nothing more than a split of the training data dedicated to optimising hyperparameters. In the words of Cawley and Talbot (2010), “model selection should be viewed as an integral part of the model fitting procedure”. Besides methodological issues, we run our own hyperparameter search on the Yelp dataset to properly disentangle the effects of KL annealing, the free bits method and verify the importance of resetting the decoder. We use the semi-supervised learning setting presented in Section 5 to evaluate the learned encoders.

### A.1 The free bits technique and variants

The *original* free bits objective (Kingma et al., 2016) is the following modification to the KL term:

$$\sum_j^K \max\left(\frac{\lambda}{K}, KL(q_j(z_j|x)||p_j(z_j))\right)$$

where indices denote components. In this formulation, each component of the multivariate normal is allowed to deviate from the prior by a small amount. Instead, in the  $\delta$ -VAE formulation, one component can use of all the  $\lambda$  free bits and the rest of the components can collapse to the prior. This is the variant called  $\delta$ , used throughout the paper:

$$\max(\lambda, KL(q(z|x)||p(z)))$$

Other modifications of the free bits technique include the use of a variable coefficient in front of the KL term (Chen et al., 2016), the target rate objective in Alemi et al. (2018), minimum desired rate (Pelsmaeker and Aziz, 2019), etc. A comparison of all these methods is out of the scope of this paper and the  $\delta$  variant satisfies our only requirement: the rate should be close to the desired rate.

### A.2 KL annealing and the original free bits method higher the rate

Our hypotheses are:

- KL annealing aims at fixing the posterior collapse and is therefore redundant with the free bits,
- KL annealing performs this role by increasing capacity inconsistently across models, making them harder to compare,
- the original free bits formulation impose the unnecessary constraint that the free bits should be balanced over all components.

To study the influence of the free bits variant as well as of KL annealing, we use the same experimental protocol as described in Section 5. To save computations, we fix  $d = 16$ . We do not perform model selection on the desired rate  $\lambda$  in order to see which methods yield the rates that are closest to the desired rate. Table 3 shows these hypotheses are correct. Therefore, all the experiments in the paper use the  $\delta$  variant without annealing.

In Li et al. (2019)’s work, the original, per-component variant of the free bits might have been chosen because it trivially maximizes a metric called *active units* (AU). However, to our knowledge, there is no evidence that this metric should be maximized, neither theoretical nor empirical.

### A.3 On the importance of resetting the decoder after pretraining

Li et al. (2019) proposed to pretrain an AE with a reconstruction loss only. Then, the parameters of the decoder are re-initialised and the (modified) KL term is added to the objective. Since it is not very clear why it would be useful, we studied the impact of this choice. Table 4 shows that it is crucial.

## B Further evidence for memorization

### B.1 Plots on other datasets

Figures 2, 3, and 4 show the reconstruction loss and the relative improvement on other datasets.

On the Yelp dataset, the penultimate token is a punctuation mark which is always followed by the end-of-sentence token, so predicting its position is equivalent to predicting the sentence length. That is why the peak at the end occurs before the last token. Moreover, on Yelp, the situation is worse with  $\lambda = 2$ : between positions 6 and 13, not only is there no improvement, but the reconstruction is higher than that of the baseline.




Figure 2: Reconstruction loss as a function of word position on the AGNews dataset. See Figure 1.




Figure 3: Reconstruction loss as a function of word position on the Amazon dataset. See Figure 1.

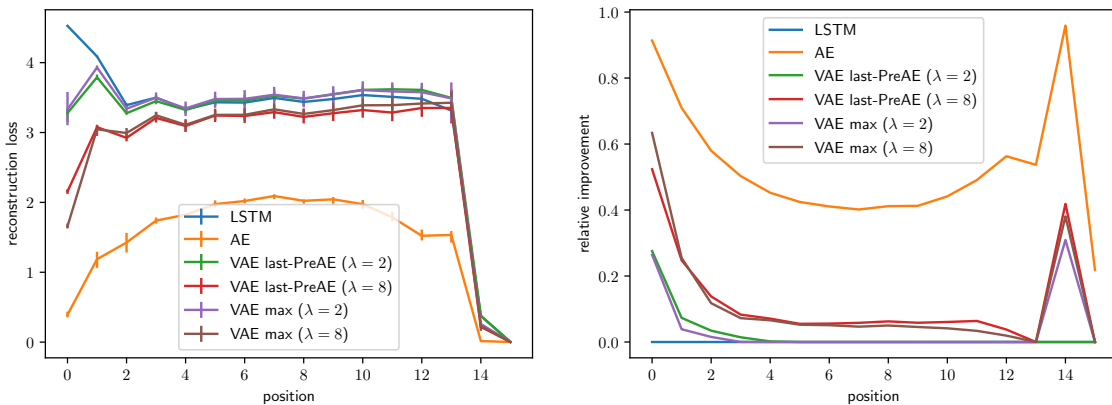


Figure 4: Reconstruction loss as a function of word position on the Yelp dataset. See Figure 1.

FB	$\lambda$	ANN.	F1(5)	F1(50)	F1(500)	F1(5000)	F1(ALL)	KL
O	2	10	53.3 $\pm$ <sub>3.3</sub> <sup>5.5</sup>	69.8 $\pm$ <sub>1.3</sub> <sup>1.8</sup>	73.6 $\pm$ <sub>1.7</sub> <sup>0.2</sup>	74.0 $\pm$ <sub>1.8</sub> <sup>0.1</sup>	73.6 $\pm$ <sub>1.1</sub>	5.27 $\pm$ <sub>0.47</sub>
O	2	0	51.8 $\pm$ <sub>6.7</sub> <sup>4.8</sup>	62.7 $\pm$ <sub>3.8</sub> <sup>2.5</sup>	67.0 $\pm$ <sub>5.6</sub> <sup>0.4</sup>	67.5 $\pm$ <sub>5.8</sub> <sup>0.1</sup>	66.9 $\pm$ <sub>2.7</sub>	2.58 $\pm$ <sub>0.46</sub>
$\delta$	2	10	51.7 $\pm$ <sub>4.7</sub> <sup>4.6</sup>	64.5 $\pm$ <sub>6.7</sub> <sup>1.9</sup>	68.3 $\pm$ <sub>7.3</sub> <sup>0.4</sup>	69.1 $\pm$ <sub>6.7</sub> <sup>0.2</sup>	68.4 $\pm$ <sub>3.3</sub>	2.5 $\pm$ <sub>0.24</sub>
$\delta$	2	0	58.7 $\pm$ <sub>3.2</sub> <sup>5.5</sup>	74.0 $\pm$ <sub>4.4</sub> <sup>2.7</sup>	78.1 $\pm$ <sub>4.1</sub> <sup>0.3</sup>	78.6 $\pm$ <sub>4.3</sub> <sup>0.1</sup>	78.6 $\pm$ <sub>1.9</sub>	2.27 $\pm$ <sub>0.02</sub>
O	8	10	60.0 $\pm$ <sub>8.7</sub> <sup>6.0</sup>	77.5 $\pm$ <sub>2.2</sub> <sup>1.2</sup>	80.8 $\pm$ <sub>4.1</sub> <sup>0.3</sup>	81.2 $\pm$ <sub>4.2</sub> <sup>0.1</sup>	81.2 $\pm$ <sub>2.1</sub>	10.67 $\pm$ <sub>0.44</sub>
O	8	0	60.2 $\pm$ <sub>4.7</sub> <sup>7.3</sup>	77.7 $\pm$ <sub>2.6</sub> <sup>2.0</sup>	81.4 $\pm$ <sub>2.2</sub> <sup>0.3</sup>	81.7 $\pm$ <sub>2.2</sub> <sup>0.1</sup>	81.5 $\pm$ <sub>0.9</sub>	9.48 $\pm$ <sub>0.08</sub>
$\delta$	8	10	57.6 $\pm$ <sub>4.2</sub> <sup>7.6</sup>	76.3 $\pm$ <sub>1.1</sub> <sup>1.4</sup>	80.3 $\pm$ <sub>3.0</sub> <sup>0.3</sup>	80.8 $\pm$ <sub>2.9</sub> <sup>0.1</sup>	80.3 $\pm$ <sub>1.0</sub>	8.21 $\pm$ <sub>0.07</sub>
$\delta$	8	0	60.4 $\pm$ <sub>3.6</sub> <sup>4.1</sup>	80.0 $\pm$ <sub>3.0</sub> <sup>1.3</sup>	82.7 $\pm$ <sub>0.9</sub> <sup>1.0</sup>	83.3 $\pm$ <sub>2.3</sub> <sup>0.1</sup>	83.5 $\pm$ <sub>0.8</sub>	8.12 $\pm$ <sub>0.02</sub>

Table 3:  $\delta$ -VAE-style free bits with no KL annealing delivers the best SSL performance and the KL value closest to the desired rate. *Ann.*: 0: no annealing, 10: anneal for 10 epochs; *FB*: free bits type; *F1(n)*: F1-score in the  $n$  data-regime; *KL*: rate obtained after training.

RESET.	$\lambda$	F1(5)	F1(50)	F1(500)	F1(5000)	F1(ALL)	KL
N	2	51.0 $\pm$ <sub>5.6</sub> <sup>4.2</sup>	61.3 $\pm$ <sub>9.2</sub> <sup>2.0</sup>	65.6 $\pm$ <sub>9.5</sub> <sup>0.5</sup>	66.2 $\pm$ <sub>9.5</sub> <sup>0.1</sup>	65.2 $\pm$ <sub>4.9</sub>	2.36 $\pm$ <sub>0.15</sub>
Y	2	58.7 $\pm$ <sub>3.2</sub> <sup>5.5</sup>	74.0 $\pm$ <sub>4.4</sub> <sup>2.7</sup>	78.1 $\pm$ <sub>4.1</sub> <sup>0.3</sup>	78.6 $\pm$ <sub>4.3</sub> <sup>0.1</sup>	78.6 $\pm$ <sub>1.9</sub>	2.27 $\pm$ <sub>0.02</sub>
N	8	57.4 $\pm$ <sub>2.4</sub> <sup>5.6</sup>	73.4 $\pm$ <sub>7.3</sub> <sup>1.5</sup>	77.2 $\pm$ <sub>6.6</sub> <sup>0.3</sup>	77.5 $\pm$ <sub>6.7</sub> <sup>0.1</sup>	77.4 $\pm$ <sub>2.6</sub>	8.23 $\pm$ <sub>0.08</sub>
Y	8	60.4 $\pm$ <sub>3.6</sub> <sup>4.1</sup>	80.0 $\pm$ <sub>3.0</sub> <sup>1.3</sup>	82.7 $\pm$ <sub>0.9</sub> <sup>1.0</sup>	83.3 $\pm$ <sub>2.3</sub> <sup>0.1</sup>	83.5 $\pm$ <sub>0.8</sub>	8.12 $\pm$ <sub>0.02</sub>

Table 4: Resetting the decoder brings very noticeable gains on all data-regimes and with different rates. Yelp dataset,  $\delta$ -VAE free bits, no KL annealing. For columns interpretations, see Table 3.

Dataset	Splits size	Label	$ \mathcal{Y} $	$H[Y]$	NLL
AGNews	110/10/10	Topic	4	1.39	128.77 $\pm$ <sub>0.21</sub>
Amazon	100/10/10	Sent.	5	1.61	82.90 $\pm$ <sub>0.10</sub>
Yahoo	100/10/10	Topic	10	2.30	81.91 $\pm$ <sub>0.36</sub>
Yelp	100/10/10	Sent.	2	0.67	34.60 $\pm$ <sub>0.28</sub>

Table 5: Datasets characteristics.  $|\mathcal{Y}|$ : number of different labels.  $H[Y]$ : entropy of labels. NLL: mean negative log-likelihood of LSTM baseline models (std. over 3 runs). Splits size: train/valid/test sizes in thousands.

## B.2 Tracing back reconstruction gains to words

If words in a document were independently modeled, any improvement in reconstruction at a certain position would indicate that information about the word in that position were encoded in the latent variable. However, words are far from being independently predicted, so how can we trace back the information to the encoder?

First, any latent information related to the first word should not yield any improvements on the prediction of the second word, because the decoder is recurrent and trained using teacher forcing, i.e., conditioned on the true first word, so that information would be redundant. However, information related to the second word in the latent variable can help the decoder predict the first word. Therefore, gains in position  $i$  can only be attributed to information pertaining to the words in positions

$\geq i$ .

Second, the correlation between words in two positions decreases as the distance between these words grow. In effect, information pertaining to the second word yields more gains on the second word than on the first word. From these two facts, we conclude that gains for a position  $i$  mostly comes from information about the word in position  $i$  itself.

## B.3 Reconstruction and memorization

To study the concrete impact of this observation for generation, we encode and decode test documents using the *last-PreAE* variant.<sup>9</sup> Then, we compute the ratio of documents for which the first word in the sources and in the reconstructions match and similarly, how often the sources and their reconstructions have the same number of words.

We compare these with scores obtained by a baseline model that outputs the most frequent first word given the label and the most common document length given the label. This baseline mimicks the behavior of a hypothetical VAE which would encode the labels of the documents (topic or sentiment) perfectly and nothing more.

Results in Table 6 show that with the *last-PreAE* the first words are reconstructed with much higher accuracy than if the latent vector only encoded the label. On the last two datasets, it recovers the first

<sup>9</sup> $\lambda = 8$ ,  $d = 16$ , beam search with beam size of 5.

Dataset	<i>last-PreAE</i>		Clf. given label	
	1st (%)	Len. (%)	1st (%)	Len. (%)
AGNews	29.6±1.1	3.6±0.1	12.9	4.8
Amazon	42.4±2.3	13.0±1.6	14.0	0
Yahoo	56.6±1.0	17.1±1.1	11.3	4.9
Yelp	53.0±0.5	33.7±1.7	14.1	9.7

Table 6: The latent variables encode more information than the label alone, in particular, information that allows to retrieve the first word and the document length with high accuracy.

words on more than half of the documents whereas the baseline only recovers the first words between 11.3 and 14.1% of the time. Accurate encoding of the number of words seems less systematic than the encoding of the first few words. For example, on AGNews, the sentence length is recovered less often than our baselines. The encoding of the sentence length is more pronounced on datasets with small documents like Yahoo and Yelp.

## C Training procedure

### C.1 Grid search

The target rates  $\lambda$  are chosen to be higher than the entropy of the labels of the documents (Table 5) as we assume that the latent variable should at least capture the annotated label. Indeed,  $\lambda = 2$  nats is enough to store the labels of all datasets without any loss, except Yahoo which has an entropy of 2.3 whereas  $\lambda = 8$  nats suffices to capture much more information than needed to store the labels on all datasets. Moreover, these rates are chosen to be much smaller than the reconstruction loss of the baselines because of the technical difficulty of increasing the rate without degrading the log-likelihood explained above.

The latent vector dimension  $d$  is either 4 or 16. Recall that our representations are evaluated on downstream tasks with very limited data in some cases (as little as 5 examples per class), so we need a small enough dimension of latent vector to be able to learn. We suppose that  $d = 4$  will be favored for the 5 or 50 examples per class regime while  $d = 16$  could be more efficient above this, but we leave this choice to the model selection procedure.

### C.2 Constant hyperparameters

All the runs are trained using SGD with a learning rate of 0.5 and gradients are clipped when their norms are higher than 5. We use the following

early stopping scheme: at every epoch, if there has not been improvements on the validation error for two epochs in a row, the learning rate is halved. Once it has been halved four times, the training stops.

All the LSTMs have hidden state size of 512 and use a batch size of 64. No dropout is applied to the encoders. The LSTM decoders use dropout ( $p = 0.5$ ) both on embeddings and on the hidden states (before the linear transformation that gives logits). Similarly, dropout is applied to the representation before the linear transformation that gives the logits for the Unigram decoder. Word embeddings are initialized randomly and learned.

### C.3 Computing infrastructure and average runtime

We performed training and evaluations of the models on a cluster containing a hundred of GPUs with various specifications (NVIDIA Tesla k80, Titan X, Titan Xp, etc.). Given that all the datasets have roughly 100000 training examples (cf. Table 5) and that neural networks are trained with BPTT (Werbos, 1990), the training time mostly depends on the average sentence length and the vocabulary size. Pretraining schemes (*PreAE*, *PreUni* and *PreLM*) require the training of two models. Roughly, the training time of a single model (pretraining or final) varied between 1 hour and 6 hours.

To be more specific, the best baselines and our best variants use pretraining phases (*PreAE* from Li et al. (2019) and *PreUni*, respectively). *PreUni* is faster because the first training phase uses a non-recurrent decoder, and the second training phase does not backpropagate and does not update the encoder. However, *PreAE* does not require pretraining for each desired target rate  $\lambda$ , unlike our approach. Overall, the approaches have comparable runtimes.

Some of our models offer an interesting compromise: *BoW-max-LSTM* with no pretraining and a simpler architecture is probably the fastest, yet outperform the *PreAE* baselines.

## D Related work

### D.1 Related models

The models that we use are similar to already proposed models.

The NVDM model of Miao et al. (2016) is precisely *BoW-max-Uni*.

Zhao et al. (2017) proposed to use two reconstruction losses: the regular reconstruction loss given by the recurrent decoder and an auxiliary loss computed from a unigram decoder. In comparison, our *Uni* models are trained in two steps: the encoder is trained jointly with the unigram decoder, then the decoder is thrown away and we train a recurrent decoder using the fixed encoder. This way, one decoder cannot dominate the other and we do not have to deal with an additional hyperparameter to weight the two losses.

Instead of using an auxiliary *loss*, we have an auxiliary *decoder* that is only used for the purpose of training the encoder. This method was presented by De Fauw et al. (2019) for training generative models of image. There is a slight difference: they use a feedforward auxiliary decoder to produce different probability distributions for all the pixels, whereas our unigram probability distribution is the same for all words of a document. This modification allows us to deal with varying lengths of documents.

Finally, the *PreLM* training procedure is related to large LM pretraining in the spirit of contextualized embeddings (Peters et al., 2018) and its successors. Note, however, two differences. Firstly, we do not use external data and stick to each individual training set, because the goal is not to evaluate transfer learning abilities. Secondly, we do not fine-tune the entire encoder, but only learn the linear transformations  $L_1$  and  $L_2$  that produce the variational parameters, to make sure that the VAE objective will have no impact on the extraction of features.

## D.2 Methods and evaluations

In their analysis of the semi-amortized VAE, Kim et al. (2018) use several saliency measures (defined as expectations of gradients) to determine which words influence the latent variable, or are influenced by it. Using these measures, they noticed that the beginning of the sentence and the end-of-sentence token have a large influence on the variable. Our method is very similar, but slightly simpler and directly interpretable in terms of quantity of information (in nats).

Ficler and Goldberg (2017) learn *LSTM-LMs* conditioned on labels that describe high-level properties of texts. Among others, they want to verify that generated texts exhibit the same properties as the conditioning labels. For instance, when the

*LSTM-LM* is conditioned on positive sentiment value, the generated texts should also exhibit a positive sentiment. To check that the conditioning variables and the generated texts are consistent, they use the following procedure. First, they extract information about the various documents using heuristics or with the help of annotators. Then, they learn *LSTM-LMs* conditioned on these labels. Finally, they quantify the ratio of generated samples which have the same labels than the conditioning labels, either by applying the same heuristics again to the generated samples or by asking human annotators once more. Our evaluation in Section 6 is similar; we simply replace the heuristics and the human annotators with classifiers learned on ground-truth data.

McCoy et al. (2019) trained autoencoders with different combinations of encoders and decoders (unidirectional, bidirectional or tree-structured) and decomposed the representations learned by the encoders using tensor product representations (Smolensky, 1990). They find that decoders “largely dictate” the way information is encoded. This is in line with our own conclusions. An important difference between our works is that they study how information is encoded in sequence-to-sequence models without capacity limitations, whereas in our study, the VAE objective puts severe constraints on the capacity.

## E Semi-supervised learning experiments

### E.1 Model selection

For a given dataset in a given data-regime, we want a measure of the performance of our models that abstracts away from i) hyperparameters for the VAEs, ii) hyperparameters for the downstream task classifiers, iii) subsampling of the dataset and iv) parameter initialisation of the VAEs. As is usually done by practitioners, we optimize over the hyperparameters of the VAEs and the classifiers, eliminating i) and ii) as sources of variance. We can study the robustness of the models by looking at the variance induced by the choice of the subsample and the initialisation of the parameters.

On a given dataset and in a given data-regime, for a given model, we note  $F_{ij}^{H_M, H_C}$  the F1-score obtained on the test set on the subsample using seed  $i$ , the parameter initialisation using seed  $j$ , VAE hyperparameters  $H_M$  and classifier hyperparameters  $H_C$ . We use repeated stratified K-fold cross-validation (Moss et al., 2018) to compute a valida-



tion error  $F_{ij}^{\widehat{H}_M, \widehat{H}_C}$ . For all training folds, we train logistic regression classifiers with  $L_2$  regularisation and a grid-search on  $H_C \in \{0.01, 0.1, 1, 10, 100\}$ . We select the best classifier hyperparameter:

$$H_C^* = \arg \max_{H_C} F_{ij}^{\widehat{H}_M, \widehat{H}_C}$$

Then, the best VAE hyperparameter is chosen by averaging over the  $s = 3$  random seeds and picking the best classifier hyperparameter,

$$H_M^* = \arg \max_{H_M} \frac{1}{s} \sum_{i=1}^s F_{ij}^{\widehat{H}_M, \widehat{H}_C^*}$$

Having optimised the hyperparameters, we compute the test set F1-score:

$$F_{ij} = F_{ij}^{H_M^*, H_C^*}$$

## E.2 Decomposing the variances of the scores

For a given model, dataset and data-regime, after optimisation of the hyperparameters of the VAE and the classifier, we collect several F1-scores  $F_{ij}$  which depend on the seed used to subsample the dataset  $i$  and the seed used to initialise the model parameters  $j$ . We posit a linear model with one random-effect factor, the initialisation seed, and where replicates are obtained by varying the subsampling seed:

$$F_{ij} = \mu + \alpha_j + \epsilon_{ij}$$

Assuming that  $\alpha_j$  and  $\epsilon_{ij}$  are independent random variables with null expectations, we can decompose the variance as

$$\begin{aligned} \text{Var}(F_{ij}) &= \mathbb{E}[(F_{ij} - \mu)^2] \\ &= \mathbb{E}[(\alpha_j + \epsilon_{ij})^2] \\ &= \mathbb{E}[\alpha_j^2] + \mathbb{E}[\epsilon_{ij}^2] \\ &= \text{Var}(\alpha_j) + \text{Var}(\epsilon_{ij}) \end{aligned}$$

This is the basis of the method of analysis of variance (ANOVA) and is often used to test hypotheses (for instance, that the effect  $\mathbb{E}[\alpha_j]$  is significant) (Oehlert, 2010). The two estimates of  $\sigma_{init}^2$  and  $\sigma^2$  are usually denoted  $MS_T$  and  $MS_E$ .

In our case, we are only interested in estimating roughly what variability is due to the model initialisation and what is due to the subsampling of the dataset.

Note that we could treat the two sources of variance  $i$  and  $j$  symmetrically by adding a term  $\beta_i$ ,

but we would need to report 3 standard deviations (that of  $\alpha_j$ ,  $\beta_i$  and  $\epsilon_{ij}$ ) to get the full picture. The most important estimate is  $\sigma_{init}$ . It quantifies the inherent robustness of the model to different initialisations. The effect of the subsampling is specific to the dataset, therefore, it is less relevant to our analysis.

## E.3 What is the representation of a document?

VAEs are mostly used for generating samples but are also sometimes used as feature extractors for SSL. In the latter case, it is not clear what the representation of a datapoint is: the mean of the approximate posterior  $\mu$  or the noisy samples  $Z \sim \mathcal{N}(\mu, I\sigma^2)$ ? Kingma et al. (2014) feed noisy samples  $z$  in the classifiers but in the literature of VAEs applied to language modeling, it is more common to use  $\mu$  without explanation or even mention.<sup>10</sup>

If we are interested purely in downstream task performance, the mean should perform best, as the samples are just noisy versions of the mean vector (it is still not completely straightforward as the noise could play a regularizing role). However, in order to evaluate what information is *effectively* transmitted to the decoder, we should use the samples. The performance of downstream task classifiers using the mean does not tell us *at all* whether the latent variable is used by the decoder to reconstruct the input. The following experiment illustrates this fact.

We train the original VAE architecture on the Yelp dataset, both with and without the *PreAE*, using the original ELBo objective ( $\lambda = 0$ ). As expected, the KL term collapses to 0. Then, we train a classifier using the procedure explained above using 5000 examples per class. We expect that its performance will be close to random chance, regardless of whether samples or the mean parameter are used as inputs. However, Table 7 shows that this is not the case. Using samples, we do get random chance predictions from the classifiers, whereas using means, the performance is remarkably high (as high as 81.5 of F1 using pretraining). The reason is that the KL term never *completely* collapses to 0. Therefore,  $\mu$  can be almost zero while still encoding a lot of information about its inputs. However, when the KL term is close to 0, the variance of the

<sup>10</sup>For instance, Li et al. (2019) and Fu et al. (2019) do not mention what representation they use but their code uses the mean; Long et al. (2019) report using a concatenation of the mean and the variance vectors.

PreAE	F1		KL
	$z$	$\mu$	
No	49.5	64.7	$1e^{-4}$
Yes	49.6	81.5	$2e^{-4}$

Table 7: When the KL collapses, the performances of classifiers trained on the mean  $\mu$  vs on samples  $z \sim \mathcal{N}(\mu, I\sigma^2)$  are very different, especially for pretrained models.  $z$  does not contain any information while  $\mu$  is very predictive of the label.

samples is close to 1, so no information is transmitted to the decoder. This tendency is exacerbated with the *PreAE* runs, for which the means encode remnants of the pretraining phase.

This experiment shows that it is crucial to report what representation ( $z$  or  $\mu$ ) is analyzed and to cautiously interpret the results. Therefore, for the purpose of analysing representations for text generation, we feed  $z$  as inputs to the classifiers.

#### E.4 Recurrent and *BoW* encoders work around max-pooling

It is counter-intuitive that *BoW-max-LSTM* improves over *LSTM-max-LSTM* (with or without *PreAE*). Indeed, taking into account word order should allow the LSTM encoder to do better inference than the *BoW* encoder, for example, by handling negation or parsing more complicated discourse structure (Pang et al., 2002).

LSTM encoders are more powerful, but it can lead them to learn undesirable behaviors. We noticed that some components of the hidden states consistently reach their maximum values at fixed positions, regardless of the inputs (i.e., for some components  $j^*$ ,  $\arg \max_i h_i^{j^*} \approx K$ ). These positions  $K$  are often early positions in the sentence. For instance, with  $\lambda = 8$ ,  $d = 16$ , *LSTM-max-LSTM-PreAE* has 70 components out of 512 that are selected on 80% of the documents on the same position on Yelp (68 on the first word, 2 on the second) and 78 on Amazon (57 on the first word, 21 on the second). In other words, some components of  $r$  act like memory slots assigned to fixed positions in the sentence. This is probably achieved through counting mechanisms (Shi et al., 2016; Suzgun et al., 2019). The decoder is also an LSTM and can count, so it can also extract the relevant components at each position to retrieve the corresponding words.

For *BoW* encoders, it is less clear. It is possi-

ble that on some datasets, capitalized words could take especially high values on some components, in order to be consistently represented after max pooling. However, we have not explored the issue further.

## F Qualitative analysis

For our qualitative analysis, we take a look at the reconstruction samples (which were also used in Section 6). We focus on the *PreUni* models which lower memorization the most with the *LSTM-max* and *BoW-max* encoders, and compare them to the two best baselines. We use only one seed and one  $z$  sample per model and per source sentence, but use two decoding strategies (beam search and greedy decoding).

In general, and for the reasons explained above, the rate  $\lambda = 8$  is chosen too small to recover exactly the source. Indeed, this rate is an order of magnitude less than the negative log-likelihood of *LSTM-LM* baselines: above 80 for all datasets except on Yelp where it is around 34.60 nats (cf. Table 5). Since the NLL is an upper-bound on the entropy of the data, it gives a crude over estimate of the information content of the average document. On Yelp, where the NLL is much smaller (around 34.60 nats), we hope to obtain good paraphrases for simple and frequent sentences. On the other datasets, we can not hope to reconstruct the sentences correctly but merely to control the generation by producing sentences which have the same labels as the source sentences. For this reason, we cherry-pick source sentences that look quite generic, because they are more probable and therefore, should be easy to reconstruct correctly.

Results are presented in Tables 10, 11, 12 and 13. Overall, we do observe less memorization of the first words and more correct sentiment or topic. Between our two models, on Amazon and Yelp (sentiment labels), it seems that *LSTM-max-LSTM-PreUni* might perform better than *BoW-max-LSTM-PreUni* because of its ability to handle negation, probably thanks to the recurrent encoder. It also seem more on topic on AGNews. On Yelp, it is able to paraphrase generic, small sentences. Therefore, we recommend this model as a future baseline.

We have already seen that it is very hard to classify sentences based on their first three words on the Amazon and Yahoo dataset, and that the baseline methods will learn representations that are not predictive of labels. However, the Yahoo dataset is

	$r$	Dec.	Pre.	5	50	500	5000	All
						$F1 \pm \sigma_{\text{init}}$		
AGNews	last	LSTM	-	$59.6 \pm 5.1$	$71.7 \pm 1.0$	$73.6 \pm 0.1$	$73.7 \pm 0.1$	$73.6 \pm 5.4$
	last	LSTM	AE	$65.8 \pm 3.3$	$81.0 \pm 0.7$	$82.8 \pm 0.3$	$83.1 \pm 0.1$	$83.4 \pm 0.3$
	max	LSTM	-	$27.3 \pm 2.4$	$30.8 \pm 3.4$	$33.1 \pm 0.9$	$33.8 \pm 0.4$	$34.6 \pm 2.4$
	max	LSTM	AE	$55.7 \pm 18.7$	$75.1 \pm 1.3$	$81.9 \pm 0.3$	$82.5 \pm 0.1$	$83.3 \pm 0.4$
	max	LSTM	-	$72.7 \pm 2.0$	$81.2 \pm 0.6$	$82.2 \pm 0.2$	$82.3 \pm 0.1$	$83.1 \pm 0.3$
	max	Uni	-	$71.6 \pm 5.5$	$80.4 \pm 0.8$	$81.8 \pm 0.5$	$82.4 \pm 0.1$	$83.9 \pm 0.3$
	last	Uni	-	$54.8 \pm 5.2$	$61.7 \pm 0.8$	$62.9 \pm 0.4$	$63.0 \pm 0.3$	$59.3 \pm 40.9$
	max	Uni	-	$71.8 \pm 4.8$	$81.4 \pm 0.5$	$82.5 \pm 0.1$	$82.5 \pm 0.1$	$83.1 \pm 0.5$
	avg	LSTM	LM	$70.8 \pm 4.3$	$81.2 \pm 1.2$	$82.6 \pm 1.3$	$82.8 \pm 0.9$	$83.5 \pm 0.1$
	Amazon	last	LSTM	-	$18.9 \pm 1.7$	$20.9 \pm 1.2$	$22.5 \pm 0.7$	$23.3 \pm 0.4$
last		LSTM	AE	$20.0 \pm 2.2$	$24.7 \pm 0.7$	$27.2 \pm 0.4$	$27.7 \pm 0.3$	$28.1 \pm 1.0$
max		LSTM	-	$19.8 \pm 0.7$	$20.4 \pm 1.1$	$22.2 \pm 0.6$	$23.0 \pm 0.3$	$23.7 \pm 0.5$
max		LSTM	AE	$22.3 \pm 2.6$	$30.5 \pm 0.9$	$33.4 \pm 0.4$	$34.1 \pm 0.3$	$34.0 \pm 1.6$
max		LSTM	-	$21.0 \pm 2.6$	$34.6 \pm 0.7$	$38.3 \pm 0.4$	$39.0 \pm 0.1$	$38.9 \pm 0.7$
max		Uni	-	$21.8 \pm 3.1$	$32.8 \pm 0.8$	$36.9 \pm 0.4$	$38.0 \pm 0.2$	$38.2 \pm 0.5$
last		Uni	-	$24.0 \pm 3.0$	$31.2 \pm 0.6$	$35.1 \pm 0.4$	$36.1 \pm 0.2$	$36.8 \pm 0.9$
max		Uni	-	$25.4 \pm 3.2$	$32.8 \pm 1.0$	$36.1 \pm 0.4$	$36.9 \pm 0.2$	$37.9 \pm 0.2$
avg		LSTM	LM	$21.8 \pm 3.8$	$35.3 \pm 0.8$	$40.2 \pm 0.4$	$41.1 \pm 0.2$	$40.0 \pm 0.4$
Yahoo		last	LSTM	-	$10.9 \pm 0.9$	$12.1 \pm 0.6$	$13.9 \pm 0.4$	$14.1 \pm 0.2$
	last	LSTM	AE	$20.7 \pm 0.7$	$32.2 \pm 0.8$	$36.1 \pm 0.2$	$36.7 \pm 0.1$	$37.2 \pm 0.7$
	max	LSTM	-	$9.9 \pm 1.0$	$13.0 \pm 0.6$	$14.6 \pm 0.3$	$14.9 \pm 0.1$	$15.7 \pm 0.5$
	max	LSTM	AE	$20.8 \pm 1.3$	$31.3 \pm 0.7$	$35.6 \pm 0.3$	$36.3 \pm 0.1$	$36.6 \pm 0.7$
	max	LSTM	-	$23.4 \pm 2.1$	$36.7 \pm 1.1$	$41.1 \pm 0.2$	$41.6 \pm 0.1$	$42.6 \pm 0.2$
	max	Uni	-	$24.9 \pm 1.3$	$33.2 \pm 0.7$	$37.3 \pm 0.1$	$37.9 \pm 0.1$	$38.9 \pm 1.7$
	last	Uni	-	$24.5 \pm 3.8$	$30.8 \pm 1.7$	$34.4 \pm 0.3$	$35.1 \pm 0.1$	$37.1 \pm 2.3$
	max	Uni	-	$24.1 \pm 2.9$	$35.0 \pm 1.2$	$39.1 \pm 0.1$	$39.5 \pm 0.1$	$40.1 \pm 0.7$
	avg	LSTM	LM	$21.9 \pm 2.3$	$36.1 \pm 0.8$	$39.9 \pm 0.2$	$40.4 \pm 0.1$	$41.7 \pm 0.3$
	Yelp	last	LSTM	-	$49.9 \pm 4.5$	$55.6 \pm 2.3$	$57.9 \pm 1.1$	$59.5 \pm 0.2$
last		LSTM	AE	$59.3 \pm 5.4$	$80.0 \pm 1.3$	$82.7 \pm 1.0$	$83.3 \pm 0.1$	$67.9 \pm 0.1$
max		LSTM	-	$61.6 \pm 8.2$	$71.4 \pm 2.3$	$76.0 \pm 0.9$	$76.5 \pm 0.1$	$78.0 \pm 1.7$
max		LSTM	AE	$59.9 \pm 10.4$	$78.7 \pm 2.4$	$82.9 \pm 0.3$	$83.3 \pm 0.1$	$84.1 \pm 0.7$
max		LSTM	-	$67.1 \pm 10.1$	$79.3 \pm 2.8$	$83.4 \pm 0.3$	$83.9 \pm 0.1$	$85.0 \pm 0.2$
max		Uni	-	$62.3 \pm 4.6$	$76.7 \pm 1.7$	$80.4 \pm 0.2$	$80.9 \pm 0.1$	$83.1 \pm 0.5$
last		Uni	-	$65.0 \pm 8.0$	$74.1 \pm 2.0$	$78.5 \pm 0.3$	$79.1 \pm 0.1$	$81.6 \pm 0.5$
max		Uni	-	$59.9 \pm 7.2$	$77.3 \pm 1.2$	$81.1 \pm 0.3$	$81.5 \pm 0.1$	$83.3 \pm 0.4$
avg		LSTM	LM	$63.6 \pm 7.4$	$81.0 \pm 1.6$	$83.2 \pm 0.7$	$83.8 \pm 0.1$	$84.4 \pm 0.5$

Table 8: Using *BoW* encoders, *Uni* decoders or *PreLM* pretraining, the learned representations are more predictive of the labels (sentiment or topic) of the documents.

Dataset	F1(All)	F1(3)	Ratio
AGNews	89.0	71.9	0.808
Amazon	48.9	29.7	0.607
Yahoo	63.0	19.1	0.303
Yelp	96.5	82.4	0.854

Table 9: Performance of bag-of-word classifiers when using all words as features versus only the first three words. Ratios of performance vary a lot across datasets.

especially challenging and our methods also struggle on it. We hypothesize that it is because only a few words per sentences are correlated with the labels. Indeed, there are many sentences of the

form “what do you think about X?” or “what is the difference between X and Y?” where only X and Y are correlated with the label and moreover often out-of-vocabulary. There might be no benefits for the model to diminish the reconstruction loss in priority on these words. By comparison, the other “hard” dataset (Amazon) is “easier”, because the sentiment is often indicated by frequent adjectives like “horrible” or “good” and other frequent verbs.

Model/“source”	Dec.	Sample
source		Michael Owen heads England ’s winner in the World Cup qualifier against Azerbaijan .
LSTM-last-LSTM-PreAE	beam	Michael Owen will be sidelined for the rest of the season with a knee injury .
LSTM-last-LSTM-PreAE	greedy	Michael Owen has been charged with a rib injury and a new team for the first time in the last two weeks .
LSTM-max-LSTM-PreAE	beam	American Bode Miller won the World Cup super @-@ G with a 6 @-@ 3 victory over the United States in the World Cup of Hockey .
LSTM-max-LSTM-PreAE	greedy	American Bode Miller won the World Cup of Hockey on Sunday , beating the United States by a record @-@ setting victory over the United States .
BoW-max-LSTM-PreUni	beam	England coach Sven @-@ Goran Eriksson says he will not be able to win the World Cup qualifier against Wales .
BoW-max-LSTM-PreUni	greedy	England captain David Beckham has been named the England captain for the 2006 World Cup qualifiers against Wales .
LSTM-max-LSTM-PreUni	beam	England coach Sven @-@ Goran Eriksson says he will be fit for the World Cup qualifier against Wales next month .
LSTM-max-LSTM-PreUni	greedy	England captain David Beckham has been named the first World Cup qualifier in the World Cup qualifier against Wales .
source		New Athlon 64 processors will compete with Intel ’s Pentium 4 Extreme Edition .
LSTM-last-LSTM-PreAE	beam	IBM ’s dual @-@ core Opteron processor will be available in the next three years .
LSTM-last-LSTM-PreAE	greedy	A new chipset for mobile phones will be available in the next three years .
LSTM-max-LSTM-PreAE	beam	New version of Windows Server 2003 .
LSTM-max-LSTM-PreAE	greedy	New version of the Linux operating system is designed to integrate Linux and Linux .
BoW-max-LSTM-PreUni	beam	Hewlett @-@ Packard Co . , the world ’s largest computer maker , has unveiled a new version of its iPod digital music player , the company said .
BoW-max-LSTM-PreUni	greedy	Hewlett @-@ Packard Co . , the world ’s largest maker of digital music player , on Tuesday unveiled a new version of its popular PlayStation 2 game console , which will be available in the next few years .
LSTM-max-LSTM-PreUni	beam	Intel has unveiled a new version of its Pentium 4 Extreme Edition processor , which will be available for the first time .
LSTM-max-LSTM-PreUni	greedy	Intel has unveiled a new version of its Pentium M processor , which is designed to help the company ’s new processor @-@ based processors .
source		Nortel said it expects revenue for the third quarter to fall short of expectations .
LSTM-last-LSTM-PreAE	beam	Coca @-@ Cola Co .
LSTM-last-LSTM-PreAE	greedy	research ) is expected to announce a new deal with the company to buy the company .
LSTM-max-LSTM-PreAE	beam	Nortel Networks Corp.
LSTM-max-LSTM-PreAE	greedy	General Electric Co. said on Thursday it will buy the company for \$ 500 million in cash and stock .
BoW-max-LSTM-PreUni	beam	Ford Motor Co . , the world ’s largest maker of photographic film , said on Thursday it expects to cut its full @-@ year earnings forecast , citing strong demand for its flagship database products .
BoW-max-LSTM-PreUni	greedy	Coca @-@ Cola Co. on Thursday said third @-@ quarter earnings rose 29 percent , helped by strong sales of its soft drinks and business software .
LSTM-max-LSTM-PreUni	beam	Nortel Networks Corp. , the world ’s largest maker of equipment , said on Thursday that its third @-@ quarter profit rose 12 percent , helped by a rebound in the value of its assets .
LSTM-max-LSTM-PreUni	greedy	Shares of Nortel Networks Corp. fell nearly 8 percent on Thursday after the company said it expects its earnings for the third quarter , citing a decline in its third @-@ quarter earnings .

Table 10: Cherry-picked AGNews samples. Baselines are the first two models, our models are the two last. In the first example, the first baseline copies “Michael Owen” and complete with generic suffixes; the second baseline is about Hockey instead of soccer. Our baselines do not copy the beginning while correctly identifying the topic of England and the World Cup’s qualifier. Similar comments can be made on the two other examples.

Model/“source”	Dec.	Sample
source		I loved the book , but was a little bit UNK with the ending .
LSTM-last-LSTM-PreAE	beam	I enjoyed the book , but the story line was not as good as the first one .
LSTM-last-LSTM-PreAE	greedy	I enjoyed the book and the story line . I was very disappointed in the book .
LSTM-max-LSTM-PreAE	beam	I was very disappointed in the quality of the book , and the content of the book is very poor .
LSTM-max-LSTM-PreAE	greedy	I was disappointed in the quality of the book , but the book is not as good as the original .
BoW-max-LSTM-PreUni	beam	I liked the story and the story line . It was a little slow at times but overall a good read .
BoW-max-LSTM-PreUni	greedy	I liked the story and the story line . It was a little slow but the ending was a little predictable .
LSTM-max-LSTM-PreUni	beam	The book was a little slow , but the story line was good . I enjoyed it .
LSTM-max-LSTM-PreUni	greedy	The book was a little slow and the story line was very good . I was very disappointed .
source		This movie wasn ’t as good as the original but I still enjoy watching it .
LSTM-last-LSTM-PreAE	beam	This movie was a little slow at times , but it was a good movie .
LSTM-last-LSTM-PreAE	greedy	This movie was a little slow and the plot was not good . I would not recommend it to anyone .
LSTM-max-LSTM-PreAE	beam	This movie was not as good as I thought it would be . I was very disappointed .
LSTM-max-LSTM-PreAE	greedy	The movie was not as good as the first one . I was disappointed in the quality of the movie .
BoW-max-LSTM-PreUni	beam	Not as good as I thought it would be . I wouldn ’t watch it again .
BoW-max-LSTM-PreUni	greedy	I didn ’t like this movie . I thought it was going to be a good movie but I wouldn ’t watch it again .
LSTM-max-LSTM-PreUni	beam	I was expecting a little more from this movie . It was a little slow and boring .
LSTM-max-LSTM-PreUni	greedy	The movie was good , but the acting was not very good . I was expecting a little more from the movie .
source		This movie is horrible . The story , the acting , the directing . Just horrible .
LSTM-last-LSTM-PreAE	beam	This is a great movie . I love it . It is a great family movie .
LSTM-last-LSTM-PreAE	greedy	This movie is great . It is a great movie and I love it .
LSTM-max-LSTM-PreAE	beam	This is the worst movie I have ever seen . It was not worth the time to watch .
LSTM-max-LSTM-PreAE	greedy	This was a good movie . The acting was good , but the story line was not very good .
BoW-max-LSTM-PreUni	beam	This movie is not worth the money . The acting is poor and the acting is poor .
BoW-max-LSTM-PreUni	greedy	The movie is very poor , the acting is poor . The acting is poor .
LSTM-max-LSTM-PreUni	beam	This movie is a waste of time and money . It was a waste of time and money .
LSTM-max-LSTM-PreUni	greedy	This movie is a waste of time . It was a waste of time and money .
source		This book is very bad and does not give a real idea of the sport of UNK
LSTM-last-LSTM-PreAE	beam	this is a great book for those who want to learn a little more about the history of the history of the history
LSTM-last-LSTM-PreAE	greedy	this book is a great book for the price , but the book is a little too short for my taste
LSTM-max-LSTM-PreAE	beam	This book is a must have for anyone who is interested in the history of the Catholic Church
LSTM-max-LSTM-PreAE	greedy	This book is a must have for anyone who is interested in the field of the world of the New Testament
BoW-max-LSTM-PreUni	beam	This is a good book , but it does not have a lot of information in it .
BoW-max-LSTM-PreUni	greedy	This book is a good book for the beginner , but it does not have a lot of information in it .
LSTM-max-LSTM-PreUni	beam	This is a good book , but it does not have enough information on how to build your own .
LSTM-max-LSTM-PreUni	greedy	This is a good book , but the book is not a good book for the money .

Table 11: Cherry-picked Amazon samples. The first two examples are moderately positive reviews. LSTM-max-LSTM-PreUNI recreate roughly the same sentiment overall with different beginnings of sentences. On the third example, our models do copy the beginning of the source sentence but do not make mistakes on the sentiment as the baselines do. On the last, all models fail to capture the very negative sentiment, but our models at least moderate the positivity.

Model/"source"	Dec.	Sample
source		What is your favorite book ?
LSTM-last-LSTM-PreAE	beam	what is a good name for a girl ?
LSTM-last-LSTM-PreAE	greedy	what is a good name for a girl ?
LSTM-max-LSTM-PreAE	beam	What is your favorite book ?
LSTM-max-LSTM-PreAE	greedy	What is your favorite book ?
BoW-max-LSTM-PreUni	beam	What is your favorite sport ?
BoW-max-LSTM-PreUni	greedy	What is your favorite movie ?
LSTM-max-LSTM-PreUni	beam	What is your favorite book ?
LSTM-max-LSTM-PreUni	greedy	What is your favorite Christmas movie ?
source		how can i get money ?
LSTM-last-LSTM-PreAE	beam	i want to know ?
LSTM-last-LSTM-PreAE	greedy	i want to know ?
LSTM-max-LSTM-PreAE	beam	how do i get money ?
LSTM-max-LSTM-PreAE	greedy	do you have any money ?
BoW-max-LSTM-PreUni	beam	where can i go to watch free online ?
BoW-max-LSTM-PreUni	greedy	where can i get free online games ?
LSTM-max-LSTM-PreUni	beam	where can i get a job ?
LSTM-max-LSTM-PreUni	greedy	how do i get a job in USA ?
source		What countries have nuclear weapons ?
LSTM-last-LSTM-PreAE	beam	What are the pros and cons ?
LSTM-last-LSTM-PreAE	greedy	What are the pros and cons of smoking ?
LSTM-max-LSTM-PreAE	beam	which countries in the world ?
LSTM-max-LSTM-PreAE	greedy	which country has the highest world cup ?
BoW-max-LSTM-PreUni	beam	how many countries are there in the world ?
BoW-max-LSTM-PreUni	greedy	what are the countries that will be able to be the most effective government in the world ?
LSTM-max-LSTM-PreUni	beam	Why are the colors of the Earth ?
LSTM-max-LSTM-PreUni	greedy	what are the three different countries in the U.S. ?
source		how to print all webpage content ?
LSTM-last-LSTM-PreAE	beam	how to create a website ?
LSTM-last-LSTM-PreAE	greedy	how to find a website ?
LSTM-max-LSTM-PreAE	beam	how can i learn english language ?
LSTM-max-LSTM-PreAE	greedy	how can i watch free online online ?
BoW-max-LSTM-PreUni	beam	how do you get a free copy of the internet ?
BoW-max-LSTM-PreUni	greedy	how to get the free internet explorer ?
LSTM-max-LSTM-PreUni	beam	how do i get a copy of my computer in the internet ?
LSTM-max-LSTM-PreUni	greedy	how do i get the power to open a computer in the internet ?

Table 12: Cherry-picked Yahoo samples. There isn't a model that clearly stands out, but we can rule out LSTM-last-LSTM-PreAE. This dataset is more difficult (see main text).

Model/“source”	Dec.	Sample
source		amazing place .
LSTM-last-LSTM-PreAE	beam	amazing customer service .
LSTM-last-LSTM-PreAE	greedy	amazing customer service .
LSTM-max-LSTM-PreAE	beam	amazing food .
LSTM-max-LSTM-PreAE	greedy	amazing food .
BoW-max-LSTM-PreUni	beam	this place is amazing .
BoW-max-LSTM-PreUni	greedy	this place is amazing .
LSTM-max-LSTM-PreUni	beam	this place is amazing .
LSTM-max-LSTM-PreUni	greedy	this place is amazing .
source		definitely going back soon !
LSTM-last-LSTM-PreAE	beam	definitely coming back !
LSTM-last-LSTM-PreAE	greedy	definitely recommend to anyone !
LSTM-max-LSTM-PreAE	beam	definitely coming back again !
LSTM-max-LSTM-PreAE	greedy	highly recommend them to anyone !
BoW-max-LSTM-PreUni	beam	i will definitely be back !
BoW-max-LSTM-PreUni	greedy	i will definitely be back !
LSTM-max-LSTM-PreUni	beam	i will be back !
LSTM-max-LSTM-PreUni	greedy	i will be back !
source		not worth the risk .
LSTM-last-LSTM-PreAE	beam	not the best .
LSTM-last-LSTM-PreAE	greedy	not the best .
LSTM-max-LSTM-PreAE	beam	not worth the money .
LSTM-max-LSTM-PreAE	greedy	not worth the money .
BoW-max-LSTM-PreUni	beam	worth the wait .
BoW-max-LSTM-PreUni	greedy	it was worth the wait .
LSTM-max-LSTM-PreUni	beam	not worth the hassle .
LSTM-max-LSTM-PreUni	greedy	it 's not worth the money .
source		overall , a huge disappointment .
LSTM-last-LSTM-PreAE	beam	pizza was good too .
LSTM-last-LSTM-PreAE	greedy	pizza was good too .
LSTM-max-LSTM-PreAE	beam	ok , nothing special .
LSTM-max-LSTM-PreAE	greedy	nothing special , but the food was bland .
BoW-max-LSTM-PreUni	beam	wow .
BoW-max-LSTM-PreUni	greedy	great experience .
LSTM-max-LSTM-PreUni	beam	what a disappointment .
LSTM-max-LSTM-PreUni	greedy	what a disappointment .

Table 13: Cherry-picked Yelp samples. On small and typical sentences, our last variant LSTM-max-LSTM-PreUni can produce paraphrases. On the other hand, BoW-max-LSTM-PreUni fails on the two negative examples, probably because it lacks the ability to deal with negation. The baseline models also fail to capture the sentiment on the last example, and copy the beginning on the first three examples.