arXiv:1803.03382v3 [cs.LG] 13 Mar 2018

Fast Decoding in Sequence Models Using Discrete Latent Variables

Łukasz Kaiser [1] **Aurko Roy** [1] **Ashish Vaswani** [1] **Niki Parmar** [1] **Samy Bengio** [1] **Jakob Uszkoreit** [1]
Noam Shazeer [1]

Abstract

Autoregressive sequence models based on deep neural networks, such as RNNs, Wavenet and the Transformer attain state-of-the-art results on many tasks. However, they are difficult to parallelize and are thus slow at processing long sequences. RNNs lack parallelism both during training and decoding, while architectures like WaveNet and Transformer are much more parallelizable during training, yet still operate sequentially during decoding. Inspired by (van den Oord et al., 2017), we present a method to extend sequence models using discrete latent variables that makes decoding much more parallelizable. We first autoencode the target sequence into a shorter sequence of discrete latent variables, which at inference time is generated autoregressively, and finally decode the output sequence from this shorter latent sequence in parallel. To this end, we introduce a novel method for constructing a sequence of discrete latent variables and compare it with previously introduced methods. Finally, we evaluate our model end-to-end on the task of neural machine translation, where it is an order of magnitude faster at decoding than comparable autoregressive models. While lower in BLEU than purely autoregressive models, our model achieves higher scores than previously proposed non-autogregressive translation models.

1. Introduction

Neural networks have been applied successfully to a variety of tasks involving natural language. In particular, recurrent neural networks (RNNs) with long short-term memory (LSTM) cells (Hochreiter & Schmidhuber, 1997) in a sequence-to-sequence configuration (Sutskever et al., 2014) have proven successful at tasks including machine translation (Sutskever et al., 2014; Bahdanau et al., 2014; Cho et al., 2014), parsing (Vinyals et al., 2015), and many others. RNNs are inherently sequential, however, and thus tend to be slow to execute on modern hardware optimized for parallel execution. Recently, a number of more parallelizable sequence models were proposed and architectures such as WaveNet (van den Oord et al., 2016), ByteNet (Kalchbrenner et al., 2016) and the Transformer (Vaswani et al., 2017) can indeed be trained faster due to improved parallelism.

When actually generating sequential output, however, their autoregressive nature still fundamentally prevents these models from taking full advantage of parallel computation. When generating a sequence $y_1 \dots y_n$ in a canonical order, say from left to right, predicting the symbol y_t first requires generating all symbols $y_1 \dots y_{t-1}$ as the model predicts

$$P(y_t | y_{t-1}\ y_{t-2} \dots\ y_1).$$

During training, the ground truth is known so the conditioning on previous symbols can be parallelized. But during decoding, this is a fundamental limitation as at least n sequential steps need to be made to generate $y_1 \dots y_n$.

To overcome this limitation, we propose to introduce a sequence of discrete latent variables $l_1 \dots l_m$, with $m < n$, that summarizes the relevant information from the sequence $y_1 \dots y_n$. We will still generate $l_1 \dots l_m$ autoregressively, but it will be much faster as $m < n$ (in our experiments we mostly use $m = \frac{n}{8}$). Then, we reconstruct each position in the sequence $y_1 \dots y_n$ from $l_1 \dots l_m$ in parallel.

For the above strategy to work, we need to autoencode the target sequence $y_1 \dots y_n$ into a shorter sequence $l_1 \dots l_m$. Autoencoders have a long history in deep learning (Hinton & Salakhutdinov, 2006; Salakhutdinov & Hinton, 2009a; Vincent et al., 2010; Kingma & Welling, 2013). Autoencoders mostly operate on continuous representations, either by imposing a bottleneck (Hinton & Salakhutdinov, 2006), requiring them to remove added noise (Vincent et al., 2010), or adding a variational component (Kingma & Welling, 2013). In our case though, we prefer the sequence $l_1 \dots l_m$ to be discrete, as we use standard autoregressive models to predict it. Despite some success (Bowman et al., 2016; Yang et al., 2017), predicting continuous latent representations does not work as well as the discrete case in our setting.

[1]Google Brain, Mountain View, USA. Correspondence to: Łukasz Kaiser <lukaszkaiser@google.com>, Aurko Roy <aurkor@google.com>.

However, using discrete latent variables can be challenging when training models end-to-end. Three techniques recently have shown how to successfuly use discrete variables in deep models: the Gumbel-Softmax (Jang et al., 2016; Maddison et al., 2016), VQ-VAE (van den Oord et al., 2017) and improved semantic hashing (Kaiser & Bengio, 2018). We compare all these techniques in our setting and introduce another one: decomposed vector quantization (DVQ) which performs better than VQ-VAE for large latent alphabet sizes.

Using either DVQ or improved semantic hashing, we are able to create a neural machine translation model that achieves good BLEU scores on the standard benchmarks while being an order of magnitude faster at decoding time than autoregressive models. A recent paper (Gu et al., 2017) reported similar gain for neural machine translation. But their techniques are hand-tuned for translation and require training with reinforcement learning. Our latent variables are learned and the model is trained end-to-end, so it can be applied to any sequence problem. Despite being more generic, our model outperforms the hand-tuned technique from (Gu et al., 2017) yielding better BLEU. To summarize, our main contributions are:

(1) A method for fast decoding for autoregressive models.

(2) An improved discretization technique: the DVQ.

(3) The resulting Latent Transformer model, achieving good results on translation while decoding much faster.

2. Discretization Techniques

In this section we introduce various *discretization bottlenecks* used to train discrete autoencoders for the target sequence. We will use the notation from (van den Oord et al., 2017) where the target sequence y is passed through an encoder, enc, to produce a continuous latent representation $\mathrm{enc}(y) \in R^D$, where D is the dimension of the latent space. Let K be the size of the discrete latent space and let $[K]$ denote the set $\{1, 2, \ldots, K\}$. The continuous latent $\mathrm{enc}(y)$ is subsequently passed through a discretization bottleneck to produce a discrete latent representation $z_d\left(y\right) \in [K]$, and an input $z_q(y)$ to be passed to the decoder dec. For integers i, m we will use $\tau_m(i)$ to denote the binary representation of i using m bits, with the inverse operation, i.e. conversion from binary to decimal denoted by τ_m^{-1}.

2.1. Gumbel-Softmax

A discretization technique that has recently received a lot of interest is the Gumbel-Softmax trick proposed by (Jang et al., 2016; Maddison et al., 2016). In this case one simply projects the encoder output $\mathrm{enc}(y)$ using a learnable projection $W \in R^D \times K$ to get the logits $l = W z_e(y)$ with the discrete code $z_d\left(y\right)$ being defined as

$$z_d\left(y\right) = \arg\max_{i \in [K]} l_i. \tag{1}$$

The decoder input $z_q(y)$ during evaluation and inference is computed using an embedding $e \in R^{K \times D}$ where $z_q(y) = e_j$. For training, the Gumbel-Softmax trick is used by generating samples $g_1, g_2, \ldots, g_K$ i.i.d samples from the Gumbel distribution: $g_i \sim -\log\left(-\log u\right)$, where $u \sim \mathcal{U}(0, 1)$ are uniform samples. Then as in (Jang et al., 2016; Maddison et al., 2016), one computes the log-softmax of l to get $w \in R^K$:

$$w_i = \frac{\exp((l_i + g_i)/\tau)}{\sum_i \exp((l_i + g_i)/\tau}, \tag{2}$$

with the input to the encoder $z_q(y)$ being simply the matrix-vector product we. Note that the Gumbel-Softmax trick makes the model differentiable and thus it can be trained using backpropagation.

For low temperature τ the vector w is close to the 1-hot vector representing the maximum index of l, which is what is used during evaluation and testing. But at higher temperatures, it is an approximation (see Figure 1 in Jang et al. (2016)).

2.2. Improved Semantic Hashing

Another discretization technique proposed by (Kaiser & Bengio, 2018) that has been recently explored stems from *semantic hashing* (Salakhutdinov & Hinton, 2009b). The main idea behind this technique is to use a simple rounding bottleneck after squashing the encoder state $z_e(y)$ using a saturating sigmoid. Recall the saturating sigmoid function from (Kaiser & Sutskever, 2016; Kaiser & Bengio, 2016):

$$\sigma'(x) = \max(0, \min(1, 1.2\sigma(x) - 0.1)). \tag{3}$$

During training, a Gaussian noise $\eta \sim \mathcal{N}(0, 1)^D$ is added to $z_e(y)$ which is then passed through a saturating sigmoid to get the vector $f_e(y)$:

$$f_e(y) = \sigma'(z_e(y) + \eta). \tag{4}$$

To compute the discrete latent representation, the binary vector $g_e(y)$ is constructed via rounding, i.e.:

$$g_e(y)_i = \begin{cases} 1 & \text{if } f_e(y)_i > 0.5 \\ 0 & \text{otherwise,} \end{cases} \tag{5}$$

with the discrete latent code $z_d\left(y\right)$ corresponding to $\tau^{-1}_{\log K}(g(y))$. The input to the decoder during training is computed using two embedding spaces $e_1, e_2 \in R^{K \times D}$, with $z_q(y) = f_e(y)e_1 + (1 - f_e(y))e_2$, while during testing and evaluation one uses $z_q(y) = g_e(y)e_1 + (1 - g_e(y))e_2$.

2.3. Vector Quantization

The Vector Quantized - Variational Autoencoder (VQ-VAE) discretization bottleneck method was proposed in (van den Oord et al., 2017). Note that vector quantization based methods have a long history of being used successfully in various Hidden Markov Model (HMM) based machine learning models (see e.g., (Huang & Jack, 1989; Lee et al., 1989)). In VQ-VAE, the encoder output $\mathrm{enc}(y) \in R^D$ is passed through a discretization bottleneck using a nearest-neighbor lookup on embedding vectors $e \in R^{K \times D}$.

More specifically, the decoder input $z_q(y)$ is defined as

$$z_q(y) = e_k, \quad k = \arg\min_{j \in [K]} \|\mathrm{enc}(y) - e_j\|_2 . \tag{6}$$

The corresponding discrete latent $z_d\,(y)$ is then the index k of the embedding vector closest to $\mathrm{enc}(y)$ in ℓ_2 distance. Let l_r be the reconstruction loss of the decoder dec given $z_q(y)$, (e.g., the cross entropy loss); then the model is trained to minimize

$$L = l_r + \beta \|\mathrm{enc}(y) - \mathrm{sg}\,(z_q(y))\|_2 , \tag{7}$$

where $\mathrm{sg}\,(.)$ is the stop gradient operator. The embeddings are trained using an exponential moving average (EMA) of the encoder output $\mathrm{enc}(y)$ over a mini-batch of targets $\{y_1, \ldots, y_l, \ldots\}$:

$$e_j \leftarrow \lambda e_j + (1-\lambda) \frac{\sum_l \mathbb{1}\left[z_q(y_l) = e_l\right] \mathrm{enc}(y_l)}{\sum_l \mathbb{1}\left[z_q(y_l) = e_l\right]}, \tag{8}$$

where $\mathbb{1}[.]$ is the indicator function and λ is a suitable decay parameter.

2.4. Decomposed Vector Quantization

When the size of the discrete latent space K is large, then an issue with the approach of Section 2.3 is *index collapse*, where only a few of the embedding vectors get trained due to a rich getting richer phenomena. In particular, if an embedding vector e_j is close to a lot of encoder outputs $\mathrm{enc}(y_1), \ldots, \mathrm{enc}(y_i)$, then it receives the strongest signal to get even closer via the EMA update of Equation (8). Thus only a few of the embedding vectors will end up actually being used. To circumvent this issue, we propose two variants of decomposing VQ-VAE that make more efficient use of the embedding vectors for large values of K.

2.4.1. Sliced Vector Quantization

The main idea behind this approach is to break up the encoder output $\mathrm{enc}(y)$ into n_d smaller slices

$$\mathrm{enc}^1(y) \circ \mathrm{enc}^2(y) \cdots \circ \mathrm{enc}^{n_d}(y), \tag{9}$$

where each $\mathrm{enc}^i(y)$ is a D/n_d dimensional vector and $\circ$ denotes the concatenation operation. Corresponding to each $\mathrm{enc}^i(y)$ we have an embedding space $e^i \in R^{K' \times D/n_d}$, where $K' = 2^{(\log_2 K)/n_d}$. Note that the reason for the particular choice of K' is information theoretic: using an embedding space of size K from Section 2.3 allows us to express discrete codes of size $\log_2 K$. In the case when we have n_d different slices, we want the total expressible size of the discrete code to be still $\log_2 K$ and so K' is set to $2^{(\log_2 K)/n_d}$. We now compute nearest neighbors for each subspace as:

$$z_q^i(y) = e^i_{k_i}, \quad k_i = \arg\min_{j \in [K]} \left\|\mathrm{enc}^i(y) - e^i_j\right\|_2 , \tag{10}$$

with the decoder input being $z_q(y) = z_q^1(y) \circ \cdots \circ z_q^{n_d}(y)$.

The training objective L is the same as in Section 2.3, with each embedding space e^i trained individually via EMA updates from $\mathrm{enc}^i(y)$ over a mini-batch of targets $\{y_1, \ldots, y_l, \ldots\}$:

$$e^i_j \leftarrow \lambda e^i_j + (1-\lambda) \frac{\sum_l \mathbb{1}\left[z_q^i(y_l) = e^i_l\right] \mathrm{enc}^i(y_l)}{\sum_l \mathbb{1}\left[z_q^i(y_l) = e^i_l\right]}, \tag{11}$$

where $\mathbb{1}[.]$ is the indicator function as before, and λ is a suitable decay parameter.

Then the discrete latent code $z_d\,(y)$ is now defined as

$$z_d\,(y) = \tau^{-1}_{\log_2 K}\left(\tau_{\log_2 K'}(k_1) \circ \cdots \circ \tau_{\log_2 K'}(k_{n_d})\right). \tag{12}$$

Observe that when $n_d = 1$, the sliced Vector Quantization reduces to the VQ-VAE of (van den Oord et al., 2017). On the other hand, when $n_d = \log_2 K$, sliced DVQ is equivalent to improved semantic hashing of Section 2.2 loosely speaking: the individual table size K' for each slice $\mathrm{enc}^i(y)$ is 2, and it gets rounded to 0 or 1 depending on which embedding is closer. However, the rounding bottleneck in semantic hashing of Section 2.2 proceeds via a saturating sigmoid and thus strictly speaking, the two techniques are different.

Note that similar decomposition approaches to vector quantization in the context of HMMs have been studied in the past under the name *multiple code-books*, see for instance (Huang et al., 1989; Rogina & Waibel, 1994; Peinado et al., 1996). The approach of sliced Vector Quantization has also been studied more recently in the context of clustering, under the name of Product or Cartesian Quantization in (Jegou et al., 2011; Norouzi & Fleet, 2013). A more recent work (Shu & Nakayama, 2018) explores a similar decomposed quantization approach coupled with the Gumbel-Softmax trick to learn compressed word embeddings.

2.4.2. PROJECTED VECTOR QUANTIZATION

Another natural way to decompose Vector Quantization is to use a set of fixed randomly initialized projections

$$\left\{\pi^i \in R^{D\times D/n_d} \mid i \in [n_d]\right\} \tag{13}$$

to project the encoder output $\mathrm{enc}(y)$ into a R^{D/n_d}-dimensional subspace. For $\mathrm{enc}^i(y) = \pi^i(\mathrm{y}) \in R^{D/n_d}$ we have an embedding space $e^i \in R^{K'\times D/n_d}$, where $K' = 2^{(\log_2 K)/n_d}$ as before. The training objective, embeddings update, the input $z_q(y)$ to the decoder, and the discrete latent representation $z_d(y)$ is computed exactly as in Section 2.4.1. Note that when $n_d = 1$, projected Vector Quantization reduces to the VQ-VAE of (van den Oord et al., 2017) with an extra encoder layer corresponding to the projections π^i. Similarly, when $n_d = \log_2 K$, projected DVQ is equivalent to improved semantic hashing of Section 2.2 with the same analogy as in Section 2.4.1, except the encoder now has an extra layer.

3. Latent Transformer

Using the discretization techniques from Section 2 we can now introduce the Latent Transformer (LT) model. Given an input-output pair $(x, y) = (x_1, \ldots x_k, y_1, \ldots y_n)$ the LT will make use of the following components.

- The function $\mathrm{ae}(y, x)$ will autoencode y into a shorter sequence $l = l_1, \ldots, l_m$ of discrete latent variables using the discretization bottleneck from Section 2.
- The latent prediction model $\mathrm{lp}(l, x)$ (a Transformer) will autoregressively predict l based on x.
- The decoder $\mathrm{ad}(l, x)$ is a parallel model that will decode y from l and the input sequence x.

The functions $\mathrm{ae}(y, x)$ and $\mathrm{ad}(l, x)$ together form an autoencoder of the targets y that has additional access to the input sequence x. For the autoregressive latent prediction we use a Transformer (Vaswani et al., 2017), a model based on multiple self-attention layers that was originally introduced in the context of neural machine translation. In this work we focused on the autoencoding functions and did not tune the Transformer: we used all the defaults from the baseline provided by the Transformer authors (6 layers, hidden size of 512 and filter size of 4096) and only varied parameters relevant to ae and ad, which we describe below. The three components above give rise to two losses:

- The autoencoder reconstruction loss l_r coming from comparing $\mathrm{ad}(\mathrm{ae}(y, x), x)$ to y.
- The latent prediction loss l_{lp} that comes from comparing $l = \mathrm{ae}(y, x)$ to the generated $\mathrm{lp}(l, x)$.

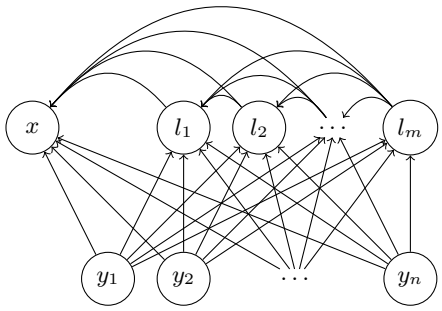

Figure 1. Dependence structure of the Latent Transformer in the form of a graphical model. We merged all inputs $x_1 \ldots x_k$ into a single node for easier visibility and we draw an arrow from node a to b if the probability of a depends on the generated b.

We train the LT model by minimizing $l_r + l_{lp}$. Note that the final outputs y are generated only depending on the latents l but not on each other, as depicted in Figure 1. In an autoregressive model, each y_i would have a dependence on all previous $y_j, j < i$, as is the case for ls in Figure 1.

The function $\mathrm{ae}(y, x)$**.** The autoencoding function $\mathrm{ae}(y, x)$ we use is a stack of residual convolutions followed by an attention layer attending to x and a stack of strided convolutions. We first apply to y a 3-layer block of 1-dimensional convolutions with kernel size 3 and padding with 0s on both sides (SAME-padding). We use ReLU nonlinearities between the layers and layer-normalization (Ba et al., 2016). Then, we add the input to the result, forming a residual block. Next we have an encoder-decoder attention layer with dot-product attention, same as in (Vaswani et al., 2017), with a residual connection. Finally, we process the result with a convolution with kernel size 2 and stride 2, effectively halving the size of s. We do this strided processing c times so as to decrease the length $C = 2^c$ times (later $C = \frac{n}{m}$). The result is put through the discretization bottleneck of Section 2. The final compression function is given by $\mathrm{ae}(y, x) = z_q(y)$ and the architecture described above is depicted in Figure 2.

The function $\mathrm{ad}(l, x)$**.** To decode from the latent sequence $l = \mathrm{ae}(y, x)$, we use the function $\mathrm{ad}(l, x)$. It consists of c steps that include up-convolutions that double the length, so effectively it increases the length $2^c = C = \frac{n}{m}$ times. Each step starts with the residual block, followed by an encoder-decoder attention to x (both as in the ae function above). Then it applies an up-convolution, which in our case is a feed-forward layer (equivalently a kernel-1 convolution) that doubles the internal dimension, followed by a reshape to twice the length. The result after the c steps is then passed to a self-attention decoder, same as in the Tranformer model (Vaswani et al., 2017).

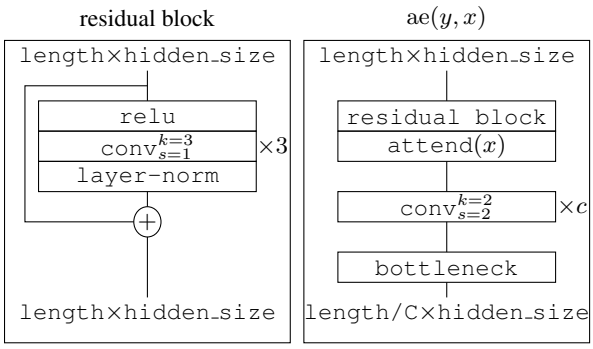

Figure 2. Architecture of $\mathrm{ae}(y, x)$. We write $\mathtt{conv}^{k=a}_{s=b}$ to denote a 1D convolution layer with kernel size a and stride b.

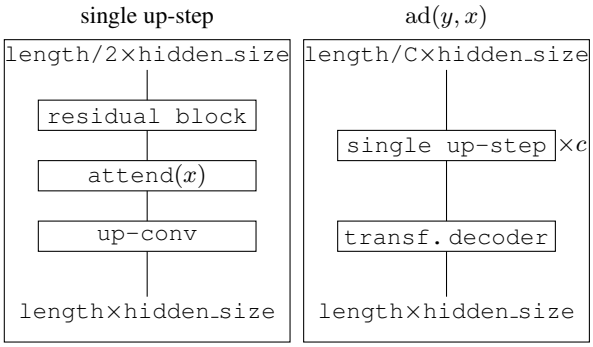

Figure 3. Architecture of $\mathrm{ad}(l, x)$. We write upconv to denote a 1D up-convolution layer.

Note that at the beginning of training (for the first 10K steps), we give the true targets y to the transformer-decoder here, instead of the decompressed latents l. This pre-training ensures that the self-attention part has reasonable gradients that are then back-propagated to the convolution stack and then back to the ae function and the discretization bottleneck of Section 2.

4. Related Work

Neural Machine Translation. Machine translation using deep neural networks achieved great success with sequence-to-sequence models (Sutskever et al., 2014; Bahdanau et al., 2014; Cho et al., 2014) that used recurrent neural networks (RNNs) with LSTM cells (Hochreiter & Schmidhuber, 1997). The basic sequence-to-sequence architecture is composed of an RNN encoder which reads the source sentence one token at a time and transforms it into a fixed-sized state vector. This is followed by an RNN decoder, which generates the target sentence, one token at a time, from the state vector. While a pure sequence-to-sequence recurrent neural network can already obtain good translation results (Sutskever et al., 2014; Cho et al., 2014), it suffers from the fact that the whole input sentence needs to be encoded into a single fixed-size vector. This clearly manifests itself in the degradation of translation quality on longer sentences and was overcome in (Bahdanau et al., 2014) by using a neural model of attention. Convolutional architectures have been used to obtain good results in word-level neural machine translation starting from (Kalchbrenner & Blunsom, 2013) and later in (Meng et al., 2015). These early models used a standard RNN on top of the convolution to generate the output, which creates a bottleneck and hurts performance. Fully convolutional neural machine translation without this bottleneck was first achieved in (Kaiser & Bengio, 2016) and (Kalchbrenner et al., 2016). The model in (Kaiser & Bengio, 2016) (Extended Neural GPU) used a recurrent stack of gated convolutional layers, while the model in (Kalchbrenner et al., 2016) (ByteNet) did away with recursion and used left-padded convolutions in the decoder. This idea, introduced in WaveNet (van den Oord et al., 2016), significantly improves efficiency of the model. The same technique was improved in a number of neural translation models recently, including (Gehring et al., 2017) and (Kaiser et al., 2017). Instead of convolutions, one can use stacked self-attention layers. This was introduced in the Transformer model (Vaswani et al., 2017) and has significantly improved state-of-the-art in machine translation while also improving the speed of training. Thus, we use the Transformer model as a baseline in this work.

Autoencoders and discretization bottlenecks. Variational autoencoders were first introduced in (Kingma & Welling, 2013; Rezende et al., 2014), however training them for discrete latent variable models has been challenging. The NVIL estimator of (Mnih & Gregor, 2014) proposes using a single sample objective to optimize the variational lower bound, while VIMCO (Mnih & Rezende, 2016) proposes using a muliti-sample objective of (Burda et al., 2015) which further speeds up convergence by using multiple samples from the inference network. There have also been several discretization bottlenecks proposed recently that have been used successfully in various learning tasks, see Section 2 for a more detailed description of the techniques directly relevant to this work.

Non-autoregressive Neural Machine Translation. Much of the recent state of the art models in Neural Machine Translation are auto-regressive, meaning that the model consumes previously generated tokens to predict the next one. A recent work that attempts to speed up decoding by training a non-autotregressive model is (Gu et al., 2017). The approach of (Gu et al., 2017) is to use the self-attention Transformer model of (Vaswani et al., 2017), together with the REINFORCE algorithm (Williams, 1992) to model the *fertilities* of words to tackle the multi-modality problem in translation. However, the main drawback of this

work is the need for extensive fine-tuning to make policy gradients work for REINFORCE, as well as the issue that this approach only works for machine translation and is not generic, so it cannot be directly applied to other sequence learning tasks.

Graphical models. The core of our approach to fast decoding consists of finding a sequence l of latent variables such that we can predict the output sequence y in parallel from l and the input x. In other words, we assume that each token y_i is conditionally independent of all other tokens y_j ($j \neq i$) given l and x: $y_i \perp\!\!\!\perp y_j \mid l, x$. Our autoencoder is thus learning to create a one-layer graphical model with m variables ($l_1 \dots l_m$) that can then be used to predict $y_1 \dots y_n$ independently of each other.

5. Experiments

We train the Latent Transformer with the base configuration to make it comparable to both the autoregressive baseline (Vaswani et al., 2017) and to the recent non-autoregressive NMT results (Gu et al., 2017). We used around 33K subword units as vocabulary and implemented our model in TensorFlow (Abadi et al., 2015). The implementation of various discretization bottlenecks including DVQ of Section 2.4 is available as open-source at [1] while the implementation of the Latent Prediction model is available at [2].

For non-autoregressive models, it is beneficial to generate a number of possible translations and re-score them with an autoregressive model. This can be done in parallel, so it is still fast, and it improves performance. This is called *noisy parallel decoding* in (Gu et al., 2017) and we include results both with and without it. The best BLEU scores obtained by different methods are summarized in Table 1. As you can see, our method with re-scoring almost matches the baseline autoregressive model without beam search.

To get a better understanding of the non-autoregressive models, we focus on performance without rescoring and investigate different variants of the Latent Transformer. We include different discretization bottlenecks, and report the final BLEU scores together with decoding speeds in Table 2. The LT is slower in non-batch mode than the simple NAT baseline of (Gu et al., 2017), which might be caused by system differences (our code is in TensorFlow and has not been optimized, while their implementation is in Torch). Latency at higher batch-size is much smaller, showing that the speed of the LT can still be significantly improved with batching.

[1] https://github.com/tensorflow/tensor2tensor/blob/master/tensor2tensor/layers/discretization.py

[2] https://github.com/tensorflow/tensor2tensor/blob/master/tensor2tensor/models/research/transformer_vae.py

Model	BLEU
Baseline Transformer [1]	27.3
Baseline Transformer [2]	23.5
Baseline Transformer [2] (no beam-search)	22.7
NAT+FT (no NPD) [2]	17.7
LT without rescoring $\left(\frac{n}{m} = 8\right)$	19.8
NAT+FT (NPD rescoring 10) [2]	18.7
LT rescornig top-10 $\left(\frac{n}{m} = 8\right)$	21.0
NAT+FT (NPD rescoring 100) [2]	19.2
LT rescornig top-100 $\left(\frac{n}{m} = 8\right)$	22.5

Table 1. BLEU scores (the higher the better) on the WMT English-German translation task on the `newstest2014` test set. The acronym NAT corresponds to the Non-Autoregressive Transformer, while FT denotes an additional Fertility Training and NPD denotes Nosiy Parallel Decoding, all of them from (Gu et al., 2017). The acronym LT denotes the Latent Transformer from Section 3. Results reported for LT are from this work, the others are from (Gu et al., 2017) [2] except for the first baseline Transformer result which is from (Vaswani et al., 2017) [1].

The choice of the discretization bottleneck seems to have a small impact on speed and both DVQ and improved semantic hashing yield good BLEU scores, while VQ-VAE fails in this context (see below for a discussion).

6. Discussion

Since the discretization bottleneck is critical to obtaining good results for fast decoding of sequence models, we focused on looking into it, especially in conjunction with the size K of the latent vocabulary, the dimension D of the latent space, and the number of decompositions n_d for DVQ.

An issue with the VQ-VAE of (van den Oord et al., 2017) that motivated the introduction of DVQ in Section 2.3 is *index collapse*, where only a few embeddings are used and subsequently trained. This can be visualized in the histogram of Figure 4, where the x-axis corresponds to the possible values of the discrete latents (in this case $\{1, \dots, K\}$), and the y-axis corresponds to the training progression of the model (time steps increase in a downward direction). On the other hand, using the DVQ from Section 2.4.1 with $n_d = 2$ leads to a much more balanced use of the available discrete latent space, as can be seen from Figure 5. We also report the percentage of available latent code-words used for different settings of n_d in Table 3; the usage of the code-words is maximized for $n_d = 2$.

The other variables for DVQ are the choice of the decomposition, and the number n_d of decompositions. For the projected DVQ, we use fixed projections $\{\pi^i\}$ initialized using the Glorot initializer (Glorot & Bengio, 2010). We also found that the optimal number of decompositions for our choice of latent vocabulary size $\log_2 K = 14$ and 16

Model	BLEU	Latency	
		$b = 1$	$b = 64$
Baseline (no beam-search)	22.7	408 ms	-
NAT	17.7	39 ms	-
NAT+NPD=10	18.7	79 ms	-
NAT+NPD=100	19.2	257 ms	-
LT, Improved Semhash	19.8	105 ms	8 ms
LT, VQ-VAE	2.78	148 ms	7 ms
LT, s-DVQ	19.8	177 ms	7 ms
LT, p-DVQ	19.7	182 ms	8 ms

Table 2. BLEU scores and decode times on the WMT English-German translation task on the `newstest2014` test set for different variants of the LT with $\frac{n}{m} = 8$ and $D = 512$ and $n_d = 2$ with s-DVQ and p-DVQ representing sliced and projected DVQ respectively. The LT model using improved semantic hashing from Section 2.2 uses $\log_2 K = 14$, while the one using VQ-VAE and DVQ from Sections 2.3 and 2.4 uses $\log_2 K = 16$. For comparison, we include the baselines from (Gu et al., 2017). We report the time to decode per sentence averaged over the whole test set as in (Gu et al., 2017); decoding is implemented in Tensorflow on a Nvidia GeForce GTX 1080. The batch size used during decoding is denoted by b and we report both $b = 1$ and $b = 64$.

was $n_d = 2$, with $n_d = 1$ (i.e., regular VQ-VAE) performing noticeably worse (see Table 2 and Figure 4). Setting higher values of n_d led to a decline in performance, possibly because the expressive power ($\log_2 K'$) was reduced for each decomposition, and the model also ended up using fewer latents (see Table 3).

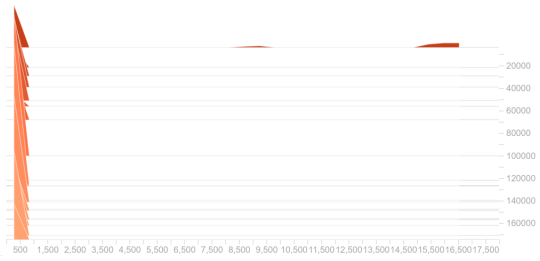

Figure 4. Histogram of discrete latent usage for VQ-VAE from (van den Oord et al., 2017), or equivalently sliced DVQ with $n_d = 1$ and $\log_2 K = 14$. The x-axis corresponds to the different possible discrete latents, while the y-axis corresponds to the progression of training steps (time increases in a downwards direction). Notice *index collapse* in the vanilla VQ-VAE where only a few latents ever get used.

Another important point about LT is that it allows making different trade-offs by tuning the $\frac{n}{m}$ fraction of the length of the original output sequence to the length of the latent sequence. As $\frac{n}{m}$ increases, so does the parallelism and

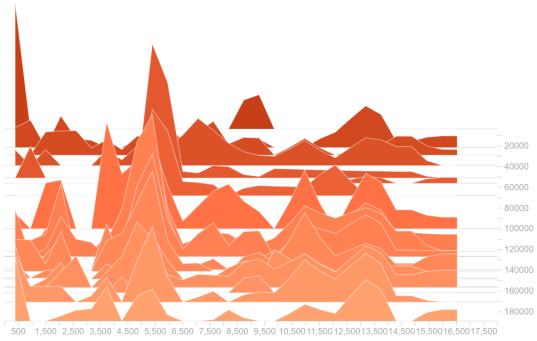

Figure 5. Histogram of discrete latent usage for sliced DVQ with $n_d = 2$ and $\log_2 K = 14$. The x-axis corresponds to the different possible discrete latents, while the y-axis corresponds to the progression of training steps (time increases in a downwards direction). Notice the diversity of latents used in this case.

n_d	Percentage of latents used
1	5%
2	74.5%
4	15.6%
8	31.2%

Table 3. Percentage of latent codewords used by the Decomposed Vector Quantization (DVQ) of Section 2.4.1 for $\log_2 K = 16$ and $D = 512$ after $500,000$ steps. Note that when $n_d = 1$, i.e. for vanilla VQ-VAE, only 5% of the available 2^{16} discrete latents (roughly 3000) are used. The latent usage is maximized for $n_d = 2$.

decoding speed, but the latents need to encode more and more information to be able to decode the outputs in parallel. To study this tradeoff, we measure the reconstruction loss (the perplexity of the reconstructed y vs the original) for different $\frac{n}{m}$ and varying the number of bits in the latent variables. The results, presented in Table 4, show clearly that reconstruction get better, as expected, if the latent state has more bits or is used to compress less subword units.

7. Conclusions

Autoregressive sequence models based on deep neural networks were made successful due to their applications in

$\frac{n}{m}$	$\log_2 K = 8$	$\log_2 K = 14$
2	1.33	0.64
4	2.04	1.26
8	2.44	1.77

Table 4. Log-perplexities of autoencoder reconstructions on the development set (`newstest2013`) for different values of $\frac{n}{m}$ and numbers of bits in latent variables (LT trained for 250K steps).

machine translation (Sutskever et al., 2014) and have since yielded state-of-the-art results on a number of tasks. With models like WaveNet and Transformer, it is possible to train them fast in a parallel way, which opened the way to applications to longer sequences, such as WaveNet for sound generation (van den Oord et al., 2016) or Transformer for long text summarization (Liu et al., 2018) and image generation (Vaswani et al., 2018). The key problem appearing in these new applications is the slowness of decoding: it is not practical to wait for minutes to generate a single example. In this work, while still focusing on the original problem of machine translation, we lay the groundwork for fast decoding for sequence models in general. While the latent transformer does not yet recover the full performance of the autoregressive model, it is already an order of magnitude faster and performs better than a heavily hand-tuned, task-specific non-autoregressive model.

In the future, we plan to improve both the speed and the accuracy of the latent transformer. A simple way to improve speed that we did not yet try is to use the methods from this work in a hierarchical way. As illustrated in Figure 1, the latents are still generated autoregressively which takes most of the time for longer sentences. In the future, we will apply the LT model to generate the latents in a hierarchical manner, which should result in further speedup. To improve the BLEU scores, on the other hand, we intend to investigate methods related to Gibbs sampling or even make the model partially autoregressive. For example, one could generate only the odd-indexed outputs, $y_1 y_3 y_5 \ldots$, based on the latent symbols l, and then generate the even-indexed ones based on both the latents and the odd-indexed outputs. We believe that including such techniques has the potential to remove the gap between fast-decoding models and purely autoregressive ones and will lead to many new applications.

References

Abadi, Martín, Agarwal, Ashish, Barham, Paul, Brevdo, Eugene, Chen, Zhifeng, Citro, Craig, Corrado, Greg, Davis, Andy, Dean, Jeffrey, Devin, Matthieu, Ghemawat, Sanjay, Goodfellow, Ian, Harp, Andrew, Irving, Geoffrey, Isard, Michael, Jia, Yangqing, Jozefowicz, Rafal, Kaiser, Lukasz, Kudlur, Manjunath, Levenberg, Josh, Mané, Dan, Monga, Rajat, Moore, Sherry, Murray, Derek, Olah, Chris, Schuster, Mike, Shlens, Jonathon, Steiner, Benoit, Sutskever, Ilya, Talwar, Kunal, Tucker, Paul, Vanhoucke, Vincent, Vasudevan, Vijay, Viégas, Fernanda, Vinyals, Oriol, Warden, Pete, Wattenberg, Martin, Wicke, Martin, Yu, Yuan, and Zheng, Xiaoqiang. Tensorflow: Large-scale machine learning on heterogeneous distributed systems, 2015. URL http://download.tensorflow.org/paper/whitepaper2015.pdf.

Ba, Jimmy Lei, Kiros, Jamie Ryan, and Hinton, Geoffrey E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. URL http://arxiv.org/abs/1607.06450.

Bahdanau, Dzmitry, Cho, Kyunghyun, and Bengio, Yoshua. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473, 2014. URL http://arxiv.org/abs/1409.0473.

Bowman, Samuel R., Vilnis, Luke, Vinyals, Oriol, Dai, Andrew M., Józefowicz, Rafal, and Bengio, Samy. Generating sentences from a continuous space. In *Proceedings of the SIGNLL'16*, pp. 10–21, 2016. URL https://arxiv.org/abs/1511.06349.

Burda, Yuri, Grosse, Roger, and Salakhutdinov, Ruslan. Importance weighted autoencoders. *CoRR*, abs/1509.00519, 2015. URL http://arxiv.org/abs/1509.00519.

Cho, Kyunghyun, van Merrienboer, Bart, Gulcehre, Caglar, Bougares, Fethi, Schwenk, Holger, and Bengio, Yoshua. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *CoRR*, abs/1406.1078, 2014. URL http://arxiv.org/abs/1406.1078.

Gehring, Jonas, Auli, Michael, Grangier, David, Yarats, Denis, and Dauphin, Yann N. Convolutional sequence to sequence learning. *CoRR*, abs/1705.03122, 2017. URL http://arxiv.org/abs/1705.03122.

Glorot, Xavier and Bengio, Yoshua. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 249–256, 2010.

Gu, Jiatao, Bradbury, James, Xiong, Caiming, Li, Victor O.K., and Socher, Richard. Non-autoregressive neural machine translation. *CoRR*, abs/1711.02281, 2017. URL http://arxiv.org/abs/1711.02281.

Hinton, Geoffrey E. and Salakhutdinov, Ruslan. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.

Hochreiter, Sepp and Schmidhuber, Jürgen. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

Huang, XD and Jack, MA. Unified techniques for vector quantization and hidden markov modeling using semi-continuous models. In *Acoustics, Speech, and Signal Processing, 1989. ICASSP-89., 1989 International Conference on*, pp. 639–642. IEEE, 1989.

Huang, XD, Hon, Hsiao-Wuen, and Lee, Kai-Fu. Multiple codebook semi-continuous hidden markov models for speaker-independent continuous speech recognition. 1989.

Jang, Eric, Gu, Shixiang, and Poole, Ben. Categorical reparameterization with gumbel-softmax. *CoRR*, abs/1611.01144, 2016. URL http://arxiv.org/abs/1611.01144.

Jegou, Herve, Douze, Matthijs, and Schmid, Cordelia. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33 (1):117–128, 2011.

Kaiser, Łukasz and Bengio, Samy. Can active memory replace attention? In *Advances in Neural Information Processing Systems*, pp. 3781–3789, 2016. URL https://arxiv.org/abs/1610.08613.

Kaiser, Łukasz and Bengio, Samy. Discrete autoencoders for sequence models. *CoRR*, abs/1801.09797, 2018. URL http://arxiv.org/abs/1801.09797.

Kaiser, Łukasz and Sutskever, Ilya. Neural GPUs learn algorithms. In *International Conference on Learning Representations (ICLR)*, 2016. URL https://arxiv.org/abs/1511.08228.

Kaiser, Łukasz, Gomez, Aidan N., and Chollet, Francois. Depthwise separable convolutions for neural machine translation. *CoRR*, abs/1706.03059, 2017. URL http://arxiv.org/abs/1706.03059.

Kalchbrenner, Nal and Blunsom, Phil. Recurrent continuous translation models. In *Proceedings EMNLP 2013*, pp. 1700–1709, 2013. URL http://www.aclweb.org/anthology/D13-1176.

Kalchbrenner, Nal, Espeholt, Lasse, Simonyan, Karen, van den Oord, Aaron, Graves, Alex, and Kavukcuoglu, Koray. Neural machine translation in linear time. *CoRR*, abs/1610.10099, 2016. URL http://arxiv.org/abs/1610.10099.

Kingma, Diederik P. and Welling, Max. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013. URL http://arxiv.org/abs/1312.6114.

Lee, K-F, Hon, H-W, Hwang, M-Y, Mahajan, Sanjoy, and Reddy, Raj. The sphinx speech recognition system. In *Acoustics, Speech, and Signal Processing, 1989. ICASSP-89., 1989 International Conference on*, pp. 445–448. IEEE, 1989.

Liu, Peter J., Saleh, Mohammad, Pot, Etienne, Goodrich, Ben, Sepassi, Ryan, Kaiser, Lukasz, and Shazeer, Noam. Generating wikipedia by summarizing long sequences. *CoRR*, abs/1801.10198, 2018. URL http://arxiv.org/abs/1801.10198.

Maddison, Chris J., Mnih, Andriy, and Teh, Yee Whye. The concrete distribution: A continuous relaxation of discrete random variables. *CoRR*, abs/1611.00712, 2016. URL http://arxiv.org/abs/1611.00712.

Meng, Fandong, Lu, Zhengdong, Wang, Mingxuan, Li, Hang, Jiang, Wenbin, and Liu, Qun. Encoding source language with convolutional neural network for machine translation. In *ACL*, pp. 20–30, 2015. URL https://arxiv.org/abs/1503.01838.

Mnih, Andriy and Gregor, Karol. Neural variational inference and learning in belief networks. *CoRR*, abs/1402.0030, 2014. URL http://arxiv.org/abs/1402.0030.

Mnih, Andriy and Rezende, Danilo. Variational inference for monte carlo objectives. In *International Conference on Machine Learning*, pp. 2188–2196, 2016.

Norouzi, Mohammad and Fleet, David J. Cartesian k-means. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pp. 3017–3024. IEEE, 2013.

Peinado, Antonio M, Segura, José C, Rubio, Antonio J, Garcia, Pedro, and Pérez, José L. Discriminative codebook design using multiple vector quantization in hmm-based speech recognizers. *IEEE transactions on speech and audio processing*, 4(2):89–95, 1996.

Rezende, Danilo Jimenez, Mohamed, Shakir, and Wierstra, Daan. Stochastic backpropagation and approximate inference in deep generative models. *CoRR*, abs/1401.4082, 2014. URL http://arxiv.org/abs/1401.4082.

Rogina, Ivica and Waibel, Alex. Learning state-dependent stream weights for multi-codebook hmm speech recognition systems. In *Acoustics, Speech, and Signal Processing, 1994. ICASSP-94., 1994 IEEE International Conference on*, volume 1, pp. I–217. IEEE, 1994.

Salakhutdinov, Ruslan and Hinton, Geoffrey E. Deep Boltzmann machines. In *Proceedings of AISTATS'09*, pp. 448–455, 2009a.

Salakhutdinov, Ruslan and Hinton, Geoffrey E. Semantic hashing. *Int. J. Approx. Reasoning*, 50(7):969–978, 2009b.

Shu, Raphael and Nakayama, Hideki. Compressing word embeddings via deep compositional code learning. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=BJRZzFlRb.

Sutskever, Ilya, Vinyals, Oriol, and Le, Quoc V. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*, pp. 3104–3112, 2014. URL http://arxiv.org/abs/1409.3215.

van den Oord, Aäron, Vinyals, Oriol, and Kavukcuoglu, Koray. Neural discrete representation learning. *CoRR*, abs/1711.00937, 2017. URL http://arxiv.org/abs/1711.00937.

van den Oord, Aäron, Dieleman, Sander, Zen, Heiga, Simonyan, Karen, Vinyals, Oriol, Graves, Alexander, Kalchbrenner, Nal, Senior, Andrew, and Kavukcuoglu, Koray. WaveNet: A generative model for raw audio. *CoRR*, abs/1609.03499, 2016. URL http://arxiv.org/abs/1609.03499.

Vaswani, Ashish, Shazeer, Noam, Parmar, Niki, Uszkoreit, Jakob, Jones, Llion, Gomez, Aidan N., Kaiser, Lukasz, and Polosukhin, Illia. Attention is all you need. *CoRR*, 2017. URL http://arxiv.org/abs/1706.03762.

Vaswani, Ashish, Parmar, Niki, Uszkoreit, Jakob, Shazeer, Noam, and Kaiser, Lukasz. Image transformer, 2018. URL https://openreview.net/forum?id=r16Vyf-0-.

Vincent, Pascal, Larochelle, Hugo, Lajoie, Isabelle, Bengio, Yoshua, and Manzagol, Pierre-Antoine. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11:3371–3408, 2010.

Vinyals, Kaiser, Koo, Petrov, Sutskever, and Hinton. Grammar as a foreign language. In *Advances in Neural Information Processing Systems*, 2015. URL http://arxiv.org/abs/1412.7449.

Williams, Ronald J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pp. 5–32. Springer, 1992.

Yang, Zichao, Hu, Zhiting, Salakhutdinov, Ruslan, and Berg-Kirkpatrick, Taylor. Improved variational autoencoders for text modeling using dilated convolutions. In *Proceedings of ICML'17*, pp. 3881–3890, 2017.