

---

# Mixture Models for Diverse Machine Translation: Tricks of the Trade

---

Tianxiao Shen<sup>\*1</sup> Myle Ott<sup>\*2</sup> Michael Auli<sup>2</sup> Marc’Aurelio Ranzato<sup>2</sup>

## Abstract

Mixture models trained via EM are among the simplest, most widely used and well understood latent variable models in the machine learning literature. Surprisingly, these models have been hardly explored in text generation applications such as machine translation. In principle, they provide a latent variable to control generation and produce a diverse set of hypotheses. In practice, however, mixture models are prone to degeneracies—often only one component gets trained or the latent variable is simply ignored. We find that disabling dropout noise in responsibility computation is critical to successful training. In addition, the design choices of parameterization, prior distribution, hard versus soft EM and online versus offline assignment can dramatically affect model performance. We develop an evaluation protocol to assess both quality and diversity of generations against multiple references, and provide an extensive empirical study of several mixture model variants. Our analysis shows that certain types of mixture models are more robust and offer the best trade-off between translation quality and diversity compared to variational models and diverse decoding approaches.<sup>1</sup>

## 1. Introduction

Machine translation (MT) is a challenging task not only because of the large and structured output space, but also because it is inherently a one-to-many mapping. There are often many plausible and semantically equivalent translations due to information asymmetry between different languages, e.g., translating from a language without grammatical gender to a language that has grammatical gender leads to two valid translation options, as well as different translation

styles such as formal/informal, literal/not literal, etc. This raises the question of how to model such *multi-modal* output distributions and how to evaluate these models.

Our first contribution is a better evaluation protocol that uses multiple references during evaluation to measure both the *quality* of translation and *diversity* of a generated hypothesis set. The second contribution of this paper is an in-depth empirical analysis of mixture models for machine translation, although we conjecture that the findings are general and might apply to other text generation tasks, such as dialogue, summarization, image captioning, etc.

Conditional mixture models, also known as mixture of experts (MoE) (Jacobs et al., 1991), are in principle well suited to generating diverse hypotheses which can be achieved through different mixture components. However, they have been largely overlooked in favor of models with richer latent structure (Zhang et al., 2016; Kaiser et al., 2018). There has been some previous work on mixture models for sequence to sequence learning (Shazeer et al., 2017; He et al., 2018), but these did not evaluate generations in terms of both quality and diversity, or they focus on a particular model variant. There is a lack of consensus whether mixture models are competitive with more complex models that rely on approximate Bayesian inference, whether they are plagued by the same “posterior collapse” degeneracy as variational models (Bowman et al., 2016), how model configurations affect performance and which one works best in practice.

This work considers all the major design choices involved in the construction of mixture models, including hard versus soft EM training, different parameterizations of mixture components, the choice of conditional prior, update frequency of responsibilities (also called membership weights), and how regularization noise is injected. We experiment on the large scale WMT English to German benchmark with a state-of-the-art model architecture and the results demonstrate intricate dependencies between these design choices. They also reveal that some ingredients are key to successful training of mixture models.

First, we show that mixture models are prone to degeneracies when trained with dropout noise, but that this can be mitigated by turning off dropout in the computation of responsibilities. The key to the specialization of experts is to make consistent use of them, and even a small amount of

---

<sup>\*</sup>Equal contribution <sup>1</sup>MIT CSAIL <sup>2</sup>Facebook AI Research. Correspondence to: Tianxiao Shen <tianxiao@csail.mit.edu>.

regularization noise can hamper that. Second, hard mixtures yield more diverse generations than soft mixtures, similar to how K-Means tends to find centroids that are farther apart from each other compared to the means found by a mixture of Gaussians (Kearns et al., 1998). Third, employing a uniform prior encourages *all* mixture components to produce good translations for any input source sentence, which is highly desirable. Finally, using independently parameterized mixture components provides greater diversifying capacity than shared parameters; but if responsibilities are refreshed online, independent parameterization is prone to a degeneracy where only a single component is trained because of the “rich gets richer” effect. Conversely, the combination of shared parameters and offline responsibility assignment may lead to another degeneracy, in which the mixture components fail to specialize and behave the same.

We extend our evaluation to three WMT benchmark datasets for which test sets with multiple human references are available. We demonstrate that mixture models, when successfully trained, consistently outperform variational NMT (Zhang et al., 2016) and diverse decoding algorithms such as diverse beam search (Li et al., 2017; Vijayakumar et al., 2018) and biased sampling (Graves, 2013; Fan et al., 2018). Our qualitative analysis shows that different mixture components can capture consistent translation styles across examples, enabling users to control generations in an interpretable and semantically meaningful way.

## 2. Related Work

Prior studies have investigated the prediction uncertainty in machine translation. Dreyer & Marcu (2012) and Galley et al. (2015) introduced new metrics to address uncertainty at evaluation time. Ott et al. (2018a) inspected the sources of uncertainty and proposed tools to check fitting between the model and the data distributions. They also observed that modern conditional auto-regressive NMT models can only capture uncertainty to a limited extent, and they tend to oversmooth probability mass over the hypothesis space.

Recent work has explored latent variable modeling for machine translation. Zhang et al. (2016) leveraged variational inference (Kingma & Welling, 2014; Bowman et al., 2016) to augment an NMT system with a single Gaussian latent variable. This work was extended by Schulz et al. (2018), who considered a sequence of latent Gaussian variables to represent each target word. Kaiser et al. (2018) proposed a similar model, but with groups of discrete multinomial latent variables. In their qualitative analysis, Kaiser et al. (2018) showed that the latent codes do affect the output predictions in interesting ways, but their focus was on speeding up regular decoding rather than producing a diverse set of hypotheses. None of these works analyzed and quantified diversity introduced by such latent variables.

The most relevant work is by He et al. (2018), who propose to use a soft mixture model with uniform prior for diverse machine translation. However, they did not evaluate on datasets with multiple references, nor did they analyze the full spectrum of design choices for building mixture models. Moreover, they used weaker base models and did not compare to variational NMT or diverse decoding baselines, which makes their empirical analysis less conclusive. We provide a comprehensive study and shed light on the different behaviors of mixture models in a variety of settings.

Besides machine translation, there is work on latent variables for dialogue generation (Serban et al., 2017; Cao & Clark, 2017; Wen et al., 2017) and image captioning (Wang et al., 2017; Dai et al., 2017). The proposed mixture model departs from these VAE or GAN-based approaches and importantly, is much simpler. It could also be applied to other text generation tasks as well.

## 3. Mixture Models for Diverse MT

A standard neural machine translation (NMT) model has an encoder-decoder structure. The encoder maps a source sentence  $x$  to a sequence of hidden states, which are then fed to the decoder to generate an output sentence one word at a time. At each time step, the decoder additionally conditions its output on the previous outputs, resulting in an auto-regressive factorization  $p(y|x; \theta) = \prod_{t=1}^T p(y_t|y_{1:t-1}, x; \theta)$ , where  $(y_1, \dots, y_T)$  are the words that compose a target sentence  $y$ .

However, the machine translation task has inherent uncertainty, due to the existence of multiple valid translations  $y$  for a given source sentence  $x$ . With the auto-regressive factorization all uncertainty is represented in the decoder output distribution, making it difficult to search for multiple modes of  $p(y|x; \theta)$ . Indeed, widely used decoding algorithms such as beam search typically produce hypotheses of low diversity with only minor differences in the suffix (Ott et al., 2018a).

Mixture models provide an alternative approach to modeling uncertainty and generating diverse translations. While these models have primarily been explored as a means of increasing model capacity (Jacobs et al., 1991; Shazeer et al., 2017), they are also a natural way of modeling different translation styles (He et al., 2018).

Formally, given a source sentence  $x$  and reference translation  $y$ , a mixture model introduces a multinomial latent variable  $z \in \{1, \dots, K\}$ , and decomposes the marginal likelihood as:

$$p(y|x; \theta) = \sum_{z=1}^K p(y, z|x; \theta) = \sum_{z=1}^K p(z|x; \theta) p(y|z, x; \theta) \quad (1)$$

where the prior  $p(z|x; \theta)$  and likelihood  $p(y|z, x; \theta)$  are learned functions parameterized by  $\theta$ . Each value of  $z$  represents an *expert*, and the posterior probability:

$$p(z|x, y; \theta) = \frac{p(z|x; \theta)p(y|z, x; \theta)}{\sum_{z'} p(z'|x; \theta)p(y|z', x; \theta)} \quad (2)$$

can be viewed as the *responsibility* each expert takes for explaining an observation  $(x, y)$ .

**Training** Given a training set  $\{(x^{(i)}, y^{(i)})\}_{i=1}^N$ , we want to find  $\theta$  that maximizes the log likelihood. To this end, for each training example we compute the gradient:

$$\nabla_{\theta} \log p(y^{(i)}|x^{(i)}; \theta) = \sum_z p(z|x^{(i)}, y^{(i)}; \theta) \cdot \nabla_{\theta} \log p(y^{(i)}, z|x^{(i)}; \theta) \quad (3)$$

and train the model with the EM algorithm (Dempster et al., 1977) by iteratively applying the following two steps:

**E-step:** estimate the responsibilities of each expert  $r_z^{(i)} \leftarrow p(z|x^{(i)}, y^{(i)}; \theta)$  using the current parameters  $\theta$ ;

**M-step:** update  $\theta$  through each expert with gradients  $\nabla_{\theta} \log p(y^{(i)}, z|x^{(i)}; \theta)$  weighted by their responsibilities  $r_z^{(i)}$ .

There are several ways to apply the E and M steps across training examples (Neal & Hinton, 1998); §3.3 will review those considered in this work.

**Decoding** All the decoding strategies for  $p(y|x; \theta)$  of a baseline model can be applied equally to  $p(y|z, x; \theta)$  in a mixture model. We adopt the most straightforward one: generating  $K$  hypotheses by first enumerating  $z$  and then greedily decoding  $\hat{y}_t = \arg \max_y p(y|\hat{y}_{1:t-1}, z, x; \theta)$ . Notably, this decoding procedure is efficient and easily parallelizable.

**Degeneracies** Unfortunately, naïve implementations of mixture models for text generation are prone to two major types of degeneracies:

**D1:** Only one component gets trained (Eigen et al., 2014; Shazeer et al., 2017) because of the “rich gets richer” effect whereby, once a component is slightly better than others, it is always picked while the other components starve and are eventually never used.

**D2:** The latent variable is ignored, similar to the collapse of variational auto-encoders where the posterior is always equal to the prior (Bowman et al., 2016).

In both cases, the model operates like a baseline model without any benefit from the latent variable. In practice, however, the chance of these degeneracies is heavily affected by a number of design decisions, which we describe in the following subsections.

### 3.1. Model Variants

Ideally, we would like different experts to specialize on different translation styles, so they can generate diverse hypotheses. Moreover, we want all of them to work well with any source sentence, so they will produce high quality translations. To this end, we explore several variants of the canonical mixture model, which vary by whether they use hard (h) or soft (s) assignments of responsibilities, and whether they use a learned prior (lp) or uniform prior (up).

The specialization of experts implies that the responsibility  $p(z|x, y; \theta)$  for explaining a particular translation  $y$  should be large for only one  $z$ , i.e., only one element in the sum  $\sum_z p(y, z|x; \theta)$  dominates. To encourage this, a *hard mixture model* directly optimizes for  $\max_z p(y, z|x; \theta)$  by assigning full responsibility for each training example to the expert with the largest joint probability. Training proceeds via hard-EM, where the M-step remains unchanged and the E-step becomes:

**E-step (hard):** estimate the responsibilities of each expert  $r_z^{(i)} \leftarrow \mathbb{1}[z = \arg \max_{z'} p(y^{(i)}, z'|x^{(i)}; \theta)]$  using the current parameters  $\theta$ .

This can also be seen as maximizing the marginal likelihood  $p(y|x; \theta)$  while minimizing the gap between the sum and the max, thus finding a balance between the two terms (Kearns et al., 1998).

To encourage all experts to generate good hypotheses for any source sentence, we may set the prior  $p(z|x; \theta)$  to be uniform. This can prevent the model from collapsing into only one working expert with extreme  $p(z|x; \theta)$  value. This also aligns with our simple decoding strategy that generates a single hypothesis from each expert.

The choices of soft (s) versus hard (h) mixture model (M) and learned prior (lp) versus uniform prior (up) give us four model variants with different loss functions:

$$\begin{aligned} \mathcal{L}_{\text{sMlp}}(\theta) &= \mathbb{E}_{(x,y) \sim \text{data}} \left[ -\log \sum_z p(z|x; \theta)p(y|z, x; \theta) \right] \\ \mathcal{L}_{\text{sMup}}(\theta) &= \mathbb{E}_{(x,y) \sim \text{data}} \left[ -\log \sum_z p(y|z, x; \theta) \right] \\ \mathcal{L}_{\text{hMlp}}(\theta) &= \mathbb{E}_{(x,y) \sim \text{data}} \left[ \min_z -\log p(z|x; \theta)p(y|z, x; \theta) \right] \\ \mathcal{L}_{\text{hMup}}(\theta) &= \mathbb{E}_{(x,y) \sim \text{data}} \left[ \min_z -\log p(y|z, x; \theta) \right] \end{aligned} \quad (4)$$

where the constant  $\log K$  term is omitted for models with a uniform prior.

Among these variants, the hMup objective is also known as multiple choice learning (MCL) for an ensemble of learners where the oracle loss is minimized (Guzman-Rivera et al., 2012), and He et al. (2018) has considered the sMup objective to train a sequence-to-sequence mixture model.

### 3.2. Parameterization

Another important design decision with mixture models is the degree of parameter sharing between experts. Using *independently* parameterized experts provides them with additional capacity to differentiate from one another, but may exacerbate overfitting since the number of parameters increases linearly with the number of experts. On the other hand, *sharing* parameters among experts may help mitigate degeneracy D1, whereby low quality experts are neglected and eventually “die” during training, since by sharing parameters even unpopular experts receive some gradients.

We test different model variants using both independent and shared parameters. With independent parameterization, each expert has a different decoder network. With shared parameterization, experts use the same decoder network but the beginning-of-sentence token at the start of the target sequence is replaced with an embedded representation of the latent variable. This requires a negligible increase in parameters over the baseline model.

### 3.3. Training Schedule

We consider two schedules for alternating between the E-step and M-step during training: *online* and *offline*. In the online EM algorithm, we minimize the loss via stochastic gradient descent, effectively interleaving the E-step and M-step for each mini-batch (Lee et al., 2016). In contrast, the offline EM algorithm performs the E-step for all training examples, trains each expert to convergence with the resulting responsibilities, and repeats. In practice, for offline training we perform the M-step for only a single epoch, rather than to convergence, before re-estimating the responsibilities.

### 3.4. Regularization

Deep neural networks with a large amount of parameters are prone to overfitting, and regularization via dropout is usually adopted to achieving good generalization performance. However, the key to make experts diversify as training progresses is to make consistent use of them, and we find that even a small amount of regularization noise in the computation of responsibilities (E-step) can hamper that. Indeed, we show in §5.1 that naïve use of dropout causes mixture models to ignore the latent variable, but this degeneracy is mitigated by disabling dropout in the E-step. We explore this phenomenon in more detail in Appendix A.

## 4. Metrics

In this section, we describe the metrics we use to quantitatively assess the quality and diversity of a set of translation hypotheses. We use BLEU (Papineni et al., 2002) based on word n-gram matching to measure corpus similarity.

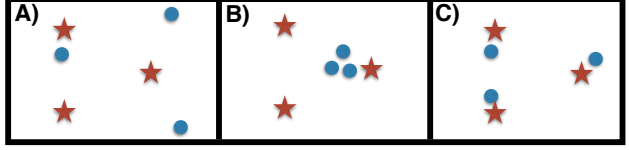


Figure 1. Toy illustration of how the metrics behave in various scenarios. Red stars represent human references, while blue dots represent system hypotheses. **A)** When there are poor hypotheses in the set, the BLEU is low and Pairwise-BLEU is low; this is the typical scenario of sampling hypotheses as well as degeneracy D1 described in §3. **B)** When hypotheses cluster closely together and in vicinity of a reference, the BLEU is high and Pairwise-BLEU is high; this is the typical scenario of beam search hypotheses and also degeneracy D2. **C)** When hypotheses match references (ideal case), the BLEU is high and Pairwise-BLEU matches human Pairwise-BLEU.

Suppose  $\{y^1, \dots, y^M\}$  are  $M$  reference translations of a source sentence  $x$ , and  $\{\hat{y}^1, \dots, \hat{y}^K\}$  are  $K$  hypotheses. Let  $\text{BLEU}\{([r_1, \dots, r_n], h)\}_{x \in \text{data}}$  denote the corpus-level BLEU for all pairs where  $h$  is a hypothesis and  $[r_1, \dots, r_n]$  is its reference list. Let  $[n]$  denote the set of  $\{1, \dots, n\}$ , and  $[y^{-i}]$  denote  $[y^1, \dots, y^{i-1}, y^{i+1}, \dots, y^M]$ . We compute the following two metrics<sup>2</sup> (Ott et al., 2018a):

- **Pairwise-BLEU:** To measure similarity among the hypotheses, we compare them with each other and compute  $\text{BLEU}\{([\hat{y}^j], \hat{y}^k)\}_{x \in \text{data}, j \in [K], k \in [K], j \neq k}$ .<sup>3</sup> The more diverse the hypothesis set, the lower the Pairwise-BLEU. Ideally, we would like a model with Pairwise-BLEU matching human Pairwise-BLEU.
- **BLEU:** We calculate human BLEU in a leave-one-out manner by computing  $\text{BLEU}\{([y^{-m}], y^m)\}_{x \in \text{data}}$  for  $m \in [M]$  and then averaging the  $M$  scores. We also use  $M - 1$  references when computing system BLEU, to be comparable with human scores, i.e. average  $\text{BLEU}\{([y^{-m}], \hat{y}^k)\}_{x \in \text{data}, k \in [K]}$  for  $m \in [M]$ . This measures the *overall quality* of a hypothesis set. If this metric scores low, it implies that some generated hypotheses have poor quality.

<sup>2</sup>See Appendix B for another diversity metric in terms of reference coverage.

<sup>3</sup>Several text GAN papers use Self-BLEU to evaluate diversity of unconditional generation (Yu et al., 2017), where each generated sentence is regarded as hypothesis against all other sentences as its reference list, i.e.  $\text{BLEU}\{([\hat{y}^{-k}], \hat{y}^k)\}_{x \in \text{data}, k \in [K]}$ . For machine translation, Pairwise-BLEU offers a more precise evaluation of diversity compared to Self-BLEU. For example, if a source sentence has two valid translations  $T_1 \neq T_2$ , system 1 provides 4 hypotheses  $H_1 = H_2 = H_3 = H_4 = T_1$ , and system 2 provides  $H_1 = H_2 = T_1, H_3 = H_4 = T_2$ , then their Self-BLEU are both 100; whereas Pairwise-BLEU for system 1 is 100 and for system 2 it is not, indicating that the latter is more diverse and more desirable (while they both have perfect quality).



Figure 1 illustrates how these metrics behave in different situations. The degeneracies outlined in §3 can be easily identified with these metrics: The first degeneracy is a situation where a single expert is responsible for all inputs (D1). This case can be identified when we measure both very low Pairwise-BLEU as well as very low BLEU, i.e., all but one latent value produce very bad generations. The second degeneracy happens when the latent variable is ignored and all experts behave almost identically (D2). This is evident when we observe good BLEU but extremely high Pairwise-BLEU (close to 100), i.e., when all latent values give good yet very similar outputs.

## 5. Experiments

**Datasets** We test mixture models and baselines on three benchmark datasets that uniquely provide multiple human references (Ott et al., 2018a; Hassan et al., 2018).

**WMT’17 English-German (En-De):** We train on all available bitext and filter sentence pairs that have source or target longer than 80 words, resulting in 4.5M sentence pairs. We use the Moses tokenizer (Koehn et al., 2007) and learn a joint source and target Byte-Pair-Encoding (Sennrich et al., 2016) with 32K types. We develop on newstest2013 and test on a 500 sentence subset of newstest2014 that has 10 reference translations (Ott et al., 2018a).

**WMT’14 English-French (En-Fr):** We borrow the setup of Gehring et al. (2017) with 36M training sentence pairs and 40K joint BPE vocabulary. We validate on newstest2012+2013, and test on a 500 sentence subset of newstest2014 with 10 reference translations (Ott et al., 2018a).

**WMT’17 Chinese-English (Zh-En):** We pre-process the training data following Hassan et al. (2018) which results in 20M sentence pairs, 48K and 32K source and target BPE vocabularies respectively. We develop on devtest2017 and report results on newstest2017 with 3 reference translations.

**Model Architecture** All models use a very similar architecture, built using the Transformer (Vaswani et al., 2017) implementation in the Fairseq toolkit (Ott et al., 2019). The encoder and decoder have 6 blocks. The number of attention heads, embedding dimension and inner-layer dimension are 8, 512, 2048 for the “base” configuration and 16, 1024, 4096 for the “big” configuration, respectively. We use the “base” configuration to compare mixture model variants in §5.1, and the “big” configuration for extended experiments in §5.2.

Mixture models with “independent” experts use independent decoders, while models with “shared” experts use a single shared decoder with an extra set of weights to embed each latent variable state. All mixture models use a shared encoder across the experts. Models that learn a prior, namely

sMlp and hMlp, have an additional module that averages the top encoder layer hidden states and predicts the conditional prior distribution  $p(z|x;\theta)$  via a one hidden layer neural network with a tanh activation in between; the number of hidden units matches the embedding dimension.

**Baselines** The first baseline we consider use the same architecture, without a latent variable, and use beam search or sampling to generate  $K$  hypotheses. We also consider the following modified versions: *Diverse Beam Search* (Vijayakumar et al., 2018) and *Biased Sampling* (Graves, 2013; Fan et al., 2018; Edunov et al., 2018). The former performs beam search sequentially and penalizes the selection of words used in previous generations. The latter improves upon straight sampling by sampling over the top- $k$  most likely words instead of all words at each step, resulting in less noisy generations. Finally, we also compare against *Variational NMT* (Zhang et al., 2016) which uses a Gaussian latent variable and variational inference for posterior approximation. We use a 512-dimensional latent space. At test time we first sample  $z$  from the conditional prior distribution and then do greedy decoding. Any additional hyperparameters for these baselines are tuned via grid search over the validation set.

**Experimental Details** Models are optimized with the Adam algorithm (Kingma & Ba, 2015) using  $\beta_1 = 0.9$ ,  $\beta_2 = 0.98$ , and  $\epsilon = 1e - 8$ . We use the same learning rate schedule as Ott et al. (2018b). We run experiments on between 8 and 128 Nvidia V100 GPUs with mini-batches of approximately 25K and 400K tokens for the experiments of §5.1 and §5.2, respectively, following Ott et al. (2018b).

### 5.1. Analysis of Mixture Models: Tricks of the Trade

In this section we compare the four variants of mixture models (§3.1), and consider for each such variant whether to share parameters among mixture components (§3.2), whether to update responsibilities once each epoch (offline mode) or after every gradient step (online mode) (§3.3), and the effect of regularization on them (§3.4). We train on the WMT’17 En-De dataset using the “base” Transformer configuration and  $K = 3$  mixture components for efficiency considerations.

**Effect of Dropout:** Naïvely optimizing the loss functions in Equation 4 with dropout noise causes mixture models to ignore the latent variable (D2). This can be seen by the very high Pairwise-BLEU (low diversity) when dropout is used in both the E-step and M-step (Table 1, E&M). Unfortunately, entirely disabling dropout exacerbates overfitting and gives quite worse translation quality, i.e., lower BLEU (Table 1, no). We provide a solution to this issue by decomposing the log-likelihood gradient as in Equation 3 and applying

	Pairwise-BLEU			BLEU		
	no	E&M	M	no	E&M	M
sMlp	53.9	98.5	57.8	58.2	66.4	64.4
sMup	55.4	98.5	61.6	58.2	67.7	64.2
hMlp	45.9	89.3	47.8	54.7	65.9	60.3
hMup	41.0	87.6	53.1	53.1	66.1	62.6

Table 1. Results on the WMT’17 En-De dataset with  $K = 3$  for mixture models with no dropout (no), dropout = 0.1 in both the E-step and M-step (E&M), or in the M-step only (M). For all model variants, dropout in the E-step causes the model to ignore the latent variable.

dropout only to the gradient computation of the experts (M-step) but not their responsibilities (E-step). This allows the model to make consistent use of the experts and helps them to diversify (Table 1, M; see also Appendix A). We adopt this dropout scheme for subsequent experiments.

**Mixture Model Exploration:** Figure 2 shows the results of different modeling choices we introduced in §3. We plot Pairwise-BLEU in descending order versus BLEU, so the right side indicates high diversity and the upper side indicates high quality. First, both Variational NMT (V-NMT) as well as soft mixture models with shared parameterization and offline responsibility updates lead to degeneracy D2 (upper left corner), where the latent variable is ignored and the generated hypotheses are almost the same, as indicated by Pairwise-BLEU close to 100. This is a well-known failure mode of VAEs for text (Bowman et al., 2016).<sup>4</sup> In the mixture model case, we surmise that training a shared mixture for a long time with the initial (random) responsibilities prevents the latent variable embeddings from specializing, and thus they provide no useful information to the decoder.

Second, mixture models trained with independently parameterized components and online responsibility assignment are prone to degeneracy D1 (lower right corner), whereby only a single expert gets trained due to the “rich gets richer” effect (Shazeer et al., 2017), and BLEU dramatically drops because of poor generations from other experts.<sup>5</sup> Parameter sharing alleviates this, as updates to one expert affect also the others.

Third, there are two settings that work consistently well: online responsibility update with shared parameters, and offline update with independent parameters. The latter yields more diverse but lower quality translations, as expected since components have more degrees of freedom to deviate from each other but also less data to train. Fourth and not surprisingly, soft assignment yields lower diversity than

<sup>4</sup>We also tried annealing the KL term for Variational NMT, but even then the latent variable is hardly used.

<sup>5</sup>We observe that the BLEU score is high for a single latent state but nearly zero for others, see Table 7 in Appendix C.

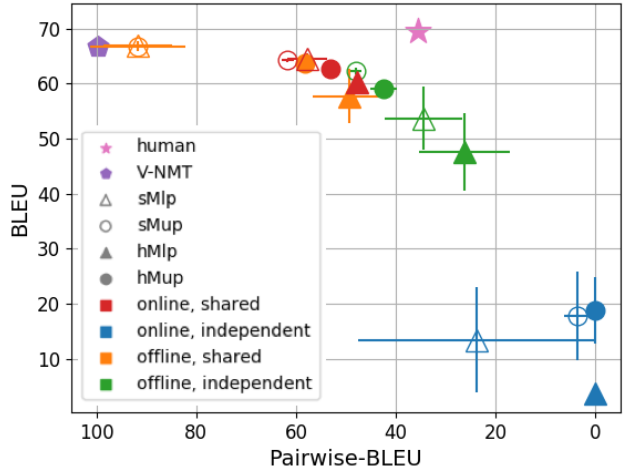


Figure 2. Comparison of mixture models with different design choices on the WMT’17 En-De dataset with  $K = 3$  mixture components (best viewed in color). Filled or empty markers show hard (h) or soft (s) mixture (M), and circle/triangle represent uniform prior (up)/learned prior (lp), respectively. Online/offline responsibility assignment and shared/independent parameterization are indicated by different colors. The star marker represents human performance, and pentagon represents Variational NMT. Error bars are computed over five random training seeds. Degeneracy D1 occurs in the lower right corner where some mixture components are hardly trained and produce poor generations. Degeneracy D2 is in the upper left, where the latent variable is ignored and different latent states produce almost the same hypotheses. We strive for configurations close to human performance, which are of high quality as well as diverse (upper right).

hard assignment (Kearns et al., 1998). Fifth, setting the prior to be uniform encourages the model to make use of all the components for each input source sentence, and thus gives better BLEU.

Overall, there are a few variants that work robustly as indicated by the small variance from random initialization, and strike a good balance between generation quality and diversity as indicated by proximity to human performance: all online-shared models, offline-shared hMup, as well as offline-independent sMup and hMup. All of these models perform well, offering slightly different trade-offs between quality and diversity.

If we also account for computational and memory cost, methods using shared parameters and hard EM are preferable, since the extra parameters are negligible and it requires only a single backward pass for the selected expert.<sup>6</sup> If we further consider simplicity of implementation, uniform prior and online responsibility update are favorable because they do not require additional model components, nor storing the

<sup>6</sup>With  $K = 10$  and the “base” Transformer architecture, hard mixture models are 2.4 times faster than the soft counterpart.

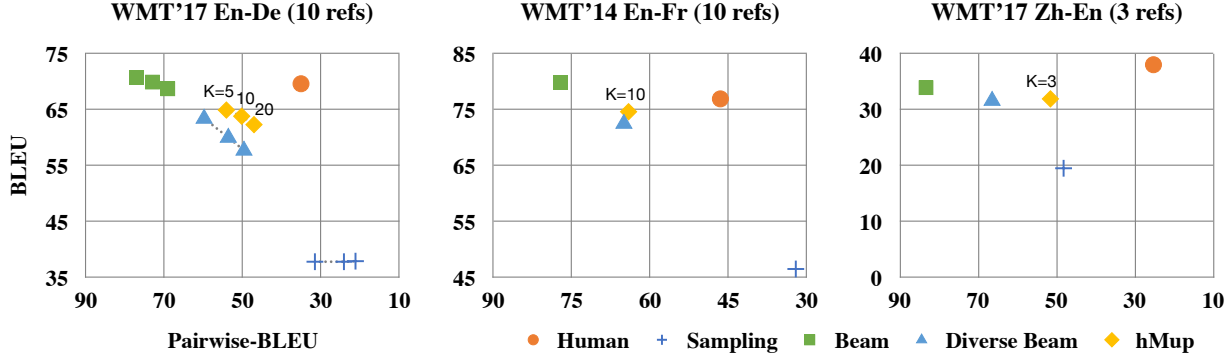


Figure 3. Comparison of the hMup mixture model with baselines on three WMT datasets. The mixture model provides the best trade-off between translation quality and diversity.

	Pairwise-BLEU			BLEU		
	en-de	en-fr	zh-en	en-de	en-fr	zh-en
sampling	24.1	32.0	48.2	37.8	46.5	19.5
beam	73.0	77.1	83.4	69.9	79.8	33.9
div-beam	53.7	64.9	66.5	60.0	72.5	31.6
hMup	50.2	64.0	51.6	63.8	74.6	31.9
human	35.5	46.5	25.3	69.6	76.9	38.0

Table 2. Results on three WMT datasets: En-De, En-Fr, Zh-En have 10, 10, 3 human references respectively. We generate the same number of hypotheses as the number of references.

responsibilities as for offline responsibility update. Taking these considerations into account, we adopt hMup (online-shared) for subsequent experiments, keeping in mind that the model variants mentioned above are likely to work similarly well.

## 5.2. Large-Scale Evaluation

We now extend our evaluation to three large benchmark datasets using the “big” Transformer configuration and a larger numbers of mixture components ( $K = 10, 10, 3$  for WMT En-De, En-Fr and Zh-En datasets respectively, to match the number of references we have on the test set).

Table 2 and Figure 3 show results of the hMup mixture model compared to other baseline approaches. The BLEU of beam search is fairly close to human, implying a remarkably high generation quality. However, it severely lacks diversity, as Pairwise-BLEU is close to 80. This suggests a scenario similar to Figure 1-B. In contrast, sampling hypotheses are very diverse, typically even more so than human references on En-De and En-Fr, but have a very poor BLEU, as illustrated in Figure 1-A. Diverse beam search increases the diversity of beam search to some extent, but pays a cost in translation quality. Overall, hMup achieves

the best trade-off between quality and diversity.

We also explore the impact of varying  $K$  on the WMT’17 En-De dataset (Figure 3, left). In general, when more hypotheses are generated, they become more diverse but of worse quality. Despite the improvement offered by hMup, there is still a large gap between the diversity of human references and system generations. The complete results of these experiments and the biased sampling baseline are provided in Appendix C.

## 5.3. Qualitative Analysis

In this section we perform a qualitative analysis with the WMT’17 Zh-En dataset. In Table 3 we show two source sentences, the corresponding reference translations, and generated hypotheses from different approaches. We see that beam search tends to produce generations that differ only in the last few words. Diverse beam search improves the diversity over beam search, but is not as diverse as hMup and may produce duplicate hypotheses (e.g., if the diversity penalty is not sufficiently high). hMup shows significant diversity in wording, word order, clause structure, etc.

To investigate whether the latent variable in hMup learns different translation styles, we examine the hypotheses generated from each latent state. For each value of  $z$ , we compute word frequencies of the corresponding generations and look for words whose frequency is significantly different as we change the value of the latent variable. We first discover that for words like *was*, *were* and *had*,  $z = 1$ ’s frequency is more than three times higher than  $z = 3$ ’s; conversely, for *has* and *says*,  $z = 3$ ’s frequency is more than twice higher than  $z = 1$ ’s. Since Chinese does not have tense unless a time phrase is explicitly stated, we speculate that when translating into English, the first latent value tends to translate with past tense whereas the third latent value tends to translate with present tense. Indeed we find that this is

Source	参与投票的成员中，58% 反对该合同交易。	自11 月份开始，俄罗斯民意也有所扭转。
References	It was rejected by 58 % of its members who voted in the ballot . Of the members who voted , 58 % opposed the contract transaction . Of the members who participated in the vote , 58 % opposed the contract .	Russian public opinion has also turned since November . Russian public opinion has started to change since November . The polls in Russian show a twist turn since the beginning of November .
Beam	Fifty-eight per cent of those voting opposed the contract deal . Fifty-eight per cent of the voting members opposed the contract deal . Fifty-eight per cent of the voting members opposed the contract transaction .	Since November , public opinion in Russia has also shifted . Since November , public opinion in Russia has also reversed . Since November , Russian public opinion has also shifted .
Diverse beam	Of the members voting , 58 per cent opposed the contract deal . Of the members voting , 58 per cent opposed the contract deal . Of the members voting , 58 per cent opposed the transaction .	Since November , public opinion in Russia has also shifted . Since November , the mood in Russia has also changed . Since November , public opinion in Russia has also been reversed .
hMup	Fifty-eight per cent of the members who voted opposed the contract deal . Of the members who voted , 58 % opposed the deal . Fifty-eight per cent of the voting members opposed the contract deal .	Since November , opinion in Russia has also reversed . Since November , the mood in Russia has also reversed . Opinion in Russia has also shifted since November .

Table 3. Examples of generations by the hMup mixture model and baselines on the WMT’17 Zh-En dataset. The mixture model shows considerable diversity compared to beam search and diverse beam search.

Source	这是一次规模巨大的作业，同时也是一次非常精密的作业。	他从不愿意与家人争吵。
Reference	This was a massive and , at the same time , very delicate operation .	He never wanted to be in any kind of altercation .
hMup	It was a huge job , and a very delicate one as well . It ’s a very large job , and it ’s a very delicate one , too . This is a huge job , but also a very delicate one .	He never liked to quarrel with his family . He never wants to quarrel with his family He never likes to argue with his family
Source	不断的恐怖袭击显然已对他造成很大打击。	我不想说这是我的最后一场比赛。
Reference	Repeat terror attacks on Turkey have clearly shaken him too .	I didn ’t want to say this was my last race .
hMup	The continuing terrorist attacks had apparently hit him hard . He is clearly already being hit hard by the continuing terrorist attacks . Repeated terrorist attacks have apparently hit him hard .	I didn ’t want to say it was my last game . I don ’t want to say it ’s my last game , he said . I don ’t want to say this is my last game .
Source	由此判断，这无疑是一场持续战。	两人2015 年缴纳了20.3% 的联邦税。
Reference	It appears that this was definitely an ongoing battle .	They paid a federal effective tax rate of 20.3 percent in 2015 .
hMup	Judging by that , it is undoubtedly a continuing battle . It is a battle that is no doubt ongoing . Judging by this , this is undoubtedly a continuing battle .	Both paid a federal tax of 20.3 per cent in 2015 . They paid a federal tax of 20.3 % in 2015 . Both paid 20.3 % of federal taxes in 2015 .

Table 4. More examples of the hMup mixture model on the WMT’17 Zh-En dataset. Different latent values learn to specialize for different translation styles consistently across examples, such as past tense vs. present tense, this vs. that, and per cent vs. %.

a consistent behavior, as seen from the first four examples in Table 4. Similarly, we find that different latent values exhibit different preferences for using `this` or `that` (see the fourth and fifth examples in Table 4), `%` or `per cent` (see the last example in Table 4 and the first example in Table 3), and so on.

## 6. Conclusion

Using large scale benchmarks and state-of-the-art architectures, we investigated 32 variants of mixture models, arising from the following design choices: use of hard versus soft EM, uniform versus learned prior, shared versus independent parameterization of mixture components, online versus offline responsibility update, and use of standard dropout regularization versus removal of this noise in responsibility computation. To the best of our knowledge this is the most extensive study of mixture models for conditional text generation to date, with machine translation as a use case. The simplicity of mixture models provides important advantages: the latent variable assignment can be explicitly enumerated

and the posterior can be computed exactly. Despite their simplicity, however, mixture models exhibit complex behaviors depending on different combinations of design choices. In particular, when instantiated with sub-optimal choices, they are prone to two typical failure modes—only one component gets trained and other components “die”, or the latent variable is ignored and all components behave the same. Our study provides insights into training deep sequence models with discrete latent variables.

Our recommended configurations enable mixture models to offer much better trade-offs between quality and diversity than variational models as well as heuristic diverse decoding approaches. In the future, we hope to broaden the scope of this work by looking at other generation tasks such as dialogue and image captioning. We would also like to investigate models with richer and more structured latent representations, and narrow the gap between model and human performance.



## Acknowledgements

We thank David Grangier for insightful discussions in the preliminary phase of this work. We also thank MIT NLP group for their helpful comments.

## References

- Bowman, S. R., Vilnis, L., Vinyals, O., Dai, A. M., Jozefowicz, R., and Bengio, S. Generating sentences from a continuous space. In *SIGNLL Conference on Computational Natural Language Learning*, 2016.
- Cao, K. and Clark, S. Latent variable dialogue models and their diversity. In *Proceedings of EACL (Short Papers)*, pp. 182–187, 2017.
- Dai, B., Fidler, S., Urtasun, R., and Lin, D. Towards diverse and natural image descriptions via a conditional gan. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2989–2998. IEEE, 2017.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pp. 1–38, 1977.
- Dreyer, M. and Marcu, D. Hyter: Meaning-equivalent semantics for translation evaluation. In *ACL*, 2012.
- Edunov, S., Ott, M., Auli, M., and Grangier, D. Understanding back-translation at scale. In *Proc. of EMNLP*, 2018.
- Eigen, D., Ranzato, M., and Sutskever, I. Learning factored representations in a deep mixture of experts. *ICLR*, 2014.
- Fan, A., Lewis, M., and Dauphin, Y. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Long Papers)*, pp. 889–898. Association for Computational Linguistics, 2018.
- Galley, M., Brockett, C., Sordoni, A., Ji, Y., Auli, M., Quirk, C., Mitchell, M., Gao, J., and Dolan, B. deltableu: A discriminative metric for generation tasks with intrinsically diverse targets. In *Proceedings of ACL (Short Papers)*, pp. 445–450, 2015.
- Gehring, J., Auli, M., Grangier, D., Yarats, D., and Dauphin, Y. N. Convolutional Sequence to Sequence Learning. In *Proc. of ICML*, 2017.
- Graves, A. Generating sequences with recurrent neural networks. *arXiv:1308.0850*, 2013.
- Guzman-Rivera, A., Batra, D., and Kohli, P. Multiple choice learning: Learning to produce multiple structured outputs. In *Advances in Neural Information Processing Systems*, pp. 1799–1807, 2012.
- Hassan, H., Aue, A., Chen, C., Chowdhary, V., Clark, J., Federmann, C., Huang, X., Junczys-Dowmunt, M., Lewis, W., Li, M., et al. Achieving human parity on automatic chinese to english news translation. *arXiv preprint arXiv:1803.05567*, 2018.
- He, X., Haffari, G., and Norouzi, M. Sequence to sequence mixture model for diverse machine translation. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pp. 583–592, 2018.
- Jacobs, R. A., Jordan, M. I., Nowlan, S., , and Hinton, G. E. Adaptive mixtures of local experts. *Neural Computation*, 3:1—12, 1991.
- Kaiser, L., Bengio, S., Roy, A., Vaswani, A., Parmar, N., Uszkoreit, J., and Shazeer, N. Fast decoding in sequence models using discrete latent variables. In *International Conference on Machine Learning*, pp. 2395–2404, 2018.
- Kearns, M., Mansour, Y., and Ng, A. Y. An information-theoretic analysis of hard and soft assignment methods for clustering. In *Learning in graphical models*, pp. 495–520. Springer, 1998.
- Kingma, D. P. and Ba, J. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014.
- Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., Dyer, C., Bojar, O., Constantin, A., and Herbst, E. Moses: Open source toolkit for statistical machine translation. In *ACL Demo Session*, 2007.
- Lee, S., Purushwalkam, S., Cogswell, M., Ranjan, V., Crandall, D., and Batra, D. Stochastic multiple choice learning for training diverse deep ensembles. In *NIPS*, 2016.
- Li, J., Monroe, W., and Jurafsky, D. A simple, fast diverse decoding algorithm for neural generation. *arXiv:1611.08562v2*, 2017.
- Neal, R. M. and Hinton, G. E. A view of the em algorithm that justifies incremental, sparse, and other variants. In *Learning in graphical models*, pp. 355–368. Springer, 1998.
- Ott, M., Auli, M., Grangier, D., and Ranzato, M. Analyzing uncertainty in neural machine translation. In *International Conference of Machine Learning*, 2018a.
- Ott, M., Edunov, S., Grangier, D., and Auli, M. Scaling neural machine translation. In *Proceedings of the Third*

- Conference on Machine Translation: Research Papers*, 2018b.
- Ott, M., Edunov, S., Baevski, A., Fan, A., Gross, S., Ng, N., Grangier, D., and Auli, M. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pp. 311–318. Association for Computational Linguistics, 2002.
- Schulz, P., Aziz, W., and Cohn, T. A stochastic decoder for neural machine translation. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 2018.
- Sennrich, R., Haddow, B., and Birch, A. Neural machine translation of rare words with subword units. In *Conference of the Association for Computational Linguistics (ACL)*, 2016.
- Serban, I. V., Sordoni, A., Lowe, R., Charlin, L., Pineau, J., Courville, A. C., and Bengio, Y. A hierarchical latent variable encoder-decoder model for generating dialogues. In *AAAI*, pp. 3295–3301, 2017.
- Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., and Dean, J. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations*, 2017.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, pp. 5998–6008, 2017.
- Vijayakumar, A. K., Cogswell, M., Selvaraju, R. R., Sun, Q., Lee, S., Crandall, D., and Batra, D. Diverse beam search for improved description of complex scenes. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Wang, L., Schwing, A., and Lazebnik, S. Diverse and accurate image description using a variational auto-encoder with an additive gaussian encoding space. In *Advances in Neural Information Processing Systems*, pp. 5756–5766, 2017.
- Wen, T.-H., Miao, Y., Blunsom, P., and Young, S. Latent intention dialogue models. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 3732–3741, 2017.
- Yu, L., Zhang, W., Wang, J., and Yu, Y. Seqgan: Sequence generative adversarial nets with policy gradient. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Zhang, B., Xiong, D., Su, J., Duan, H., and Zhang, M. Variational neural machine translation. *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2016.

## A. Effect of Dropout

In Section 5.1 we observed that it is crucial to turn off dropout during the computation of responsibilities (E-step) to avoid model collapse, see Table 1. In this section, we further investigate the impact of dropout on the responsibility computation, using hMup as an example.

We speculate that dropout noise weakens the dependency on the latent variable, causing the hard E-step to select among the latent values at random. This prevents different latent states from specializing and ultimately causes the model to ignore them.

To test this hypothesis, we show in Figure 4 the effect of dropout noise on the E-step at the beginning of training (i.e., with a randomly initialized model). On the y-axis we plot how often the optimal value of  $z$  changes after applying dropout with different rates. We see that as we increase the dropout probability, the optimal value of  $z$  is quickly corrupted—even with a small dropout probability of 0.1 we observe a 42% chance that the optimal assignment of  $z$  changes.

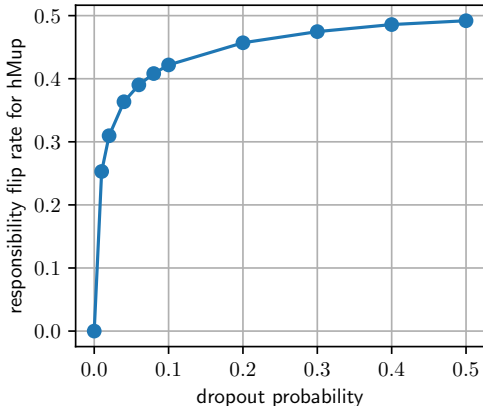


Figure 4. The effect of dropout on the E-step at the beginning of training. On the x-axis the dropout rate, and on the y-axis the fraction of times that the (hard) responsibility assignment is changed when dropout noise is turned on. Even a small amount of dropout noise causes excessive randomness of responsibility assignment, which in turn makes the model ignore the latent variable. It is therefore important to turn off dropout when estimating responsibilities. Experiments are performed on the WMT’17 En-De dataset with 2 latent categories ( $K = 2$ ) and the “base” Transformer architecture.

## B. Another Diversity Metric

In addition to Pairwise-BLEU, we also consider another metric to evaluate the diversity of a set of hypotheses, the **reference coverage**.

We pair each hypothesis to its best matching reference

(breaking ties randomly), count how many distinct references are matched to at least one hypothesis, and average this count over all sentences in the test set. A low coverage number indicates that all hypotheses are close to a few references. Instead, we would like a diverse set that covers most of the references.

Compared to Pairwise-BLEU, this metric offers more intuitive numerical values, as it ranges between 1 and the total number of available references. In the next section, we will report both values for completeness.

## C. Detailed Results

Table 5 compares two well performing mixture model variants, namely (online, shared) hMup and sMup, to several baselines on the three benchmark datasets we have considered in this work, expanding on the results reported in Table 2 and Figure 3 of the main paper. Once again, mixture models offer a good trade-off between translation quality and diversity.

In Table 6 we compare different approaches for generating diverse translations on the WMT’17 En-De dataset. We additionally compare each approach as we vary the number of desired translations ( $K$ ) (see also Figure 3, left). We observe that sampling produces diverse but low quality outputs. We can improve translation quality by restricting sampling to the top- $k$  candidates at each output position ( $k = 2$  performed best), but translation quality is still worse than hMup. Beam search produces the highest quality outputs, but with low diversity. Diverse beam search provides a reasonable balance between diversity and translation quality, but hMup produces even more diverse and better quality translations. Finally, except for unrestricted sampling, hMup covers the largest number of references among all the approaches evaluated.

We conclude with Table 7 which shows the values used to generate Figure 2, together with other metrics, such as corpus level BLEU with hypotheses generated by a fixed latent variable state throughout the whole test set. This metric is useful to detect models affected by degeneracy D1, as there will be states that yield very low corpus level BLEU because they rarely generate good hypotheses.

	Pairwise-BLEU			BLEU			#refs covered		
	en-de	en-fr	zh-en	en-de	en-fr	zh-en	en-de	en-fr	zh-en
Sampling	24.1	32.0	48.2	37.8	46.5	19.5	4.6	4.3	1.5
Beam	73.0	77.1	83.4	69.9	79.8	33.9	3.1	3.2	1.3
Diverse beam	53.7	64.9	66.5	60.0	72.5	31.6	3.7	3.5	1.4
sMup (He et al., 2018)	68.9	80.4	60.9	68.1	79.6	32.8	2.9	2.7	1.4
hMup	50.2	64.0	51.6	63.8	74.6	31.9	4.0	3.7	1.6
Human	35.5	46.5	25.3	69.6	76.9	38.0	-	-	-

Table 5. Results on three WMT datasets. Extended version of Table 2, including the results of another mixture model sMup (He et al., 2018) and the reference coverage metric. hMup and sMup provide different trade-offs between quality and diversity: the former is more diverse (lower Pairwise-BLEU), while the latter gives higher translation quality (BLEU).

	Pairwise-BLEU			BLEU			#refs covered		
	$K = 5$	10	20	5	10	20	5	10	20
Sampling	31.6	24.1	21.2	37.8	37.8	37.9	3.1	4.6	6.2
Biased sampling (top-2)	49.3	47.8	46.7	59.7	60.0	60.4	2.7	3.7	4.7
Beam	77.1	73.0	69.1	70.7	69.9	68.7	2.3	3.1	4.0
Diverse beam	59.8	53.7	49.7	63.4	60.0	57.7	2.5	3.7	4.8
hMup	54.2	50.2	47.1	64.9	63.8	62.3	2.8	4.0	5.3

Table 6. Results on the WMT’17 En-De dataset with various numbers of generations ( $K$ ). We compare: multinomial sampling (Sampling); sampling restricted to the top- $k$  candidates at each step (Biased sampling (top-2));  $k=2$  performed best); beam search with varying beam widths (Beam); diverse beam search (Vijayakumar et al., 2018) with varying number of outputs (Diverse beam; note that the number of groups  $G$  and diversity strength are tuned separately for each value of  $K$ ); and the hMup mixture model with  $K$  components (hMup).

schedule	parameterization	loss	BLEU per latent			Pairwise-BLEU	BLEU	#refs covered
			$z = 1$	2	3			
online	shared	sMlp	25.9	24.9	22.6	57.8 (4.0)	64.4 (1.0)	1.9
online	shared	sMup	25.8	25.2	22.8	61.6 (1.3)	64.2 (0.4)	1.9
online	shared	hMlp	25.5	22.6	21.5	47.8 (0.6)	60.3 (0.3)	2.1
online	shared	hMup	25.6	24.4	21.3	53.1 (1.2)	62.6 (0.6)	2.1
online	independent	sMlp	25.5	0.0	0.0	23.8 (23.8)	13.4 (9.6)	2.6
online	independent	sMup	25.8	0.6	0.0	3.5 (2.9)	17.8 (8.1)	2.6
online	independent	hMlp	26.1	0.0	0.0	0.0 (0.0)	3.7 (0.0)	2.7
online	independent	hMup	25.7	0.2	0.0	0.1 (0.1)	18.7 (6.0)	2.4
offline	shared	sMlp	25.8	25.7	25.6	91.7 (6.8)	66.8 (0.7)	1.4
offline	shared	sMup	25.7	25.5	25.3	91.7 (9.6)	66.8 (0.9)	1.4
offline	shared	hMlp	25.6	21.2	19.4	49.2 (7.3)	57.7 (5.0)	2.1
offline	shared	hMup	25.3	24.4	23.6	58.2 (1.5)	63.7 (0.6)	2.0
offline	independent	sMlp	25.7	19.4	16.4	34.4 (7.8)	53.6 (5.7)	2.2
offline	independent	sMup	25.3	23.5	22.9	48.0 (1.1)	62.3 (0.6)	2.1
offline	independent	hMlp	25.7	15.8	13.0	26.4 (9.1)	47.6 (7.0)	2.4
offline	independent	hMup	25.5	22.5	19.8	42.5 (2.6)	59.1 (1.6)	2.2

Table 7. Comparison of mixture models with different design choices on the WMT’17 En-De dataset with  $K = 3$  mixture components. See §3.1, §3.2, §3.3 for a detailed discussion about these model configurations. Pairwise-BLEU versus BLEU are plotted in Figure 2. Each configuration was run five times with different random seeds. We report the mean value of each metric (and for Pairwise-BLEU and BLEU also the standard deviation in parentheses). We also report corpus level BLEU w.r.t. one reference when greedily decoding the test set using a fixed latent variable state (columns labeled with  $z = 1, 2, 3$ ). Configurations that have values colored in red exhibit degeneracy of type D1, while configurations that have values colored in blue exhibit degeneracy of type D2.