

VARIATIONAL SMOOTHING IN RECURRENT NEURAL NETWORK LANGUAGE MODELS

Lingpeng Kong, Gabor Melis, Wang Ling, Lei Yu, Dani Yogatama

DeepMind

{lingpenk, melisgl, lingwang, leiyu, dyogatama}@google.com

ABSTRACT

We present a new theoretical perspective of data noising in recurrent neural network language models (Xie et al., 2017). We show that each variant of data noising is an instance of Bayesian recurrent neural networks with a particular variational distribution (i.e., a mixture of Gaussians whose weights depend on statistics derived from the corpus such as the unigram distribution). We use this insight to propose a more principled method to apply at prediction time and propose natural extensions to data noising under the variational framework. In particular, we propose variational smoothing with tied input and output embedding matrices and an element-wise variational smoothing method. We empirically verify our analysis on two benchmark language modeling datasets and demonstrate performance improvements over existing data noising methods.

1 INTRODUCTION

Recurrent neural networks (RNNs) are state of the art models in various language processing tasks. However, their performance heavily depends on proper regularization at training time (Melis et al., 2018b; Merity et al., 2018). The two predominant approaches to regularize RNNs are dropout (randomly zeroing out neurons; Srivastava et al., 2014) and ℓ_2 regularization (applying ℓ_2 penalty to model parameters; Hoerl & Kennard, 1970). Recently, Xie et al. (2017) proposed data noising to regularize language models. Their method is formulated as a data augmentation method that randomly replaces words with other words drawn from a proposal distribution. For example, we can use the unigram distribution which models the number of word occurrences in the training corpus, or a more sophisticated proposal distribution that takes into account the number of bigram types in the corpus. Data noising has been shown to improve perplexity on small and large corpora, and that this improvement is complementary to other regularization techniques and translates to improvements on downstream models such as machine translation.

Xie et al. (2017) derived connections between data noising and smoothing in classical n -gram language models, which we review in §2. In smoothing (Chen & Goodman, 1996), since empirical counts for unseen sequences are zero, we smooth our estimates by a weighted average of higher order and lower order n -gram models. There are various ways to choose the weights and the lower order models leading to different smoothing techniques, with Kneser-Ney smoothing widely considered to be the most effective. Xie et al. (2017) showed that the pseudocounts of a noised data correspond to a mixture of different n -gram models.

In this paper, we provide a new theoretical foundation for data noising and show that it can be understood as a form of Bayesian recurrent neural network with a particular variational distribution (§4.1 and §4.2). Our derivation relates data noising to dropout and variational dropout (Gal & Ghahramani, 2016b), and naturally leads to a data dependent ℓ_2 regularization coefficient. We use this insight to arrive at a more principled way to do prediction with data noising—i.e., by taking the mean of the variational distribution, as opposed to the mode—and propose several extensions under the variational framework in §4.3. Specifically, we show how to use variational smoothing with tied input and output embeddings and propose element-wise smoothing. In §5, we validate our analysis in language modeling experiments on the Penn Treebank (Marcus et al., 1994) and Wikitext-2 (Merity et al., 2017) datasets.

2 RECURRENT NEURAL NETWORK LANGUAGE MODELS

We consider a language modeling problem where the goal is to predict the next word x_t given previously seen context words $x_{<t} = \{x_0, x_1, \dots, x_{t-1}\}$. Let \mathbf{W} be parameters of a recurrent neural network, and $p(x_t | x_{<t}; \mathbf{W}) = \text{RNN}(x_{<t}; \mathbf{W})$. Following previous work in language modeling (Melis et al., 2018b; Merity et al., 2018), we use LSTM (Hochreiter & Schmidhuber, 1997) as our RNN function, although other variants such as GRU (Cho et al., 2014) can be used as well.

Given a training corpus $\mathcal{D} = \{x_0, \dots, x_T\}$, the likelihood we would like to maximize is:

$$\mathcal{L}(\mathbf{W}; \mathcal{D}) = \prod_{t=1}^T p(x_t | x_{<t}; \mathbf{W}).$$

Directly optimizing the (log) likelihood above often leads to overfitting. We typically augment the objective function with a regularizer (e.g., ℓ_2 regularizer where $\mathcal{R}(\mathbf{W}) = \lambda \|\mathbf{W}\|_2^2$) or use dropout by randomly zeroing out neurons (Srivastava et al., 2014).

Data noising as smoothing. Xie et al. (2017) proposed a method to regularize recurrent neural network language models by noising the data. For each input word i in $x_{<t}$ (sometimes also the corresponding output word), we replace it with another word sampled from a proposal distribution \mathcal{J} with probability γ_i . They introduced various methods on how to choose γ_i and \mathcal{J} .

For example, if $\gamma_i = \gamma$ for all $i \in \mathcal{V}$ and \mathcal{J} is the unigram distribution of words in the corpus, $\mathbb{E}_{\mathcal{J}}[p(x_t | \text{noise}_{\gamma, \mathcal{J}}(x_{<t}); \mathbf{W})]$ corresponds to a mixture of n -gram models with fixed weights (i.e., linear interpolation smoothing).

Another option is to set $\gamma_i = \gamma \frac{\text{distinct}(i, \bullet)}{\text{count}(i)}$, where $\text{distinct}(i, \bullet)$ denotes the number of distinct continuations preceded by word i (i.e., the number of bigram types that has i as the first word), and $\text{count}(i)$ denotes the number of times i appears in the corpus. For this choice of γ_i , when \mathcal{J} is the unigram distribution, the expectation corresponds to absolute discounting; whereas if $\mathcal{J}_i = \frac{\text{distinct}(\bullet, i)}{\sum_{v \in \mathcal{V}} \text{distinct}(\bullet, v)}$ and we replace both the input and output words, it corresponds to bigram Kneser-Ney smoothing. We summarize their proposed methods in Table 1.

At prediction (test) time, Xie et al. (2017) do not apply any noising and directly predict $p(x_t | x_{<t}; \mathbf{W})$. They showed that a combination of smoothing and dropout achieves the best result on their language modeling and machine translation experiments.

Name	Noised	γ_i	\mathcal{J}_x
Blank noising	x_0	γ	$\mathcal{J}_{_} = 1$
Linear interpolation	x_0	γ	unigram
Absolute discounting	x_0	$\gamma \frac{\text{distinct}(i, \bullet)}{\text{count}(i)}$	unigram
Kneser-Ney	x_0, x_1	$\gamma \frac{\text{distinct}(i, \bullet)}{\text{count}(i)}$	$\frac{\text{distinct}(\bullet, i)}{\sum_{v \in \mathcal{V}} \text{distinct}(\bullet, v)}$

Table 1: Variants of data noising techniques proposed in Xie et al. (2017) for context word x_0 and target word x_1 . Blank noising replaces an input word with a blank word, denoted by “_”.

3 BAYESIAN RECURRENT NEURAL NETWORKS

In Bayesian RNNs, we define a prior over our parameters \mathbf{W} and consider the following maximization problem:

$$\mathcal{L}(\mathcal{D}) = \int \prod_{t=1}^T p(x_t | x_{<t}; \mathbf{W}) p(\mathbf{W}) d\mathbf{w}.$$

A common prior is the standard normal distribution $p(\mathbf{W}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$.

For recurrent neural networks, the posterior over \mathbf{W} given the data \mathcal{D} is intractable. We approximate the posterior with a variational distribution $q(\mathbf{W})$, and minimize the KL divergence between the variational distribution and the true posterior by:

$$\begin{aligned} \mathbb{KL}(q(\mathbf{W})\|p(\mathbf{W} | \mathcal{D})) &= \mathbb{KL}(q(\mathbf{W})\|p(\mathbf{W})) - \int \log p(\mathcal{D} | \mathbf{W})q(\mathbf{W})d\mathbf{W} + \text{constant} \\ &= \mathbb{KL}(q(\mathbf{W})\|p(\mathbf{W})) - \sum_{t=1}^T \int \log p(x_t | x_{<t}; \mathbf{W})q(\mathbf{W})d\mathbf{W} + \text{constant} \end{aligned}$$

The integral is often approximated with Monte Carlo integration with one sample $\tilde{\mathbf{W}} \sim q(\mathbf{W})$:¹

$$\int \log p(x_t | x_{<t}; \mathbf{W})q(\mathbf{W})d\mathbf{W} \approx \log p(x_t | x_{<t}; \tilde{\mathbf{W}})q(\tilde{\mathbf{W}})d\tilde{\mathbf{W}}. \quad (1)$$

At test time, for a new sequence y_0, \dots, y_T , we can either set \mathbf{W} to be the mean of $q(\mathbf{W})$ or sample and average the results: $p(y_t | y_{<t}) = \frac{1}{S} \sum_{s=1}^S p(y_t | y_{<t}; \tilde{\mathbf{W}}_s)$, where S is the number of samples and $\tilde{\mathbf{W}}_s \sim q(\mathbf{W})$.

4 VARIATIONAL SMOOTHING

We now provide theoretical justifications for data smoothing under the variational framework. In a recurrent neural network language model, there are three types of parameters: an input (word) embedding matrix \mathbf{E} , an LSTM parameter matrix \mathbf{R} , and an output embedding matrix \mathbf{O} that produces logits for the final softmax function. We have $\mathbf{W} = \{\mathbf{E}, \mathbf{R}, \mathbf{O}\}$.

4.1 LINEAR INTERPOLATION SMOOTHING.

We first focus on the simplest data noising method—linear interpolation smoothing—and show how to extend it to other data noising methods subsequently.

Word embedding matrix \mathbf{E} . Denote the word vector corresponding to word $i \in \mathcal{V}$ in the input embedding matrix \mathbf{E} by \mathbf{e}_i . We obtain a similar effect to linear interpolation smoothing by using the following mixture of Gaussians variational distribution for \mathbf{e}_i :

$$\begin{aligned} \mathbf{e}_i \sim q(\mathbf{e}_i) &= (1 - \gamma)\mathcal{N}(\mathbf{e}_i, \sigma\mathbf{I}) + \gamma \sum_{v \in \mathcal{V}} \mathcal{U}_v \mathcal{N}(\mathbf{e}_v, \sigma\mathbf{I}) \\ &= (1 - \gamma + \gamma \mathcal{U}_i)\mathcal{N}(\mathbf{e}_i, \sigma\mathbf{I}) + \gamma \sum_{v \neq i \in \mathcal{V}} \mathcal{U}_v \mathcal{N}(\mathbf{e}_v, \sigma\mathbf{I}), \end{aligned} \quad (2)$$

where \mathcal{U}_v is the unigram probability of word v and σ is small. In other words, with probability $\gamma \mathcal{U}_v$, we replace the embedding \mathbf{e}_i with another embedding sampled from a normal distribution centered at \mathbf{e}_v .

Note that noising the input word is equivalent to choosing a different word embedding vector to be used in a standard recurrent neural network. Under the variational framework, we sample a different word embedding matrix for every sequence t at training time, since the integral is approximated with Monte Carlo by sampling from $q(\mathbf{e}_i)$ (Eq. 1).

The above formulation is related to word embedding dropout (Dai & Le, 2015; Iyyer et al., 2015; Kumar et al., 2016), although in word embedding dropout $q(\mathbf{e}_i) = (1 - \gamma)\mathcal{N}(\mathbf{e}_i, \sigma\mathbf{I}) + \gamma\mathcal{N}(\mathbf{0}, \sigma\mathbf{I})$.

KL term. For $p(\mathbf{e}_i) = \mathcal{N}(\mathbf{0}, \mathbf{I})$, we use Proposition 1 in Gal & Ghahramani (2016a) and approximate the KL divergence between a mixture of Gaussians $q(\mathbf{e}_i)$ and $p(\mathbf{e}_i)$ as:

$$\begin{aligned} \mathbb{KL}(q(\mathbf{e}_i)\|p(\mathbf{e}_i)) &\approx \lambda \sum_{v \in \mathcal{V}} \frac{\theta_v^i}{2} (\mathbf{e}_v^\top \mathbf{e}_v + \text{tr}(\sigma\mathbf{I}) - V(1 + \log 2\pi) - \log |\sigma\mathbf{I}|) \\ &= \lambda \sum_{v \in \mathcal{V}} \frac{\theta_v^i}{2} \|\mathbf{e}_v\|_2^2 + \text{constant}, \end{aligned}$$

¹In the followings, we use \tilde{x} to denote a sample from a distribution $q(x)$.

where V is the vocabulary size and θ_v^i is the mixture proportion for word i : $(1 - \gamma + \gamma \mathcal{U}_i)$ for $v = i$ and γ otherwise. In practice, the $\mathbb{K}\mathbb{L}$ term directly translates to an ℓ_2 regularizer on each word embedding vector, but the regularization coefficient is data dependent. More specifically, the regularization coefficient for word vector \mathbf{e}_i taking into account contributions from $\mathbb{K}\mathbb{L}(q(\mathbf{e}_v) \| p(\mathbf{e}_v))$ for $v \in \mathcal{V}$ is:

$$\lambda \frac{(V-1)\gamma + (1-\gamma + \gamma \mathcal{U}_i)}{2} = \lambda \frac{1 - (V-2 + \mathcal{U}_i)\gamma}{2} \quad (3)$$

In other words, the variational formulation of data smoothing results in a regularization coefficient that is a function of corpus statistics (the unigram distribution).

Other parameters. For other parameters \mathbf{R} and \mathbf{O} , we can use either simple variational distributions such as $q(\mathbf{R}) \sim \mathcal{N}(\mathbf{R}, \sigma \mathbf{I})$ and $q(\mathbf{O}) \sim \mathcal{N}(\mathbf{O}, \sigma \mathbf{I})$, which become standard ℓ_2 regularizers on these parameters; or incorporate dropout by setting $q(\mathbf{r}_i) \sim (1 - \alpha)\mathcal{N}(\mathbf{r}_i, \sigma \mathbf{I}) + \alpha\mathcal{N}(\mathbf{0}, \sigma \mathbf{I})$, where α is the dropout probability and \mathbf{r}_i is the i -th row of \mathbf{R} (Gal & Ghahramani, 2016b).

Training. In the original noising formulation (Xie et al., 2017), given a sequence $\{a, b, a, c\}$, it is possible to get a noised sequence $\{\underline{d}, b, a, c\}$, since the decision to noise at every timestep is independent of others. In the variational framework, we use Monte Carlo integration $\tilde{\mathbf{W}} \sim q(\mathbf{W})$ for each sequence to compute $p(x_t | x_{<t}; \tilde{\mathbf{W}})$. We can either sample one embedding matrix \mathbf{E} per sequence similar to Gal & Ghahramani (2016b), or sample one embedding matrix per timestep (Melis et al., 2018a). While the first method is computationally more efficient (we use it in our experiments), if we decide to noise the first a to d , we will have a noised sequence where every a is replaced by d : $\{\underline{d}, b, \underline{d}, c\}$. At training time, we go through sequence $x_{<t}$ multiple times (once per epoch) to get different noised sequences.

Predictions. For predictions, Xie et al. (2017) do not noise any input. In the variational framework, this corresponds to taking the mode of the variational distribution, since $1 - \gamma$ is almost always greater than γ . They also reported that they did not observe any improvements by sampling. Using the mode is rather uncommon in Bayesian RNNs. A standard approach is to take the mean, so we can set:

$$\bar{\mathbf{e}}_i = (1 - \gamma + \gamma \mathcal{U}_i)\mathbf{e}_i + \sum_{v \neq i \in \mathcal{V}} \gamma \mathcal{U}_v \mathbf{e}_v.$$

Extension to absolute discounting. It is straightforward to extend linear interpolation smoothing to absolute discounting by setting $\gamma_i = \gamma \frac{\text{distinct}(i, \bullet)}{\text{count}(i)}$ in Eq. 2 and Eq. 3 above.

4.2 KNESER-NEY SMOOTHING

We now consider the variational analog of Kneser-Ney data noising. Instead of smoothing towards lower order n -grams, Kneser-Ney uses models that take into account contextual diversity. For example, for bigram Kneser-Ney smoothing, we replace the unigram distribution with $\mathcal{K}_i = \frac{\text{distinct}(\bullet, i)}{\sum_{v \in \mathcal{V}} \text{distinct}(\bullet, v)}$, where $\text{distinct}(\bullet, i)$ denotes the number of distinct bigrams that end with word i . As a result, even if bigrams such as “San Francisco” and “Los Angeles” appear frequently in the corpus; “Francisco” and “Angeles” will not have high probabilities in \mathcal{K} since they often follow “San” and “Los”.

Recall that similar to absolute discounting, Kneser-Ney also uses $\gamma_i = \gamma \frac{\text{distinct}(i, \bullet)}{\text{count}(i)}$. However, if we replace an input word x_{t-1} , Xie et al. (2017) proposed to also replace the corresponding output word x_t . The intuition behind this is that since the probability of replacing an input word x_{t-1} is proportional to the number of distinct bigrams that start with x_{t-1} , when we replace the input word, we also need to replace the output word (e.g., if we replace “San” by a word sampled from \mathcal{K} , we should also replace “Francisco”).

Word embedding matrix \mathbf{E} . To get (bigram) Kneser-Ney smoothing, we use the following mixture of Gaussians variational distribution for \mathbf{e}_i :

$$\begin{aligned} \mathbf{e}_i \sim q(\mathbf{e}_i) &= (1 - \gamma_i)\mathcal{N}(\mathbf{e}_i, \sigma\mathbf{I}) + \gamma_i \sum_{v \in \mathcal{V}} \mathcal{K}_v \mathcal{N}(\mathbf{e}_v, \sigma\mathbf{I}) \\ &= (1 - \gamma_i + \gamma_i \mathcal{K}_i)\mathcal{N}(\mathbf{e}_i, \sigma\mathbf{I}) + \gamma_i \sum_{v \neq i \in \mathcal{V}} \mathcal{K}_v \mathcal{N}(\mathbf{e}_v, \sigma\mathbf{I}). \end{aligned} \quad (4)$$

Output embedding matrix \mathbf{O} . For the output embedding matrix, we also use the same variational distribution as the word embedding matrix:

$$\begin{aligned} \mathbf{o}_i \sim q(\mathbf{o}_i) &= (1 - \gamma_i)\mathcal{N}(\mathbf{o}_i, \sigma\mathbf{I}) + \gamma_i \sum_{v \in \mathcal{V}} \mathcal{K}_v \mathcal{N}(\mathbf{o}_v, \sigma\mathbf{I}) \\ &= (1 - \gamma_i + \gamma_i \mathcal{K}_i)\mathcal{N}(\mathbf{o}_i, \sigma\mathbf{I}) + \gamma_i \sum_{v \neq i \in \mathcal{V}} \mathcal{K}_v \mathcal{N}(\mathbf{o}_v, \sigma\mathbf{I}). \end{aligned} \quad (5)$$

$\mathbb{K}\mathbb{L}$ term. Following similar derivations in the previous subsection, it is straightforward to show that the approximated $\mathbb{K}\mathbb{L}$ term introduces the following regularization term to the overall objective:

$$\lambda \sum_{v \in \mathcal{V}} \frac{1 - (V - 2 + \mathcal{K}_v)\gamma_v}{2} \|\mathbf{e}_v\|_2^2 + \text{constant},$$

and similarly for the output embedding matrix \mathbf{O} .

Training. Recall that we use Monte Carlo integration $\tilde{\mathbf{W}} \sim q(\mathbf{W})$ for each sequence to compute $p(x_t | x_{<t}; \tilde{\mathbf{W}})$ at training time. Since we need to noise the output word when the input word is replaced, we sample $\tilde{\gamma}_i$ at $t - 1$ (for $x_{t-1} = i$) and decide whether we will keep the original (input and output) words or replace them. If we decide to replace, we need to sample two new words from $\mathcal{K}_v \mathcal{N}(\mathbf{e}_v, \sigma\mathbf{I})$ and $\mathcal{K}_v \mathcal{N}(\mathbf{o}_v, \sigma\mathbf{I})$ respectively. For other words k that are not the target at the current timestep (i.e., $k \neq x_t$), we can either assume that $\tilde{\mathbf{o}}_k = \mathbf{o}_k$ or alternatively sample (which can be expensive since we need to sample additional $V - 1$ times per timestep).

Predictions For predictions, Xie et al. (2017) also take the mode of the variational distribution (assuming $1 - \gamma_i$ is almost always greater than γ_i). We use the mean instead, similar to what we do in variational linear interpolation smoothing.

4.3 EXTENSIONS

Our derivations above provide insights into new methods under the variational smoothing framework. We describe three variants in the followings:

Tying input and output embeddings. Inan et al. (2017) and Press & Wolf (2017) showed that tying the input and output embedding matrices (i.e., using the same embedding matrix for both) improves language modeling performance and reduces the number of parameters significantly. We take inspirations from this work and sample both \mathbf{O} and \mathbf{E} from the same base matrix. As a result, we have fewer parameters to train due to this sharing mechanism, but we still have different samples of input and output embedding matrices per sequence.² Similar to previous results in language modeling (Inan et al., 2017; Melis et al., 2018b), our experiments demonstrate that tying improves the performance considerably.

Combining smoothing and dropout. We can combine variational smoothing and variational dropout by modifying the variational distribution to incorporate a standard Gaussian component:

$$\mathbf{e}_i \sim q(\mathbf{e}_i) = (1 - \alpha) \left\{ (1 - \gamma_i)\mathcal{N}(\mathbf{e}_i, \sigma\mathbf{I}) + \gamma_i \sum_{v \in \mathcal{V}} \mathcal{K}_v \mathcal{N}(\mathbf{e}_v, \sigma\mathbf{I}) \right\} + \alpha \mathcal{N}(0, \sigma\mathbf{I}),$$

²Alternatively, we could sample one matrix for both the input and output embeddings. We found that this approach is slightly worse in our preliminary experiments.

where α is the dropout probability. Note that in this formulation, we either drop an entire word by setting its embedding vector to zero (i.e., similar to word embedding dropout and blank noising) or choose an embedding vector from the set of words in the vocabulary.

However, it is more common to apply dropout to each dimension of the input embedding and output embedding matrix. In this formulation, we have:

$$e_{i,j} \sim q(e_{i,j}) = (1 - \alpha_{i,j}) \left\{ (1 - \gamma_i) \mathcal{N}(e_{i,j}, \sigma) + \gamma_i \sum_{v \in \mathcal{V}} \mathcal{K}_v \mathcal{N}(e_{v,j}, \sigma) \right\} + \alpha_{i,j} \mathcal{N}(0, \sigma). \quad (6)$$

At training time, while we sample $\tilde{\alpha}_{i,j}$ multiple times (once per word embedding dimension for each word), we only sample $\tilde{\gamma}_i$ once per word to ensure that when the element is not noised, we still use the same base embedding vector. We use this variant of smoothing and dropout throughout our experiments for our models.

Element-wise smoothing. The variational formulation above allows us to derive a variant of data smoothing that samples each element of the embedding vector \mathbf{e}_i independently (and similarly for \mathbf{o}_i). Consider the variational distribution in Eq. 6. At training time, if we sample both $\tilde{\alpha}_{i,j}$ and $\tilde{\gamma}_{i,j}$ multiple times (once per word embedding dimension for each word) we arrive at a new element-wise smoothing method. The main difference between this model and the previous model is that each dimension in the input (and output) embedding vector is sampled independently. As a result, the vector that is used is a combination of elements from various word vectors. Notice that the mean under this new scheme is still the same as sampling per vector, so we do not need to change anything at test time. One major drawback about this model is that it is computationally expensive since we need to sample each element of each embedding vector.

5 EXPERIMENTS

5.1 SETUP

We evaluate our approaches on two standard language modeling datasets: Penn Treebank (PTB) and Wikitext-2. We use a two-layer LSTM as our base language model. We perform non-episodic training with batch size 64 using RMSprop (Hinton, 2012) as our optimization method. We tune the RMSprop learning rate and ℓ_2 regularization hyperparameter λ for all models on a development set by a grid search on $\{0.002, 0.003, 0.004\}$ and $\{10^{-4}, 10^{-3}\}$ respectively, and use perplexity on the development set to choose the best model. We also tune γ from $\{0.1, 0.2, 0.3, 0.4\}$. We use recurrent dropout (Semeniuta et al., 2016) for \mathbf{R} and set it to 0.2, and apply (element-wise) input and output embedding dropouts for \mathbf{E} and \mathbf{O} and set it to 0.5 when $\mathbf{E}, \mathbf{O} \in \mathbb{R}^{V \times 512}$ and 0.7 when $\mathbf{E}, \mathbf{O} \in \mathbb{R}^{V \times 1024}$ based on preliminary experiments. We tie the input and output embedding matrices in all our experiments (i.e., $\mathbf{E} = \mathbf{O}$), except for the vanilla LSTM model, where we report results for both tied and untied.³

5.2 MODELS

We compare the following methods in our experiments:

- **Baseline:** a vanilla LSTM language model. We evaluate two variants: with tied input and output embeddings and without.
- **Data noising (DN):** an LSTM language model trained with data noising using linear interpolation smoothing or bigram Kneser-Ney smoothing (Xie et al., 2017).
- **Variational smoothing (VS):** an LSTM language model with variational smoothing using linear interpolation or Kneser-Ney. For both models, we use the mean of the variational distribution at test time.⁴

³Our preliminary experiments are consistent with previous work (Inan et al., 2017; Melis et al., 2018b; Merity et al., 2018) that show tying the input and output embedding matrices results in better models with fewer numbers of parameters.

⁴We discuss results using sampling at test time in §5.4.

Model	LSTM hidden size	# of params.	PTB		Wikitext-2	
			Dev	Test	Dev	Test
Vanilla LSTM (Xie et al., 2017)	512	-	84.3	80.4	-	-
Vanilla LSTM (Xie et al., 2017)	1500	-	81.6	77.5	-	-
DN: Kneser-Ney (Xie et al., 2017)	512	-	79.9	76.9	-	-
DN: Kneser-Ney (Xie et al., 2017)	1500	-	76.2	73.4	-	-
Var. dropout (Gal & Ghahramani, 2016b)	1500	-	-	73.4	-	-
Vanilla LSTM: untied	512	14M/38M	89.6	84.5	106.3	100.8
Vanilla LSTM: tied			80.0	74.0	90.6	86.6
DN: linear interpolation	512	9M/21M	79.4	73.3	88.9	84.6
DN: Kneser-Ney			75.0	70.7	86.1	82.1
VS: linear interpolation			76.3	71.2	84.0	79.6
VS: Kneser-Ney	512	9M/21M	74.5	70.6	84.9	80.9
VS: element-wise			70.5	66.8	-	-
Vanilla LSTM: untied	1024	37M/85M	90.3	85.5	97.6	91.9
Vanilla LSTM: tied			75.9	70.2	85.2	81.0
DN: linear interpolation	1024	27M/50M	75.5	70.2	84.3	80.1
DN: Kneser-Ney			71.4	67.3	81.9	78.3
VS: linear interpolation			71.7	67.8	80.5	76.6
VS: Kneser-Ney	1024	27M/50M	70.8	66.9	80.9	76.7
VS: element-wise			68.6	64.5	-	-

Table 2: Perplexity on PTB and Wikitext-2 datasets. DN and VS denote data noising and variational smoothing. The two numbers (*M/*M) in the # of params. column denote the number of parameters for PTB and Wikitext-2 respectively.

- **Variational element-wise smoothing:** for the smaller PTB dataset, we evaluate an LSTM language model that uses elementwise Kneser-Ney variational smoothing and dropout. We also use the mean at test time.

5.3 RESULTS

Our results are summarized in Table 2. Consistent with previous work on tying the input and output embedding matrices in language models (Inan et al., 2017; Melis et al., 2018b; Merity et al., 2018), we see a large reduction in perplexity (lower is better) when doing so. While our numbers are generally better, the results are also consistent with Xie et al. (2017) that show linear interpolation data noising is slightly better than vanilla LSTM with dropout, and that Kneser-Ney data noising outperforms these methods for both the medium (512) and large (1024) models.

Variational smoothing improves over data noising in all cases, both for linear interpolation and Kneser-Ney. Recall that the main differences between variational smoothing and data noising are: (1) using the mean at test time, (2) having a data dependent ℓ_2 regularization coefficient that comes from the $\mathbb{K}\mathbb{L}$ term,⁵ and (3) how each method interacts with the input and output embedding tying mechanism.⁶ Our results suggest that the choice of the proposal distribution to sample from is less important for variational smoothing. In our experiments, Kneser-Ney outperforms linear interpolation on PTB but linear interpolation is slightly better on Wikitext-2.

Element-wise variational smoothing performs the best for both small and large LSTM models on PTB. We note that this improvement comes at a cost, since this method is computationally expensive.⁷ It

⁵The data dependent ℓ_2 coefficient penalizes vectors that are sampled more often higher.

⁶In the variational framework, if we sample one matrix for the input and output embeddings, it effectively noises the output words even for linear interpolation. If we sample two matrices from the same base matrix, these matrices can be different at training time even if the parameters are tied. As described in §4.3, we use the latter in our experiments.

⁷Element-wise dropout can be implemented efficiently by sampling a mask (zero or one with some probability) and multiply the entire embedding matrix with this mask. In element-wise smoothing, we need to sample an index for each dimension and reconstruct the embedding matrix for each timestep.




Figure 1: Validation set perplexities on PTB for VS: Kneser-Ney (512 dimensions).

took about one day to train the smaller model as opposed to a couple hours without element-wise smoothing. As a result, we were unable to train this on a much bigger dataset such as Wikitext-2. Nonetheless, our results show that applying smoothing in the embedding (latent) space results in better models.

5.4 DISCUSSIONS

Sampling at test time. In order to better understand another prediction method for these models, we perform experiments where we sample at test time for both data noising and variational smoothing (instead of taking the mode or the mean). We use twenty samples and average the log likelihood. Our results agree with Xie et al. (2017) that mentioned sampling does not provide additional benefits. In our experiments, the perplexities of 512- and 1024-dimensional DN-Kneser-Ney models increase to 96.5 and 85.1 (from 75.0 and 71.4) on the PTB validation set. For VS-Kneser-Ney models, the perplexities increase to 89.8 and 78.7 (from 74.5 and 70.8). Both our results and Xie et al. (2017) suggest that introducing data dependent noise at test time is detrimental for recurrent neural network language models.

Sensitivity to γ . We evaluate the sensitivity of variational smoothing to hyperparameter γ . Figure 1 shows perplexities on the PTB validation set for a variant of our models. The model generally performs well within the range of 0.1 and 0.3, but it becomes progressively worse as we increase γ since there is too much noise.

Other applications. While we focus on language modeling in this paper, the proposed technique is applicable to other language processing tasks. For example, Xie et al. (2017) showed that data smoothing improves machine translation, so our techniques can be used in that setup as well. It is also interesting to consider how variational smoothing interacts with and/or can be applied to memory augmented language models (Tran et al., 2016; Grave et al., 2017; Yogatama et al., 2018) and state-of-the-art language models (Yang et al., 2017). We leave these for future work.

6 CONCLUSION

We showed that data noising in recurrent neural network language models can be understood as a Bayesian recurrent neural network with a variational distribution that consists of a mixture of Gaussians whose mixture weights are a function of the proposal distribution used in data noising (e.g., the unigram distribution, the Kneser-Ney continuation distribution). We proposed using the mean of the variational distribution at prediction time as a better alternative to using the mode. We combined it with variational dropout, presented two extensions (i.e., variational smoothing with tied input and output embedding matrices and element-wise smoothing), and demonstrated language modeling improvements on Penn Treebank and Wikitext-2.

REFERENCES

- Stanley F. Chen and Joshua Goodman. An empirical study of smoothing techniques for language modeling. In *Proc. of ACL*, 1996.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proc. of EMNLP*, 2014.
- Andrew M. Dai and Quoc V. Le. Semi-supervised sequence learning. In *Proc. of NIPS*, 2015.
- Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Proc. of ICML*, 2016a.
- Yarin Gal and Zoubin Ghahramani. A theoretically grounded application of dropout in recurrent neural networks. In *Proc. of NIPS*, 2016b.
- Edouard Grave, Armand Joulin, and Nicolas Usunier. Improving neural language models with a continuous cache. In *Proc. of ICLR*, 2017.
- Geoffrey Hinton. Neural networks for machine learning, 2012. Lecture 6.5.
- Sepp Hochreiter and Jurgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8): 1735–1780, 1997.
- Arthur E. Hoerl and Robert W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- Hakan Inan, Khashayar Khosravi, and Richard Socher. Tying word vectors and word classifiers: A loss framework for language modeling. In *Proc. of ICLR*, 2017.
- Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daume III. Deep unordered composition rivals syntactic methods for text classification. In *Proc. of ACL*, 2015.
- Ankit Kumar, Ozan Irsoy, Jonathan Su, James Bradbury, Robert English, Brian Pierce, Peter Ondruska, Ishaan Gulrajani, and Richard Socher. Ask me anything: Dynamic memory networks for natural language processing. In *Proc. of ICML*, 2016.
- Mitchell Marcus, Grace Kim, Mary Ann Marcinkiewicz, Robert MacIntyre, Ann Bies, Mark Ferguson, Karen Katz, and Britta Schasberger. The penn treebank: Annotating predicate argument structure. In *Proc. of the Workshop on Human Language Technology*, 1994.
- Gabor Melis, Charles Blundell, Tomas Kocisky, Karl Moritz Hermann, Chris Dyer, and Phil Blunsom. Pushing the bounds of dropout. *arXiv preprint*, 2018a.
- Gabor Melis, Chris Dyer, and Phil Blusom. On the state of the art of evaluation in neural language models. In *Proc. of ICLR*, 2018b.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. In *Proc. of ICLR*, 2017.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. Regularizing and optimizing lstm language models. In *Proc. of ICLR*, 2018.
- Ofir Press and Lior Wolf. Using the output embedding to improve language models. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 157–163. Association for Computational Linguistics, 2017. URL <http://aclweb.org/anthology/E17-2025>.
- Stanislau Semeniuta, Aliaksei Severyn, and Erhardt Barth. Recurrent dropout without memory loss. In *Proc. of COLING*, 2016.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.

Ke Tran, Arianna Bisazza, and Christof Monz. Recurrent memory networks for language modeling. In *Proc. of NAACL-HLT*, 2016.

Ziang Xie, Sida I. Wang, Jiwei Li, Daniel Levy, Aiming Nie, Dan Jurafsky, and Andrew Y. Ng. Data noising as smoothing in neural network language models. In *Proc. of ICLR*, 2017.

Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, and William W Cohen. Breaking the softmax bottleneck: A high-rank rnn language model. *arXiv preprint arXiv:1711.03953*, 2017.

Dani Yogatama, Yishu Miao, Gabor Melis, Wang Ling, Adhiguna Kuncoro, Chris Dyer, and Phil Blunsom. Memory architectures in recurrent neural network language models. In *Proc. of ICLR*, 2018.