

N-GRAM LANGUAGE MODELING USING RECURRENT NEURAL NETWORK ESTIMATION

Ciprian Chelba, Mohammad Norouzi, Samy Bengio

Google

{ciprianchelba,mnorouzi,bengio}@google.com

ABSTRACT

We investigate the effective memory depth of RNN models by using them for n -gram language model (LM) smoothing.

Experiments on a small corpus (UPenn Treebank, one million words of training data and 10k vocabulary) have found the LSTM cell with dropout to be the best model for encoding the n -gram state when compared with feed-forward and vanilla RNN models. When preserving the sentence independence assumption the LSTM n -gram matches the LSTM LM performance for $n = 9$ and slightly outperforms it for $n = 13$. When allowing dependencies across sentence boundaries, the LSTM 13-gram almost matches the perplexity of the unlimited history LSTM LM.

LSTM n -gram smoothing also has the desirable property of improving with increasing n -gram order, unlike the Katz or Kneser-Ney back-off estimators. Using multinomial distributions as targets in training instead of the usual one-hot target is only slightly beneficial for low n -gram orders.

Experiments on the One Billion Words benchmark show that the results hold at larger scale: while LSTM smoothing for short n -gram contexts does not provide significant advantages over classic N-gram models, it becomes effective with long contexts ($n > 5$); depending on the task and amount of data it can match fully recurrent LSTM models at about $n = 13$. This may have implications when modeling short-format text, e.g. voice search/query LMs.

Building LSTM n -gram LMs may be appealing for some practical situations: the state in a n -gram LM can be succinctly represented with $(n - 1) * 4$ bytes storing the identity of the words in the context and batches of n -gram contexts can be processed in parallel. On the downside, the n -gram context encoding computed by the LSTM is discarded, making the model more expensive than a regular recurrent LSTM LM.

1 INTRODUCTION

A statistical language model (LM) estimates the prior probability values $P(W)$ for strings of words W in a vocabulary \mathcal{V} whose size is usually in the tens or hundreds of thousands. Typically the string W is broken into sentences, or other segments such as utterances in automatic speech recognition which are assumed to be conditionally independent; the independence assumption has certain advantages in practice but is not strictly necessary.

Applying the chain rule to a sentence $W = w_1, w_2, \dots, w_n$ we get:

$$P(W) = \prod_{k=1}^n P(w_k | w_1, w_2, \dots, w_{k-1}) \quad (1)$$

Since the parameter space of $P(w_k | w_1, w_2, \dots, w_{k-1})$ is too large, the language model is forced to put the *context* $W_{k-1} = w_1, w_2, \dots, w_{k-1}$ into an *equivalence class* determined by a function $\Phi(W_{k-1})$. As a result:

$$P(W) \cong \prod_{k=1}^n P(w_k | \Phi(W_{k-1})) \quad (2)$$

Research in language modeling consists of finding appropriate equivalence classifiers Φ and methods to estimate $P(w_k|\Phi(W_{k-1}))$.

1.1 PERPLEXITY AS A MEASURE OF LANGUAGE MODEL QUALITY

A commonly used quality measure for a given model M is related to the entropy of the underlying source and was introduced under the name of *perplexity* (PPL) Jelinek (1997):

$$PPL(W, M) = \exp\left(-\frac{1}{N} \sum_{k=1}^N \ln [P_M(w_k|W_{k-1})]\right) \quad (3)$$

To give intuitive meaning to perplexity, it represents the average number of guesses the model needs to make in order to ascertain the identity of the next word, when running over the test word string $W = w_1 \dots w_N$ from left to right. It can be easily shown that the perplexity of a language model that uses the uniform probability distribution over words in the vocabulary \mathcal{V} equals the size of the vocabulary; a good language model should of course have lower perplexity, and thus the vocabulary size is an upper bound on the perplexity of a given language model.

Very likely, not all words in the test data are part of the language model vocabulary. It is common practice to map all words that are out-of-vocabulary to a distinguished *unknown word* symbol, and report the out-of-vocabulary (OOV) rate on test data—the rate at which one encounters OOV words in the test sequence W —as yet another language model performance metric besides perplexity. Usually the unknown word is assumed to be part of the language model vocabulary—*open vocabulary* language models—and its occurrences are counted in the language model perplexity calculation in Eq. (3). A situation less common in practice is that of *closed vocabulary* language models where all words in the test data will always be part of the vocabulary \mathcal{V} .

1.2 SMOOTHING

Since the language model is meant to assign non-zero probability to unseen strings of words (or equivalently, ensure that the cross-entropy of the model over an arbitrary test string is not infinite), a desirable property is that:

$$P(w_k|\Phi(W_{k-1})) > \epsilon > 0, \forall w_k, W_{k-1} \quad (4)$$

also known as the *smoothing* requirement.

There are currently two dominant approaches for building LMs:

1.2.1 n -GRAM LANGUAGE MODELS

The most widespread paradigm in language modeling makes a Markov assumption and uses the $(n-1)$ -gram equivalence classification, that is, defines

$$\Phi(W_{k-1}) \doteq w_{k-n+1}, w_{k-n+2}, \dots, w_{k-1} = h \quad (5)$$

A large body of work has accumulated over the years on various smoothing methods for n -gram LMs. The two most popular smoothing techniques are probably Kneser & Ney (1995) and Katz (1987), both making use of *back-off* to balance the specificity of long contexts with the reliability of estimates in shorter n -gram contexts. Goodman (2001) provides an excellent overview that is highly recommended to any practitioner of language modeling.

Approaches that depart from the nested features used in back-off n -gram LMs have shown excellent results at the cost of increasing the number of features and parameters stored by the model, e.g. Pelemans et al. (2016).

1.2.2 NEURAL LANGUAGE MODELS

Neural networks (NNLM) have emerged in recent years as an alternative to estimating and storing n -gram LMs. Words (or some other modeling unit) are represented using an embedding vector $E(w) \in \mathbb{R}^d$. A simple NNLM architecture makes the Markov assumption and feeds the concatenated embedding vectors for the words in the n -gram context to one or more layers each consisting

of an affine transform followed by a non-linearity (typically \tanh); the output of the last such layer is then fed to the output layer consisting again of an affine transform but this time followed by an exponential non-linearity that is normalized to guarantee a proper probability over the vocabulary. This is commonly named a feed-forward architecture for an n -gram LM (FF-NNLM), first introduced by Bengio et al. (2001).

An alternative is the recurrent NNLM architecture that feeds the embedding of each word $E(w_k)$ one at a time, advancing the state $S \in \mathbb{R}^s$ of a *recurrent cell* and producing a new output $U \in \mathbb{R}^u$:

$$\begin{aligned} [S_k, U_k] &= RNN(S_{k-1}, E(w_k)) \\ S_0 &= 0 \end{aligned} \tag{6}$$

This provides a representation for the context W_{k-1} that can be directly plugged into Eq. 2:

$$\Phi(W_{k-1}) = U_{k-1}(W_{k-1}) \tag{7}$$

Similar to the FF-NNLM architecture, the output U of the recurrent cell is then fed to a soft-max layer consisting of an affine transform O followed by an exponential non-linearity properly normalized over the vocabulary.

The recurrent cell $RNN(\cdot)$ can consist of one or more simple affine/non-linearity layers, often called a *vanilla RNN* architecture, see Mikolov et al. (2010). The *LSTM* cell due to Hochreiter & Schmidhuber (1997) has proven very effective at modeling long range dependencies and has become the state-of-the-art architecture for language modeling using RNNs, see Józefowicz et al. (2016).

In this work we approximate unlimited history (R)NN models with n -gram models in an attempt to identify the order n at which they become equivalent from a perplexity point of view. This is a promising direction in a few ways:

- the training data can be reduced to n -gram sufficient statistics, and the target distribution presented to the NN n -gram LM in a given context can be a multinomial pmf instead of the one-hot encoding used in on-line training for (R)NN LMs;
- unlike many LSTM LM implementations, back-propagation through time for the LSTM n -gram need not be truncated at the beginning of segments used to batch the training data;
- the state in a n -gram LM can be succinctly represented with $(n - 1) * 4$ bytes storing the identity of the words in the context; this is in stark contrast with the state $S \in \mathbb{R}^s$ for an RNN LM, where $s = 1024$ or higher, making the n -gram LM much easier to use in decoders such as for ASR/SMT;
- similar to Brants et al. (2007), batches of n -gram contexts can be processed in parallel to estimate a *sharded* (R)NN n -gram model; this is particularly attractive because it allows scaling both the amount of training data and the NNLM size significantly (100X).

2 METHOD

As mentioned in the previous section, the Markov assumption made by n -gram models allows us to present to the NN *multinomial* training targets specifying the full distribution in a given n -gram context instead of the usual *one-hot* target specifying the predicted word occurring in a given context instance. In addition, when using *multinomial* targets we can either weight each training sample by the context count or simply present each context token encountered in the training data along with the conditional multinomial pmf computed from the entire training set.

We thus have three main training regimes:

- context-weighted multinomial targets
- multinomial targets (context count $count(h) = 1$)
- one-hot targets (context count $count(h) = 1$, word count $count(h, w) = 1$)

The loss function optimized in training is the cross-entropy between the model pmf $P(w|h; \theta)$ in some n -gram context h and the relative frequency $f(w|h; \mathcal{T})$ in the training data \mathcal{T} (or development

\mathcal{D} , or test data \mathcal{E}) is computed as:

$$H(P; \mathcal{T}) = -1/T \sum_h \text{count}(h) \sum_w f(w|h; \mathcal{T}) \log P(w|h; \theta) \quad (8)$$

where T is the length of the training data \mathcal{T} and $P(\cdot; \theta)$ is the n -gram model being evaluated/trained as parameterized by θ .

The baseline back-off n -gram models (Katz, interpolated Kneser-Ney) are trained by making a sentence independence assumption. As a result n -gram contexts at the beginning of the sentence are padded to the left to reach the full context length. The same n -gram counting strategy is used when preparing the data for the various NN n -gram LMs that we experimented with. Since RNN LMs are normally trained and evaluated without making this independence assumption by passing the LM state across sentence boundaries, we also evaluated the impact of resetting the RNN LM state at sentence beginning.

We next detail the various flavors of NN LM implementations we experimented with.

For all NN LMs we represent context words using an embedding vector $E(w) \in \mathbb{R}^d$. Unless otherwise stated, all models are trained to minimize the cross-entropy on training data in Eq. (8), using Adagrad (Duchi et al. (2011)) and gradient norm clipping (Pascanu et al. (2012)); the model parameters are initialized by sampling from a truncated normal distribution of zero mean and a given standard deviation.

Training proceeds for a fixed number of epochs for every given point on the grid of hyper-parameters explored for a given model type; the best performing model (parameter values) on development data \mathcal{D} is retained as the final one to be evaluated on test data \mathcal{E} in order to report the model perplexity.

All models were implemented using TensorFlow, see Abadi & et al. (2015b).

2.1 FEED FORWARD n -GRAM LM

Each word w in the n -gram context $h = w_{k-n+1} \dots w_{k-1}$ is embedded using the mapping $E(w)$; the resulting vectors are concatenated to form a $d \cdot (n - 1)$ dimensional vector that is first fed into a dropout layer Srivastava et al. (2014) and then into an affine layer followed by a tanh non-linearity. The output of this so-called “hidden” layer is again fed into a dropout layer and then followed by an affine layer O whose output is of the same dimensionality as the vocabulary. An exponential “soft-max” layer converts the activations produced by the last affine layer into probabilities over the vocabulary.

To summarize:

$$\begin{aligned} X &= \text{concat}(E(w_{k-n+1}), \dots, E(w_{k-1})) \\ D(X) &= \text{dropout}(X; P_{\text{keep}}) \\ Y &= \tanh(H \cdot D(X) + H_{\text{bias}}) \\ D(Y) &= \text{dropout}(Y; P_{\text{keep}}) \\ P(\cdot | w_{k-n+1} \dots w_{k-1}) &= \exp(O \cdot D(Y) + O_{\text{bias}}) \end{aligned} \quad (9)$$

The parameters of the model are the embedding matrix $E \in \mathbb{R}^{d \times V}$, the keep probability for dropout layers P_{keep} , the affine input layer parameterized by $H \in \mathbb{R}^{s \times (n-1) \cdot d}$, $H_{\text{bias}} \in \mathbb{R}^s$ and the output one parameterized by $O \in \mathbb{R}^{V \times s}$, $O_{\text{bias}} \in \mathbb{R}^V$.

The hyper-parameters controlling the training are: number of training epochs, n -gram order, dimensionality of the model parameters d, s , keep probability value, gradient norm clipping value, standard deviation for the initializer and the Adagrad learning rate and initial accumulator value.

2.2 “VANILLA” RECURRENT n -GRAM LM

Each word w in the n -gram context $h = w_{k-n+1} \dots w_{k-1}$ is embedded using the mapping $E(w)$ followed by dropout and then fed in left-to-right order into the RNN cell in Eq. (6). The final output of the RNN cell is then fed first into a dropout layer and then into an affine layer followed by exponential “soft-max”.

Assuming that we encode the context $h = w_{k-n+1} \dots w_{k-1}$ with an RNN cell defined as follows (using the running index $l = k - n + 1 \dots k$ to traverse the context and a dropout layer on the embedding $D(E(w_l)) = dropout(E(w_l), P_{keep})$):

$$\begin{aligned} [S_l, U_l] &= \tanh(R \cdot [S_{l-1}, D(E(w_l))] + R_{bias}) \\ S_{k-n} &= 0 \end{aligned} \tag{10}$$

we pick the last output U_{k-1} and feed it into a dropout layer followed by an affine layer and soft-max output:

$$\begin{aligned} D(U_{k-1}) &= dropout(U_{k-1}; P_{keep}) \\ P(\cdot | w_{k-n+1} \dots w_{k-1}) &= \exp(O \cdot D(U_{k-1}) + O_{bias}) \end{aligned} \tag{11}$$

The parameters of the model are the embedding matrix $E \in \mathbb{R}^{d \times V}$, the keep probability for dropout layers P_{keep} , the RNN affine layer parameterized by $R \in \mathbb{R}^{(d+s) \times 2 \cdot s}$ and $R_{bias} \in \mathbb{R}^{2 \cdot s}$ and the output one parameterized by $O \in \mathbb{R}^{V \times s}$ and $O_{bias} \in \mathbb{R}^V$. Note that we choose to use the same dimensionality s for both $S, U \in \mathbb{R}^s$.

The hyper-parameters controlling the training are the same as in the previous section.

2.3 LSTM RECURRENT n -GRAM LM

Finally, we replace the “vanilla” RNN cell defined above with a multilayer LSTM cell with dropout. Since this was the most effective model, we experimented with a few options:

- forward context encoding: context words $h = w_{k-n+1} \dots w_{k-1}$ are fed in left-to-right order in the LSTM cell; the LSTM cell output after the last context word w_{k-1} is then fed into the output layer;
- reverse context encoding: context words $h = w_{k-n+1} \dots w_{k-1}$ are fed in left-to-right order in the LSTM cell; the LSTM cell output after the first context word w_{k-n+1} is then fed into the output layer;
- stacked output for either of the above: we concatenate the output vectors along the way and feed that into the output layer;
- bidirectional context encoding: we encode the context twice, forward and reverse order respectively, using two separate LSTM cells; the two outputs are then concatenated and fed to the output layer;
- forward context encoding with incremental loss with/out exponential decay as a function of the context length

The last item above deserves a more detailed explanation. It is possible that the LSTM encoder would benefit from incremental error back-propagation along the n -gram context instead of just one back-propagation step at the end of the context. As such, we modify the loss function to be the cumulative cross-entropy between the relative frequency and the model output distribution *at each step in the for loop feeding the n -gram context into the LSTM cell instead of just the last one*. Thus amounts to targetting a mix of $1 \dots n$ -gram target distributions; to have better control over the contribution of different n -gram orders to the loss function, we weigh each loss function by an exponential term $exp(-decay \cdot (n - 1 - l))$. The $decay > 0$ value controls how fast the contribution to the loss function from lower n -gram orders decays; note that the highest order $l = n - 1$ has weight 1.0 so a very large value $decay = \infty$ restores the regular training loss function. For this training regime we only implemented one-hot targets: the amount of data that needs to be fed to the TensorFlow graph would increase significantly for incremental multinomial targets.

The hyper-parameters controlling the training are: number of training epochs, n -gram order, embedding dimensionality d , LSTM cell output dimensionality s and number of layers, keep probability value, gradient norm clipping value, standard deviation for the initializer. To match the fully recurrent LSTM LM implemented by the UPenn Treebank TensorFlow tutorial, we estimated all of our LSTM n -gram models using gradient descent with variable learning rate: initially the learning rate is constant for a few iterations after which it follows a linear decay schedule. The hyper-parameters

controlling this schedule were not optimized but rather we used the same values as in the RNN LM tutorial provided with Abadi & et al. (2015a) or the implementation in Józefowicz (2016), respectively.

Perhaps a bit of a technicality but it is worth pointing out a major difference between error back-propagation through time (BPTT) as implemented in either of the above and the error back-propagation in the LSTM/RNN n -gram LM: Abadi & et al. (2015a) and Józefowicz (2016) implement BPTT by *segmenting* the training data into non-overlapping segments (of length 35 or 20, respectively)¹. The error BPTT does not cross the left boundary of such segments, whereas the LSTM state is of course copied forward. As a result, the first word in a segment is not really contributing to training, and the immediately following ones have a limited effect. This is in contrast to error back-propagation for the LSTM/RNN n -gram LM: the n -gram window slides over the training/test data, and error back-propagation covers the entire n -gram context; the LSTM cell state and output computed for a given n -gram context are discarded once the output distribution is computed.

3 EXPERIMENTS

3.1 UPENN TREEBANK CORPUS

For our initial set of experiments we used the same data set as in Abadi & et al. (2015a), with exactly the same training/validation/test set partition and vocabulary. The training data consists of about one million words, and the vocabulary contains ten thousand words; the validation/test data contains 73760/82430 words, respectively (including the end-of-sentence token). The out-of-vocabulary rate on validation/test data is 5.0/5.8%, respectively.

As an initial batch of experiments we trained and evaluated back-off n -gram models using Katz and interpolated Kneser-Ney smoothing. We also used the medium setting in Abadi & et al. (2015a) as an LSTM/RNN LM baseline; since the baseline n -gram models are trained under a sentence independence assumption, we also ran the LSTM/RNN LM baseline by resetting the LSTM state at each sentence beginning. The results are presented in Table 1.

Model	Order	Test PPL
n-gram, baseline		
Katz, back-off	5	167
Katz, back-off	9	182
Interpolated Kneser-Ney, back-off	5	143
Interpolated Kneser-Ney, back-off	9	143
LSTM RNN, baseline		
LSTM (medium setting)	reset state at <S>	95
LSTM (medium setting)	∞	84

Table 1: UPenn Treebank: baseline back-off n -gram and LSTM perplexity values.

As expected Kneser-Ney (KN) is better than Katz, and it does not improve with the n -gram order past a certain value, in this case $n = 5$. This behavior is due to the fact that the n -gram hit ratio on test data (number of test n -grams that were observed in training) decreases dramatically with the n -gram order: the percentage of covered n -grams² for $n = 1 \dots 9$ is 100, 81, 42, 18, 8.6, 5.0, 3.3, 2.5, 2.0, respectively.

The medium setting for the LSTM LM in Abadi & et al. (2015a) performs significantly better than the KN baseline. Resetting the state at sentence beginning degrades PPL significantly by 13% relative.

We then trained and evaluated various NN-smoothed n -gram LMs, as described in Section 2. The results are presented in Table 2. The best model among the ones considered is by far the LSTM

¹We have evaluated the impact of reducing the segment length dramatically, e.g. 4 instead of 35. Much to our surprise, the LSTM PPL increased modestly, from 84 to 88, see the before last row in Table 1; for the One Billion Words experiments using a segment of length 5 did not change PPL at all.

²For the hit ratio calculation the n -grams are not padded to the left of sentence beginning; if we are to count hit ratios using padded n -grams, the values are: 100, 81, 44.7, 24.0, 16.5, 13.7, 12.5, 11.8, 11.5, respectively.

Model	Order	Test PPL	
		multinomial	one-hot
<i>n</i>-gram, baseline			
Interpolated Kneser-Ney, back-off	5	143	
Feed-fwd <i>n</i>-gram			
Feed-fwd <i>n</i> -gram	5	127	128
Feed-fwd <i>n</i> -gram	9	125	126
Feed-fwd <i>n</i> -gram	13	125	127
“Vanilla” RNN <i>n</i>-gram			
RNN <i>n</i> -gram	9	127	131
LSTM RNN <i>n</i>-gram			
LSTM <i>n</i> -gram, forward context encoding	5	103	106
LSTM <i>n</i> -gram, forward context encoding	9	94	93
LSTM <i>n</i> -gram, forward context encoding	13	91	90
LSTM <i>n</i> -gram, reversed context encoding	9	102	107
LSTM <i>n</i> -gram, bidirectional context encoding	9	100	102
incremental LSTM <i>n</i> -gram with decay, <i>decay</i> = 2.0	13	—	91
LSTM RNN, baseline			
LSTM (medium setting)	reset at <S>	—	95

Table 2: UPenn Treebank: perplexity values for neural network smoothed *n*-gram LM.

n-gram. The most significant experimental result is that the LSTM *n*-gram can match and even outperform the fully recurrent LSTM LM as we increase the order *n*: *n* = 9 matches the LSTM LM performance, decreasing the LM perplexity by 34% relative over the Kneser-Ney baseline. LSTM *n*-gram smoothing also has the desirable property of improving with the *n*-gram order, unlike the Katz or Kneser-Ney back-off estimators, which can be credited to better feature extraction from the *n*-gram context.

Multinomial targets can slightly outperform the one-hot ones although the difference is shrinking as we increase the *n*-gram order. Weighting the contribution of each context to the loss function by its count did not work; we suspect this is because on-line training does not work well with the Zipf distribution on context counts.

Among the various flavors of LSTM models we experimented with, the forward context encoding performs best. The incremental LSTM *n*-gram with a fairly large decay (*decay* = 2.0) is slightly better but we do not consider the difference to be statistically significant (it also entails significantly more computation, we need to perform *n* − 1 back-propagation steps for each input *n*-gram).

To compare with the LSTM RNN LM that does not reset state at sentence beginning we also trained LSTM *n*-gram models (forward context encoding only) that straddle the sentence beginning. The results are presented in Table 3. Again, we notice that for a large enough order the LSTM *n*-gram LM comes very close to matching the fully recurrent LSTM baseline.

Model	Order	Test PPL	
		multinomial	one-hot
LSTM RNN <i>n</i>-gram			
LSTM <i>n</i> -gram, forward context encoding, straddling <S>	5	102	104
LSTM <i>n</i> -gram, forward context encoding, straddling <S>	9	91	95
LSTM <i>n</i> -gram, forward context encoding, straddling <S>	13	87	91
LSTM RNN, baseline			
LSTM (medium setting)	∞	—	84

Table 3: UPenn Treebank: perplexity values for neural network smoothed *n*-gram LM when straddling the sentence beginning boundary.

3.2 ONE BILLION WORDS BENCHMARK

In a second set of experiments we used the corpus in Chelba et al. (2013), the same as in Józefowicz et al. (2016). For the baseline LSTM model we used the single machine implementation provided by Józefowicz (2016); the LSTM n -gram variant was implemented as a minor tweak on this code-base and is thus different from the one used in the UPenn Treebank experiments in Section 3.1.

We experimented with the LSTM configuration in Table 3 of Józefowicz et al. (2016) for both baseline LSTM and n -gram variant, which are also the default settings in Józefowicz (2016): embedding and projection layer dimensionality was 128, one layer with state dimensionality of 2048. Training used Adagrad with gradient clipping by global norm (10.0) and dropout (probability 0.1); back-propagation at the output soft-max layer is done using importance sampling as described in Józefowicz et al. (2016) with a set 8192 “negative” samples. An additional set of experiments investigated the benefits of adding one more layer to both baseline and n -gram LSTM.

The results are presented in Tables 4-5; unlike the UPenn Treebank experiments, we did not tune the hyper-parameters for the n -gram LSTM and instead just used the same ones as for the LSTM baseline; as a result the perplexity values for the n -gram LSTM may be slightly suboptimal.

Similar to the UPenn Treebank experiments, we examined the effect of resetting state at sentence boundaries. As expected PPL did not change significantly because the sentences in the training and test data were randomized, see Chelba et al. (2013); in fact modeling the sentence independence explicitly is slightly beneficial.

We observe that on large amounts of data LSTM smoothing for short n -gram contexts does not provide significant advantages over classic back-off n -gram models. This may have implications for short-format text, e.g. voice search/query LMs. On the other hand, LSTM smoothing becomes very effective with long contexts ($n > 5$) approaching the fully recurrent LSTM model perplexity at about $n = 13$.

Training times are significantly different between the LSTM baseline and the n -gram variant, with the latter being about an order of magnitude slower due to the fact that the LSTM state is recomputed and discarded for every new training sample.

Model	Order	Test PPL
Training Target		one-hot
n-gram, baseline		
Interpolated Kneser-Ney, back-off	5	68
LSTM RNN n-gram		
LSTM n -gram, forward context encoding	5	70
LSTM n -gram, forward context encoding	9	54
LSTM n -gram, forward context encoding	13	49
LSTM RNN, baseline		
LSTM	reset at <S>	48
2-layer LSTM RNN n-gram		
LSTM n -gram, forward context encoding	5	68
LSTM n -gram, forward context encoding	9	51
LSTM n -gram, forward context encoding	13	46
2-layer LSTM RNN, baseline		
LSTM	reset at <S>	43

Table 4: One Billion Words Benchmark: perplexity values for neural network smoothed n -gram LM when enforcing the sentence independence.

4 CONCLUSIONS AND FUTURE WORK

We investigated the effective memory depth of (R)NN models by using them for word-level n -gram LM smoothing. The LSTM cell with dropout was by far the best (R)NN model for encoding the n -gram state.

Model	Order	Test PPL
Training Target		one-hot
LSTM RNN n-gram		
LSTM n -gram, forward context encoding, straddling $\langle S \rangle$	5	70
LSTM n -gram, forward context encoding, straddling $\langle S \rangle$	9	54
LSTM n -gram, forward context encoding, straddling $\langle S \rangle$	13	49
LSTM RNN, baseline		
LSTM	∞	49
2-layer LSTM RNN n-gram		
LSTM n -gram, forward context encoding, straddling $\langle S \rangle$	5	68
LSTM n -gram, forward context encoding, straddling $\langle S \rangle$	9	51
LSTM n -gram, forward context encoding, straddling $\langle S \rangle$	13	46
2-layer LSTM RNN, baseline		
LSTM	∞	43

Table 5: One Billion Words Benchmark: perplexity values for neural network smoothed n -gram LM when straddling the sentence beginning boundary.

When preserving the sentence independence assumption the LSTM n -gram matches the LSTM LM performance for $n = 9$ and slightly outperforms it for $n = 13$. When allowing dependencies across sentence boundaries, the LSTM 13-gram almost matches the perplexity of the unlimited history LSTM LM.

We can thus conclude that the memory of LSTM LMs seems to be about 9-13 previous words which is not a trivial depth but not that large either.

Compared to standard n -gram smoothing methods LSTMs have excellent statistical properties: they improve with the n -gram order well beyond the point where Katz or Kneser-Ney back-off smoothing methods saturate, proving that they are able to extract richer features from the same context. Using multinomial targets in training is only slightly beneficial in this setting, although the advantage over one-hot diminishes with increasing n -gram order.

Experiments on the One Billion Words benchmark confirm that n -gram LSTMs can match the performance of fully recurrent LSTMs at larger amounts of data.

Building LSTM n -gram LMs is attractive due to the fact that the state in a n -gram LM can be succinctly represented on $4 \cdot (n - 1)$ bytes storing the identity of the context words. This is in stark contrast with the state $H \in \mathbb{R}^h$ for an RNN LM, where $h = 1024$ or higher, making the n -gram LM easier to use in decoders such as for ASR/SMT. The LM requests in the decoder can be batched, making the RNN LM operation more efficient on GPUs.

On the downside, the LSTM encoding for the n -gram context is discarded and cannot be re-used; caching it for frequent LM states is possible.

ACKNOWLEDGMENTS

We would like to thank Oriol Vinyals and Rafał Józefowicz for support with the baseline implementation of LSTM LMs for UPenn Treebank in Abadi & et al. (2015a) and One Billion Words Benchmark in Józefowicz (2016), respectively. We would also like to thank Maxim Krikun for thorough code reviews and useful discussions.

REFERENCES

- M. Abadi and et al. Recurrent neural networks tutorial (language modeling), 2015a. URL <https://www.tensorflow.org/versions/r0.11/tutorials/recurrent/index.html>. UPenn Treebank language modeling.
- M. Abadi and et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015b. URL <http://tensorflow.org/>. Software available from tensorflow.org.

- Yoshua Bengio, Réjean Ducharme, and Pascal Vincent. A neural probabilistic language model. 2001. URL http://www.iro.umontreal.ca/~lisa/pointeurs/nips00_lm.ps.
- T. Brants, A. C. Papat, P. Xu, F. J. Och, and J. Dean. Large language models in machine translation. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pp. 858–867, 2007. URL <http://www.aclweb.org/anthology/D/D07/D07-1090>.
- Ciprian Chelba, Tomas Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, and Phillipp Koehn. One billion word benchmark for measuring progress in statistical language modeling. *CoRR*, abs/1312.3005, 2013. URL <http://arxiv.org/abs/1312.3005>.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *J. Mach. Learn. Res.*, 12:2121–2159, July 2011. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=1953048.2021068>.
- Joshua Goodman. A bit of progress in language modeling, extended version. Technical report, Microsoft Research, 2001.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 1997.
- Frederick Jelinek. *Information Extraction From Speech And Text*, chapter 8, pp. 141–142. MIT Press, 1997.
- Rafal Józefowicz. Single machine implementation of LSTM language model on One Billion Words benchmark using synchronized gradient updates, 2016. URL <https://github.com/rafaljozefowicz/lm>.
- Rafal Józefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. Exploring the limits of language modeling. *CoRR*, abs/1602.02410, 2016. URL <http://arxiv.org/abs/1602.02410>.
- S. Katz. Estimation of probabilities from sparse data for the language model component of a speech recognizer. In *IEEE Transactions on Acoustics, Speech and Signal Processing*, volume 35, pp. 400–01, 1987.
- R. Kneser and H. Ney. Improved backing-off for m-gram language modeling. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, volume 1, pp. 181–184, 1995.
- Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocký, and Sanjeev Khudanpur. Recurrent neural network based language model. In *Interspeech*, volume 2, pp. 3, 2010.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. Understanding the exploding gradient problem. *CoRR*, abs/1211.5063, 2012. URL <http://arxiv.org/abs/1211.5063>.
- Joris Pelemans, Noam Shazeer, and Ciprian Chelba. Sparse non-negative matrix language modeling. *Transactions of the Association for Computational Linguistics*, 4:329–342, 2016. ISSN 2307-387X. URL <https://transacl.org/ojs/index.php/tacl/article/view/561>.
- 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. URL <http://jmlr.org/papers/v15/srivastava14a.html>.