

Factored Neural Language Models

Andrei Alexandrescu

Department of Comp. Sci. Eng.
University of Washington
andrei@cs.washington.edu

Katrin Kirchhoff

Department of Electrical Engineering
University of Washington
katrin@ee.washington.edu

Abstract

We present a new type of neural probabilistic language model that learns a mapping from both words and explicit word features into a continuous space that is then used for word prediction. Additionally, we investigate several ways of deriving continuous word representations for unknown words from those of known words. The resulting model significantly reduces perplexity on sparse-data tasks when compared to standard backoff models, standard neural language models, and factored language models.

1 Introduction

Neural language models (NLMs) (Bengio et al., 2000) map words into a continuous representation space and then predict the probability of a word given the continuous representations of the preceding words in the history. They have previously been shown to outperform standard back-off models in terms of perplexity and word error rate on medium and large speech recognition tasks (Xu et al., 2003; Emami and Jelinek, 2004; Schwenk and Gauvain, 2004; Schwenk, 2005). Their main drawbacks are computational complexity and the fact that only distributional information (word context) is used to generalize over words, whereas other word properties (e.g. spelling, morphology etc.) are ignored for this purpose. Thus, there is also no principled way of handling out-of-vocabulary (OOV) words.

Though this may be sufficient for applications that use a closed vocabulary, the current trend of porting systems to a wider range of languages (esp. highly-inflected languages such as Arabic) calls for dynamic dictionary expansion and the capability of assigning probabilities to newly added words without having seen them in the training data. Here, we introduce a novel type of NLM that improves generalization by using vectors of word features (stems, affixes, etc.) as input, and we investigate deriving continuous representations for unknown words from those of known words.

2 Neural Language Models




Figure 1: *NLM architecture.* Each word in the context maps to a row in the matrix M . The output is next word's probability distribution.

A standard NLM (Fig. 1) takes as input the previous $n - 1$ words, which select rows from a continuous word representation matrix M . The next layer's input \mathbf{i} is the concatenation of the rows in M corresponding to the input words. From here, the network is a standard multi-layer perceptron with hidden layer $\mathbf{h} = \tanh(\mathbf{i} * W_{ih} + \mathbf{b}_h)$ and output layer $\mathbf{o} = \mathbf{h} * W_{ho} + \mathbf{b}_o$, where $\mathbf{b}_{h,o}$ are the biases on the respective layers. The vector \mathbf{o} is normalized by the softmax function $f_{softmax}(o_i) = \frac{e^{o_i}}{\sum_{k=1}^{|V|} e^{o_k}}$. Back-propagation (BKP) is used to learn model parame-

ters, including the M matrix, which is shared across input words. The training criterion maximizes the regularized log-likelihood of the training data.

3 Generalization in Language Models

An important task in language modeling is to provide reasonable probability estimates for n-grams that were not observed in the training data. This generalization capability is becoming increasingly relevant in current large-scale speech and NLP systems that need to handle unlimited vocabularies and domain mismatches. The smooth predictor function learned by NLMs can provide good generalization if the test set contains n-grams whose individual words have been seen in similar context in the training data. However, NLMs only have a simplistic mechanism for dealing with words that were not observed at all: OOVs in the test data are mapped to a dedicated class and are assigned the singleton probability when predicted (i.e. at the output layer) and the features of a randomly selected singleton word when occurring in the input. In standard backoff n-gram models, OOVs are handled by reserving a small fixed amount of the discount probability mass for the generic OOV word and treating it as a standard vocabulary item. A more powerful backoff strategy is used in factored language models (FLMs) (Bilmes and Kirchhoff, 2003), which view a word as a vector of word features or “factors”: $w = \langle f_1, f_2, \dots, f_k \rangle$ and predict a word jointly from previous words and their factors: A generalized backoff procedure uses the factors to provide probability estimates for unseen n-grams, combining estimates derived from different backoff paths. This can also be interpreted as a generalization of standard class-based models (Brown et al., 1992). FLMs have been shown to yield improvements in perplexity and word error rate in speech recognition, particularly on sparse-data tasks (Vergyri et al., 2004) and have also outperformed backoff models using a linear decomposition of OOVs into sequences of morphemes. In this study we use factors in the input encoding for NLMs.

4 Factored Neural Language Models

NLMs define word similarity solely in terms of their context: words are assumed to be close in the contin-

uous space if they co-occur with the same (subset of) words. But similarity can also be derived from word shape features (affixes, capitalization, hyphenation etc.) or other annotations (e.g. POS classes). These allow a model to generalize across classes of words bearing the same feature. We thus define a *factored neural language model* (FNLM) (Fig. 2) which takes as input the previous $n - 1$ vectors of factors. Different factors map to disjoint row sets of the matrix. The h and o layers are identical to the standard NLM’s. Instead of predicting the probabilities for




Figure 2: FNLM architecture. Input vectors consisting of word and feature indices are mapped to rows in M . The final multiplicative layer outputs the word probability distribution.

all words at the output layer directly, we first group words into classes (obtained by Brown clustering) and then compute the conditional probability of each word given its class: $P(w_t) = P(c_t) \times P(w_t | c_t)$.

This is a speed-up technique similar to the hierarchical structuring of output units used by (Morin and Bengio, 2005), except that we use a “flat” hierarchy. Like the standard NLM, the network is trained to maximize the log-likelihood of the data. We use BKP with cross-validation on the development set and L2 regularization (the sum of squared weight values penalized by a parameter λ) in the objective function.

5 Handling Unknown Factors in FNLMs

In an FNLM setting, a subset of a word’s factors may be known or can be reliably inferred from its shape although the word itself never occurred in the training data. The FNLM can use the continuous representation for these known factors directly in the input. If unknown factors are still present, new continuous representations are derived for them from those of known factors of the same type. This is done by averaging over the continuous vectors of a selected subset of the words in the training data, which places the new item in the center of the region occupied by

the subset. For example, proper nouns constitute a large fraction of OOVs, and using the mean of the rows in M associated with words with a proper noun tag yields the “average proper noun” representation for the unknown word. We have experimented with the following strategies for subset selection: NULL (the null subset, i.e. the feature vector components for unknown factors are 0), ALL (average of all known factors of the same type); TAIL (averaging over the least frequently encountered factors of that type up to a threshold of 10%); and LEAST, i.e. the representation of the single least frequent factors of the same type. The prediction of OOVs themselves is unaffected since we use a factored encoding only for the input, not for the output (though this is a possibility for future work).

6 Data and Baseline Setup

We evaluate our approach by measuring perplexity on two different language modeling tasks. The first is the LDC CallHome Egyptian Colloquial Arabic (ECA) Corpus, consisting of transcriptions of phone conversations. ECA is a morphologically rich language that is almost exclusively used in informal spoken communication. Data must be obtained by transcribing conversations and is therefore very sparse. The present corpus has 170K words for training ($|V| = 16026$), 32K for development (dev), 17K for evaluation (eval97). The data was preprocessed by collapsing hesitations, fragments, and foreign words into one class each. The corpus was further annotated with morphological information (stems, morphological tags) obtained from the LDC ECA lexicon. The OOV rates are 8.5% (development set) and 7.7% (eval97 set), respectively.

Model	ECA ($\cdot 10^2$)		Turkish ($\cdot 10^2$)	
	dev	eval	dev	eval
baseline 3gram	4.108	4.128	6.385	6.438
hand-optimized FLM	4.440	4.327	4.269	4.479
GA-optimized FLM	4.325	4.179	6.414	6.637
NLM 3-gram	4.857	4.581	4.712	4.801
FNLM-NULL	5.672	5.381	9.480	9.529
FNLM-ALL	5.691	5.396	9.518	9.555
FNLM-TAIL 10%	5.721	5.420	9.495	9.540
FNLM-LEAST	5.819	5.479	10.492	10.373

Table 1: Average probability (scaled by 10^2) of known words with unknown words in order-2 context

The second corpus consists of Turkish newspa-

per text that has been morphologically annotated and disambiguated (Hakkani-Tür et al., 2002), thus providing information about the word root, POS tag, number and case. The vocabulary size is 67510 (relatively large because Turkish is highly agglutinative). 400K words are used for training, 100K for development (11.8% OOVs), and 87K for testing (11.6% OOVs). The corpus was preprocessed by removing segmentation marks (titles and paragraph boundaries).

7 Experiments and Results

We first investigated how the different OOV handling methods affect the average probability assigned to words with OOVs in their context. Table 1 shows that average probabilities increase compared to the strategy described in Section 3 as well as other baseline models (standard backoff trigrams and FLM, further described below), with the strongest increase observed for the scheme using the least frequent factor as an OOV factor model. This strategy is used for the models in the following perplexity experiments.

We compare the perplexity of word-based and factor-based NLMs with standard backoff trigrams, class-based trigrams, FLMs, and interpolated models. Evaluation was done with (the “w/unk” column in Table 2) and without (the “no unk” column) scoring of OOVs, in order to assess the usefulness of our approach to applications using closed vs. open vocabularies. The baseline Model 1 is a standard back-off 3-gram using modified Kneser-Ney smoothing (model orders beyond 3 did not improve perplexity). Model 2 is a class-based trigram model with Brown clustering (256 classes), which, when interpolated with the baseline 3-gram, reduces the perplexity (see row 3). Model 3 is a 3-gram word-based NLM (with output unit clustering). For NLMs, higher model orders gave improvements, demonstrating their better scalability: for ECA, a 6-gram (w/o unk) and a 5-gram (w/unk) were used; for Turkish, a 7-gram (w/o unk) and a 5-gram (w/unk) were used. Though worse in isolation, the word-based NLMs reduce perplexity considerably when interpolated with Model 1. The FLM baseline is a hand-optimized 3-gram FLM (Model 5); we also tested an FLM optimized with a genetic algorithm as de-

#	Model	ECA dev		ECA eval		Turkish dev		Turkish eval	
		no unk	w/unk	no unk	w/unk	no unk	w/unk	no unk	w/unk
1	Baseline 3-gram	191	176	183	172	827	569	855	586
2	Class-based LM	221	278	219	269	1642	1894	1684	1930
3	1) & 2)	183	169	178	167	790	540	814	555
4	Word-based NLM	208	341	204	195	1510	1043	1569	1067
5	1) & 4)	178	165	173	162	758	542	782	557
6	Word-based NLM	202	194	204	192	1991	1369	2064	1386
7	1) & 6)	175	162	173	160	754	563	772	580
8	hand-optimized FLM	187	171	178	166	827	595	854	614
9	1) & 8)	182	167	174	163	805	563	832	581
10	genetic FLM	190	188	181	188	761	1181	776	1179
11	1) & 10)	183	166	175	164	706	488	720	498
12	factored NLM	189	173	190	175	1216	808	1249	832
13	1) & 12)	169	155	168	155	724	487	744	500
14	1) & 10) & 12)	165	155	165	154	652	452	664	461

Table 2: *Perplexities for baseline backoff LMs, FLMs, NLMs, and LM interpolation*

scribed in (Duh and Kirchhoff, 2004) (Model 6). Rows 7-10 of Table 2 display the results. Finally, we trained FNLMs with various combinations of factors and model orders. The combination was optimized by hand on the dev set and is therefore most comparable to the hand-optimized FLM in row 8. The best factored NLM (Model 7) has order 6 for both ECA and Turkish. It is interesting to note that the best Turkish FNLM uses only word factors such as morphological tag, stem, case, etc. but not the actual words themselves in the input. The FNLM outperforms all other models in isolation except the FLM; its interpolation with the baseline (Model 1) yields the best result compared to all previous interpolated models, for both tasks and both the unk and no/unk condition. Interpolation of Model 1, FLM and FNLM yields a further improvement. The parameter values of the (F)NLMs range between 32 and 64 for d , 45-64 for the number of hidden units, and 362-1024 for C (number of word classes at the output layer).

8 Conclusion

We have introduced FNLMs, which combine neural probability estimation with factored word representations and different ways of inferring continuous word features for unknown factors. On sparse-data Arabic and Turkish language modeling task FNLMs were shown to outperform all comparable models (standard backoff 3-gram, word-based NLMs) except FLMs in isolation, and all models when interpolated with the baseline. These conclusions apply

to both open and closed vocabularies.

Acknowledgments

This work was funded by NSF under grant no. IIS-0326276 and DARPA under Contract No. HR0011-06-C-0023. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of these agencies.

References

- Y. Bengio, R. Ducharme, and P. Vincent. 2000. A neural probabilistic language model. In *NIPS*.
- J.A. Bilmes and K. Kirchhoff. 2003. Factored language models and generalized parallel backoff. In *HLT-NAACL*.
- P. F. Brown, V. J. Della Pietra, P. V. deSouza, J. C. Lai, and R. L. Mercer. 1992. Class-based n-gram models of natural language. *Computational Linguistics*, 18(4).
- K. Duh and K. Kirchhoff. 2004. Automatic learning of language model structure. In *COLING 2004*.
- A. Emami and F. Jelinek. 2004. Exact training of a neural syntactic language model. In *ICASSP 2004*.
- D. Hakkani-Tür, K. Oflazer, and G. Tür. 2002. Statistical morphological disambiguation for agglutinative languages. *Journal of Computers and Humanities*, 36(4).
- F. Morin and Y. Bengio. 2005. Hierarchical probabilistic neural network language model. In *AISTATS*.
- H. Schwenk and J.L. Gauvain. 2004. Neural network language models for conversational speech recognition. In *ICSLP 2004*.
- H. Schwenk. 2005. Training neural network language models on very large corpora. In *HLT/EMNLP*.
- D. Vergyri, K. Kirchhoff, K. Duh, and A. Stolcke. 2004. Morphology-based language modeling for arabic speech recognition. In *ICSLP*.
- P. Xu, A. Emami, and F. Jelinek. 2003. Training connectionist models for the structured language model. In *EMNLP 2003*.